AVERAGE-CASE ACCELERATION FOR BILINEAR GAMES AND NORMAL MATRICES

Carles Domingo-Enrich
Computer Science Department
Courant Institute of Mathematical Sciences
New York University
New York, NY 10012, USA
cd2754@nyu.edu

Fabian Pedregosa
Google Research
pedregosa@google.com

Damien Scieur
Samsung SAIT AI Lab & Mila
Montreal, Canada
damien.scieur@gmail.com

ABSTRACT

Advances in generative modeling and adversarial learning have given rise to renewed interest in smooth games. However, the absence of symmetry in the matrix of second derivatives poses challenges that are not present in the classical minimization framework. While a rich theory of average-case analysis has been developed for minimization problems, little is known in the context of smooth games. In this work we take a first step towards closing this gap by developing average-case optimal first-order methods for a subset of smooth games. We make the following three main contributions. First, we show that for zero-sum bilinear games the average-case optimal method is the optimal method for the minimization of the Hamiltonian. Second, we provide an explicit expression for the optimal method corresponding to normal matrices, potentially non-symmetric. Finally, we specialize it to matrices with eigenvalues located in a disk and show a provable speed-up compared to worst-case optimal algorithms. We illustrate our findings through numerical simulations with a varying degree of mismatch with our assumptions.

1 INTRODUCTION

The traditional analysis of optimization algorithms is a worst-case analysis (Nemirovski, 1995; Nesterov, 2004). This type of analysis provides a complexity bound for any input from a function class, no matter how unlikely. However, since hard-to-solve inputs might rarely occur in practice, the worst-case complexity bounds might not be representative of the observed running time.

A more representative analysis is given by the average-case complexity, averaging the algorithm's complexity over all possible inputs. This analysis is standard for analyzing, e.g., sorting (Knuth, 1997) and cryptography algorithms (Katz & Lindell, 2014). Recently, a line of work (Berthier et al., 2020; Pedregosa & Scieur, 2020; Lacotte & Pilanci, 2020; Paquette et al., 2020) focused on optimal methods for the optimization of quadratics, specified by a symmetric matrix. While worst-case analysis uses bounds on the matrix eigenvalues to yield upper and lower bounds on convergence, average-case analysis relies on the expected distribution of eigenvalues and provides algorithms with sharp optimal convergence rates. While the algorithms developed in this context have been shown to be efficient for minimization problems, these have not been extended to smooth games.

A different line of work considers algorithms for smooth games but studies *worst-case* optimal methods (Azizian et al., 2020). In this work, we combine average-case analysis with smooth games, and develop novel average-case optimal algorithms for finding the root of a linear system determined by a (potentially non-symmetric) normal matrix. We make the following **main contributions**:

1. Inspired by the problem of finding equilibria in smooth games, we develop average-case optimal algorithms for finding the root of a non-symmetric affine operator, both under a normality assumption (Thm. 4.1), and under the extra assumption that eigenvalues of the operator are supported in a disk (Thm. 4.2). The proposed method shows a polynomial speedup compared to the worst-case optimal method, verified by numerical simulations.
2. We make a novel connection between average-case optimal methods for optimization, and average-case optimal methods for bilinear games. In particular, we show that solving the Hamiltonian using an average-case optimal method is optimal (Theorem 3.1) for bilinear games. This result complements (Azizian et al., 2020), who proved that Polyak Heavy Ball algorithm on the Hamiltonian is asymptotically worst-case optimal for bilinear games.

2 Average-case analysis for normal matrices

In this paper we consider the following class of problems.

Definition 1. *Let $\boldsymbol{A} \in \mathbb{R}^{d\times d}$ be a real matrix and $\boldsymbol{x}^\star \in \mathbb{R}^d$ a vector. The non-symmetric (affine) operator (**NSO**) problem is defined as:*

$$\textit{Find } \boldsymbol{x} \ : \ F(\boldsymbol{x}) \stackrel{\textit{def}}{=} \boldsymbol{A}(\boldsymbol{x}-\boldsymbol{x}^\star) = \boldsymbol{0}\,. \tag{NSO}$$

This problem generalizes that of minimization of a convex quadratic function f, since we can cast the latter in this framework by setting the operator $F = \nabla f$. The set of solutions is an affine subspace that we will denote $\mathcal{X}^\star$. We will find convenient to consider the distance to this set, defined as

$$\operatorname{dist}(\boldsymbol{x},\, \mathcal{X}^\star) \stackrel{\text{def}}{=} \min_{\boldsymbol{v}\in\mathcal{X}^\star} \|\boldsymbol{x}-\boldsymbol{v}\|^2, \quad \text{with } \mathcal{X}^\star = \{\boldsymbol{x}\in\mathbb{R}^d \,|\, \boldsymbol{A}(\boldsymbol{x}-\boldsymbol{x}^\star)=\boldsymbol{0}\}\,. \tag{1}$$

In this paper we will develop *average-case* optimal methods. For this, we consider $\boldsymbol{A}$ and $\boldsymbol{x}^\star$ to be random vectors, and a random initialization $\boldsymbol{x}_0$. This induces a probability distribution over NSO problems, and we seek to find methods that have an optimal *expected* suboptimality w.r.t. this distribution. Denoting $\mathbb{E}_{(\boldsymbol{A},\boldsymbol{x}^\star,\boldsymbol{x}_0)}$ the expectation over these random problems, we have that average-case optimal methods they verify the following property at each iteration t

$$\min_{\boldsymbol{x}_t} \mathbb{E}_{(\boldsymbol{A},\boldsymbol{x}^\star,\boldsymbol{x}_0)} \operatorname{dist}(\boldsymbol{x}_t,\, \mathcal{X}^\star) \quad \text{s.t. } \boldsymbol{x}_i \in \boldsymbol{x}_0 + \mathbf{span}(\{F(\boldsymbol{x}_j)\}_{j=0}^{i-1}),\ \forall i\in[1:t]. \tag{2}$$

The last condition on $\boldsymbol{x}_t$ stems from restricting the class of algorithms to first-order methods. The class of first-order methods encompasses many known schemes such as gradient descent with momentum, or full-matrix AdaGrad. However, methods such as Adam (Kingma & Ba, 2015) or diagonal AdaGrad (Duchi et al., 2011) are *not* in this class, as the diagonal re-scaling creates iterates $\boldsymbol{x}_t$ outside the span of previous gradients. Although we will focus on the distance to the solution, the results can be extended to other convergence criteria such as $\|F(\boldsymbol{x}_t)\|^2$.

Finally, note that the expectations in this paper are on the problem instance and *not* on the randomness of the algorithm.

2.1 Orthogonal residual polynomials and first-order methods

The analysis of first-order methods simplifies through the use of polynomials. This section provides the tools required to leverage this connection.

Definition 2. *A **residual polynomial** is a polynomial P that satisfies $P(0)=1$.*

Proposition 2.1. *(Hestenes et al., 1952) If the sequence $(\boldsymbol{x}_t)_{t\in\mathbb{Z}_+}$ is generated by a first-order method, then there exist residual polynomials P_t, each one of degree at most t, verifying*

$$\boldsymbol{x}_t - \boldsymbol{x}^\star = P_t(\boldsymbol{A})(\boldsymbol{x}_0-\boldsymbol{x}^\star). \tag{3}$$

As we will see, optimal average-case method are strongly related to orthogonal polynomials. We first define the inner product between polynomials, where we use z^* for the complex conjugate of $z\in\mathbb{C}$.

Definition 3. *For $P, Q \in \mathbb{R}[X]$, we define the inner product $\langle \cdot, \cdot \rangle_\mu$ for a measure μ over $\mathbb{C}$ as*

$$\langle P, Q \rangle_\mu \stackrel{def}{=} \int_{\mathbb{C}} P(\lambda) Q(\lambda)^* \, \mathrm{d}\mu(\lambda) \,. \tag{4}$$

Definition 4. *A sequence of polynomials $\{P_i\}$ is **orthogonal** (resp. **orthonormal**) w.r.t. $\langle \cdot, \cdot \rangle_\mu$ if*

$$\langle P_i, P_i \rangle_\mu > 0 \;\; (\textit{resp.} = 1); \qquad \langle P_i, P_j \rangle_\mu = 0 \;\; \textit{if } i \neq j.$$

2.2 Expected Spectral Distribution

Following (Pedregosa & Scieur, 2020), we make the following assumption on the problem family.

Assumption 1. *$\boldsymbol{x}_0 - \boldsymbol{x}^\star$ is independent of $\boldsymbol{A}$, and $\mathbb{E}_{(\boldsymbol{x}_0, \boldsymbol{x}^\star)}[(\boldsymbol{x}_0 - \boldsymbol{x}^\star)(\boldsymbol{x}_0 - \boldsymbol{x}^\star)^\top] = \frac{R^2}{d} \boldsymbol{I}_d$.*

We will also require the following definitions to characterize difficulty of a problem class. Let $\{\lambda_1, \ldots, \lambda_d\}$ be the eigenvalues of a matrix $\boldsymbol{A} \in \mathbb{R}^{d \times d}$. We define the **empirical spectral distribution** of $\boldsymbol{A}$ as the probability measure

$$\hat{\mu}_{\boldsymbol{A}}(\lambda) \stackrel{\text{def}}{=} \tfrac{1}{d} \textstyle\sum_{i=1}^d \delta_{\lambda_i}(\lambda) \,, \tag{5}$$

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. Note that with this definition, $\int_{\mathcal{D}} \mathrm{d}\hat{\mu}_{\boldsymbol{A}}(\lambda)$ corresponds to the proportion of eigenvalues in $\mathcal{D}$.

When $\boldsymbol{A}$ is a matrix-valued random variable, $\mu_{\boldsymbol{A}}$ is a measure-valued random variable. As such, we can define its **expected spectral distribution**

$$\mu_{\boldsymbol{A}} \stackrel{\text{def}}{=} \mathbb{E}_{\boldsymbol{A}}[\hat{\mu}_{\boldsymbol{A}}] \,, \tag{6}$$

which by the Riesz representation theorem is the measure that verifies $\int f \, \mathrm{d}\mu = \mathbb{E}_{\boldsymbol{A}}[\int f \, \mathrm{d}\mu_{\boldsymbol{A}}]$ for all measureable f. Surprisingly, the expected spectral distribution is the only required characteristic to design optimal algorithms in the average-case.

2.3 Expected error of first-order methods

In this section we provide an expression for the expected convergence in terms of the residual polynomial and the expected spectral distribution introduced in the previous section. To go further in the analysis, we have to assume that $\boldsymbol{A}$ is a normal matrix.

Assumption 2. *The (real) random matrix $\boldsymbol{A}$ is normal, that is, it verifies $\boldsymbol{A}\boldsymbol{A}^\top = \boldsymbol{A}^\top \boldsymbol{A}$.*

Normality is equivalent to $\boldsymbol{A}$ having the spectral decomposition $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^*$, where $\boldsymbol{U}$ is unitary, i.e., $\boldsymbol{U}^*\boldsymbol{U} = \boldsymbol{U}\boldsymbol{U}^* = \mathbf{I}$. We now have everything to write the expected error of a first-order algorithm applied to (NSO).

Theorem 2.1. *Consider the application of a first-order method associated to the sequence of polynomials $\{P_t\}$ (Proposition 2.1) on the problem (NSO). Let μ be the expected spectral distribution of $\boldsymbol{A}$. Under Assumptions 1 and 2, we have*

$$\mathbb{E}[\mathrm{dist}(\boldsymbol{x}_t, \mathcal{X}^\star)] = R^2 \int_{\mathbb{C} \setminus \{0\}} |P_t|^2 \, \mathrm{d}\mu \,, \tag{7}$$

Before designing optimal algorithms for certain specific distributions, we compare our setting with the average-case accelerating for minimization problems of Pedregosa & Scieur (2020), who proposed optimal *optimization* algorithms in the average-case.

2.4 Difficulties of First-Order Methods on Games and Related Work

This section compares our contribution with the existing framework of average-case optimal methods for quadratic minimization problems.

Definition 5. *Let $\boldsymbol{H} \in \mathbb{R}^{d\times d}$ be a random symmetric positive-definite matrix and $\boldsymbol{x}^\star \in \mathbb{R}^d$ a random vector. These elements determine the following* ***random quadratic minimization problem***

$$\min_{\boldsymbol{x}\in\mathbb{R}^d} \left\{ f(\boldsymbol{x}) \stackrel{def}{=} \frac{1}{2}(\boldsymbol{x}-\boldsymbol{x}^\star)^\top \boldsymbol{H}(\boldsymbol{x}-\boldsymbol{x}^\star) \right\}. \tag{OPT}$$

As in our paper, Pedregosa & Scieur (2020) find deterministic optimal first-order algorithms in *expectation* w.r.t. the matrix $\boldsymbol{H}$, the solution $\boldsymbol{x}^\star$, and the initialization $\boldsymbol{x}_0$. Since they work with problem (OPT), their problem is equivalent to (NSO) with the matrix $\boldsymbol{A} = \boldsymbol{H}$. However, they have the *stronger* assumption that the matrix is *symmetric*, which implies being normal. The normality assumption is restrictive in the case of game theory, as they do not always naturally fit such applications. However, this set is expressive enough to consider interesting cases, such as bilinear games, and our experiments show that our findings are also consistent with non-normal matrices.

Using orthogonal residual polynomials and spectral distributions, they derive the explicit formula of the expected error. Their result is similar to Theorem 2.1, but the major difference is the domain of the integral, a real positive line in convex optimization, but a shape in the complex plane in our case. This shape plays a crucial role in the rate of converge of first-order algorithms, as depicted in the work of Azizian et al. (2020); Bollapragada et al. (2018).

In the case of optimization methods, they show that optimal schemes in the average-case follow a simple three-term recurrence arising from the three-term recurrence for residual orthogonal polynomials for the measure $\lambda\mu(\lambda)$. Indeed, by Theorem 2.1 the optimal method corresponds to the residual polynomials minimizing $\langle P, P\rangle_\mu$, and the following result holds:

Theorem 2.2. *(Fischer, 1996, §2.4) When μ is supported in the real line, the residual polynomial of degree t minimizing $\langle P, P\rangle_\mu$ is given by the degree t residual orthogonal polynomial w.r.t. $\lambda\mu(\lambda)$.*

However, the analogous result does not hold for general measures in $\mathbb{C}$, and hence our arguments will make use of the following Theorem 2.3 instead, which links the residual polynomial of degree at most t that minimizes $\langle P, P\rangle_\mu$ to the sequence of orthonormal polynomials for μ.

Theorem 2.3. *[Theorem 1.4 of Assche (1997)] Let μ be a positive Borel measure in the complex plane. The minimum of the integral $\int_{\mathbb{C}} |P(\lambda)|^2 \,\mathrm{d}\mu(\lambda)$ over residual polynomials P of degree lower or equal than t is uniquely attained by the polynomial*

$$P^\star(\lambda) = \frac{\sum_{k=0}^t \phi_k(\lambda)\phi_k(0)^*}{\sum_{k=0}^t |\phi_k(0)|^2}, \quad \textit{with optimal value} \int_{\mathbb{C}} |P^\star(\lambda)|^2 \,\mathrm{d}\mu(\lambda) = \frac{1}{\sum_{k=0}^t |\phi_k(0)|^2}, \tag{8}$$

where $(\phi_k)_k$ is the orthonormal sequence of polynomials with respect to the inner product $\langle\cdot,\cdot\rangle_\mu$.

In the next sections we consider cases where the optimal scheme is identifiable.

3 Average-case Optimal Methods for Bilinear Games

We consider the problem of finding a Nash equilibrium of the zero-sum minimax game given by

$$\min_{\boldsymbol{\theta}_1} \max_{\boldsymbol{\theta}_2} \ell(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \stackrel{\text{def}}{=} (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_1^\star)^\top \boldsymbol{M}(\boldsymbol{\theta}_2 - \boldsymbol{\theta}_2^\star). \tag{9}$$

Let $\boldsymbol{\theta}_1, \boldsymbol{\theta}_1^\star \in \mathbb{R}^{d_1}, \boldsymbol{\theta}_2, \boldsymbol{\theta}_2^\star \in \mathbb{R}^{d_2}, \boldsymbol{M} \in \mathbb{R}^{d_1\times d_2}$ and $d \stackrel{\text{def}}{=} d_1 + d_2$. The vector field of the game (Balduzzi et al., 2018) is defined as $F(\boldsymbol{x}) = \boldsymbol{A}(\boldsymbol{x} - \boldsymbol{x}^\star)$, where

$$F(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) = \begin{bmatrix} \nabla_{\boldsymbol{\theta}_1}\ell(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \\ -\nabla_{\boldsymbol{\theta}_2}\ell(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & \boldsymbol{M} \\ -\boldsymbol{M}^\top & 0 \end{bmatrix}}_{=\boldsymbol{A}} \left(\underbrace{\begin{bmatrix} \boldsymbol{\theta}_1 \\ \boldsymbol{\theta}_2 \end{bmatrix}}_{=\boldsymbol{x}} - \underbrace{\begin{bmatrix} \boldsymbol{\theta}_1^\star \\ \boldsymbol{\theta}_2^\star \end{bmatrix}}_{=\boldsymbol{x}^\star} \right) = \boldsymbol{A}(\boldsymbol{x} - \boldsymbol{x}^\star). \tag{10}$$

As before, $\mathcal{X}^\star$ denotes the set of points $\boldsymbol{x}$ such that $F(\boldsymbol{x}) = 0$, which is equivalent to the set of Nash equilibrium. If $\boldsymbol{M}$ is sampled independently from $\boldsymbol{x}_0, \boldsymbol{x}^\star$ and $\boldsymbol{x}_0 - \boldsymbol{x}^\star$ has covariance $\frac{R^2}{d}\boldsymbol{I}_d$, Assumption 1 is fulfilled. Since $\boldsymbol{A}$ is skew-symmetric, it is in particular normal and Assumption 2 is also satisfied.

We now show that the optimal average-case algorithm to solve bilinear problems is Hamiltonian gradient descent with momentum, described below in its general form. Contrary to the methods in Azizian et al. (2020), the method we propose is *anytime* (and not only asymptotically) average-case optimal.

Optimal average-case algorithm for bilinear games.

Initialization. $\boldsymbol{x}_{-1} = \boldsymbol{x}_0 = \left(\boldsymbol{\theta}_{1,0},\ \boldsymbol{\theta}_{2,0}\right)$, sequence $\{h_t, m_t\}$ given by Theorem 3.1.
Main loop. For $t \geq 0$,

$$\begin{aligned} \boldsymbol{g}_t &= F(\boldsymbol{x}_t - F(\boldsymbol{x}_t)) - F(\boldsymbol{x}_t) \qquad \left(= \tfrac{1}{2}\nabla\|F(\boldsymbol{x}_t)\|^2 \text{ by (12)}\right) \\ \boldsymbol{x}_{t+1} &= \boldsymbol{x}_t - h_{t+1}\boldsymbol{g}_t + m_{t+1}(\boldsymbol{x}_{t-1} - \boldsymbol{x}_t) \end{aligned} \tag{11}$$

The quantity $\frac{1}{2}\|F(\boldsymbol{x})\|^2$ is commonly known as the Hamiltonian of the game (Balduzzi et al., 2018), hence the name *Hamiltonian gradient descent*. Indeed, $\boldsymbol{g}_t = \nabla\left(\frac{1}{2}\|F(\boldsymbol{x})\|^2\right)$ when F is affine:

$$\begin{aligned} F(\boldsymbol{x} - F(\boldsymbol{x})) - F(\boldsymbol{x}) &= \boldsymbol{A}(\boldsymbol{x} - \boldsymbol{A}(\boldsymbol{x} - \boldsymbol{x}^\star) - \boldsymbol{x}^\star) - \boldsymbol{A}(\boldsymbol{x} - \boldsymbol{x}^\star) = -\boldsymbol{A}(\boldsymbol{A}(\boldsymbol{x} - \boldsymbol{x}^\star)) \\ &= \boldsymbol{A}^\top(\boldsymbol{A}(\boldsymbol{x} - \boldsymbol{x}^\star)) = \nabla\left(\frac{1}{2}\|\boldsymbol{A}(\boldsymbol{x} - \boldsymbol{x}^\star)\|^2\right) = \nabla\left(\frac{1}{2}\|F(\boldsymbol{x})\|^2\right). \end{aligned} \tag{12}$$

The following theorem shows that (11) is indeeed the optimal average-case method associated to the minimization problem $\min_{\boldsymbol{x}}\left(\frac{1}{2}\|F(\boldsymbol{x})\|^2\right)$, as the following theorem shows.

Theorem 3.1. *Suppose that Assumption 1 holds and that the expected spectral distribution of $\boldsymbol{M}\boldsymbol{M}^\top$ is absolutely continuous with respect to the Lebesgue measure. Then, the method (11) is average-case optimal for bilinear games when h_t, m_t are chosen to be the coefficients of the average-case optimal minimization of $\frac{1}{2}\|F(\boldsymbol{x})\|^2$.*

How to find optimal coefficients? Since $\frac{1}{2}\|F(\boldsymbol{x})\|^2$ is a quadratic problem, the coefficients $\{h_t,\ m_t\}$ can be found using the average-case framework for quadratic minimization problems of (Pedregosa & Scieur, 2020, Theorem 3.1).

Proof sketch. When computing the optimal polynomial $\boldsymbol{x}_t = P_t(\boldsymbol{A})(\boldsymbol{x}_0 - \boldsymbol{x}^\star)$, we have that the residual orthogonal polynomial P_t behaves differently if t is even or odd.

- **Case 1: t is even.** In this case, we observe that the polynomial $P_t(\boldsymbol{A})$ can be expressed as $Q_{t/2}(-\boldsymbol{A}^2)$, where $(Q_t)_{t\geq 0}$ is the sequence of orthogonal polynomials w.r.t. the expected spectral density of $-\boldsymbol{A}^2$, whose eigenvalues are real and positive. This gives the recursion in (11).
- **Case 2: t is odd.** There is no residual orthogonal polynomial of degree t for t odd. Instead, odd iterations do correspond to the intermediate computation of $\boldsymbol{g}_t$ in (11), but not to an actual iterate.

3.1 Particular case: $\mathbf{M}$ with i.i.d. components

We now show the optimal method when the entries of $\boldsymbol{M}$ are i.i.d. sampled. For simplicity, we order the players such that $d_1 \leq d_2$.

Assumption 3. *Assume that each component of $\boldsymbol{M}$ is sampled iid from a distribution of mean 0 and variance σ^2, and we take $d_1, d_2 \to \infty$ with $\frac{d_1}{d_2} \to r < 1$.*

In such case, the spectral distribution of $\frac{1}{d_2}\boldsymbol{M}\boldsymbol{M}^\top$ tends to the Marchenko-Pastur law, supported in $[\ell, L]$ and with density:

$$\rho_{MP}(\lambda) \stackrel{\text{def}}{=} \frac{\sqrt{(L-\lambda)(\lambda-\ell)}}{2\pi\sigma^2 r\lambda}, \quad \text{where } L \stackrel{\text{def}}{=} \sigma^2(1+\sqrt{r})^2, \ell \stackrel{\text{def}}{=} \sigma^2(1-\sqrt{r})^2. \tag{13}$$

Proposition 3.1. *When $\boldsymbol{M}$ satisfies Assumption 3, the optimal parameter of scheme (11) are*

$$h_t = -\frac{\delta_t}{\sigma^2\sqrt{r}}, \quad m_t = 1 + \rho\delta_t, \quad \text{where } \rho = \frac{1+r}{\sqrt{r}}, \quad \delta_t = (-\rho - \delta_{t-1})^{-1}, \quad \delta_0 = 0. \tag{14}$$

Proof. By Theorem 3.1, the problem reduces to finding the optimal average-case algorithm for the problem $\min_{\boldsymbol{x}} \frac{1}{2}\|F(\boldsymbol{x})\|^2$. Since the expected spectral distribution of $\frac{1}{d_2}\boldsymbol{M}\boldsymbol{M}^\top$ is the Marchenko-Pastur law, we can use the optimal algorithm from (Pedregosa & Scieur, 2020, Section 5). □

4 General average-case optimal method for normal operators

In this section we derive general average-case optimal first-order methods for normal operators. First, we need to assume the existence of a three-term recurrence for residual orthogonal polynomials (Assumption 4). As mentioned in subsection 2.4, for general measures in the complex plane, the existence of a three-term recurrence of orthogonal polynomials is not ensured. In Proposition B.3 in Appendix B we give a sufficient condition for its existence, and in the next subsection we will show specific examples where the residual orthogonal polynomials satisfy the three-term recurrence.

Assumption 4 (Simplifying assumption)**.** *The sequence of residual polynomials $\{\psi_t\}_{t\geq 0}$ orthogonal w.r.t. the measure μ, defined on the complex plane, admits the three-term recurrence*

$$\psi_{-1} = 0, \quad \psi_0 = 1, \quad \psi_t(\lambda) = (a_t + b_t\lambda)\psi_{t-1}(\lambda) + (1 - a_t)\psi_{t-2}(\lambda). \tag{15}$$

Under Assumption 4, Theorem 4.1 shows that the optimal algorithm can also be written as an average of iterates following a simple three-terms recurrence.

Theorem 4.1. *Under Assumption 4 and the assumptions of Theorem 2.1, the following algorithm is optimal in the average case, with $\boldsymbol{y}_{-1} = \boldsymbol{y}_0 = \boldsymbol{x}_0$:*

$$\begin{aligned} \boldsymbol{y}_t &= a_t\boldsymbol{y}_{t-1} + (1 - a_t)\boldsymbol{y}_{t-2} + b_t F(\boldsymbol{y}_{t-1}) \\ \boldsymbol{x}_t &= \frac{B_t}{B_t + \beta_t}\boldsymbol{x}_{t-1} + \frac{\beta_t}{B_t + \beta_t}\boldsymbol{y}_t\,, \quad \beta_t = \phi_t^2(0), \quad B_t = B_{t-1} + \beta_{t-1}, \quad B_0 = 0\,. \end{aligned} \tag{16}$$

where $(\phi_k(0))_{k\geq 0}$ can be computed using the three-term recurrence (upon normalization). Moreover, $\mathbb{E}_{(\boldsymbol{A},\boldsymbol{x}^\star,\boldsymbol{x}_0)} \operatorname{dist}(\boldsymbol{x}_t, \mathcal{X}^\star)$ converges to zero at rate $1/B_t$.

Remark. Notice that it is not immediate that (16) fulfills the definition of first-order algorithms stated in (2), as $\boldsymbol{y}_t$ is clearly a first-order method but $\boldsymbol{x}_t$ is an average of the iterates $\boldsymbol{y}_t$. Using that F is an affine function we see that $\boldsymbol{x}_t$ indeed fulfills (2).

Remark. Assumption 4 is needed for the sequence $(\boldsymbol{y}_t)_{t\geq 0}$ to be computable using a three-term recurrence. However, for some distribution, the associated sequence of orthogonal polynomials may admit another recurrence that may not satisfy Assumption 4.

4.1 Circular spectral distributions

In random matrix theory, the circular law states that if $\boldsymbol{A}$ is an $n \times n$ matrix with i.i.d. entries of mean C and variance R^2/n, as $n \to \infty$ the spectral distribution of $\boldsymbol{A}$ tends to the uniform distribution on $D_{C,R}$. In this subsection we apply Theorem 4.1 to a class of spectral distributions specified by Assumption 5, which includes the uniform distribution on $D_{C,R}$. Even though the random matrices with i.i.d entries are not normal, we will see in section 6 that the empirical results for such matrices are consistent with our theoretical results under the normality assumption.

Assumption 5. *Assume that the expected spectral distribution $\mu_{\boldsymbol{A}}$ is supported in the complex plane on the disk $D_{C,R}$ of center $C \in \mathbb{R}, C > 0$ and radius $R < C$. Moreover, assume that the spectral density is circularly symmetric, i.e. there exists a probability measure μ_R supported on $[0, R]$ such for all f measurable and $r \in [0, R]$, $\mathrm{d}\mu_{\boldsymbol{A}}(C + re^{i\theta}) = \frac{1}{2\pi}\,\mathrm{d}\theta\,\mathrm{d}\mu_R(r)$.*

Proposition 4.1. *If μ satisfies Assumption 5, the sequence of orthonormal polynomials is $(\phi_t)_{t\geq 0}$,*

$$\phi_t(\lambda) = \frac{(\lambda - C)^t}{K_{t,R}}, \text{ where } K_{t,R} = \sqrt{\textstyle\int_0^R r^{2t}\,\mathrm{d}\mu_R(r)}\,. \tag{17}$$

Example. The uniform distribution in $D_{C,R}$ is to $\mathrm{d}\mu_R = \frac{2r}{R^2}\,\mathrm{d}r$, and $K_{t,R} = R^t/\sqrt{t+1}$.

From Proposition 4.1, the sequence of residual polynomials is given by $\phi_t(\lambda)/\phi_t(0) = \left(1 - \frac{\lambda}{C}\right)^t$, which implies that Assumption 4 is fulfilled with $a_t = 1, b_t = -\frac{1}{C}$. Thus, by Theorem 4.1 we have

Theorem 4.2. *Given an initialization $\boldsymbol{x}_0(\boldsymbol{y}_0 = \boldsymbol{x}_0)$, if Assumption 5 is fulfilled with $R < C$ and the assumptions of Theorem 2.1 hold, then the average-case optimal first-order method is*

$$\begin{aligned} \boldsymbol{y}_t &= \boldsymbol{y}_{t-1} - \tfrac{1}{C}F(\boldsymbol{y}_{t-1}), \quad \beta_t = C^{2t}/K_{t,R}^2, \quad B_t = B_{t-1} + \beta_{t-1}, \\ \boldsymbol{x}_t &= \frac{B_t}{B_t + \beta_t}\boldsymbol{x}_{t-1} + \frac{\beta_t}{B_t + \beta_t}\boldsymbol{y}_t. \end{aligned} \tag{18}$$

Moreover, $\mathbb{E}_{(\boldsymbol{A},\boldsymbol{x}^\star,\boldsymbol{x}_0)}\,\mathrm{dist}(\boldsymbol{x}_t, \mathcal{X}^\star)$ converges to zero at rate $1/B_t$.

We now compare Theorem 4.2 with worst-case methods studied in Azizian et al. (2020). They give a worst-case convergence lower bound of $(R/C)^{2t}$ on the quantity $\mathrm{dist}(\boldsymbol{z}_t, \mathcal{X}^\star)$ for first-order methods $(\boldsymbol{z}_t)_{t\geq 0}$ on matrices with eigenvalues in the disk $D_{C,R}$. By the classical analysis of first-order methods, this rate is achievable by gradient descent with stepsize $1/C$, i.e. the iterates $\boldsymbol{y}_t$ defined in (18). However, by equation (79) in Proposition D.3 we have that under slight additional assumptions (those of Proposition 5.2), $\lim_{t\to\infty} \mathbb{E}\left[\mathrm{dist}(\boldsymbol{x}_t, \mathcal{X}^\star)\right]/\mathbb{E}\left[\mathrm{dist}(\boldsymbol{y}_t, \mathcal{X}^\star)\right] = 1 - \frac{R^2}{C^2}$ holds. That is, the average-case optimal algorithm outperforms gradient descent by a constant factor depending on the conditioning R/C.

5 Asymptotic behavior

The recurrence coefficients of the average-case optimal method typically converges to limiting values when $t \to \infty$, which gives an "average-case asymptotically optimal first-order method" with constant coefficients. For the case of symmetric operators with spectrum in $[\ell, L]$, Scieur & Pedregosa (2020) show that under mild conditions, the asymptotically optimal algorithm is the Polyak momentum method with coefficients depending only on ℓ and L. For bilinear games, since the average-case optimal algorithm is the average-case optimal algorithm of an optimization algorithm, we can make use of their framework to obtain the asymptotic algorithm (see Theorem 3 of Scieur & Pedregosa (2020)).

Proposition 5.1. *Assume that the expected spectral density $\mu_{\boldsymbol{M}\boldsymbol{M}^\top}$ of $\boldsymbol{M}\boldsymbol{M}^\top$ is supported in $[\ell, L]$ for $0 < \ell < L$, and strictly positive in this interval. Then, the asymptotically optimal algorithm for bilinear games is the following version of Polyak momentum:*

$$\begin{aligned} \boldsymbol{g}_t &= F(\boldsymbol{x}_t - F(\boldsymbol{x}_t)) - F(\boldsymbol{x}_t) \\ \boldsymbol{x}_{t+1} &= \boldsymbol{x}_t + \left(\tfrac{\sqrt{L}-\sqrt{\ell}}{\sqrt{L}+\sqrt{\ell}}\right)^2 (\boldsymbol{x}_{t-1} - \boldsymbol{x}_t) - \left(\tfrac{2}{\sqrt{L}+\sqrt{\ell}}\right)^2 \boldsymbol{g}_t \end{aligned} \tag{19}$$

Notice that the algorithm in (19) is the worst-case optimal algorithm from Proposition 4 of Azizian et al. (2020). For the case of circularly symmetric spectral densities with support on disks, we can also compute the asymptotically optimal algorithm.

Proposition 5.2. *Suppose that the assumptions of Theorem 4.2 hold with $\mu_R \in \mathcal{P}([0, R])$ fulfilling $\mu_R([r, R]) = \Omega((R - r)^\kappa)$ for r in $[r_0, R]$ for some $r_0 \in [0, R)$ and for some $\kappa \in \mathbb{Z}$. Then, the average-case asymptotically optimal algorithm is, with $\boldsymbol{y}_0 = \boldsymbol{x}_0$:*

$$\begin{aligned} \boldsymbol{y}_t &= \boldsymbol{y}_{t-1} - \tfrac{1}{C}F(\boldsymbol{y}_{t-1}), \\ \boldsymbol{x}_t &= \left(\tfrac{R}{C}\right)^2 \boldsymbol{x}_{t-1} + \left(1 - \left(\tfrac{R}{C}\right)^2\right) \boldsymbol{y}_t. \end{aligned} \tag{20}$$

Moreover, the convergence rate for this algorithm is asymptotically the same one as for the optimal algorithm in Theorem 4.2. Namely, $\lim_{t\to\infty} \mathbb{E}\left[\mathrm{dist}(\boldsymbol{x}_t, \mathcal{X}^\star)\right] B_t = 1$.

The condition on μ_R simply rules out cases in which the spectral density has exponentially small mass around 1. It is remarkable that in algorithm (20) the averaging coefficients can be expressed so simply in terms of the quantity R/C. Notice also that while the convergence rate of the algorithm

is slower than the convergence rate for the optimal algorithm by definition, both rates match in the limit, meaning that the asymptotically optimal algorithm also outperforms gradient descent by a constant factor $1 - \frac{R^2}{C^2}$ in the limit $t \to \infty$.

6 EXPERIMENTS

We compare some of the proposed methods on settings with varying degrees of mismatch with our assumptions.

Bilinear Games. We consider min-max bilinear problems of the form (10), where the entries of $\boldsymbol{M}$ are generated i.i.d. from a standard Gaussian distribution. We vary the ratio $r = d/n$ parameter for $d = 1000$ and compare the average-case optimal method of Theorems 3.1 and 5.1, the asymptotic worst-case optimal method of (Azizian et al., 2020) and extragradient (Korpelevich, 1976). In all cases, we use the convergence-rate optimal step-size assuming knowledge of the edges of the spectral distribution.

The spectral density for these problems is displayed in the first row of Figure 1 and the benchmark results on the second row. Average-case optimal methods always outperform other methods, and the largest gain is in the ill-conditioned regime $(r \approx 1)$.

Circular Distribution. For our second experiment we choose $\boldsymbol{A}$ as a matrix with iid Gaussian random entries, therefore the support of the distribution of its eigenvalue is a disk. Note that $\boldsymbol{A}$ does not satisfy the normality assumption of Assumption 2. Figure 1 (third row) compares the average-case optimal methods from Theorems 4.2 and 5.2 on two datasets with different levels of conditioning. Note that the methods converge despite the violation of Assumption 2, suggesting a broader applicability than the one proven in this paper. We leave this investigation for future work.

7 DISCUSSION AND FUTURE RESEARCH DIRECTIONS

In this paper, we presented a general framework for the design of optimal algorithms in the average-case for affine operators F, whose underlying matrix is possibly non-symmetric. However, our approach presents some limitations, the major one being the restriction to normal matrices. Fortunately, the numerical experiments above suggests that this assumption can be relaxed. Developing a theory without that assumption is left for future work. Another avenue for future work is to analyze the nonlinear-case in which the non-symmetric operator $\boldsymbol{A}$ is non-linear, as well as the case in which it is accessed through a stochastic estimator (as done by (Loizou et al., 2020) for the worst-case analysis).

ACKNOWLEDGEMENTS

C. Domingo-Enrich has been partially funded by "la Caixa" Foundation (ID 100010434), under agreement LCF/BQ/AA18/11680094, and partially funded by the NYU Computer Science Department.

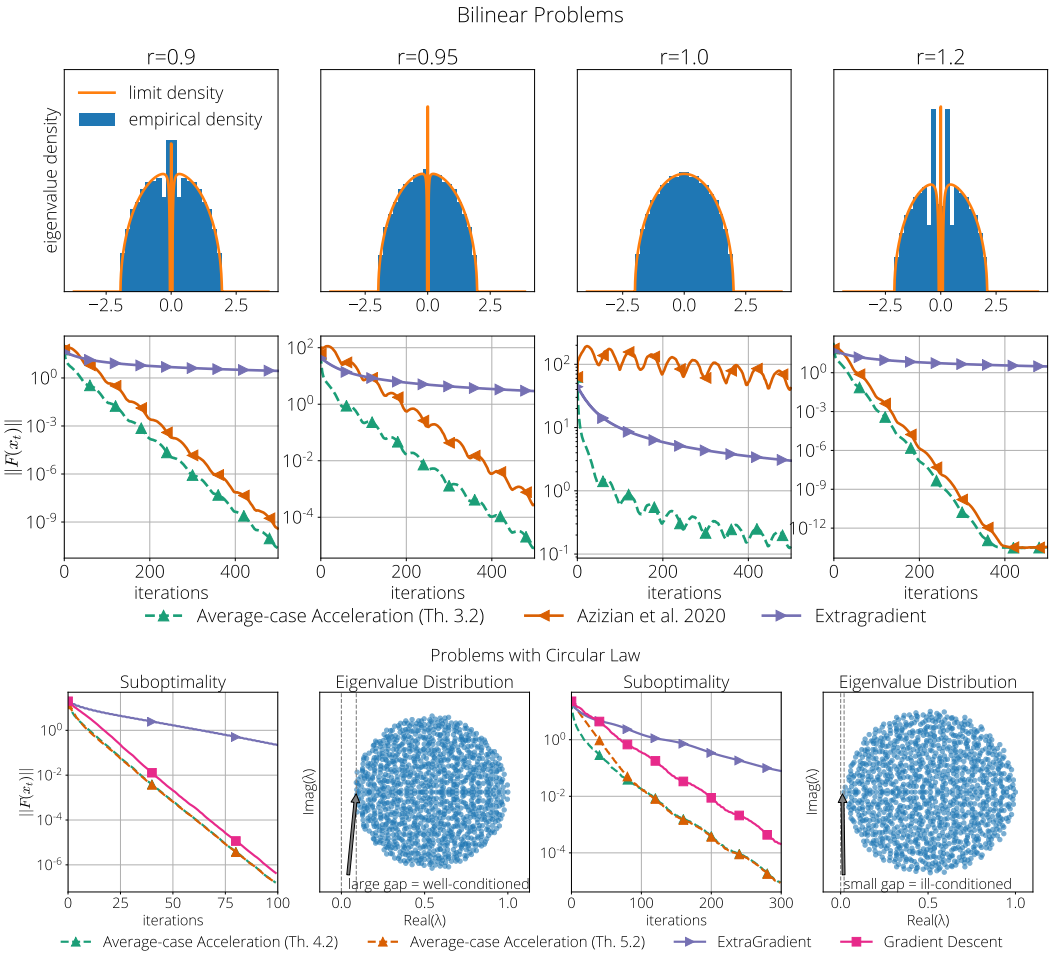

Figure 1: **Benchmarks and spectral density for different games.** *Top row*: spectral density associated with bilinear games for varying values of the ratio parameter $r = n/d$ (the x-axis represents the imaginary line). *Second row*: Benchmarks. Average-case optimal methods always outperform other methods, and the largest gain is in the ill-conditioned regime ($r \approx 1$). *Third row*. Benchmarks (columns 1 and 3) and eigenvalue distribution of a design matrix generated with iid entries for two different degrees of conditioning. Depite the normality assumption not being satisfied, we still observe an improvement of average-case optimal methods vs worst-case optimal ones.

REFERENCES

Walter Van Assche. Orthogonal polynomials in the complex plane and on the real line. In *Fields Institute Communications*, volume 14, pp. 211–245, 1997.

Waïss Azizian, Damien Scieur, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. Accelerating smooth games by manipulating spectral shapes. In *Proceedings of Machine Learning Research*, 2020.

David Balduzzi, Sébastien Racanière, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In *Proceedings of the International Conference on Machine Learning*, 2018.

Raphaël Berthier, Francis Bach, and Pierre Gaillard. Accelerated gossip in networks of given dimension using Jacobi polynomial iterations. *SIAM Journal on Mathematics of Data Science*, 2 (1):24–47, 2020.

Raghu Bollapragada, Damien Scieur, and Alexandre d'Aspremont. Nonlinear acceleration of momentum and primal-dual algorithms. *arXiv preprint arXiv:1810.04539*, 2018.

John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.

Bernd Fischer. *Polynomial Based Iteration Methods for Symmetric Linear Systems*. Vieweg+Teubner Verlag, 1996.

Magnus R Hestenes, Eduard Stiefel, et al. Methods of conjugate gradients for solving linear systems. *Journal of research of the National Bureau of Standards*, 1952.

Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC press, 2014.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Donald Knuth. *The art of computer programming*, volume 3. Pearson Education, 1997.

GM Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12, 1976.

Jonathan Lacotte and Mert Pilanci. Optimal randomized first-order methods for least-squares problems. *Proceedings of the 37th International Conference on Machine Learning*, 2020.

Nicolas Loizou, Hugo Berard, Alexia Jolicoeur-Martineau, Pascal Vincent, Simon Lacoste-Julien, and Ioannis Mitliagkas. Stochastic Hamiltonian gradient methods for smooth games. *arXiv preprint arXiv:2007.04202*, 2020.

Arkadi Nemirovski. *Information-based complexity of convex programming*. Lecture Notes, 1995.

Yurii Nesterov. *Introductory Lectures on Convex Optimization*. Springer, 2004.

Courtney Paquette, Bart van Merriënboer, and Fabian Pedregosa. Halting time is predictable for large models: A universality property and average-case analysis. *arXiv preprint arXiv:2006.04299*, 2020.

Fabian Pedregosa and Damien Scieur. Average-case acceleration through spectral density estimation. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.

Damien Scieur and Fabian Pedregosa. Universal average-case optimality of Polyak momentum. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.