

A Robust Framework for Analyzing Gradient-Based Dynamics in Bilinear Games

Working Paper**Author(s):**

Anagnostides, Ioannis; Penna, Paolo

Publication date:

2020-10

Permanent link:

<https://doi.org/10.3929/ethz-b-000456977>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

arXiv

A Robust Framework for Analyzing Gradient-Based Dynamics in Bilinear Games

Ioannis Anagnostides
ianagnost@student.ethz.ch

Paolo Penna
paolo.penna@inf.ethz.ch

Abstract

In this work, we establish a frequency-domain framework for analyzing gradient-based algorithms in linear minimax optimization problems; specifically, our approach is based on the Z-transform, a powerful tool applied in Control Theory and Signal Processing in order to characterize linear discrete-time systems. We employ our framework to obtain the first tight analysis of stability of Optimistic Gradient Descent/Ascent (OGDA), a natural variant of Gradient Descent/Ascent that was shown to exhibit last-iterate convergence in bilinear games by Daskalakis et al. [DISZ18]. Importantly, our analysis is considerably simpler and more concise than the existing ones.

Moreover, building on the intuition of OGDA, we consider a general family of gradient-based algorithms that augment the memory of the optimization through multiple historical steps. We reduce the convergence – to a saddle-point – of the dynamics in bilinear games to the stability of a polynomial, for which efficient algorithmic schemes are well-established. As an immediate corollary, we obtain a broad class of algorithms – that contains OGDA as a special case – with a last-iterate convergence guarantee to the space of Nash equilibria of the game.

1 Introduction

The fundamental problem of converging to a Nash equilibrium in multi-agent systems has been a topic of prolific research in several fields, including Mathematics, Economics, Algorithmic Game Theory, Optimization [vNMR44, Nas50, Sio58, DGP06, NRTV07, Nes05] and more recently Machine Learning in the context of Generative Adversarial Networks (GANs) [GPAM⁺14, ACB17] and multi-agent reinforcement learning [HW98]. In particular, finding a Nash equilibrium in a two-player *zero-sum* game with objective function $f : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$ is tantamount to computing a pair of strategies $(\mathbf{x}^*, \mathbf{y}^*) \in \mathcal{X} \times \mathcal{Y}$ such that

$$f(\mathbf{x}^*, \mathbf{y}) \leq f(\mathbf{x}^*, \mathbf{y}^*) \leq f(\mathbf{x}, \mathbf{y}^*), \quad \forall (\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}, \quad (1)$$

where \mathcal{X} and \mathcal{Y} represent the space of strategies of each player respectively. In words, an equilibrium – or a *saddle-point* – is a pair of strategies such that neither player can improve her utility by a unilateral deviation. In this context, the celebrated min-max theorem implies that if $f(\mathbf{x}, \mathbf{y})$ is a convex-concave function, then an equilibrium always exists; yet, it does not inform us on whether natural learning algorithms can converge to a minimax pair with a reasonable amount of computational resources.

This question has given rise to profound research, commencing from the analysis of fictitious play by J. Robinson [Rob51], and leading to the development of the no-regret paradigm [CBL06, BM05, AHK12]. However, despite the intrinsic appeal of this framework, these results usually suffer from a lack of last-iterate guarantees, given that a regret-based – or time-average – analysis cannot distinguish between a self-stabilizing system and one with recurrent cycles. Indeed, it has been recognized that limit cycles persist in broad classes of no-regret schemes, such as Mirror Descent [MPP18, PPP17, PP16, PS14]. It is important to point out that a time-average guarantee is insufficient for two main reasons. First, practical applications pose considerable impediments, and implementations of such algorithms could potentially require enormous memory resources; e.g., training GANs through fictitious play would necessitate the storage of millions of parameters at each time step [GLL⁺18, GXC⁺18]. Thus, it is strongly motivated to design algorithms that perform the optimization step through a limited time horizon. Moreover, the stability of the system is the primary desideratum from a control-theoretic perspective, and subsequently of fundamental importance within Game Theory which endeavors to analyze and control the behavior of rational agents.

These shortcomings of traditional learning algorithms – such as Gradient Descent/Ascent (GDA) – were tackled by Daskalakis et al. [DISZ18] in the context of training GANs. Specifically, the authors established a pointwise convergence guarantee for Optimistic Gradient Descent/Ascent (henceforth OGDA) – a natural variant of GDA that augments the memory of the optimization with the previously observed gradient – for bilinear games, i.e. the objective function is linear in the strategies of both players. This result has subsequently engendered a considerable body of work along several directions [MLZ⁺19, LS19, MOP20, DP19, DP18]. Our study follows this long line of research and addresses one such question: Can we extend the last-iterate guarantee of OGDA to a broad class of first-order learning algorithms?

More precisely, building on the intuition of OGDA, we consider a general family of gradient-based dynamics that augment the memory with multiple historical steps. This consideration is strongly motivated by the no-regret literature, where learning through the previously observed costs lies at the heart of the optimization. Naturally, the main challenge resides primarily in the analysis, and elementary techniques appear to be of no use. In this context, our results are established based on the Z-transform, a powerful tool for deriving a frequency-domain representation of linear dynamics. The main contribution of our paper is to reduce the convergence to a Nash equilibrium – through the considered class of algorithms – to the stability of a polynomial,

for which efficient algorithmic schemes are well-established in Mathematics and Control Theory. We illustrate the power of our framework in the analysis of OGDA, providing a remarkably simple and precise characterization.

Related Work Our work follows the line of research initiated by Daskalakis et al. [DISZ18]; their main contribution was showing through an inductive argument that a natural variant of GDA exhibits last-iterate convergence in *unconstrained* bilinear games. Their approach was inspired by a technique previously introduced and applied in the fields of Online Learning and Convex Optimization [RS13, SALS15, SL14, WA18], namely *optimism*. More precisely, this technique is mainly employed when the future cost functions are smooth – or predictable; e.g., they derive from a regularized algorithm such as Follow-The-Regularized-Leader (FTRL). In this case, it has been shown that optimism can substantially accelerate the dynamics.

Moreover, Liang and Stokes [LS19] established that the convergence’s rate of OGDA is linear, assuming that the matrix of the game is square and non-singular. This result was also obtained by Mokhtari et al. [MOP20], while they additionally extended the last-iterate guarantee of OGDA in convex-concave minimax problems. In a closely related direction, [MOP20] also considered a variant of OGDA with slightly perturbed coefficients. Their approach was founded on the observation that OGDA can be examined as an approximate version of the proximal point method [KT98].

The more challenging *constrained* case has also been recently addressed in a series of works. First, Mertikopoulos et al. [MLZ⁺19] established a convergence guarantee for the projected version of OGDA and Extra-Gradient methods through the Variational Inequality framework. Furthermore, the optimistic variant of multiplicative weights was analyzed by Daskalakis and Panageas [DP19]. A no-regret analysis of Optimistic Mirror Descent can be found in [KHSC18]. For a characterization of the stationary points of GDA and OGDA beyond convex-concave settings we refer to [DP18]. Finally, several algorithms have been proposed specifically for solving saddle-point problems; we refer to [ADLH19, SA19, MJS19] and references thereof.

Our Contributions In the first part of our work (Section 3), we analyze Optimistic Gradient Descent/Ascent [DISZ18] in bilinear games, strengthening and improving prior results along several lines. Specifically, we derive a concise expression for the characteristic equation of the induced dynamical system (Proposition 3.1); subsequently, we obtain the first – to the best of our knowledge – tight analysis of stability with respect to the learning rate, providing the exact region of convergence in Theorem 3.1. Our characterization also implies the learning rate that leads to the most rapid convergence (see Appendix A.1), a result which is of clear significance in practical implementations of the algorithm. Our analysis is considerably simpler than the existing ones, and is established through a frequency-domain representation of the dynamics. In particular, we apply the Z-transform, a powerful tool employed in Signal Processing and Control Theory to analyze linear discrete-time systems.

Moreover, we employ the Z-transform in order to study a general class of gradient-based optimization algorithms in bilinear games (Section 4). More precisely, we first identify sufficient conditions under which the limit points of the dynamics are Nash equilibria (Proposition 4.1). Next, we reduce the stability of the dynamics to the stability of a polynomial (Theorem 4.1), for which efficient algorithmic schemes are well-documented in the literature of Mathematics and Control Theory. As an immediate consequence, we show that Optimistic Gradient Descent/Ascent is an instance of a broad class of algorithms that exhibit an analogous convergence guarantee.

Our results strongly supplement the work of [MOP20], where the authors studied a variant of OGDA with slightly perturbed coefficients. The class of algorithms we consider allows for

arbitrary coefficients and importantly, incorporates multiple historical gradient steps – within a certain bounded time horizon. Naturally, historical methods are of great importance within the literature of Game Theory and the no-regret framework, and providing robust paradigms for their analysis constitutes a crucial endeavor. We also believe that our approach could be of independent interest in the realm of Optimization. Our only hypothesis is that the matrix of the game is square and non-singular, a normative assumption (e.g. see [LS19, MOP20]) made to simplify the analysis.

2 Preliminaries

Optimistic Gradient Descent/Ascent Consider a continuously differentiable function $f : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$ that represents the objective function of the game, with $f(\mathbf{x}, \mathbf{y})$ the payoff of player \mathcal{X} to player \mathcal{Y} under strategies $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$ respectively. Throughout this paper, we focus exclusively on the *unconstrained* case, i.e. $\mathcal{X} = \mathbb{R}^n$ and $\mathcal{Y} = \mathbb{R}^m$. Perhaps the most natural optimization algorithm for solving the induced saddle-point problem is by performing simultaneously Gradient Descent on \mathbf{x} and Gradient Ascent on \mathbf{y} ; formally, if $\eta > 0$ denotes some positive constant – typically referred to as the *learning rate*, GDA can be described as follows:

$$\begin{aligned}\mathbf{x}_t &= \mathbf{x}_{t-1} - \eta \nabla_{\mathbf{x}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}), \\ \mathbf{y}_t &= \mathbf{y}_{t-1} + \eta \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}).\end{aligned}\tag{2}$$

However, there are very simple examples where the system of equations (2) diverges; for instance, when $f(x, y) = xy$ with $x, y \in \mathbb{R}$ and $(x_0, y_0) \neq (0, 0)$, GDA is known to diverge for any learning rate $\eta > 0$. This inadequacy has motivated *optimistic* variants of GDA that incorporate some prediction on the next iteration’s gradient through the regularization term (recall that Gradient Descent can be viewed as an instance of FTRL with Euclidean regularizer [Sha12]). With OGDA we refer to the optimization variant that arises when the prediction of the next iteration’s gradient is simply the previously observed gradient; this yields the following update rules:

$$\begin{aligned}\mathbf{x}_t &= \mathbf{x}_{t-1} - 2\eta \nabla_{\mathbf{x}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) + \eta \nabla_{\mathbf{x}} f(\mathbf{x}_{t-2}, \mathbf{y}_{t-2}), \\ \mathbf{y}_t &= \mathbf{y}_{t-1} + 2\eta \nabla_{\mathbf{y}} f(\mathbf{x}_{t-1}, \mathbf{y}_{t-1}) - \eta \nabla_{\mathbf{y}} f(\mathbf{x}_{t-2}, \mathbf{y}_{t-2}).\end{aligned}\tag{3}$$

We remark that in our analysis of OGDA (Section 3) we allow η to take negative values, although this consideration might appear counter-intuitive. In a *bilinear* game the objective function is linear in the strategies of both players, i.e. $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A \mathbf{y}$; here, A denotes an $n \times m$ real matrix which will be referred to as the matrix of the game. We say that the dynamics are *stable* if the strategies of both players converge.

2.1 The Z-transform

The purpose of this subsection is to review the Z-transform, and certain of its properties that will be employed in the forthcoming parts of our study. Specifically, this tool is widely applied in Mathematics and Signal Processing in order to convert a discrete-time signal into a frequency-domain representation; one should view it as the discrete-time analog of the Laplace transform [Mur42]. Importantly for this work, the Z-transform can greatly simplify the analysis and the characterization of discrete-time systems. Indeed, in Control Theory the designer predominantly regulates the system’s behavior in the frequency domain; for an introduction to the basic principles of Control Theory we refer to [KCY05, YLA13]. We should also point out that

the Z-transform has numerous applications in Combinatorics, where the *unilateral* expression is known as *ordinary generating function* (OGF) of a sequence; e.g. see [FS09, Knu11].

Consider a discrete-time signal $x_n \in \mathbb{R}$; we will assume that x_n is *causal*, that is $x_n = 0$ for every $n < 0$. For causal signals, the Z-transform coincides with the unilateral – or *one-sided* – Z-transform, and is defined as

$$X(z) = \mathcal{Z}\{x_n\} = \sum_{n=-\infty}^{+\infty} x_n z^{-n} = \sum_{n=0}^{+\infty} x_n z^{-n}, \quad (4)$$

where we assume that the frequency parameter z resides in the region of convergence:

$$\text{ROC} = \left\{ z \in \mathbb{C}^* : \left| \sum_{n=0}^{+\infty} x_n z^{-n} \right| < +\infty \right\}. \quad (5)$$

In our setting, the dynamics involve multi-dimensional signals that represent the players' strategies throughout the evolution of the game. In this case, the Z-transform is defined by applying the definition (4) in each coordinate or dimension of the signal. Moreover, a crucial feature of the – unilateral – Z-transform is the following time delay property:

$$\mathcal{Z}\{x_{n+k}\} = z^k \left(X(z) - \sum_{i=0}^{k-1} x_i z^{-i} \right), \quad (6)$$

for some $k \in \mathbb{N}$ and x_0, x_1, \dots, x_{k-1} the initial conditions. This property greatly simplifies the analysis of linear discrete-time systems, as it will be illustrated in the forthcoming sections. We also remark the following special cases of (6) that will be used in the analysis of OGDA:

- $\mathcal{Z}\{x_{n+1}\} = z(X(z) - x_0)$
- $\mathcal{Z}\{x_{n+2}\} = z^2(X(z) - x_0 - x_1 z^{-1})$

Of course, it is easy to see that these properties also extend for the multi-dimensional definition. In this context, we state a fundamental theorem in Signal Processing that connects the asymptotic behavior of a causal signal with its Z-transform.

Theorem 2.1. *Let x_n a causal discrete-time signal with $X(z) = \mathcal{Z}\{x_n\}$ a rational function; then, x_n converges for $n \rightarrow \infty$ if and only if every pole of $X(z)$ resides inside the unit circle¹, except from at most one pole at $z = 1$.*

Note that the poles of a rational function in reduced form – the denominator has no common factors with the nominator – coincide with the roots of the polynomial in the denominator. As a result, the convergence of the dynamics can be reduced to analyzing the roots of a polynomial. For a comprehensive study on the Z-transform and applications in discrete-time systems we refer to [Oga95, RZ52, EA05] and references thereof.

Transfer Function Consider a linear discrete-time system that maps an input signal x_t to an output signal y_t through the following difference equation:

$$a_n y_{t+n} + \dots + a_0 y_t = b_m x_{t+m} + \dots + b_0 x_t, \quad (7)$$

for some real coefficients $a_0, \dots, a_n, b_0, \dots, b_m$. The transfer function of the system $H(z)$ is defined as

¹Recall that the unit circle in the complex plane is defined as $\mathbb{U}_1 = \{z \in \mathbb{C} : |z| = 1\}$.

$$H(z) = \frac{Y(z)}{X(z)} = \frac{b_0 + b_1z + \dots + b_mz^m}{a_0 + a_1z + \dots + a_nz^n}. \quad (8)$$

Notice that in the definition of the transfer function, the initial conditions are taken to be zero.

Notation We use lowercase boldface for time domain vectors and capital boldface for frequency domain vectors. We represent with t the discrete time, while subscripts mainly refer to time indices. We use \mathbf{I}_k and $\mathbf{0}_{k \times \ell}$ to refer to the identity matrix of size $k \times k$ and the zero matrix of size $k \times \ell$ respectively; when $k = \ell$ we simply write $\mathbf{0}_k$ instead of $\mathbf{0}_{k \times k}$. For a matrix S , we denote with $\|S\|$ the *spectral norm* of S , i.e. the square root of the maximum eigenvalue of $S^T S$. Finally, a polynomial is said to be stable if all of its roots reside within the unit circle in the complex plane.

3 Analysis of OGDA

In this section, we employ the Z-transform in order to analyze OGDA in bilinear and unconstrained zero-sum games; that is, $f(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{A} \mathbf{y}$, $\mathcal{X} = \mathbb{R}^n$, and $\mathcal{Y} = \mathbb{R}^m$. We will allow arbitrary initial conditions, represented with $\mathbf{x}_0, \mathbf{x}_1 \in \mathbb{R}^n$ and $\mathbf{y}_0, \mathbf{y}_1 \in \mathbb{R}^m$ respectively. In this setting, the update rules of OGDA (3) can be described for any $t \geq 0$ as follows:

$$\begin{aligned} \mathbf{x}_{t+2} &= \mathbf{x}_{t+1} - 2\eta \mathbf{A} \mathbf{y}_{t+1} + \eta \mathbf{A} \mathbf{y}_t, \\ \mathbf{y}_{t+2} &= \mathbf{y}_{t+1} + 2\eta \mathbf{A}^T \mathbf{x}_{t+1} - \eta \mathbf{A}^T \mathbf{x}_t. \end{aligned} \quad (9)$$

These equations can be expressed more concisely in matrix form as

$$\begin{pmatrix} \mathbf{x}_{t+2} \\ \mathbf{y}_{t+2} \end{pmatrix} = \begin{pmatrix} \mathbf{I}_n & -2\eta \mathbf{A} \\ 2\eta \mathbf{A}^T & \mathbf{I}_m \end{pmatrix} \begin{pmatrix} \mathbf{x}_{t+1} \\ \mathbf{y}_{t+1} \end{pmatrix} + \begin{pmatrix} \mathbf{0}_n & \eta \mathbf{A} \\ -\eta \mathbf{A}^T & \mathbf{0}_m \end{pmatrix} \begin{pmatrix} \mathbf{x}_t \\ \mathbf{y}_t \end{pmatrix}. \quad (10)$$

In correspondence to the last expression, let us introduce the following matrices:

$$B = \begin{pmatrix} \mathbf{I}_n & -2\eta \mathbf{A} \\ 2\eta \mathbf{A}^T & \mathbf{I}_m \end{pmatrix}, C = \begin{pmatrix} \mathbf{0}_n & \eta \mathbf{A} \\ -\eta \mathbf{A}^T & \mathbf{0}_m \end{pmatrix}, \mathbf{w}_t = \begin{pmatrix} \mathbf{x}_t \\ \mathbf{y}_t \end{pmatrix}. \quad (11)$$

With this notation, Equation (10) can be re-written as

$$\mathbf{w}_{t+2} = B \mathbf{w}_{t+1} + C \mathbf{w}_t. \quad (12)$$

Let us denote with $\mathbf{W}(z) = \mathcal{Z}\{\mathbf{w}_t\}$ the Z-transform of \mathbf{w}_t ; if we transfer Equation (12) in the frequency domain and apply the time delay property (6) it follows that

$$\begin{aligned} z^2(\mathbf{W}(z) - \mathbf{w}_0 - \mathbf{w}_1 z^{-1}) &= zB(\mathbf{W}(z) - \mathbf{w}_0) + C\mathbf{W}(z) \iff \\ (z^2 \mathbf{I}_{n+m} - zB - C)\mathbf{W}(z) &= (z^2 \mathbf{I}_{n+m} - zB)\mathbf{w}_0 + z\mathbf{w}_1, \end{aligned}$$

where we assume that z resides in the region of convergence of $\mathbf{W}(z)$. Therefore, it is easy to see that $\mathbf{W}(z)$ can be expressed as a rational function in every dimension; indeed, recall that for a non-singular and square matrix S

$$S^{-1} = \frac{\text{adj}(S)}{\det(S)}, \quad (13)$$

where if M_{ij} denotes the (i, j) -minor of S - the determinant of the matrix that results from removing the i^{th} row and the j^{th} column from S , $\text{adj}(S)_{ij} = ((-1)^{i+j} M_{ij})^T$. As a result, each entry of matrix $\text{adj}(z^2 \mathbf{I}_{n+m} - zB - C)$ is a polynomial of z and hence, $\mathbf{W}(z)$ can be expressed - in each dimension - as the ratio of two polynomials. Note that this property holds for any sequence that satisfies a linear difference equation. Thus, the characteristic equation of the system can be expressed as

$$\det(z^2 \mathbf{I}_{n+m} - zB - C) = \begin{vmatrix} (z^2 - z) \mathbf{I}_n & (2z - 1) \eta A \\ -(2z - 1) \eta A^T & (z^2 - z) \mathbf{I}_m \end{vmatrix} = 0. \quad (14)$$

With the assumption that $m = n$, this determinant can be expanded as

$$\begin{vmatrix} (z^2 - z) \mathbf{I}_n & (2z - 1) \eta A \\ -(2z - 1) \eta A^T & (z^2 - z) \mathbf{I}_m \end{vmatrix} = \det((z^2 - z)^2 \mathbf{I}_n + (2z - 1)^2 \eta^2 A A^T), \quad (15)$$

where we employed a well-known identity for the determinant of a block matrix. Let $\alpha(z)$ represent the characteristic polynomial of matrix $\eta^2 A A^T$ - or equivalently $\eta^2 A^T A$; given that $z = 1/2$ is not a root of Equation (14) it follows that

$$(14) \iff \det\left(\left(\frac{z^2 - z}{2z - 1}\right)^2 \mathbf{I}_n + \eta^2 A A^T\right) = 0 \iff \alpha\left(-\left(\frac{z^2 - z}{2z - 1}\right)^2\right) = 0. \quad (16)$$

As a result, we have derived a concise expression for the characteristic equation of the system as a function of the spectrum of matrix A and the learning rate η .

Proposition 3.1. *Consider a square matrix A , and some learning rate η ; if $\alpha(z)$ represents the characteristic polynomial of matrix $\eta^2 A A^T$, the characteristic equation of OGDA in bilinear games (9) can be expressed as*

$$\chi(z) = \alpha\left(-\left(\frac{z^2 - z}{2z - 1}\right)^2\right) = 0. \quad (17)$$

This expression provides a complete characterization of the stability of the system, as well as the rate of convergence. To be more precise, let \mathcal{P} the set of the roots of $\chi(z) = 0$. For a non-singular matrix A it follows that $1 \notin \mathcal{P}$; hence, we know from Theorem 2.1 that the dynamics converge if and only if $|z| < 1, \forall z \in \mathcal{P}$. Moreover, assuming that the system is stable, the rate of convergence is determined by the root with the largest norm; i.e. if λ is the rate of convergence, it follows that $\lambda = \max\{|z| : z \in \mathcal{P}\}$. We should also remark that the region of convergence of $\mathbf{W}(z)$ can be expressed as

$$\text{ROC} = \{z \in \mathbb{C}^* : |z| > \max\{|z'| : z' \in \mathcal{P}\}\}. \quad (18)$$

Through this prism, the stability criterion for the system is equivalent to $e^{j\theta} \in \text{ROC}$ - for any θ . Having established Proposition 3.1, the remainder of the analysis reduces to trivial algebraic calculations (see Appendix A.1). In particular, considering the non-trivial case of $\gamma = \|A\| \neq 0$, we prove the following theorem:

Theorem 3.1. *Consider a square and non-singular matrix A . For any learning rate η such that $|\eta| < 1/(\sqrt{3}\gamma)$ OGDA (9) converges from any initial state, whilst for η such that $|\eta| \geq 1/(\sqrt{3}\gamma)$ the behavior of the dynamics is inherently unstable ².*

²The instability is implied in the sense that at least under certain initial states the dynamics will not converge.

A surprising consequence of this theorem is that the stability of the system and indeed, the rate of convergence depends on η^2 ; thus, the dynamics could converge with negative learning rate (see Figure 1), a very counter-intuitive property in the context of Gradient Descent.

We provide a simple example to illustrate the main implications of Theorem 3.1. In particular, we consider the objective function $f(x, y) = xy$ for $x, y \in \mathbb{R}$. We exhibit the qualitative behavior of OGDA and the impact of the learning rate in Figures 1 and 2.

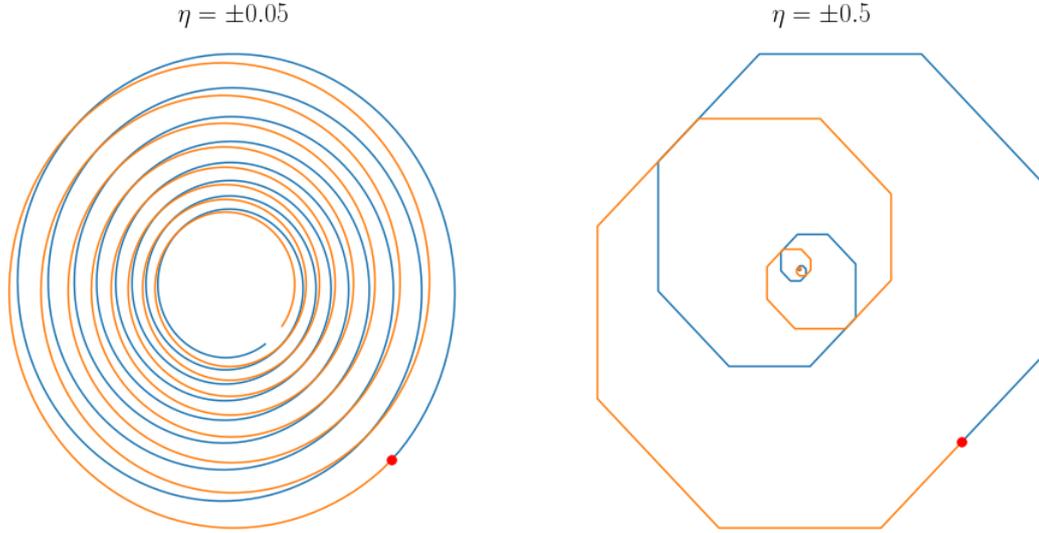


Figure 1: The behavior of the dynamics for different values of learning rate. Note that the blue spiral corresponds to a positive η , while the orange spiral to the corresponding negative. The red point represents the initial state of the system. It is important to point out that Proposition 3.1 implicitly provides the learning rate that leads to the most rapid convergence, establishing an exact characterization of the convergence's rate with respect to η . In our simple example, the optimal learning rate is $\eta = 0.5$, illustrated in the rightmost image.

Finally, we conclude this section with a characterization of the limit points of OGDA; specifically, it is easy to establish the following claim.

Proposition 3.2. *If OGDA converges, the limit points are Nash equilibria* ³.

Proof. We first remark that a pair of strategies $(\mathbf{x}^*, \mathbf{y}^*)$ is a Nash equilibrium of the bilinear game if and only if $A\mathbf{y}^* = \mathbf{0}$ and $A^T\mathbf{x}^* = \mathbf{0}$. Assuming that the dynamics converge, we denote with $\mathbf{x}_\infty = \lim_{t \rightarrow +\infty} \mathbf{x}_t$ and $\mathbf{y}_\infty = \lim_{t \rightarrow +\infty} \mathbf{y}_t$. It follows from Equation (9) that

$$\lim_{t \rightarrow +\infty} (\mathbf{x}_{t+1} - \mathbf{x}_t) = \mathbf{0} = -2\eta A\mathbf{y}_\infty + \eta A\mathbf{y}_\infty \implies A\mathbf{y}_\infty = \mathbf{0}, \quad (19)$$

$$\lim_{t \rightarrow +\infty} (\mathbf{y}_{t+1} - \mathbf{y}_t) = \mathbf{0} = -2\eta A^T\mathbf{x}_\infty + \eta A^T\mathbf{x}_\infty \implies A^T\mathbf{x}_\infty = \mathbf{0}. \quad (20)$$

□

We should point that if A is square and non-singular the unique Nash equilibrium of the game is at $(\mathbf{x}^*, \mathbf{y}^*) = (\mathbf{0}, \mathbf{0})$; nonetheless, Proposition 3.2 holds for any arbitrary matrix A . In the following section we provide a general characterization of the limit points for a broad class of linear optimization algorithms, extending the result of Proposition 4.1.

³This proposition posits that $\eta \neq 0$; otherwise, the limit points coincide with the initial state

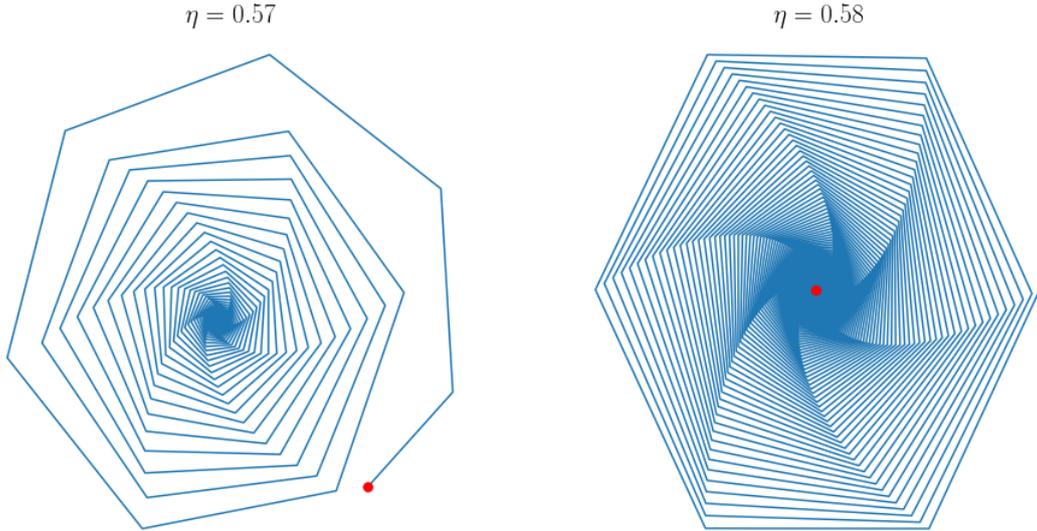


Figure 2: The behavior of the dynamics near the critical value of the learning rate. In particular, according to Theorem 3.1 the system converges only if $\eta < 1/\sqrt{3} \simeq 0.577$. Thus, we can indeed verify the transition of the system from stability (leftmost image) to divergent behavior (rightmost image). We also remark that the convergence of the system becomes gradually slower as the learning rate exceeds the optimal value of $\eta = 0.5$ and advances towards the critical point.

4 Historical Methods

In this section, we analyze a general class of linear gradient-based optimization algorithms. More precisely, consider some constant parameter k that corresponds to the time horizon of each player. Note that for Gradient Descent/Ascent this parameter would be $k = 1$ as the step of each player depends only on the current state of the system, while the optimistic variant we analyzed in the previous section has a $k = 2$ time horizon since it incorporates an additional historical gradient. More broadly, we consider the following update rules for $t \geq 0$:

$$\begin{aligned}
 \mathbf{x}_{t+k} &= \sum_{i=1}^k p_i \mathbf{x}_{t+k-i} - \eta \sum_{i=1}^k q_i \nabla_{\mathbf{x}} f(\mathbf{x}_{t+k-i}, \mathbf{y}_{t+k-i}), \\
 \mathbf{y}_{t+k} &= \sum_{i=1}^k p_i \mathbf{y}_{t+k-i} + \eta \sum_{i=1}^k q_i \nabla_{\mathbf{y}} f(\mathbf{x}_{t+k-i}, \mathbf{y}_{t+k-i}),
 \end{aligned} \tag{21}$$

where $\mathbf{x}_0, \dots, \mathbf{x}_{k-1} \in \mathbb{R}^n$ and $\mathbf{y}_0, \dots, \mathbf{y}_{k-1} \in \mathbb{R}^m$ the initial conditions of the system, and $p_1, \dots, p_k, q_1, \dots, q_k \in \mathbb{R}$ some real parameters. These dynamics will be referred to as Historical Gradient Descent/Ascent (henceforth HGDA). Let us elaborate on Equation (21) for the case of bilinear games. Specifically, one should distinguish the following two main components in the dynamics:

Smoothness Term The first component of the HGDA dynamics is the smoothness or exploitation term, and determines how the previously followed strategies within the time horizon regularize the forthcoming action. For instance, in OGD (or indeed GD) this particular term for the x player would take the form $\mathcal{S}\{\mathbf{x}_t\} = \mathbf{x}_{t+1} - \mathbf{x}_t$. In general, the smoothness term can be expressed as

$$\mathcal{S}\{\mathbf{x}_t\} = \mathbf{x}_{t+k} - p_1\mathbf{x}_{t+k-1} - \cdots - p_k\mathbf{x}_t. \quad (22)$$

The smoothness term can be captured through the transfer function $S(z)$ of system (22); in particular, it follows that

$$S(z) = z^k - p_1z^{k-1} - p_2z^{k-2} - \cdots - p_k. \quad (23)$$

Gradient Term The second component of the HGDA dynamics is the gradient or exploration term, and determines how each player explores the space of strategies through the previously observed gradients. Notice that for bilinear games $\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{y}) = \mathbf{A}\mathbf{y}$ and $\nabla_{\mathbf{y}}f(\mathbf{x}, \mathbf{y}) = A^T\mathbf{x}$; thus, the gradient term follows after applying a linear transformation to the opponent's previous strategies. Let us focus on the weighted average of the opponent's strategies within the time horizon. For OGDA, this weighted average would be $\mathcal{G}\{\mathbf{y}_t\} = 2\mathbf{y}_{t+1} - \mathbf{y}_t$; in general, it can be expressed as

$$\mathcal{G}\{\mathbf{y}_t\} = q_1\mathbf{y}_{t+k-1} + \cdots + q_k\mathbf{y}_t. \quad (24)$$

Similarly, the transfer function of system (24) can be written as

$$G(z) = q_1z^{k-1} + q_2z^{k-2} + \cdots + q_k. \quad (25)$$

As a result, HGDA in bilinear games obtains the following form:

$$\begin{aligned} \mathcal{S}\{\mathbf{x}_t\} &= -\eta\mathbf{A}\mathcal{G}\{\mathbf{y}_t\}, \\ \mathcal{S}\{\mathbf{y}_t\} &= \eta A^T\mathcal{G}\{\mathbf{x}_t\}. \end{aligned} \quad (26)$$

First, we provide sufficient conditions under which the study of the induced dynamical system (26) is meaningful from a game-theoretic standpoint. More precisely, we prove the following proposition:

Proposition 4.1. *Consider the HGDA dynamics (26) such that $S(1) = 0$ and $G(1) \neq 0$; then, if the dynamics converge, the limit points are Nash equilibria⁴.*

Proof. First, we have that

$$S(1) = 0 \iff \sum_{i=1}^k p_i = 1, \quad (27)$$

$$G(1) \neq 0 \iff \sum_{i=1}^k q_i = q_0 \neq 0. \quad (28)$$

Recall that a pair of strategies $(\mathbf{x}^*, \mathbf{y}^*)$ is a Nash equilibrium if and only if $\mathbf{A}\mathbf{y}^* = \mathbf{0}$ and $A^T\mathbf{x}^* = \mathbf{0}$. Let $\mathbf{x}_\infty = \lim_{t \rightarrow +\infty} \mathbf{x}_t$ and $\mathbf{y}_\infty = \lim_{t \rightarrow +\infty} \mathbf{y}_t$; it follows that

$$\lim_{t \rightarrow +\infty} \mathcal{S}\{\mathbf{x}_t\} = \mathbf{x}_\infty \left(1 - \sum_{i=1}^k p_i \right) = \mathbf{0}, \quad (29)$$

⁴We again assume that $\eta \neq 0$.

$$\lim_{t \rightarrow +\infty} \mathcal{G}\{\mathbf{y}_t\} = \mathbf{y}_\infty \sum_{i=1}^k q_i = q_0 \mathbf{y}_\infty. \quad (30)$$

Therefore, it follows from (26) that $\mathbf{0} = -\eta q_0 A \mathbf{y}_\infty \iff A \mathbf{y}_\infty = \mathbf{0}$, since we have assumed that $\eta \neq 0$. Analogously, we can show that $A^T \mathbf{x}_\infty = \mathbf{0}$. \square

As a result, we have reduced finding the Nash equilibrium to the stability of the dynamical system (26). It is also easy to establish the following proposition:

Proposition 4.2. *The HGDA dynamics (26) can be transformed in the z-space as*

$$\begin{pmatrix} S(z)\mathbf{I}_n & G(z)\eta A \\ -G(z)\eta A^T & S(z)\mathbf{I}_m \end{pmatrix} \mathbf{W}(z) = \mathcal{Q}(z; \mathbf{w}_0, \mathbf{w}_1, \dots, \mathbf{w}_{k-1}), \quad (31)$$

where $\mathcal{Q}(z) \in \mathbb{R}^{n+m}$ is a polynomial of z in each dimension that depends on the initial state of the system.

This proposition follows directly after transferring (26) to the z-space and applying the generalized time delay property (6). Moreover, Proposition 4.2 implies that the characteristic equation of HGDA in bilinear games can be expressed as

$$\begin{vmatrix} S(z)\mathbf{I}_n & G(z)\eta A \\ -G(z)\eta A^T & S(z)\mathbf{I}_m \end{vmatrix} = 0. \quad (32)$$

It is clear that every common root of $G(z)$ and $S(z)$ will also satisfy Equation (32), yielding an additional constraint on the structure of the optimization algorithm. In particular, let \mathcal{P} denote the set of all common roots of $G(z)$ and $S(z)$, and consider the following polynomial:

$$P(z) = \prod_{\rho \in \mathcal{P}} (z - \rho). \quad (33)$$

In this context, we proceed with the following factorization: $G(z) = P(z)G'(z)$ and $S(z) = P(z)S'(z)$; notice that the polynomials G' and S' do not have common roots. Applying this factorization yields that

$$(32) \iff (P(z))^{n+m} \begin{vmatrix} S'(z)\mathbf{I}_n & G'(z)\eta A \\ -G'(z)\eta A^T & S'(z)\mathbf{I}_m \end{vmatrix} = 0. \quad (34)$$

As a result, we are able to reduce the stability of Historical Gradient Descent/Ascent in bilinear games to the stability of a polynomial, as stated in the following theorem:

Theorem 4.1. *Consider a square and non-singular matrix A , and let $x^n + a_1 x^{n-1} + \dots + a_n$ the characteristic polynomial of matrix $\eta^2 A A^T$; then, converging to a Nash equilibrium through HGDA in bilinear games reduces to the stability of the following polynomial:*

$$(-S^2(z))^n + a_1 (-S^2(z))^{n-1} (G^2(z)) + \dots + a_n (G^2(z))^n, \quad (35)$$

where $S(z)$ and $G(z)$ are the transfer functions of the smoothness and the gradient term respectively with $S(1) = 0$ and $G(1) \neq 0$.

Proof. First, we have that

$$\begin{vmatrix} S'(z)\mathbf{I}_n & G'(z)\eta A \\ -G'(z)\eta A^T & S'(z)\mathbf{I}_m \end{vmatrix} = \det((S'(z))^2 \mathbf{I}_n + (G'(z))^2 \eta^2 A A^T). \quad (36)$$

Let \mathcal{R} denote the set of roots of $G'(z) = 0$. If $\alpha(z)$ represents the characteristic polynomial of matrix $\eta^2 AA^T$, it follows that for $z \notin \mathcal{R}$

$$(34) \iff (P(z))^{2n} \alpha \left(- \left(\frac{S'(z)}{G'(z)} \right)^2 \right) = 0. \quad (37)$$

Moreover, for $z \notin \mathcal{R}$ we obtain that

$$(37) \iff (-S^2(z))^n + a_1(-S^2(z))^{n-1}(G^2(z)) + \dots + a_n(G^2(z))^n = 0. \quad (38)$$

Finally, consider some $\rho \in \mathcal{R}$. It is clear that ρ satisfies (34) if and only if $\rho \in \mathcal{P}$; as a result, Equation (34) is equivalent to Equation (38) without the constraint $z \notin \mathcal{R}$, and the theorem follows directly from Proposition 4.1. \square

Singular Case Our reduction (Theorem 4.1) is only applicable for a square and non-singular matrix. Indeed, let us assume that $\det(A) = 0$. It follows that $z = 1$ is a (multiple) root of Equation (35) and hence, the induced polynomial is not stable. However, recall from Theorem 2.1 that the dynamics could converge with a factor of $(z - 1)$ in the denominator; in fact, this term determines the steady state solution – or the limit points – of the dynamics. Yet, handling the singular case appears to be technically challenging in the z -space, as one would have to argue about the multiplicity of the factor $(z - 1)$ in the adjoint matrix $\text{adj}(S^2(z)\mathbf{I}_n + G^2(z)\eta^2 AA^T)$ (recall (13)).

Stability of a Polynomial The problem of characterizing the stability of a polynomial has been thoroughly investigated and understood in Mathematics and Control Theory. First, although deriving closed-form solutions for arbitrary polynomials is impossible (Abel’s Theorem), numerical methods yield very accurate approximations [Leh61]; in this context, Theorem 4.1 also implies the rate of convergence – up to some small error. Moreover, the literature of Control Theory provides efficient algorithmic schemes for characterizing the stability of an arbitrary polynomial. Most notably for discrete-time systems, we mention the Bistritz criterion and the Jury test. The details of these methods extend beyond the scope of our work; we refer the interested reader to [VK77, Bis84, PJ93, Bis02] and references thereof.

Finally, as a concrete example we consider the set of HGDA algorithms with $S(z) = z(z - 1)P(z)$ and $G(z) = (2z - 1)P(z)$, where $P(z)$ is a stable polynomial; notice that $S(1) = 0$ and $G(1) \neq 0$, as required by Proposition 4.1. It should be clear that OGDAs are an instance of this class for $P(z) = 1$. In this case, the convergence of the dynamics is implied – for any $\eta < 1/(\sqrt{3}\gamma)$ – by Theorem 3.1 and Theorem 4.1. As a result, we have identified an infinite class of optimization algorithms with a convergence guarantee to the Nash equilibrium of the game.

5 Concluding Remarks

The main contribution of our paper is twofold. First, we provided an exact characterization of Optimistic Gradient Descent/Ascent in bilinear games through a remarkably simple analysis. More importantly, we introduced a novel class of optimization algorithms – namely Historical Gradient Descent/Ascent, and we established an algorithmic procedure for evaluating the stability of the dynamics in bilinear games. Our results have been founded on the Z-transform, a robust tool for characterizing linear discrete-time systems.

We believe that the stability of OGDA – and variants thereof – can be better understood through the lens of Control Theory. In particular, the optimistic term should be construed as a regulator that stabilizes the system. An interesting avenue for future research would be to stabilize Gradient Descent/Ascent via explicit techniques from Control Theory; e.g. the celebrated method of PID control. Moreover, our work has left open the case of a singular or non-square matrix, although this direction may be technically challenging. Finally, our frequency-domain framework can be directly applied to analyze and control continuous-time dynamics through the Laplace transform.

References

- [ACB17] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *CoRR*, abs/1701.07875, 2017.
- [ADLH19] Leonard Adolphs, Hadi Daneshmand, Aurélien Lucchi, and Thomas Hofmann. Local saddle point optimization: A curvature exploitation approach. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019*, volume 89 of *Proceedings of Machine Learning Research*, pages 486–495. PMLR, 2019.
- [AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(6):121–164, 2012.
- [Bis84] Y. Bistritz. Zero location with respect to the unit circle of discrete-time linear system polynomials. *Proceedings of the IEEE*, 72(9):1131–1142, 1984.
- [Bis02] Y. Bistritz. Zero location of polynomials with respect to the unit-circle unhampered by nonessential singularities. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 49(3):305–314, 2002.
- [BM05] Avrim Blum and Yishay Mansour. From external to internal regret. In Peter Auer and Ron Meir, editors, *Learning Theory*, pages 621–636. Springer Berlin Heidelberg, 2005.
- [CBL06] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [DGP06] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing, STOC '06*, page 71–78. Association for Computing Machinery, 2006.
- [DISZ18] Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net, 2018.
- [DP18] Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*, pages 9256–9266, 2018.

- [DP19] Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019*, volume 124 of *LIPICs*, pages 27:1–27:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [EA05] Refaat El Attar. *Lecture notes on Z-transform*. Morrisville, NC : Lulu Press, 1st ed edition, 2005.
- [FS09] Philippe Flajolet and Robert Sedgewick. *Analytic Combinatorics*. Cambridge University Press, USA, 1 edition, 2009.
- [GLL⁺18] Paulina Grnarova, Kfir Y. Levy, Aurélien Lucchi, Thomas Hofmann, and Andreas Krause. An online learning approach to generative adversarial networks. In *6th International Conference on Learning Representations, ICLR 2018*. OpenReview.net, 2018.
- [GPAM⁺14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [GXC⁺18] Hao Ge, Yin Xia, Xu Chen, Randall Berry, and Ying Wu. Fictitious GAN: training gans with historical models. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision - ECCV 2018 - 15th European Conference, 2018, Proceedings, Part I*, volume 11205 of *Lecture Notes in Computer Science*, pages 122–137. Springer, 2018.
- [HW98] Junling Hu and Michael P. Wellman. Multiagent reinforcement learning: Theoretical framework and an algorithm. In *Proceedings of the Fifteenth International Conference on Machine Learning, ICML '98*, page 242–250. Morgan Kaufmann Publishers Inc., 1998.
- [KCY05] Kiam Heong Ang, G. Chong, and Yun Li. Pid control system analysis, design, and technology. *IEEE Transactions on Control Systems Technology*, 13(4):559–576, 2005.
- [KHSC18] Ehsan Asadi Kangarshahi, Ya-Ping Hsieh, Mehmet Fatih Sahin, and Volkan Cevher. Let’s be honest: An optimal no-regret framework for zero-sum games. *CoRR*, abs/1802.04221, 2018.
- [Knu11] Donald E. Knuth. *The Art of Computer Programming: Combinatorial Algorithms, Part 1*. Addison-Wesley Professional, 1st edition, 2011.
- [KT98] Alexander Kaplan and Rainer Tichatschke. Proximal point methods and nonconvex optimization. *J. Global Optimization*, 13(4):389–406, 1998.
- [Leh61] D. H. Lehmer. A machine method for solving polynomial equations. *J. ACM*, 8(2):151–162, April 1961.
- [LS19] Tengyuan Liang and James Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In Kamalika Chaudhuri and

- Masashi Sugiyama, editors, *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019*, volume 89 of *Proceedings of Machine Learning Research*, pages 907–915. PMLR, 2019.
- [MJS19] Eric V. Mazumdar, Michael I. Jordan, and S. Shankar Sastry. On finding local nash equilibria (and only local nash equilibria) in zero-sum games. *CoRR*, abs/1901.00838, 2019.
- [MLZ⁺19] Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. In *7th International Conference on Learning Representations, ICLR 2019*. OpenReview.net, 2019.
- [MOP20] Aryan Mokhtari, Asuman E. Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In Silvia Chiappa and Roberto Calandra, editors, *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020*, volume 108 of *Proceedings of Machine Learning Research*, pages 1497–1507. PMLR, 2020.
- [MPP18] Panayotis Mertikopoulos, Christos H. Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 2703–2717. SIAM, 2018.
- [Mur42] F. J. Murray. Review: David vernon widder, the laplace transform. *Bulletin of the American Mathematical Society*, 48:642–646, 09 1942.
- [Nas50] John F. Nash. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Science*, 36(1):48–49, January 1950.
- [Nes05] Yu Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, May 2005.
- [NRTV07] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, USA, 2007.
- [Oga95] Katsuhiko Ogata. *Discrete-Time Control Systems (2nd Ed.)*. Prentice-Hall, Inc., USA, 1995.
- [PJ93] K. Premaratne and E.I. Jury. On the bistriz tabular form and its relationship with the schur-cohn minors and inner determinants. *Journal of the Franklin Institute*, 330(1):165 – 182, 1993.
- [PP16] Christos Papadimitriou and Georgios Piliouras. From nash equilibria to chain recurrent sets: Solution concepts and topology. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, ITCS ’16*, page 227–235. Association for Computing Machinery, 2016.
- [PPP17] Gerasimos Palaiopanos, Ioannis Panageas, and Georgios Piliouras. Multiplicative weights update with constant step-size in congestion games: Convergence, limit cycles and chaos. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 5872–5882, 2017.

- [PS14] Georgios Piliouras and Jeff S. Shamma. Optimization despite chaos: Convex relaxations to complex limit sets via poincaré recurrence. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, pages 861–873. SIAM, 2014.
- [Rob51] Julia Robinson. An iterative method of solving a game. *Annals of Mathematics*, 54(2):296–301, 1951.
- [RS13] Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3066–3074. Curran Associates, Inc., 2013.
- [RZ52] J. R. Ragazzini and L. A. Zadeh. The analysis of sampled-data systems. *Transactions of the American Institute of Electrical Engineers, Part II: Applications and Industry*, 71(5):225–234, 1952.
- [SA19] Florian Schäfer and Anima Anandkumar. Competitive gradient descent. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019*, pages 7623–7633, 2019.
- [SALS15] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2989–2997. Curran Associates, Inc., 2015.
- [Sha12] Shai Shalev-Shwartz. Online learning and online convex optimization. *Found. Trends Mach. Learn.*, 4(2):107–194, 2012.
- [Sio58] Maurice Sion. On general minimax theorems. *Pacific J. Math.*, 8(1):171–176, 1958.
- [SL14] Jacob Steinhardt and Percy Liang. Adaptivity and optimism: An improved exponentiated gradient algorithm. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1593–1601. PMLR, 2014.
- [VK77] A. Vieira and T. Kailath. On another approach to the schur-cohn criterion. *IEEE Transactions on Circuits and Systems*, 24(4):218–220, 1977.
- [vNMR44] John von Neumann, Oskar Morgenstern, and Ariel Rubinstein. *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press, 1944.
- [WA18] Jun-Kun Wang and Jacob D. Abernethy. Acceleration through optimistic no-regret dynamics. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*, pages 3828–3838, 2018.
- [YLA13] Se Young Yoon, Zongli Lin, and Paul E. Allaire. *Introduction to Control Theory*, pages 149–184. Springer London, London, 2013.

A Omitted Proofs

A.1 Proof of Theorem 3.1

Proof. Let Λ be the set of eigenvalues of matrix $\eta^2 AA^T$ – or equivalently $\eta^2 A^T A$ – and let $\lambda \in \Lambda$; notice that $\lambda > 0$, given that $\eta^2 AA^T$ is symmetric and positive definite. Proposition 3.1 implies that λ induces a root z in the characteristic equation of the system (9) such that

$$\left(\frac{z^2 - z}{2z - 1}\right)^2 = -\lambda \iff \frac{z^2 - z}{2z - 1} = \pm\sqrt{\lambda}j, \quad (39)$$

where j denotes the imaginary unit. This equation yields the following quadratic equations:

$$z^2 - z(1 + 2\sqrt{\lambda}j) + \sqrt{\lambda}j = 0, \quad (40)$$

$$z^2 - z(1 - 2\sqrt{\lambda}j) - \sqrt{\lambda}j = 0. \quad (41)$$

It is easy to see that a complex number z satisfies Equation (40) if and only if its conjugate \bar{z} is a solution to Equation (41); hence, it suffices to restrict our attention to the former quadratic equation, given that the induced solutions have the same norm. In particular, Equation (40) has the following solutions:

$$z_{1,2}(\lambda) = \frac{1 + 2\sqrt{\lambda}j \pm \sqrt{1 - 4\lambda}}{2}. \quad (42)$$

We first consider the case where $\lambda \in (0, \frac{1}{4}]$; specifically, it follows that

$$|z_1(\lambda)| = \frac{\sqrt{2 + 2\sqrt{1 - 4\lambda}}}{2}, \quad (43)$$

$$|z_2(\lambda)| = \frac{\sqrt{2 - 2\sqrt{1 - 4\lambda}}}{2}. \quad (44)$$

As a result, for all $\lambda \in (0, \frac{1}{4}]$ it holds that $|z_1(\lambda)|, |z_2(\lambda)| < 1$ (see Figure 3). Next, we investigate the case where $\lambda > 1/4$; in particular, we obtain that

$$|z_1(\lambda)| = \sqrt{2\lambda + \sqrt{\lambda(4\lambda - 1)}}, \quad (45)$$

$$|z_2(\lambda)| = \sqrt{2\lambda - \sqrt{\lambda(4\lambda - 1)}}. \quad (46)$$

It is easy to see that $|z_2| < 1$ and hence, the region of convergence is determined by the solutions of the inequality $|z_1(\lambda)| < 1$, for all $\lambda \in \Lambda$. Note that for the $\lambda > 1/4$ that satisfies $|z_1(\lambda)| = 1$ the system oscillates, given that $z_1(\lambda) \neq 1$ ($\iff \lambda \neq 0$). Moreover, it follows that $|z_1(\lambda)| < 1 \iff \lambda < 1/3$ (see Figure 3). As a result, the system is stable if and only if every eigenvalue of $\eta^2 AA^T$ is smaller than $1/3$, or equivalently $\eta^2 \gamma^2 < 1/3$, completing the proof. \square

Optimal Learning Rate Our analysis directly provides the learning rate that optimizes the rate of convergence. Specifically, let $\lambda_{min} = \min\{\lambda : \lambda \in \Lambda\}$ and $\lambda_{max} = \max\{\lambda : \lambda \in \Lambda\}$. Note that the rate of convergence depends on the norm of the $z_1(\lambda)$ branch of solutions (the blue curve in Figure 3), since $|z_1(\lambda)| \geq |z_2(\lambda)|, \forall \lambda > 0$. Indeed, the rate of convergence depends on the root with the largest norm. We consider the following two cases:

- Case $\lambda_{min} = \lambda_{max}$ (i.e. all the eigenvalues of AA^T are equal): The optimal learning rate is simply $\eta_{opt} = 1/(2\sqrt{\lambda_{max}}) = 1/(2\gamma)$; for instance, see our example in Figure 1.
- Case $\lambda_{min} \neq \lambda_{max}$: Consider that we begin with a very small learning rate, and we gradually increment its value. Then, the solutions $|z_1(\lambda)|$, for $\lambda \in \Lambda$, will progressively traverse the blue curve in Figure 3. It is easy to see that the rate of convergence - i.e. the $\max\{|z_1(\lambda)| : \lambda \in \Lambda\}$ - will decrease (the dynamics will accelerate), until we reach the learning rate η_{opt} that satisfies $|z_1(\eta_{opt}^2 \lambda_{min})| = |z_1(\eta_{opt}^2 \lambda_{max})|$. Naturally, it follows that $\eta_{opt}^2 \lambda_{min} < 1/4$ and $\eta_{opt}^2 \lambda_{max} > 1/4$ and thus, we can obtain the optimal learning rate η_{opt} through Equations (43) and (45).

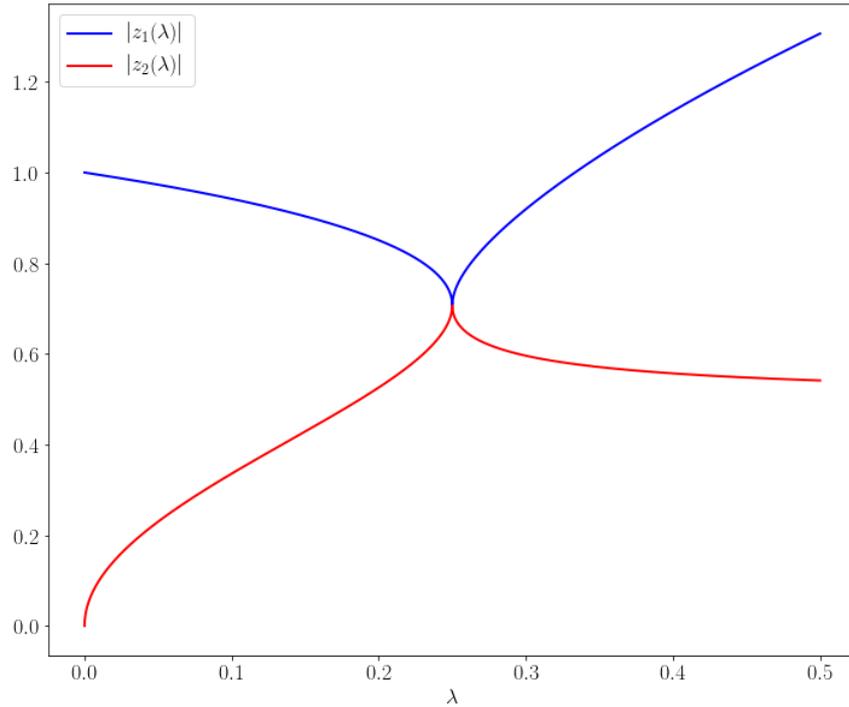


Figure 3: The norm of the solutions of Equation (40) with respect to λ .