
Universal Average-Case Optimality of Polyak Momentum

Damien Scieur^{*1} Fabian Pedregosa^{*2}

Abstract

Polyak momentum (PM), also known as the heavy-ball method, is a widely used optimization method that enjoys an asymptotic optimal worst-case complexity on quadratic objectives. However, its remarkable empirical success is not fully explained by this optimality, as the worst-case analysis –contrary to the average-case– is not representative of the expected complexity of an algorithm. In this work we establish a novel link between PM and the average-case analysis. Our main contribution is to prove that *any* optimal average-case method converges in the number of iterations to PM, under mild assumptions. This brings a new perspective on this classical method, showing that PM is asymptotically both worst-case and average-case optimal.

1. Introduction

Polyak momentum (PM), also known as the heavy-ball method, is a widely used optimization method. Originally developed to solve linear equations (Frankel, 1950; Rutishauser, 1959), it was generalized to smooth functions and popularized in the optimization community by Boris Polyak (Polyak, 1964; 1987). This method has seen a renewed interest in recent years, as its stochastic variant which replaces the gradient with a stochastic estimate is effective on deep learning problems (Sutskever et al., 2013; Zhang et al., 2020).

PM also enjoys a locally optimal rate of convergence for strongly convex and twice differentiable objectives. As is common within the optimization literature, this optimality is relative to the *worst-case* analysis, that provides complexity bounds for *any* input from a function class, no matter how unlikely. Despite its widespread use, the worst-case is not representative of the typical behavior of opti-

mization methods. The simplex method, for example, has a worst-case exponential complexity, that becomes polynomial when considering the average-case (Spielman & Teng, 2004).

A more representative analysis of the typical behavior is given by the *average-case* complexity, which averages the algorithm’s complexity over all possible inputs. The average-case analysis is standard for analyzing sorting (Knuth, 1997) and cryptography (Katz & Lindell, 2014) algorithms, to name a few. However, little is known of the average-complexity of optimization algorithms, whose analysis depends on the often unknown probability distribution over the inputs.

The recent work of Pedregosa & Scieur (2020); Lacotte & Pilanci (2020) overcame this dependency on the input probability distribution through the use of random matrix theory techniques. In the same papers, the authors noticed the convergence of some optimal average-case methods to PM, as the number of iterations grows (see Figure 1). This is rather surprising given their crucial differences. For instance, average-case optimal methods use knowledge of the full spectral distribution, while PM only requires knowledge of its edges (i.e., smallest and largest eigenvalue). Since this convergence was only shown on specific methods, it raises the question on whether this is a spurious phenomenon or if this holds more generally:

Conjecture

As the number of iterations grows, all average-case optimal methods converge to Polyak momentum.

The **main contribution** of this paper is to give a positive answer to this conjecture. The main, but not so restrictive assumption, is that the probability density function of the eigenvalues is non-zero on the interval containing its support. With this we can show the previously unknown property that PM is *asymptotically* optimal under the average-case analysis, bringing a new perspective on the remarkable empirical performance of this classical method. Furthermore, this statement is *universal*, i.e., independent of the probability distribution over the inputs.

^{*}Equal contribution ¹Samsung SAIT AI Lab, Montreal ²Google Research. Correspondence to: Damien Scieur <damien.scieur@gmail.com>.




Figure 1. Convergence of optimal average-case methods to Polyak Momentum. For the Marchenko-Pastur and uniform distribution of eigenvalues (left), we construct the method that has optimal average-case complexity and plot the momentum (middle) and steps-size (right) parameters. For the two methods considered, the momentum and step-size parameters converge as the number of iterations grows to those of Polyak momentum, displayed here as a straight line.

1.1. Related work

This work draws from the fields of optimization, complexity analysis and orthogonal polynomials, of which we comment on the most closely related ideas.

Average-case analysis. The average-case analysis has a long history in computer science and numerical analysis. Often it is used to justify the superior performances of algorithms such as Quicksort (Hoare, 1962) and the simplex method in linear programming (Spielman & Teng, 2004). Despite this rich history, it’s challenging to transfer these ideas into continuous optimization due to the ill-defined notion of a typical continuous optimization problem.

In the context of optimization, Pedregosa & Scieur (2020) derived a framework for analyzing the average-case gradient-based methods and developed methods that are non-asymptotic optimal algorithms with respect to the average-case. Such average-case analysis finds applications in various domains. For instance, Lacotte & Pilanci (2020) use this framework to derive optimal average-case algorithms to minimize least-squares with random matrix sketching. Prior to this stream of papers, Berthier et al. (2018) use methods based on Jacobi polynomials to design average-case optimal gossip methods, but without generalizing the framework.

In the numerical analysis literature, Deift & Trogdon (2019) have recently developed an average-case complexity of conjugate gradient.

Asymptotics or orthonormal polynomials. A key ingredient of the proof are asymptotics or orthonormal polynomials. This is a vast subject with applications in stochastic processes (Grenander & Szegő, 1958), random matrix

theory (Deift, 1999) and numerical integration (Mhaskar, 1997) to name a few. The monograph of (Lubinsky, 2000) discusses all results used in this paper.

Notation. Throughout the paper we denote vectors in lowercase boldface (\mathbf{x}), matrices in uppercase boldface letters (\mathbf{H}), and polynomials in uppercase latin letter (P, Q). We will sometimes omit integration variable, with the understanding that $\int \varphi d\mu$ is a shorthand for $\int \varphi(\lambda) d\mu(\lambda)$.

2. Average-Case Analysis of Gradient-Based Methods

The goal of the average-case analysis is to quantify the expected error $\mathbb{E} \|\mathbf{x}_t - \mathbf{x}^*\|^2$, where \mathbf{x}_t is the t -th update of some optimization method and the expectation is taken over all possible problem instances. To make this analysis tractable, and following (Pedregosa & Scieur, 2020), we consider quadratic optimization problems of the form

$$\min_{\mathbf{x} \in \mathbb{R}^d} \left\{ f(\mathbf{x}) \stackrel{\text{def}}{=} \frac{1}{2} (\mathbf{x} - \mathbf{x}^*)^\top \mathbf{H} (\mathbf{x} - \mathbf{x}^*) \right\}, \quad (\text{OPT})$$

where $\mathbf{H} \in \mathbb{R}^{d \times d}$ is a *random* symmetric positive-definite matrix and \mathbf{x}^* is a *random* d -dimensional vector which is a solution of (OPT).

Remark 1. Problem (OPT) subsumes the quadratic minimization problem $\min_{\mathbf{x}} \mathbf{x}^\top \mathbf{H} \mathbf{x} + \mathbf{b}^\top \mathbf{x} + c$ but the notation above will be more convenient for our purposes.

Remark 2. The expectation in $\mathbb{E} \|\mathbf{x}_t - \mathbf{x}^*\|^2$ is over the inputs and not over any randomness of the algorithm, as is common in the stochastic literature. In this paper we only consider deterministic algorithms.

We consider in this paper the class of *first order methods*, which build \mathbf{x}_t using a pre-defined linear combination of an initial guess and previous gradients:

$$\mathbf{x}_t \in \mathbf{x}_0 + \text{span}\{\nabla f(\mathbf{x}_0), \dots, \nabla f(\mathbf{x}_t)\}. \quad (1)$$

This wide class includes most gradient-based optimization methods, such as gradient descent and momentum. However, it excludes quasi-Newton methods, preconditioned gradient descent or Adam (to cite a few), as the preconditioning allows the iterates to go outside span.

2.1. Tools of the trade: orthogonal polynomials and spectral densities

Average-case optimal methods rely on two key concepts that we now introduce: *residual orthogonal polynomials* and the *expected spectral distribution*.

2.1.1. ORTHOGONAL (RESIDUAL) POLYNOMIALS

This section defines orthogonal polynomials and residual polynomials.

Definition 1. Let α be a non-decreasing function such that $\int Q d\alpha$ is finite for all polynomials Q . We will say that the sequence of polynomials P_0, P_1, \dots is orthogonal with respect to $d\alpha$ if P_i has degree i and

$$\int_{\mathbb{R}} P_i P_j d\alpha \begin{cases} = 0 & \text{if } i \neq j \\ > 0 & \text{if } i = j \end{cases}. \quad (2)$$

Furthermore, if they verify $P_i(0) = 1$ for all i , we call these *residual orthogonal polynomials*.

Residual orthogonal polynomials verify a three-term recurrence (Fischer, 1996, §2.4), that is, there exists a sequence of real values a_t, b_t such that

$$P_t(\lambda) = (a_t + b_t \lambda) P_{t-1}(\lambda) + (1 - a_t) P_{t-2}(\lambda), \quad (3)$$

where $P_0(\lambda) = 1$ and $P_1(\lambda) = 1 + b_1 \lambda$.

2.1.2. EXPECTED SPECTRAL DISTRIBUTION

The expected spectral distribution and the extreme eigenvalues of the matrix \mathbf{H} play similar roles in the case of, respectively, average-case and worst-case optimal methods. They measure the problem's difficulty and define the optimal method's parameters.

Definition 2 (Empirical/Expected Spectral Measure). Let \mathbf{H} be a random matrix with eigenvalues $\{\lambda_1, \dots, \lambda_d\}$. The *empirical spectral measure* of \mathbf{H} , called $\mu_{\mathbf{H}}$, is the probability measure

$$\mu_{\mathbf{H}}(\lambda) \stackrel{\text{def}}{=} \frac{1}{d} \sum_{i=1}^d \delta_{\lambda_i}(\lambda), \quad (4)$$

where δ_{λ_i} is the Dirac delta, a distribution equal to zero everywhere except at λ_i and whose integral over the entire real line is equal to one.

Since \mathbf{H} is random, the empirical spectral measure $\mu_{\mathbf{H}}$ is a random measure. Its expectation over \mathbf{H} ,

$$\mu \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{H}}[\mu_{\mathbf{H}}], \quad (5)$$

is called the *expected spectral distribution*.

Example 1 (Marchenko-Pastur density and large least squares problems). Consider a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$, where each entry is an iid random variable with mean zero and variance σ^2 . Then it is known that the expected spectral distribution of $\mathbf{H} = \frac{1}{n} \mathbf{A}^\top \mathbf{A}$ converges to the Marchenko-Pastur distribution (Marchenko & Pastur, 1967) as n and $d \rightarrow \infty$ at a rate in which the asymptotic ratio $d/n \rightarrow r$ is finite. The Marchenko-Pastur distribution $d\mu_{\text{MP}}$ is defined as

$$\max\{1 - \frac{1}{r}, 0\} \delta_0(\lambda) + \frac{\sqrt{(L - \lambda)(\lambda - \ell)}}{2\pi\sigma^2 r \lambda} 1_{\lambda \in [\ell, L]}. \quad (6)$$

Here $\ell \stackrel{\text{def}}{=} \sigma^2(1 - \sqrt{r})^2$, $L \stackrel{\text{def}}{=} \sigma^2(1 + \sqrt{r})^2$ are the extreme nonzero eigenvalues, δ_0 is a Dirac delta at zero (which disappears if $r \geq 1$) and $1_{\lambda \in [\ell, L]}$ is a rectangular window function, equal to 1 for $\lambda \in [\ell, L]$ and 0 elsewhere.

2.2. Average-case optimal methods

With these two ingredients, we can construct the method with optimal average-case complexity. We rewrite the expected error as an integral using the expected spectral density μ as a weight function.

Theorem 1. (Pedregosa & Scieur, 2020) Assume $\mathbf{x}_0, \mathbf{x}^*$ are random variables independent of \mathbf{H} , satisfying $\mathbb{E}[(\mathbf{x}_0 - \mathbf{x}^*)(\mathbf{x}_0 - \mathbf{x}^*)^\top] = R^2 \mathbf{I}$. Let \mathbf{x}_t be generated by a first-order method, associated to the polynomial P_t . Then the expected error at iteration t reads

$$\mathbb{E}\|\mathbf{x}_t - \mathbf{x}^*\|^2 = \overbrace{R^2}^{\text{initialization}} \int_{\mathbb{R}} \underbrace{P_t^2}_{\text{algorithm}} \overbrace{d\mu}^{\text{problem}}. \quad (7)$$

The optimal first order method is obtained by minimizing (7) over the space of residual polynomials of degree t . This is equivalent to finding a sequence of residual polynomials $\{P_i\}$ orthogonal w.r.t. the weight function $\lambda d\mu(\lambda)$, as shown in the following theorem.

Theorem 2. (Pedregosa & Scieur, 2020) Let a_t and b_t be the coefficients of the three-term recurrence (3) for the sequence of residual polynomials orthogonal w.r.t. $\lambda d\mu(\lambda)$. Then the following method has optimal average-case com-

plexity over the class of problems (OPT):¹

$$\begin{aligned} \mathbf{x}_1 &= \mathbf{x}_0 + b_1 \nabla f(\mathbf{x}_0), \\ \mathbf{x}_t &= \mathbf{x}_{t-1} + (1 - a_t)(\mathbf{x}_{t-2} - \mathbf{x}_{t-1}) + b_t \nabla f(\mathbf{x}_{t-1}). \end{aligned} \quad (8)$$

Due to the dependency of the coefficients a_t, b_t on the expected spectral distribution, equation (8) does not represent a single scheme, but rather a family of algorithms: each different expected spectral distribution generates a different optimal method. Below is an example of such optimal algorithm w.r.t the Marchenko-Pastur expected spectral distribution (see Example 2).

Example 2 (Marchenko-Pastur acceleration, (Pedregosa & Scieur, 2020)). *Let $d\mu$ be the density associated with the Marchenko-Pastur distribution. Then, the recurrence of the optimal average-case method associated with this distribution is*

$$\begin{aligned} \rho &= \frac{1+r}{\sqrt{r}}, \quad \delta_0 = 0; \\ \mathbf{x}_1 &= \mathbf{x}_0 - \frac{1}{(1+r)\sigma^2} \nabla f(\mathbf{x}_0); \\ \delta_t &= -(\rho + \delta_{t-1})^{-1}; \\ \mathbf{x}_t &= \mathbf{x}_{t-1} + (1 + \rho\delta_t)(\mathbf{x}_{t-2} - \mathbf{x}_{t-1}) + \frac{\delta_t}{\sigma^2\sqrt{r}} \nabla f(\mathbf{x}_{t-1}). \end{aligned}$$

The coefficients come from the orthogonal polynomials w.r.t. $\lambda d\mu(\lambda)$, which is a shifted Chebyshev polynomials of the second kind.

2.3. Polyak Momentum and worst-case optimality

The Polyak momentum algorithm (Polyak, 1964) has an optimal worst-case convergence rate over the class of first order methods with constant coefficients (Polyak, 1987; Scieur et al., 2017). The method requires knowledge of the smallest and largest eigenvalue of the Hessian \mathbf{H} (denoted ℓ and L respectively) and iterates as follows:

$$\begin{aligned} \mathbf{x}_1 &= \mathbf{x}_0 - \frac{2}{L+\ell} \nabla f(\mathbf{x}_0) \\ \mathbf{x}_{t+1} &= \mathbf{x}_t + \left(\frac{\sqrt{L}-\sqrt{\ell}}{\sqrt{L}+\sqrt{\ell}} \right)^2 (\mathbf{x}_{t-1} - \mathbf{x}_t) - \left(\frac{2}{\sqrt{L}+\sqrt{\ell}} \right)^2 \nabla f(\mathbf{x}_t) \end{aligned} \quad (\text{PM})$$

Remark 3. *Unlike the Marchenko-Pastur accelerated method of Example 2, coefficients of this method are constant in the iterations. Furthermore, these coefficients only depend on the edges of the spectral distribution and not on the full density.*

¹Throughout the paper, we will color-code **momentum** and **step-size** parameters.

3. All Roads Lead to Polyak Momentum

Main result

Theorem 3. *Assume the density function $d\mu$ is strictly positive in the interval $[\ell, L]$ with $\ell > 0$. Then the parameters of the optimal average-case method (8) converge to those of (PM):*

$$\lim_{t \rightarrow \infty} (1 - a_t) = - \underbrace{\left(\frac{\sqrt{L} - \sqrt{\ell}}{\sqrt{L} + \sqrt{\ell}} \right)^2}_{= (\text{PM}) \text{ momentum}}, \quad \text{and} \quad (9)$$

$$\lim_{t \rightarrow \infty} b_t = - \underbrace{\left(\frac{2}{\sqrt{L} + \sqrt{\ell}} \right)^2}_{= (\text{PM}) \text{ step-size}}. \quad (10)$$

The key insight of the proof is to cast the three-term recurrence of residual *orthogonal* polynomials into orthonormal polynomials² in the interval $[-1, 1]$. Once this is done, we will use asymptotic properties of these polynomials. The proof is split into three steps.

- Step 0 introduces notation and some known results.
- Step 1 writes the coefficients of optimal average-case methods in terms of properties of a class of orthonormal polynomials in the $[-1, 1]$ interval.
- Step 2 computes the limits of the expressions derived in the previous step by using known asymptotic properties of orthonormal polynomials.

Step 0: Definitions. In the classical theory of orthogonal polynomials, the weight function associated with orthogonal polynomials is defined in the interval $[-1, 1]$. However, in our case the spectral densities are instead defined in $[\ell, L]$. To translate results from one setting to the other we define the following linear mapping from $[\ell, L]$ to $[-1, 1]$:

$$m(\lambda) = \frac{L+\ell}{L-\ell} - \frac{2}{L-\ell} \lambda. \quad (11)$$

For notational convenience, we will also use the shorthand $m_0 \stackrel{\text{def}}{=} m(0)$. We also define Q_i as the i -th degree orthonormal polynomial with respect to the density function $d\nu : [-1, 1] \mapsto \mathbb{R}$, where $d\nu$ satisfies

$$d\nu(m(\lambda)) = \lambda d\mu(\lambda). \quad (12)$$

Similar to the three-term recurrence of residual orthogonal polynomials, orthonormal polynomials also verify a similar three-term recurrence with coefficients α_t, β_t :

$$\alpha_t Q_t(\xi) = (\xi - \beta_t) Q_{t-1}(\xi) - \alpha_{t-1} Q_{t-2}(\xi). \quad (13)$$

²A sequence Q_1, Q_2, \dots of orthogonal polynomials with respect to $d\omega$ is orthonormal if $\int Q_i^2 d\omega = 1$.

Step 1: Parameters of optimal method and orthonormal polynomials. Theorem 2 links the coefficient a_t, b_t of the optimal average-case method to the residual orthogonal polynomial. Instead, this step relates these coefficients to a different *orthonormal* polynomial in the interval $[-1, 1]$. This eases the usage of the rich theory that studies asymptotic properties of orthonormal polynomials. This relationship is stated in the following lemma:

Lemma 1. *Let a_t, b_t be the parameters associated with optimal average-case method (Theorem 2). These coefficients verify the following identity,*

$$1 - a_t = -\frac{\alpha_{t-1}}{\alpha_t} \frac{Q_{t-2}(m_0)}{Q_t(m_0)}, \quad \text{and} \quad (14)$$

$$b_t = -\frac{2}{\alpha_t(L-\ell)} \frac{Q_{t-1}(m_0)}{Q_t(m_0)}. \quad (15)$$

Proof. We first note that $\{Q_i(m(\lambda))\}_i$ is a sequence of orthogonal polynomials with respect to the weight function $\lambda d\mu(\lambda)$. Indeed, if $i \neq j$, using the substitution $\xi = m(\lambda)$ gives

$$\int_{\ell}^L Q_i(m(\lambda)) Q_j(m(\lambda)) \underbrace{\lambda d\mu(\lambda)}_{=d\nu(m(\lambda))}, \quad (16)$$

$$= -\frac{L-\ell}{2} \int_{-1}^1 Q_i Q_j d\nu = 0, \quad (17)$$

where the last identity follows by orthogonality. Since orthogonality is preserved after multiplication by a scalar, the polynomial $P_t(\lambda) \stackrel{\text{def}}{=} Q_t(m(\lambda))/Q_t(m_0)$ is also orthogonal with respect to the weight function $\lambda d\mu(\lambda)$. The normalization $1/Q_t(m_0)$ ensure the polynomial to be a residual polynomial. Note that $Q_t(m_0)$ cannot be zero and so the normalization is well defined. This is because the roots of orthogonal polynomials are always contained in their support, and by construction m_0 is outside $[-1, 1]$.

Using Theorem 2, the coefficients of the optimal average-case method can be derived from the three-term recurrence of this polynomial. Indeed, starting from the three-term recurrence of Q_i (13), we obtain for P_t

$$\begin{aligned} P_t(\lambda) &= (m - \beta_{t-1}) \frac{Q_{t-1}(m(\lambda))}{\alpha_t Q_t(m_0)} - \alpha_{t-1} \frac{Q_{t-2}(m(\lambda))}{\alpha_t Q_t(m_0)} \\ &= \underbrace{\frac{1}{\alpha_t} \left(\frac{L+\ell}{L-\ell} - \beta_{t-1} - \frac{2}{L-\ell} \lambda \right) \frac{Q_{t-1}(m_0)}{Q_t(m_0)}}_{=(a_t+b_t\lambda)} P_{t-1} \\ &\quad - \underbrace{\frac{\alpha_{t-1}}{\alpha_t} \frac{Q_{t-2}(m_0)}{Q_t(m_0)}}_{=-(1-a_t)} P_{t-2}(\lambda), \end{aligned}$$

where in the last line we used the definition of m and the identity $P_i = Q_i(m(\lambda))/Q_i(m_0)$ for $i = t-1$ and $i = t-2$. Finally, matching the coefficients of this recurrence with (3) yields the identity in the Lemma. \square

Step 2: Asymptotics of orthonormal polynomials. This step uses known result on asymptotics of orthonormal polynomials to compute the limit $t \rightarrow \infty$ of expressions derived in the previous step.

We use the following theorem on the asymptotic ratio between two successive orthonormal polynomials.

Theorem 4 (Rakhmanov (1983);³ Ratio Asymptotics). *Let $\{Q_i\}$ be a sequence of orthonormal polynomials with respect to a weight function strictly positive in $] -1, 1[$ and zero elsewhere. Then we have the following limit for the ratio of polynomials evaluated outside of the support,*

$$\lim_{t \rightarrow \infty} \frac{Q_t(\xi)}{Q_{t-1}(\xi)} = \xi + \sqrt{\xi^2 - 1} \quad \text{for } \xi \notin [-1, 1]. \quad (18)$$

We can use this result to compute the limit of the ratio $Q_{t-1}(m_0)/Q_t(m_0)$, that appears in (14), as $m_0 > 1$ (and thus is not in the interval $[-1, 1]$):

$$\lim_{t \rightarrow \infty} \frac{Q_{t-1}(m_0)}{Q_t(m_0)} \stackrel{(18)}{=} \left(\frac{L+\ell}{L-\ell} + \sqrt{\left(\frac{L+\ell}{L-\ell} \right)^2 - 1} \right)^{-1} \quad (19)$$

$$= \frac{\sqrt{L} - \sqrt{\ell}}{\sqrt{L} + \sqrt{\ell}}. \quad (20)$$

The other dependency of Eq. (14) on the iteration t is through the coefficients α_t, β_t . To compute the limits of these we use the following known asymptotics:⁴

Theorem 5 (Máté et al. (1985); Limits of recurrence coefficients). *Under the same assumptions as Theorem 4, the limits of the coefficients α_t, β_t in the orthonormal three-term recurrence is*

$$\lim_{t \rightarrow \infty} \alpha_t = \frac{1}{2}, \quad \lim_{t \rightarrow \infty} \beta_t = 0. \quad (21)$$

Using this last theorem together with (19), we have

$$\begin{aligned} \lim_{t \rightarrow \infty} (1 - a_t) &\stackrel{(14)}{=} - \left(\lim_{t \rightarrow \infty} \frac{\alpha_{t-1}}{\alpha_t} \right) \left(\lim_{t \rightarrow \infty} \frac{Q_{t-2}(m_0)}{Q_t(m_0)} \right) \\ &\stackrel{(19,21)}{=} - \left(\frac{\sqrt{L} - \sqrt{\ell}}{\sqrt{L} + \sqrt{\ell}} \right)^2, \end{aligned}$$

which is the claimed limit.

To conclude the proof, we compute the same limit for the step-size b_t :

$$\lim_{t \rightarrow \infty} b_t \stackrel{(14)}{=} -\frac{2}{L-\ell} \left(\lim_{t \rightarrow \infty} \alpha_t^{-1} \right) \left(\lim_{t \rightarrow \infty} \frac{Q_{t-1}(m_0)}{Q_t(m_0)} \right) \quad (22)$$

$$\stackrel{(19,21)}{=} - \left(\frac{2}{\sqrt{L} + \sqrt{\ell}} \right)^2. \quad (23)$$

³The original version of this theorem was stated for monic orthogonal polynomials but is valid for polynomials with other normalizations like orthonormal, see for instance (Lubinsky, 2000; Denisov, 2004).

⁴It can be shown that the last two theorems are equivalent (Nevai, 1979, Theorem 13), although for simplicity we will treat them as independent results.

4. Asymptotic Expected Convergence Rates

The previous section showed convergence of the method's parameters to PM, but said nothing about its rate of convergence. This section fills this gap by providing the asymptotic convergence of the expected convergence rate $\mathbb{E}\|\mathbf{x}_t - \mathbf{x}^*\|^2$. More precisely, we show that the expected convergence rate converges to the rate of convergence of Polyak, and that this convergence rate is *independent* of the probability distribution.

Theorem 4.1. *Under the same assumptions of Theorem 3, the asymptotic expected rate of convergence of the optimal method converges to the worst-case rate of convergence,*

$$\lim_{t \rightarrow \infty} \sqrt[t]{\mathbb{E} \left[\frac{\|\mathbf{x}_t - \mathbf{x}^*\|^2}{\|\mathbf{x}_0 - \mathbf{x}^*\|^2} \right]} = \left(\frac{\sqrt{L} - \sqrt{\ell}}{\sqrt{L} + \sqrt{\ell}} \right)^2. \quad (24)$$

Proof. Let P_t be the residual orthogonal polynomial w.r.t. $\lambda d\mu(\lambda)$. Pedregosa & Scieur (2020, Theorem 3.1) showed that the expected rate of convergence for average-case optimal methods admits the following simple form⁵

$$\mathbb{E}\|\mathbf{x}_t - \mathbf{x}^*\|^2 = R^2 \int_{\mathbb{R}} P_t d\mu. \quad (25)$$

This form is particularly convenient for us, as we can then use the three-term recurrence to obtain a recurrence of this expression. Let $r_t = \int_{\mathbb{R}} P_t d\mu$. After using the recurrence over P_t ,

$$\begin{aligned} r_t &= \int_{\mathbb{R}} (a_t + \lambda b_t) P_{t-1}(\lambda) + (1 - a_t) P_{t-2}(\lambda) d\mu(\lambda) \\ &= a_t \underbrace{\int_{\mathbb{R}} P_{t-1} d\mu}_{=r_{t-1}} + (1 - a_t) \underbrace{\int_{\mathbb{R}} P_{t-2} d\mu}_{=r_{t-2}}, \end{aligned} \quad (26)$$

where in the last identity we have used the orthogonality between P_t and $P_0(\lambda) = 1$ w.r.t. $\lambda d\mu(\lambda)$. In all, we have that the convergence rate r_t is described by the recurrence

$$r_t = a_t r_{t-1} + (1 - a_t) r_{t-2}, \quad r_0 = 1, \quad r_1 = \int_{\mathbb{R}} P_1 d\mu. \quad (27)$$

Replacing a_t by its asymptotic value a_∞ , which reads

$$a_\infty \stackrel{\text{def}}{=} 1 + \left(\frac{\sqrt{L} - \sqrt{\ell}}{\sqrt{L} + \sqrt{\ell}} \right)^2, \quad (28)$$

we obtain a different recurrence, which we denote \tilde{r}_t with the same limit:

$$\tilde{r}_t = a_\infty \tilde{r}_{t-1} + (1 - a_\infty) \tilde{r}_{t-2}. \quad (29)$$

⁵In particular, the result shows that the integral of the squared polynomial is equal to the integral of the polynomial if the polynomial is optimal w.r.t. $\lambda d\mu(\lambda)$.

This recurrence is easier to analyze than the previous one because the coefficients are constant. It's a property of linear recurrence relations with constant coefficients that there exists constants c_1, c_2 such that

$$\tilde{r}_t = c_1(1 - a_\infty)^t + c_2. \quad (30)$$

Since the method is convergent, we must have $r_\infty = 0$, which implies $c_2 = 0$. Taking limits of the t -th root we recover the desired rate:

$$\begin{aligned} \lim_{t \rightarrow \infty} \sqrt[t]{r_t} &= \lim_{t \rightarrow \infty} \sqrt[t]{\tilde{r}_t} = \lim_{t \rightarrow \infty} \sqrt[t]{c_1 \left(\frac{\sqrt{L} - \sqrt{\ell}}{\sqrt{L} + \sqrt{\ell}} \right)^{2t}} \\ &= \lim_{t \rightarrow \infty} \sqrt[t]{c_1} \left(\frac{\sqrt{L} - \sqrt{\ell}}{\sqrt{L} + \sqrt{\ell}} \right)^2 = \left(\frac{\sqrt{L} - \sqrt{\ell}}{\sqrt{L} + \sqrt{\ell}} \right)^2. \end{aligned} \quad (31)$$

□

5. Discussion and Simulations: Speed of Convergence to PM

The main result (Theorem 3) shows that, asymptotically, any average-case optimal method converge towards Polyak momentum. This could be interpreted as evidence against average-case optimal methods, as average-case optimal methods are not “essentially different” from PM. However, simulations show other dynamics at play.

In Figure 2 we plot the speed of convergence of the parameters of the optimal average-method for the Marchenko-Pastur distribution with different ratios $r = \frac{d}{n}$ (and hence condition number) and for the uniform distribution with different intervals. We see a clear effect of the condition number on the speed of convergence. The more ill-conditioned the problem, the slower the convergence of the optimal method to PM, implying that PM behaves sub-optimally for a larger number of iterations. This observation is consistent with the results of (Pedregosa & Scieur, 2020), who showed important speedups in the ill-conditioned regime.

6. Conclusion and Perspectives

In this work, we've shown that optimal average-case methods for minimizing quadratics converge to PM under mild assumptions on the expected spectral distribution. This universality over the probability measure is somewhat surprising, as Polyak momentum method only depends on the edges of the spectrum, while on the other hand optimal average-case methods depend on the whole spectrum.

A potential area for future work is the analysis of the rate of convergence of optimal method to Polyak momentum algorithm. It seems the convergence of the step-size and momentum parameters are bounded polynomially in the number of iterations. This observation indicates the potential

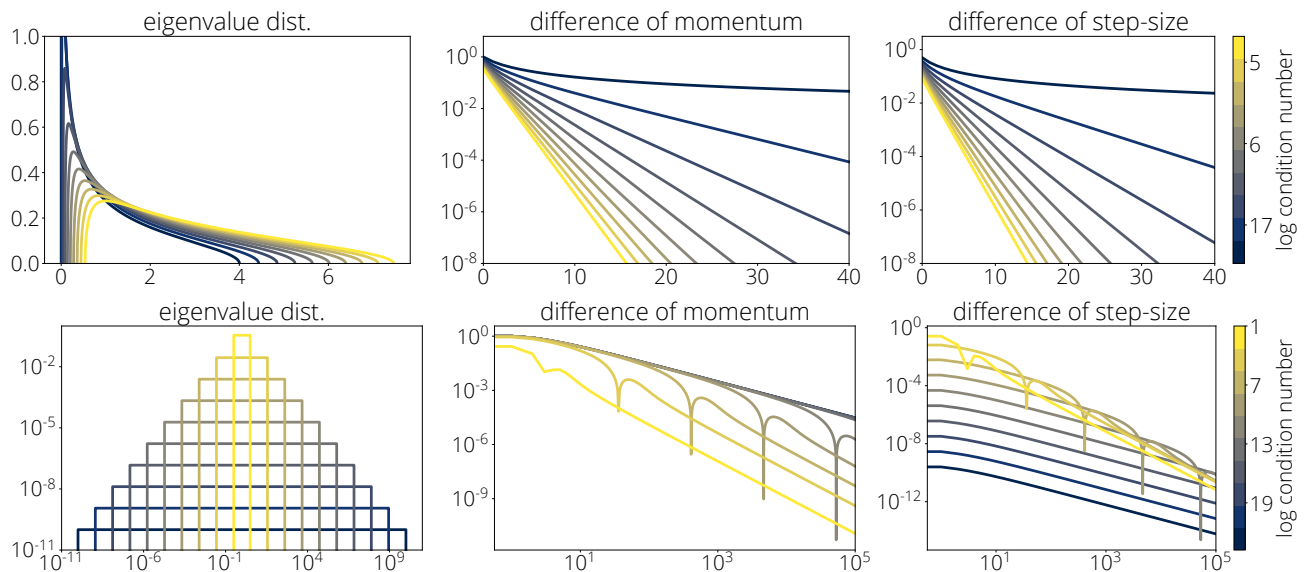


Figure 2. Speed of convergence to Polyak Momentum. For different parametrizations of the Marchenko-Pastur (*top line*) and uniform (*bottom line*) distributions, we plot the absolute difference between the average-case optimal momentum parameter (*middle*) and average-case optimal step-size (*right*) and the momentum and step-size of the Polyak method. The plots show a high anti-correlation between the speed of convergence of optimal average-case methods to PM and problem conditioning: for well-conditioned problems (small condition number) the parameters converge faster to PM than for ill-conditioned (large condition number) problems. Thus, in a regime where we perform only a few iterations, Polyak momentum may not be the best choice.

benefit of optimal methods over PM in the case where we perform a small number of iteration, typical in machine-learning problems.

A second research direction is the study of optimal polynomials on the *complex* plane. In this case, we are no longer solving the optimization problem (OPT). Instead, we aim to solve the linear system $Ax = b$, where the matrix A is non-symmetric, with potentially complex eigenvalues. This has implication in the study of optimal algorithm in game theory (Azizian et al., 2020) or in the acceleration of primal-dual algorithms (Bollapragada et al., 2018).

Finally, our results are only valid in the strongly convex regime ($\ell > 0$), ruling out the important case $r = 1$ in the Marchenko-Pastur distribution, which corresponds to large least squares problems with a square matrix. After the first version of this paper appeared, Paquette et al. (2020) derived an average-case analysis for gradient descent and showed a gap between the asymptotic average-case and worst-case convergence rate. The development of average-case optimal methods and the study of their asymptotic limits in this regime remains an open problem.

Acknowledgements

We would like to thank our colleagues Reza Babanejad, Simon Lacoste-Julien, Remi Lepriol, Nicolas Loizou, Adam Ibrahim, Nicolas Leroux and Courtney Paquette for their insightful discussions and relevant remarks. We also thank Francis Bach and Raphaël Berthier for their useful remarks and pointers.

References

- Azizian, W., Scieur, D., Mitliagkas, I., Lacoste-Julien, S., and Gidel, G. Accelerating smooth games by manipulating spectral shapes. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- Berthier, R., Bach, F., and Gaillard, P. Accelerated Gossip in Networks of Given Dimension using Jacobi Polynomial Iterations. *arXiv preprint arXiv:1805.08531*, 2018.
- Bollapragada, R., Scieur, D., and d’Aspremont, A. Non-linear acceleration of momentum and primal-dual algorithms. *arXiv preprint arXiv:1810.04539*, 2018.
- Deift, P. *Orthogonal polynomials and random matrices: a Riemann-Hilbert approach*, volume 3. American Mathematical Soc., 1999.
- Deift, P. and Trogon, T. The conjugate gradient algorithm on well-conditioned Wishart matrices is almost deterministic. *arXiv preprint arXiv:1901.09007*, 2019.
- Denisov, S. On Rakhmanov’s theorem for Jacobi matrices. *Proceedings of the American Mathematical Society*, 2004.
- Fischer, B. *Polynomial based iteration methods for symmetric linear systems*. Springer, 1996.
- Frankel, S. P. Convergence rates of iterative treatments of partial differential equations. *Mathematical Tables and Other Aids to Computation*, 1950.
- Grenander, U. and Szegő, G. *Toeplitz forms and their applications*. American Mathematical Soc., 1958.
- Hoare, C. A. R. Quicksort. *The Computer Journal*, 1962.
- Katz, J. and Lindell, Y. *Introduction to modern cryptography*. CRC press, 2014.
- Knuth, D. E. *The art of computer programming*, volume 3. Pearson Education, 1997.
- Lacotte, J. and Pilanci, M. Optimal randomized first-order methods for least-squares problems. *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.
- Lubinsky, D. Asymptotics of orthogonal polynomials: some old, some new, some identities. *Acta Applicandae Mathematica*, 2000.
- Marchenko, V. A. and Pastur, L. A. Distribution of eigenvalues for some sets of random matrices. *Matematicheskii Sbornik*, 1967.
- Máté, A., Nevai, P., and Totik, V. Asymptotics for the ratio of leading coefficients of orthonormal polynomials on the unit circle. *Constructive Approximation*, 1985.
- Mhaskar, H. N. *Introduction to the theory of weighted polynomial approximation*. World Scientific, 1997.
- Nevai, P. G. *Orthogonal polynomials*, volume 213. American Mathematical Soc., 1979.
- Paquette, C., van Merriënboer, B., and Pedregosa, F. Halting Time is Predictable for Large Models: A Universality Property and Average-case Analysis. *arXiv preprint arXiv:2006.04299*, 2020.
- Pedregosa, F. and Scieur, D. Average-case Acceleration through spectral density estimation. *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.
- Polyak, B. T. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 1964.
- Polyak, B. T. *Introduction to optimization*. Optimization Software, Inc. Publications Division, New York, 1987.
- Rakhmanov, E. A. On the asymptotics of the ratio of orthogonal polynomials. II. *Mathematics of the USSR-Sbornik*, 1983.
- Rutishauser, H. Theory of gradient methods. In *Refined iterative methods for computation of the solution and the eigenvalues of self-adjoint boundary value problems*. Springer, 1959.
- Scieur, D., Roulet, V., Bach, F., and d’Aspremont, A. Integration methods and optimization algorithms. In *Advances in Neural Information Processing Systems*, 2017.
- Spielman, D. A. and Teng, S.-H. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 2004.
- Sutskever, I., Martens, J., Dahl, G., and Hinton, G. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, 2013.
- Zhang, A., Lipton, Z., Li, M., and Smola, A. J. Dive into deep learning. <https://d2l.ai>, 2020.