

---

# Accelerating Smooth Games by Manipulating Spectral Shapes

---

Waïss Azizian<sup>1,†</sup> Damien Scieur<sup>2</sup> Ioannis Mitliagkas<sup>3,‡</sup> Simon Lacoste-Julien<sup>2,3,‡</sup> Gauthier Gidel<sup>3</sup>  
<sup>1</sup>École Normale Supérieure, Paris <sup>2</sup>Mila & SAIT AI Lab, Montreal <sup>3</sup>Mila & DIRO, Université de Montréal

## Abstract

We use matrix iteration theory to characterize acceleration in smooth games. We define the *spectral shape* of a family of games as the set containing all eigenvalues of the Jacobians of standard gradient dynamics in the family. Shapes restricted to the real line represent well-understood classes of problems, like minimization. Shapes spanning the complex plane capture the added numerical challenges in solving smooth games. In this framework, we describe gradient-based methods, such as extragradient, as transformations on the spectral shape. Using this perspective, we propose an optimal algorithm for bilinear games. For smooth and strongly monotone operators, we identify a continuum between convex minimization, where acceleration is possible using Polyak’s momentum, and the worst case where gradient descent is optimal. Finally, going beyond first-order methods, we propose an accelerated version of consensus optimization.

## 1 Introduction

Recent successes of multi-agent formulations in various areas of deep learning (Goodfellow et al., 2014; Pfau and Vinyals, 2016) have caused a surge of interest in the theoretical understanding of first-order methods for the solution of differentiable multi-player games (Palaniappan and Bach, 2016; Gidel et al., 2019a; Balduzzi et al., 2018; Mescheder et al., 2017, 2018; Mazumdar et al., 2019). This exploration hinges on a key question:

*How fast can a first-order method be?*

In convex minimization, Nesterov (1983, 2004) answered this question with lower bounds for the rate

---

<sup>†</sup>Work done during an internship at Mila.

<sup>‡</sup>Canada CIFAR AI Chair.

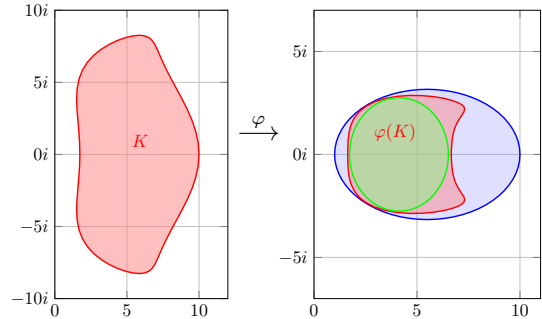


Figure 1: Transformation of the spectral shape  $K$  (in red from left to right) by the extragradient operator  $\varphi : \lambda \mapsto \lambda(1 - \eta\lambda)$ . Any ellipse (e.g. in blue) that contains the transformed red shape  $\varphi(K)$  provides an upper convergence bound using extragradient with Polyak momentum (with step-size and momentum that depends on the ellipse parameters). Any ellipse included in it (e.g. in green) provides a lower bound. See §3.3.

of convergence and an accelerated, momentum-based algorithm matching that optimal lower bound.

The dynamics of numerical methods is often described by a vector field,  $F$ , and summarized in the spectrum of its Jacobian. In minimization problems, the eigenvalues of the Jacobian lie on the real line. On strongly convex problems, the *condition number* (the dynamic range of eigenvalues) is at the heart of Nesterov’s upper and lower bound results, characterizing the hardness of a minimization problem.

Our understanding of differentiable games is nowhere close to this point. There, the eigenvalues of the Jacobian at the solution are distributed on the complex plane, suggesting a richer, more complex set of dynamics (Mescheder et al., 2017; Balduzzi et al., 2018). Some old papers (Korpelevich, 1976; Tseng, 1995) and many recent ones (Nemirovski, 2004; Chen et al., 2014; Palaniappan and Bach, 2016; Mescheder et al., 2017; Gidel et al., 2019a,b; Daskalakis et al., 2018; Mokhtari et al., 2019; Azizian et al., 2019) suggest new methods and provide better upper bounds.

All of the above work relies on bounding the magnitude or the real part of the eigenvalues of submatrices of the Jacobian. This coarse-grain approach can be oblivious

to the dependence of upper and lower bounds on the exact distribution of eigenvalues on the complex plane. More importantly, the questions of acceleration and optimality have not been answered for smooth games.

In this paper, we take a different approach. We use matrix iteration theory to characterize acceleration in smooth games. Our analysis framework revolves around the *spectral shape* of a family of games, defined as the set containing all eigenvalues of the Jacobians of natural gradient dynamics in the family (cf. §3.2). This fine-grained analysis framework can capture the dependence of upper and lower bounds on the specific shape of the spectrum. Critically, it allows us to establish acceleration in specific families of smooth games.

**Contributions.** Our main contribution is a geometric interpretation of the conditioning of a game (via its *spectral shape* as illustrated in Fig. 1, and discussed with more details in §3.3). Our result links the “hardness” of a game to the distribution of the eigenvalues of its Jacobian of the game at the optimum. Using our framework, we make the following contributions.

1. We show a reduction from bilinear games to games with *real* eigenvalues, where acceleration is possible through momentum. We provide lower bounds and design an optimal algorithm for this class.
2. Showing that acceleration persists even if there is an “imaginary perturbation”, we propose an accelerated version of extragradient (EG) for bilinear games.
3. We accelerate *consensus optimization* (CO), a *cheap* second-order method. We combine it with momentum to achieve a nearly-accelerated rate, improving the best rate previously known for this method.

**Organisation.** We recall the definition of the *asymptotic convergence factor* in §4 and use it to show that acceleration is *not possible* for the general class of smooth and strongly monotone games. In §5 we show that bilinear games or games with a “small imaginary perturbation” can be accelerated. Finally, in §6 we improve the rate of CO by using momentum.

## 2 Related work

**Matrix iteration theory.** There is extensive literature on iterative methods for linear systems, due to their countless applications. An important line of work considers the design of iterative methods through the lens of approximation problems by polynomials on the complex plane. Eiermann and Niethammer (1983) then used complex analysis tools to define, for a given compact set, its *asymptotic convergence factor*: it is the optimal asymptotic convergence rate a first-order method can achieve for all linear systems with spectrum in the set. Niethammer and Varga (1983) bring tools from summability theory to analyze multi-step iterative

methods in this framework and provide optimal methods, in particular, the momentum method for ellipses. Eiermann et al. (1985) continued in this direction, summarizing and improving the previous results. Finally Eiermann et al. (1989) study how polynomial transformations of the spectrum help compute the asymptotic convergence factor and the optimal method for a given set, potentially yielding faster convergence.

**Acceleration and lower bounds.** Lower bounds of convergence are standard in convex optimization (Nesterov, 2004) but are often non-asymptotic or cast in an infinite-dimensional space. Arjevani et al. (2016); Arjevani and Shamir (2016) showed non-asymptotic lower bounds using a framework called  $p$ -SCLI close to matrix iteration theory. Ibrahim et al. (2019); Azizian et al. (2019) extended this framework to multi-player games, but they consider lower and upper-bounds on the eigenvalues of the Jacobian of the game rather than their distribution in the complex plane. Two main acceleration methods in convex optimization achieve these lower bounds, Polyak’s momentum (Polyak, 1964) and Nesterov’s acceleration Nesterov (1983). The latter is the only one that has global convergence guarantees for convex functions. Nevertheless, Polyak’s momentum still plays a crucial role in the training of large scale machine learning models Sutskever et al. (2013).

**Acceleration for games.** Recent work applied acceleration techniques to game optimization. Gidel et al. (2019b) showed that negative momentum with alternating updates converges on bilinear games, but with the same geometrical rate as EG. Chen et al. (2014) provided a version of the mirror-prox method which improves the constant but not its rate. In the context of minimax optimization, Palaniappan and Bach (2016) used Catalyst (Lin et al., 2015), a generic acceleration method, to improve the convergence of variance-reduced algorithms for min-max problems. In the context of variational inequalities, the standard assumptions on the operator are Lipschitzness and (strong) monotonicity (Tseng, 1995; Nesterov, 2003). Nemirovski (2004) provided a lower bound in  $\mathcal{O}(1/t)$  on the convergence rate for smooth monotone games, which suggests that EG is nearly optimal in the strongly monotone case. In our work, we show that acceleration is possible by substituting the smoothness and monotonicity assumptions on the operator into more precise assumptions on the *eigenvalues of its Jacobian*.

## 3 Setting and notation

We consider the problem of finding a stationary point  $\omega^* \in \mathbb{R}^d$  of a vector field  $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ , i.e.,  $F(\omega^*) = 0$ , the solution of an unconstrained *variational inequality* problem (Harker and Pang, 1990). A relevant special case is a  $n$ -player convex game, where  $\omega^*$  corresponds

to a Nash equilibrium (Von Neumann and Morgenstern, 1944; Balduzzi et al., 2018). Consider  $n$  players  $i = 1, \dots, n$  who want to minimize their loss  $l_i(\omega^{(i)}, \omega^{(-i)})$ . The notation  $\cdot^{(-i)}$  means all indexes but  $i$ . A Nash equilibrium satisfies

$$(\omega^*)^{(i)} \in \arg \min_{\omega^{(i)} \in \mathbb{R}^{d_i}} l_i(\omega^{(i)}, (\omega^*)^{(-i)}) \quad \forall i \in \{1, \dots, n\}.$$

In this situation no player can unilaterally reduce its loss. The vector field of the game is

$$F(\omega) = [\nabla_{\omega_1} l_1^T(\omega^{(1)}, \omega^{(-1)}), \dots, \nabla_{\omega_n} l_n^T(\omega^{(n)}, \omega^{(-n)})]^T.$$

### 3.1 First-order methods

To study lower bounds of convergence, we need a class of algorithms. We consider the classic definition<sup>1</sup> of first-order methods from Nemirovsky and Yudin (1983).

**Definition 1.** A *first-order method* generates

$$\omega_t \in \omega_0 + \text{Span}\{F(\omega_0), \dots, F(\omega_{t-1})\}, \quad t \geq 1.$$

This class is widely used in large-scale optimization, as it involves only gradient computation. For instance, Nesterov’s acceleration belongs to the class of first-order methods. On the contrary, this definition does not cover Adagrad (Duchi et al., 2011), that could conceptually be also considered as first-order. This is due to the diagonal re-scaling, so  $\omega_t$  can go *outside* the span of gradients. The next proposition gives a way to easily identify first-order methods that fit our definition.

**Proposition 1.** (Arjevani and Shamir, 2016) *first-order methods can be written as*

$$\omega_{t+1} = \sum_{k=0}^t \alpha_k^{(t)} F(\omega_k) + \beta_k^{(t)} \omega_k, \quad (1)$$

where  $\sum_{k=0}^t \beta_k^{(t)} = 1$ . The method is called *oblivious* if the coefficients  $\alpha_k^{(t)}$  and  $\beta_k^{(t)}$  are known in advance.

Oblivious methods allow the knowledge of “side information” on the function, like its smoothness constant. Most of first-order methods belong to this class, but it excludes for instance methods with adaptive step-sizes. We show how standard methods fit into this framework.

**Gradient method.** Consider the gradient method with time-dependant step-size:  $\omega_{t+1} = \omega_t - \eta_t F(\omega_t)$ . This is a first-order method, where  $\alpha_t^{(t)} = -\eta_t$ ,  $\beta_t^{(t)} = 1$  and all the other coefficients set to zero.

**Momentum method.** The momentum method defines iterates as  $\omega_{t+1} = \omega_t - \alpha F(\omega_t) + \beta(\omega_t - \omega_{t-1})$ . It fits into the previous framework with  $\alpha_t^{(t)} = -\alpha$ ,  $\beta_t^{(t)} = 1 + \beta$ ,  $\beta_{t-1}^{(t)} = -\beta$ .

<sup>1</sup>Technically, first-order algorithms are more generally methods that have access only to first-order oracles.

**Extragradient method.** Though slightly trickier, the extragradient method (EG) is also encompassed by this definition. The iterates of EG are defined by  $\omega_{t+1} = \omega_t - \eta F(\omega_t - \eta F(\omega_t))$  where

$$\begin{cases} \beta_t^{(t)} = 0, \beta_{t-1}^{(t)} = 1 & \text{if } t \text{ is odd (update)}, \\ \beta_t^{(t)} = 1, \beta_{t-1}^{(t)} = 0 & \text{if } t \text{ is even (extrapolation)}, \end{cases}$$

and  $\alpha_t^{(t)} = -\eta$  the step size.

The next (known) lemma shows that when  $F$  is linear, first-order methods can be written using *polynomials*.

**Lemma 1.** (e.g. Chihara, 2011) *If  $F(\omega) = A\omega + b$ ,*

$$\omega_t - \omega^* = p_t(A)(\omega_0 - \omega^*), \quad (2)$$

where  $\omega^*$  satisfies  $A\omega^* + b = 0$  and  $p_t$  is a real polynomial of degree at most  $t$  such that  $p_t(0) = 1$ .

We denote by  $\mathcal{P}_t$  the set of real polynomials of degree at most  $t$  such that  $p_t(0) = 1$ . Hence, the convergence of a first-order method can be analyzed through the sequence of polynomials  $(p_t)_t$  it defines.

### 3.2 Problem class

In the previous section, when  $F$  is the linear function  $F = Ax + b$ , the iterates  $\omega_t$  follow the relation (2) involving the polynomial  $p_t$ . Since all first-order methods can be written using polynomials (1), they follow

$$\|\omega_t - \omega^*\|_2 = \|p_t(A)(\omega_0 - \omega^*)\|_2. \quad (3)$$

This gives the rate of convergence of the method for a specific matrix  $A$ . Instead, we consider a larger class of problems. It consists of a set  $\mathcal{M}_K$  of matrices  $A$  whose eigenvalues belong to a set  $K$  on the complex plane,

$$\mathcal{M}_K := \{A \in \mathbb{R}^d : \text{Sp}(A) \subset K \subset \mathbb{C}_+\}, \quad (4)$$

where  $\text{Sp}(A)$  is the set of eigenvalues of  $A$  and  $\mathbb{C}_+$  is the set of complex numbers with positive real part. Moreover, we assume that  $d \geq 2$  to avoid trivial cases.

### 3.3 Geometric intuition

Our paper is entirely based on the study of the support  $K$  of the eigenvalues of the Jacobian of the operator  $F$ , denoted by  $\mathbf{J}_F(\omega^*)$ . Before detailing our theoretical results, we give a high-level explanation of our objectives. This geometric intuition comes from the fact that the standard assumptions made in the literature correspond to particular problem classes  $\mathcal{M}_K$ .

**Smooth and strongly convex minimization.** Consider the minimization of a twice-differentiable,  $L$ -smooth and  $\mu$ -strongly convex function  $f$ ,

$$\mu \mathbf{I} \preceq \nabla^2 f(\omega) \preceq L \mathbf{I} \quad \forall \omega \in \mathbb{R}^d.$$

There is a link between minimization problems and games, since the vector field  $F$  becomes the gradient of the objective, and its Jacobian  $\mathbf{J}_F(\omega)$  is the Hessian  $\nabla^2 f(\omega)$ . Thus, the class corresponding to the minimization of smooth, strongly convex functions is

$$\{F : \forall \omega \in \mathbb{R}^d, \text{Sp } \mathbf{J}_F(\omega) \subset [\mu, L]\}, \quad 0 < \mu \leq L.$$

**Bilinear games.** Consider the following problem,

$$\min_{x \in \mathbb{R}^d} \max_{y \in \mathbb{R}^d} x^\top A y.$$

Its Jacobian  $\mathbf{J}_F(\omega)$  is constant and skew-symmetric. It is a standard linear algebra result (see Lem. 7) to show that  $\text{Sp } \mathbf{J}_F(\omega) \in \pm[i\sigma_{\min}(A), i\sigma_{\max}(A)]$ .

**Variational inequalities.** The Lipschitz assumption

$$\|F(\omega) - F(\omega')\|_2^2 \leq L\|\omega - \omega'\|_2^2 \quad (5)$$

implies an upper bound on the magnitude of the eigenvalues of  $\mathbf{J}_F(\omega^*)$ . The strong monotonicity assumption

$$(\omega - \omega')^T (F(\omega) - F(\omega')) \geq \mu\|\omega - \omega'\|_2^2 \quad (6)$$

implies a lower bound on the real part of the eigenvalues of  $\mathbf{J}_F(\omega^*)$  (see Lem. 5 in §B) which thus belong to

$$K = \{\lambda \in \mathbb{C} : 0 < \mu \leq \Re \lambda, |\lambda| \leq L\}.$$

This set is the intersection between a circle and a half-plane, as shown in Figure 2 (left).

**Fine-grained bounds.** Nemirovski (2004) provides a lower-bound for the class of strongly monotone and Lipschitz operators (see §4.2) excluding the possibility of acceleration in that general setting. It motivates the adoption of more refined assumptions on the eigenvalues of  $\mathbf{J}_F(\omega^*)$ . We consider the class of games where these eigenvalues belong to a specified set  $K$ . Since  $\mathbf{J}_F(\omega^*)$  is real, its spectrum is symmetric w.r.t. the real axis, so we assume that  $K$  is too. For this class of problem, we have a simple method to compute lower and upper convergence bounds using a class of well studied shapes: ellipses.

**Proposition 2** (Ellipse method for lower and upper bound (Informal)). *Let  $K \subset \mathbb{C}_+$  be a compact set, then any ellipse symmetric w.r.t. the real axis that includes (resp. is included in)  $K$  provides an upper (resp. lower) convergence bound for the class of problem  $\mathcal{M}_K$  using Polyak momentum with a step-size and a momentum depending on the ellipse.*

See Appendix C.2, Thm. 6 for the precise result on ellipses. The proposition extends to any shape whose optimal algorithm (resp. lower bound) is known. This proposition, illustrated in Fig. 1, heavily relies on the fact that, the optimal method for ellipses is Polyak momentum (Niethammer and Varga, 1983).

Any first-order method can be seen as a way to transform the set  $K$ . In order to illustrate that we consider Lemma 1: since a first-order method update for a linear operator  $F = Ax + b$  can be written using a polynomial  $p$ , the eigenvalues to consider are not the ones of  $A$  but the ones of  $p(A)$ . Thus, the set of interest is  $p(K)$ .

As an example, consider EG with momentum. This consists in applying the momentum method to the transformed vector field  $\omega \mapsto F(\omega - \eta F(\omega))$ . From a spectral point of view, this is equivalent to first transforming the shape  $K$  into  $\varphi(K)$  with the extragradient mapping  $\varphi_\eta : \lambda \mapsto \lambda(1 - \eta\lambda)$ , then study the effect of momentum on  $\varphi(K)$ . This example of transformation is illustrated in Fig. 1, and this idea is used in §5.4.

## 4 Asymptotic convergence factor

We recall known results that compute lower bounds for some classes of games using the *asymptotic convergence factor* (Eiermann and Niethammer, 1983; Eiermann et al., 1985; Nevanlinna, 1993). Then, we illustrate them on two particular classes of problems.

### 4.1 Lower bounds for a class of problems

We now show how to lower bound the worst-case rate of convergence of a *specific* method over the class  $\mathcal{M}_K$  (4), with the worst possible initialisation  $\omega_0$ . We start with equation (3), but this time we pick the worst-case over all matrices  $A \in \mathcal{M}_K$ , i.e.,

$$\max_{A \in \mathcal{M}_K} \|p_t(A)(\omega_0 - \omega^*)\|_2.$$

Now, we can pick an arbitrary bad initialisation  $\omega_0$ , in particular, the one that corresponds to the largest eigenvalue of  $p_t(A)$  in magnitude. This gives

$$\begin{aligned} \exists \omega_0 : \|\omega_t - \omega^*\|_2 &\geq \max_{A \in \mathcal{M}_K} \rho(p_t(A)) \|\omega_0 - \omega^*\|_2 \\ &= \max_{\lambda \in K} |p_t(\lambda)| \|\omega_0 - \omega^*\|_2. \end{aligned} \quad (7)$$

It remains to lower bound  $\max_{\lambda \in K} |p_t(\lambda)|$  over *all possible* first-order methods. This is called the *asymptotic convergence factor*, presented in the next section.

### 4.2 Asymptotic convergence factor

Here we recall the definition of the *asymptotic convergence factor* (Eiermann and Niethammer, 1983), which gives a lower bound for the rate of convergence over matrices which belong to the class  $\mathcal{M}_k$  (4), for all possible first-order methods. We mainly follow the definition of Nevanlinna (1993) (see Rmk. 1 in §B for details).

The simplest way to lower bound  $\|\omega_t - \omega^*\|_2$  is given by minimizing (7) over all polynomials corresponding



to a first-order method. By Lemma 1, this class of polynomials is given by  $\mathcal{P}_t$ . Thus, for some  $\omega_0$ ,

$$\|\omega_t - \omega^*\| \geq \min_{p_t \in \mathcal{P}_t} \max_{\lambda \in K} |p_t(\lambda)| \cdot \|\omega_0 - \omega^*\|_2.$$

The *asymptotic convergence factor*  $\rho(K)$  for the class  $K$  is given by taking the *minimum average rate of convergence* over  $t$  for any  $t$ , i.e.,

$$\rho(K) = \inf_{t > 0} \min_{p_t \in \mathcal{P}_t} \max_{\lambda \in K} \sqrt[t]{|p_t(\lambda)|}. \quad (8)$$

This way, by construction,  $\rho(K)$  gives a lower-bound on the *worst-case* rate of convergence for the class  $\mathcal{M}_K$ . We formalize this statement in the proposition below.

**Proposition 3.** (*Nevanlinna, 1993*) *Let  $K \subset \mathbb{C}$  be a subset of  $\mathbb{C}$  symmetric w.r.t. the real axis, which does not contain 0 and such that  $\mathcal{M}_K \neq \emptyset$ . Then, any oblivious first-order method (whose coefficients only depend on  $K$ ) satisfies,*

$$\forall t \geq 0, \exists A \in \mathcal{M}_K, \exists \omega_0 : \|\omega_t - \omega^*\|_2 \geq \rho(K)^t \|\omega_0 - \omega^*\|_2.$$

However, the object  $\rho(K)$  may be complicated to obtain as it depends on the solution of a minimax problem *over a set*  $K \subset \mathbb{C}_+$ . If the set is simple enough, we can lower-bound the asymptotic rate of convergence. We start by giving the two extreme cases: when  $K$  is a segment on the real line (convex and smooth minimization) or  $K$  is a disc (monotone and smooth games).

### 4.3 Extreme cases: real segments and discs

#### Smooth and strongly convex minimization.

In the case where we are interested in lower-bounds, we can consider the restricted class of functions where  $J_F(\omega) (= \nabla^2 f(\omega))$  is constant, i.e., independent of  $\omega$ . This corresponds to quadratic minimization, and our restricted class becomes

$$\mathcal{M}_K \quad \text{where } K = [\mu, L].$$

For this specific class, where  $K$  is a segment in the real line, the solution to the subproblem associated to the *asymptotic rate of convergence* (8), i.e.,

$$\min_{p \in \mathcal{P}_t} \max_{\lambda \in [\mu, L]} |p(\lambda)| \quad (9)$$

is well-known. The optimal polynomial  $p_t^*$  is a properly scaled and translated Chebyshev polynomial of the first kind of degree  $t$  (Golub and Varga, 1961; Manteuffel, 1977). The rate of convergence of  $p_t$  evolves with  $t$ , but asymptotically converges to

$$\rho([\mu, L]) = \frac{\sqrt{L} - \sqrt{\mu}}{\sqrt{L} + \sqrt{\mu}}.$$

This is the lower bound of Nesterov (2004, Thm. 2.1.13), which corresponds to an accelerated linear rate. The condition number  $L/\mu$  appears as a square root unlike




Figure 2: **Left:** Illustration of the proof of Cor. 1. The yellow set correspond to  $K$ , the set of strongly monotone problems while the red disc is the disc of center  $\frac{1}{2}(\mu + L)$  and radius  $\frac{1}{2}(L - \mu)$  which fits inside. **Right:** Illustration of  $K_\epsilon$  of Prop. 6 with  $\epsilon = \sqrt{\mu L}$ .

for the rate of the plain gradient descent, which implies a huge (asymptotic) improvement.

In this section, we have seen that when the spectrum is constrained to be on a segment in the real line, one can expect acceleration. The next section shows that this is not the case for the class of discs.

**Discs and strongly monotone vector fields** Consider a disc with a real positive center

$$K = \{z \in \mathbb{C} : |z - c| \leq r\}, \quad \text{with } 0 < c < r.$$

This time again, the shape is simple enough to have an explicit solution for the optimal polynomials

$$p_t^*(\lambda) = \arg \min_{p_t \in \mathcal{P}_t} \max_{\lambda \in K} |p_t(\lambda)|.$$

In this case, the optimal polynomial reads  $p_t^*(\omega) = (1 - \omega/c)^t$ , and this corresponds to gradient descent with step-size  $\eta = 1/c$ . Hence, with this specific shape, gradient method is optimal (Eiermann et al., 1985, §6.2); Nevanlinna (1993, Example 3.8.2). A direct consequence of this result is a lower bound of convergence for the class of Lipschitz, strongly monotone vector fields, i.e., vector fields  $F$  that satisfies (5)-(6). For linear vector fields parameterized by the matrix  $A$  as in Lemma 1, this is included in the set

$$\mathcal{M}_K, K = \{\lambda \in \mathbb{C} : 0 < \mu \leq \Re \lambda, |\lambda| \leq L\}. \quad (10)$$

This set is the intersection between a circle and a half-plane, as shown in Figure 2 (left). Notice that the disc of center  $\frac{\mu+L}{2}$  and radius  $\frac{L-\mu}{2}$  actually fits in  $K$ , as illustrated by Fig. 2. Since this disc is *included* in  $K$ , a lower bound for the disc also gives a lower bound for  $K$ , as stated in the following corollary.

**Corollary 1.** *Let  $K$  be defined in (10). Then,*

$$\rho(K) > \frac{L-\mu}{L+\mu} = 1 - \frac{2\mu}{L+\mu}.$$

The rate of Cor. 1 is already achieved by first-order methods, without momentum or acceleration, such as EG. Thus, acceleration is *not possible* for the general class of smooth, strongly monotone games.

## 5 Acceleration in games

We present our contributions in this section. The previous section highlights a big contrast between optimization and games. In the former, acceleration is possible, but this does not generalize for the latter. Here, we explore acceleration via a sharp analysis of intermediate cases, like imaginary segments (bilinear games) or thin ellipses (perturbed acceleration), via lower and upper bounds. Since we use spectral arguments, the convergence guarantees of our algorithms are local, but lower bounds remain valid globally.

### 5.1 Local convergence of optimization methods for nonlinear vector fields

Before presenting our result, we recall the classical local convergence theorem from Polyak (1964). In this section, we are interested in finding the fixed point  $\omega^*$  of a vector field  $V$ , i.e.,  $V(\omega^*) = \omega^*$ .  $V$  here plays the role of an iterative optimization methods and defines iterates according to the fixed-point iteration

$$\omega_{t+1} = V(\omega_t). \quad (11)$$

Analysing the properties of the vector field  $V$  is usually challenging, as  $V$  can be any nonlinear function. However, under mild assumption, we can simplify the analysis by considering the linearization  $V(\omega) \approx V(\omega^*) + \mathbf{J}_V(\omega^*)(\omega - \omega^*)$ , where  $\mathbf{J}_V(\omega)$  is the Jacobian of  $V$  evaluated at  $\omega^*$ . The next theorem shows we can deduce the rate of convergence of (11) using the spectral radius of  $\mathbf{J}_V(\omega^*)$ , denoted by  $\rho(\mathbf{J}_V(\omega^*))$ .

**Theorem 1** (Polyak (1987)). *Let  $V : \mathbb{R}^d \rightarrow \mathbb{R}^d$  be continuously differentiable and let  $\omega^*$  one of its fixed-points. Assume that there exists  $\rho^* > 0$  such that,*

$$\rho(\mathbf{J}_V(\omega^*)) \leq \rho^* < 1.$$

*For  $\omega_0$  close to  $\omega^*$ , (11) converges linearly to  $\omega^*$  at a rate  $\mathcal{O}((\rho^* + \epsilon)^t)$ . If  $V$  is linear, then  $\epsilon = 0$ .*

Recent works such as Mescheder et al. (2017); Gidel et al. (2019b); Daskalakis and Panageas (2018) used this connection to study game optimization methods.

Thm. 1 can be applied directly on methods which use only the last iterate, such as gradient or EG. For methods that do not fall into this category, such as momentum, a small adjustment is required, called *system augmentation*.

Consider that  $V : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$  follows the recursion

$$\omega_{t+1} = V(\omega_t, \omega_{t-1}). \quad (12)$$

Instead we consider its *augmented operator*

$$\begin{bmatrix} \omega_t \\ \omega_{t+1} \end{bmatrix} = V_{\text{augm}}(\omega_t, \omega_{t-1}) = \begin{bmatrix} \omega_t \\ V(\omega_t, \omega_{t-1}) \end{bmatrix},$$

to which we can now apply the previous theorem. This technique is summarized in the following lemma.

**Lemma 2.** *Let  $V : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$  be continuously differentiable and let  $\omega^*$  satisfies  $V(\omega^*, \omega^*) = \omega^*$ . Assume there exists  $\rho^* > 0$  such that,  $\rho(\mathbf{J}_{V_{\text{augm}}}(\omega^*)) \leq \rho^* < 1$ . If  $\omega_0$  and  $\omega_1$  are close to  $\omega^*$ , then (11) converges linearly to  $\omega^*$  at rate  $(\rho^* + \epsilon)^t$ . If  $V$  is linear, then  $\epsilon = 0$ .*

### 5.2 Acceleration for bilinear games

For convex minimization, adding momentum results in an accelerated rate for strongly convex functions we have discuss above. For instance, if  $\text{Sp } \nabla F(\omega^*) \subset [\mu, L]$ , the Polyak's Heavy-ball method (see the full statement in Appendix C.1), Polyak (1964, Thm. 9)

$$\begin{aligned} \omega_{t+1} &= V^{\text{Polyak}}(\omega_t, \omega_{t-1}) \\ &:= \omega_t - \alpha F(\omega_t) + \beta(\omega_t - \omega_{t-1}) \end{aligned} \quad (13)$$

converges (locally) with the accelerated rate

$$\rho(\mathbf{J}_{V^{\text{Polyak}}}(\omega^*, \omega^*)) \leq \frac{\sqrt{L} - \sqrt{\mu}}{\sqrt{L} + \sqrt{\mu}}.$$

Another example are bilinear games. Most known methods converge at a rate of  $(1 - c\sigma_{\min}(A)^2/\sigma_{\max}(A)^2)^t$  for some  $c > 0$  (Daskalakis et al., 2018; Mescheder et al., 2017; Gidel et al., 2019a,b; Liang and Stokes, 2018; Abernethy et al., 2019). Using results from Eiermann et al. (1989), we show that this rate is suboptimal.

For bilinear games, the eigenvalues of the Jacobian  $\mathbf{J}_F$  are purely imaginary (see Lem. 7 in Appendix C.1), i.e.,

$$K = [i\sigma_{\min}(A), i\sigma_{\max}(A)] \cup [-i\sigma_{\min}(A), -i\sigma_{\max}(A)].$$

A method that follows strictly the vector field  $F$  does not converge, as its flow is composed by only concentric circles, thus leading to oscillations. This problem is avoided if we transform the vector field into another one with better properties. For example, the transformation

$$F^{\text{real}}(\omega) = \frac{1}{\eta}(F(\omega - \eta F(\omega)) - F(\omega)) \quad (14)$$

can be seen as a finite-difference approximation of  $\nabla(\frac{1}{2}\|F\|_2^2)$ . It is easier to find the equilibrium of  $V$  since the eigenvalues of  $\mathbf{J}_V(\omega) = -\mathbf{J}_F^2(\omega)$  are located on a real segment. Thus, we can use standard minimization methods like the Polyak Heavy-Ball method.

**Proposition 4.** *Let  $F$  be a vector field such that  $\text{Sp } \nabla F(\omega^*) \subset [ia, ib] \cup [-ia, -ib]$ , for  $0 < a < b$ . Setting  $\sqrt{\alpha} = \frac{2}{a+b}$ ,  $\sqrt{\beta} = \frac{b-a}{b+a}$ , the Polyak Heavy-Ball method (13) on the transformation (14), i.e.,*

$$\omega_{t+1} = \omega_t - \alpha F^{\text{real}}(\omega_t) + \beta(\omega_t - \omega_{t-1}).$$

*converges locally at a linear rate  $\mathcal{O}((1 - \frac{2a}{a+b})^t)$ .*

Using results from Eiermann et al. (1989), we show that this method is optimal. Indeed, for this set, we can compute explicitly  $\rho(K)$  from (8), the lower bound for the local convergence factor.

**Proposition 5.** *Let  $K = [ia, ib] \cup [-ia, -ib]$  for  $0 < a < b$ . Then,  $\rho(K) = \sqrt{\frac{b-a}{b+a}}$ .*

*Proof.* (Sketch). The transformation that we have applied, i.e.  $\lambda \mapsto -\lambda^2$ , preserves the asymptotic convergence factor  $\rho$  (up to a square root), as it satisfies the assumptions of Eiermann et al. (1989, Thm. 6).  $\square$

The difference of a square root between the lower bound and the bound on the spectral radius is explained by the fact that the method presented here queries two gradient per iteration and so one of its iterations actually corresponds to two steps of a first-order method as defined in Definition 1.

In this subsection, we showed that when the eigenvalues of the Jacobian are purely real or imaginary, acceleration is possible using momentum on the right vector field. Yet the previous subsection shows it is not the case for general smooth, strongly monotone games. The question of acceleration remains for intermediate shapes, like ellipses. The next subsection shows how to recover an accelerated rate of convergence in this case.

### 5.3 Perturbed acceleration

As we cannot compute  $\rho$  explicitly for most sets  $K$ , we focus on ellipses to answer this question. They have been well studied, and optimal methods are again based on Chebyshev polynomials (Manteuffel, 1977).

In this section we study games whose eigenvalues of the Jacobian belong to a thin ellipse. These ellipses correspond to the real segments  $[\mu, L]$  perturbed in an elliptic way, see Fig. 2 (right). Mathematically, we have for  $0 < \mu < L$  and  $\epsilon > 0$ , the equation

$$K_\epsilon = \left\{ z \in \mathbb{C} : \left( \frac{\Re z - \frac{\mu+L}{2}}{\frac{L-\mu}{2}} \right)^2 + \left( \frac{\Im z}{\epsilon} \right)^2 \leq 1 \right\}$$

When  $\epsilon = 0$  (with the convention that  $0/0 = 0$ ), Polyak momentum achieves the rate of  $1 - 2\frac{\sqrt{\mu}}{\sqrt{\mu} + \sqrt{L}}$ . However, when  $\epsilon = \frac{L-\mu}{2}$ , we showed the lower bound of  $1 - 2\frac{\mu}{\mu+L}$  in Cor. 1. To check if acceleration still persists for intermediate cases, we study the behaviour of the asymptotic convergence factor (when  $L/\mu \rightarrow +\infty$ ) as a function of  $\epsilon$ . The next proposition uses results from Niethammer and Varga (1983); Eiermann et al. (1985) to show that acceleration is still possible on  $K_\epsilon$ .

**Proposition 6.** *Define  $\epsilon(\mu, L)$  as  $\frac{\epsilon(\mu, L)}{L} = \left(\frac{\mu}{L}\right)^\theta$  with*

$\theta > 0$  and  $a \wedge b = \min(a, b)$ . Then, when  $\frac{\mu}{L} \rightarrow 0$ ,

$$\rho(K_\epsilon) = \begin{cases} 1 - 2\sqrt{\frac{\mu}{L}} + \mathcal{O}\left(\left(\frac{\mu}{L}\right)^{\theta \wedge 1}\right), & \text{if } \theta > \frac{1}{2} \\ 1 - 2(\sqrt{2} - 1)\sqrt{\frac{\mu}{L}} + \mathcal{O}\left(\frac{\mu}{L}\right), & \text{if } \theta = \frac{1}{2} \\ 1 - \left(\frac{\mu}{L}\right)^{1-\theta} + \mathcal{O}\left(\left(\frac{\mu}{L}\right)^{1 \wedge (2-3\theta)}\right), & \text{if } \theta < \frac{1}{2}. \end{cases}$$

Moreover, the momentum method is optimal for  $K_\epsilon$ . This means there exists  $\alpha > 0$  and  $\beta > 0$  (function of  $\mu, L$  and  $\epsilon$  only) such that if  $\text{Sp } \mathbf{J}_F(\omega^*) \subset K_\epsilon$ , then,  $\rho(\mathbf{J}_{V^{\text{Polyak}}}(\omega^*, \omega^*)) \leq \rho(K_\epsilon)$ .

This shows that the convergence rate interpolates continuously between the accelerated rate and the non-accelerated one. Crucially, for small perturbations, that is to say if the ellipse is thin enough, acceleration persists until  $\theta = \frac{1}{2}$  or equivalently  $\epsilon \sim \sqrt{\mu L}$ . That's why Prop. 6 plays a central role in our forthcoming analyses of accelerated EG and CO.

### 5.4 Accelerating extragradient

We now consider the acceleration of EG using momentum. Its main appealing property is its convergence on bilinear games, unlike the gradient method. On the class of bilinear problems, EG achieves a convergence rate of  $(1 - ca^2/b^2)$  for some constant  $c > 0$ .

In the previous section, we achieved an accelerated rate on bilinear games by applying momentum to the operator  $F^{\text{real}}(\omega)$  instead of  $F$ , as the Jacobian of  $F^{\text{real}}$  has real eigenvalues when  $\mathbf{J}_F(\omega^*)$  has its spectrum in  $K$ . Here we try to apply momentum to the EG operator  $F^{\text{e-g}}(\omega)$ , defined as

$$F^{\text{e-g}}(\omega) = F(\omega - \eta F(\omega)). \quad (15)$$

Unfortunately, when  $\text{Sp } \mathbf{J}_F \subset K$ , the spectrum of  $F^{\text{e-g}}(\omega^*)$  is never purely real. Using the insight from Prop. 6, we can choose  $\eta > 0$  such that we are in the first case of Prop. 6, making acceleration possible.

**Proposition 7.** *Consider the vector field  $F$ , where  $\text{Sp } \mathbf{J}_F(\omega^*) \subset [ia, ib] \cup [-ia, -ib]$  for  $0 < a < b$ . There exists  $\alpha, \beta, \eta > 0$  such that, the operator defined by*

$$\omega_{t+1} = \omega_t - \alpha F(\omega_t - \eta F(\omega_t)) + \beta(\omega_t - \omega_{t-1}),$$

*converges locally at a linear rate  $O\left(\left(1 - c\frac{a}{b} + M\frac{a^2}{b^2}\right)^t\right)$  where  $c = \sqrt{2} - 1$  and  $M$  is an absolute constant.*

One drawback is that, to achieve fast convergence on bilinear games, one has to tune the two step-sizes  $\alpha, \eta$  of EG precisely and separately. They actually differ by a factor  $\frac{b^2}{a^2}$ :  $\eta$  is roughly proportional to  $\frac{1}{a}$  while  $\alpha$  behaves like  $\frac{a}{b^2}$  (see Lem. 9 in Appendix C.4).

## 6 Beyond typical first-order methods

In the previous section, we achieved acceleration with first-order methods for specific problem classes. However, the lower bound from Cor. 1 still prevents us from doing so for the larger problem classes for smooth and strongly monotone games. To bypass this limitation, we can consider going *beyond* first-order methods. In this section, we consider two different approaches. The first one is adaptive acceleration, which is a *non-oblivious* first-order method. The second is consensus optimization, an inversion-free second order method.

### 6.1 Adaptive acceleration

In previous sections, we considered shapes whose optimal polynomial is known. This optimal polynomial lead to an optimal first-order method. However, when the shape is *unknown*, we cannot use better methods than EG with an appropriate stepsize.

Recent work in optimization analysed adaptive algorithms, such as *Anderson Acceleration* (Walker and Ni, 2011), that are adaptive to the problem constants. They can be seen as an automatic way to find the optimal combination of the previous iterates. Recent works on Anderson Acceleration extended the theory for non-quadratic minimization, by using regularisation (Scieur et al., 2016) (RNA method). The theory has also been extended to “non symmetric operators” (Bollapragada et al., 2018), and this setting fits perfectly the one of games, as  $\mathbf{J}_F(\omega^*)$  is not symmetric.

Anderson acceleration and its extension RNA are similar to quasi-Newton (Fang and Saad, 2009), but remains first-order methods. Even if they find the optimal first-order method (for linear  $F$ ), they cannot beat a lower bound similar to Cor. 1, when the number of iterations is smaller than the dimension of the problem. The next section shows how to use *cheap* second-order information to improve the convergence rate.

### 6.2 Momentum consensus optimization

CO (Mescheder et al., 2017) iterates as follow:

$$\omega_{t+1} = \omega_t - \alpha(F(\omega_t) + \tau \mathbf{J}_F^T(\omega)F(\omega)).$$

Albeit being a second-order method, each iteration requires only one Jacobian-vector multiplication. This operation can be computed efficiently by modern machine learning frameworks, with automatic differentiation and back-propagation. For instance, for neural networks, the computation time of this product or the gradient is comparable. Moreover, unlike Newton’s method, CO does *not* require a matrix inversion.

Though CO is a second-order method, its analysis can

still be reduced to our framework by considering the following transformation of the initial operator  $F(\omega)$ ,

$$F^{\text{cons.}}(\omega) = F(\omega) + \tau \nabla \left( \frac{1}{2} \|F\|^2 \right) (\omega). \quad (16)$$

Though the eigenvalues of  $\mathbf{J}_{F^{\text{cons.}}}$  are not purely real in general, their imaginary to real part ratio can be controlled by Mescheder et al. (2017, Lem. 9) as,

$$\max_{\lambda \in \text{Sp } \mathbf{J}_{F^{\text{cons.}}}(\omega^*)} \frac{|\Im \lambda|}{|\Re \lambda|} = O\left(\frac{1}{\tau}\right).$$

Therefore, if  $\tau$  increases, these eigenvalues move closer to the real axis and can be included in a thin ellipse as described by §5.3. We then show that, if  $\tau$  is large enough, this ellipse can be chosen thin enough to fall into the accelerated regime of Prop. 6 and therefore, adding momentum achieves acceleration.

**Proposition 8.** *Let  $\sigma_i$  be the singular values of  $\mathbf{J}_F(\omega^*)$ . Assume that*

$$\gamma \leq \sigma_i \leq L, \quad \tau = \frac{L}{\gamma^2}.$$

*There exists  $\alpha, \beta$ , s.t., momentum applied to  $F^{\text{cons.}}$ ,*

$$\omega_{t+1} = \omega_t - \alpha F^{\text{cons.}}(\omega_t) + \beta(\omega_t - \omega_{t-1})$$

*converges locally at a rate  $O\left(\left(1 - c\frac{\gamma}{L} + M\frac{\gamma^2}{L^2}\right)^t\right)$  where  $c = \sqrt{2} - 1$  and  $M$  is an absolute constant.*

Hence, adding momentum to CO yields an accelerated rate. The assumption on the Jacobian encompasses both strongly monotone and bilinear games. On these two classes of problems, CO is at least as fast as any oblivious first-order method as its rate roughly matches the lower bounds of Prop. 3 and 5.

Note that, choosing  $\tau$  of this order is what is done by Abernethy et al. (2019) for (non-accelerated) CO. They claim that this point of view – seeing consensus as a perturbation of gradient descent on  $\frac{1}{2}\|F\|^2$  – is justified by practice as in the experiments of Mescheder et al. (2017),  $\tau$  is set to 10.

## 7 Conclusion

This paper shows that a spectral perspective is fundamental to understand the conditioning of games. The latter is indeed linked to the geometric properties of the distribution of the spectrum of its Jacobian. In the light of this perspective, we demonstrate how several gradient-based methods transform the spectral shape of a game to achieve accelerated convergence when combined with Polyak momentum. Our main tool throughout this paper was the flexible and convenient class of ellipses; we left as future work the study of more intricate shapes, which – ideally – would fit the distribution of the eigenvalues of applications of challenging machine learning problems such as GANs.



## Acknowledgments

This research was partially supported by the Canada CIFAR AI Chair Program, the Canada Excellence Research Chair in “Data Science for Realtime Decision-making”, by the NSERC Discovery Grants RGPIN-2017-06936 and RGPIN-2019-06512, the FRQNT new researcher program 2019-NC-257943, by a Borealis AI fellowship, a Google Focused Research award and a startup grant by IVADO. Simon Lacoste-Julien is a CIFAR Associate Fellow in the Learning in Machines & Brains program.

## Bibliography

- Jacob Abernethy, Kevin A. Lai, and Andre Wibisono. Last-iterate convergence rates for min-max optimization. *arXiv*, 2019.
- Yossi Arjevani and Ohad Shamir. On the iteration complexity of oblivious first-order optimization algorithms. In *ICML*, 2016.
- Yossi Arjevani, Shai Shalev-Shwartz, and Ohad Shamir. On Lower and Upper Bounds in Smooth and Strongly Convex Optimization. *JMLR*, 2016.
- Waïss Azizian, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. A tight and unified analysis of extragradient for a whole spectrum of differentiable games. *arXiv*, 2019.
- Joseph Bak and Donald J. Newman. *Complex Analysis*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, 2010.
- David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The Mechanics of n-Player Differentiable Games. In *ICML*, 2018.
- Raghu Bollapragada, Damien Scieur, and Alexandre d’Aspremont. Nonlinear acceleration of momentum and primal-dual algorithms. *arXiv*, 2018.
- Yunmei Chen, Guanghui Lan, and Yuyuan Ouyang. Accelerated Schemes For A Class of Variational Inequalities. *arXiv*, 2014.
- Theodore S Chihara. *An introduction to orthogonal polynomials*. Courier Corporation, 2011.
- Constantinos Daskalakis and Ioannis Panageas. The Limit Points of (Optimistic) Gradient Descent in Min-Max Optimization. *arXiv*, 2018.
- Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training GANs with Optimism. In *ICLR*, 2018.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *JMLR*, 2011.
- M. Eiermann, W. Niethammer, and R. S. Varga. A study of semiiterative methods for nonsymmetric systems of linear equations. *Numerische Mathematik*, 1985.
- Michael Eiermann and Wilhelm Niethammer. On the Construction of Semiterative Methods. *SIAM Journal on Numerical Analysis*, 1983.
- Michael Eiermann, X Li, and Richard Varga. On Hybrid Semi-Iterative Methods. *Siam Journal on Numerical Analysis*, 1989.
- Francisco Facchinei and Jong-Shi Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems Vol I*. Springer Series in Operations Research and Financial Engineering, Finite-Dimensional Variational Inequalities and Complementarity Problems. Springer-Verlag, 2003.
- Haw-ren Fang and Yousef Saad. Two classes of multisecond methods for nonlinear acceleration. *Numerical Linear Algebra with Applications*, 2009.
- Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A Variational Inequality Perspective on Generative Adversarial Networks. In *ICLR*, 2019a.
- Gauthier Gidel, Reyhane Askari Hemmat, Mohammad Pezeshki, Remi Lepriol, Gabriel Huang, Simon Lacoste-Julien, and Ioannis Mitliagkas. Negative Momentum for Improved Game Dynamics. In *AISTATS*, 2019b.
- Gene H. Golub and Richard S. Varga. Chebyshev semi-iterative methods, successive overrelaxation iterative methods, and second order Richardson iterative methods. *Numerische Mathematik*, 1961.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *NeurIPS*. 2014.
- Patrick T Harker and Jong-Shi Pang. Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Mathematical programming*, 1990.
- Adam Ibrahim, Waïss Azizian, Gauthier Gidel, and Ioannis Mitliagkas. Lower Bounds and Conditioning of Differentiable Games. *arXiv*, 2019.
- G.M. Korpelevich. The extragradient method for finding saddle points and other problems., 1976.
- Peter D. Lax. *Linear Algebra and Its Applications*. John Wiley & Sons, 2007.
- Alistair Letcher, David Balduzzi, Sébastien Racanière, James Martens, Jakob N. Foerster, Karl Tuyls, and Thore Graepel. Differentiable game mechanics. *JMLR.*, 2019.

- Tengyuan Liang and James Stokes. Interaction Matters: A Note on Non-asymptotic Local Convergence of Generative Adversarial Networks. In *AISTATS*, 2018.
- Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. In *NeurIPS*, 2015.
- Thomas A. Manteuffel. The Tchebychev iteration for nonsymmetric linear systems. *Numerische Mathematik*, 1977.
- Eric V. Mazumdar, Michael I. Jordan, and S. Shankar Sastry. On Finding Local Nash Equilibria (and Only Local Nash Equilibria) in Zero-Sum Games. *arXiv:1901.00838 [cs, math, stat]*, 2019.
- Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. The numerics of GANs. In *NeurIPS*, 2017.
- Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which Training Methods for GANs do actually Converge? In *ICML*, 2018.
- Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A Unified Analysis of Extra-gradient and Optimistic Gradient Methods for Saddle Point Problems: Proximal Point Approach. *arXiv*, 2019.
- Zeev Nehari. *Conformal Mapping*. McGraw-Hill, 1952.
- Arkadi Nemirovski. Prox-Method with Rate of Convergence  $O(1/t)$  for Variational Inequalities with Lipschitz Continuous Monotone Operators and Smooth Convex-Concave Saddle Point Problems. *SIAM Journal on Optimization*, 2004.
- Arkadii Semenovich Nemirovsky and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- Y. E. Nesterov. A method for solving the convex programming problem with convergence rate  $O(1/k^2)$ . *Dokl. Akad. Nauk SSSR*, 1983.
- Yurii Nesterov. Dual Extrapolation and its Applications for Solving Variational Inequalities and Related Problems'. SSRN Scholarly Paper, Social Science Research Network, 2003.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer, 2004.
- Olavi Nevanlinna. *Convergence of Iterations for Linear Equations*. Lectures in Mathematics. ETH Zürich. Birkhäuser Basel, 1993.
- W. Niethammer and R. S. Varga. The analysis of step iterative methods for linear systems from summability theory. *Numerische Mathematik*, 1983.
- Balamurugan Palaniappan and Francis Bach. Stochastic variance reduction methods for saddle-point problems. In *NeurIPS*, 2016.
- David Pfau and Oriol Vinyals. Connecting Generative Adversarial Networks and Actor-Critic Methods. *arXiv*, 2016.
- B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 1964.
- Boris T Polyak. *Introduction to Optimization*. Optimization Software, 1987.
- Thomas Ransford. *Potential Theory in the Complex Plane*. Cambridge University Press, 1995.
- V. K. Saul'yev. *Integration of Equations of Parabolic Type by the Method of Nets*. Elsevier, 1964.
- Damien Scieur, Alexandre d'Aspremont, and Francis Bach. Regularized nonlinear acceleration. In *NeurIPS*, 2016.
- Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *ICML*, 2013.
- Paul Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 1995.
- John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. 1944.
- Homer F Walker and Peng Ni. Anderson acceleration for fixed-point iterations. *SIAM Journal on Numerical Analysis*, 2011.
- J. L. Walsh. *Interpolation and Approximation by Rational Functions in the Complex Domain*. American Mathematical Soc., 1935.
- Fuzhen Zhang, editor. *The Schur Complement and Its Applications*. Numerical Methods and Algorithms. Springer-Verlag, 2005.