Negative Momentum for Improved Game Dynamics

Gauthier Gidel* **Reyhane Askari Hemmat*** **Mohammad Pezeshki**
Rémi Le Priol **Gabriel Huang** **Simon Lacoste-Julien**[†,‡] **Ioannis Mitliagkas**[†]
Mila & DIRO, Université de Montréal [†]Canada CIFAR AI chair [‡]CIFAR fellow

Abstract

Games generalize the single-objective optimization paradigm by introducing different objective functions for different players. Differentiable games often proceed by simultaneous or alternating gradient updates. In machine learning, games are gaining new importance through formulations like generative adversarial networks (GANs) and actor-critic systems. However, compared to single-objective optimization, game dynamics is more complex and less understood. In this paper, we analyze gradient-based methods with momentum on simple games. We prove that alternating updates are more stable than simultaneous updates. Next, we show both theoretically and empirically that alternating gradient updates with a negative momentum term achieves convergence in a difficult toy adversarial problem, but also on the notoriously difficult to train saturating GANs.

1 INTRODUCTION

Recent advances in machine learning are largely driven by the success of gradient-based optimization methods for the training process. A common learning paradigm is empirical risk minimization, where a (potentially non-convex) objective, that depends on the data, is minimized. However, some recently introduced approaches require the joint minimization of several objectives. For example, actor-critic methods can be written as a bi-level optimization problem (Pfau and Vinyals, 2016) and generative adversarial networks (GANs) (Goodfellow et al., 2014) use a two-player game formulation.

*Equal contribution.

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

Games generalize the standard optimization framework by introducing different objective functions for different optimizing agents, known as *players*. We are commonly interested in finding a local *Nash equilibrium*: a set of parameters from which no player can (locally and unilaterally) improve its objective function. Games with differentiable objectives often proceed by simultaneous or alternating gradient steps on the players' objectives. Even though the dynamics of gradient based methods is well understood for minimization problems, new issues appear in multi-player games. For instance, some stable stationary points of the dynamics may not be (local) Nash equilibria (Adolphs et al., 2018).

Motivated by a decreasing trend of momentum values in GAN literature (see Fig. 1), we study the effect of two particular algorithmic choices: (i) the choice between simultaneous and alternating updates, and (ii) the choice of step-size and momentum value. The idea behind our approach is that a momentum term combined with the alternating gradient method can be used to manipulate the natural oscillatory behavior of adversarial games. We summarize our main contributions as follows:

- We prove in §5 that the alternating gradient method with negative momentum is the only setting within our study parameters (Fig. 2) that converges on a bilinear smooth game. Using a zero or positive momentum value, or doing simultaneous updates in such games fails to converge.
- We show in §4 that, for general dynamics, when the eigenvalues of the Jacobian have a large imaginary part, negative momentum can improve the local convergence properties of the gradient method.
- We confirm the benefits of negative momentum for training GANs with the notoriously ill-behaved saturating loss on both toy settings, and real datasets.

Outline. §2 describes the fundamentals of the analytic setup that we use. §3 provides a formulation for the optimal step-size, and discusses the constraints and intuition behind it. §4 presents our theoretical results

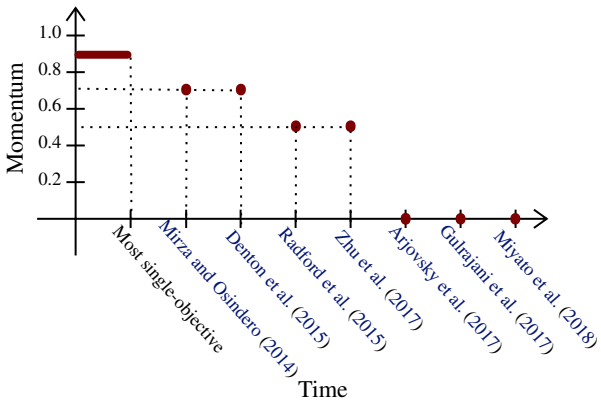

Figure 1: Decreasing trend in the value of momentum used for training GANs across time.

and guarantees on negative momentum. §5 studies the properties of alternating and simultaneous methods with negative momentum on a bilinear smooth game. §6 contains experimental results on toy and real datasets. Finally, in §7, we review some of the existing work on smooth game optimization as well as GAN stability and convergence.

2 BACKGROUND

Notation In this paper, scalars are lower-case letters (e.g., λ), vectors are lower-case bold letters (e.g., $\boldsymbol{\theta}$), matrices are upper-case bold letters (e.g., $\boldsymbol{A}$) and operators are upper-case letters (e.g., F). The spectrum of a squared matrix $\boldsymbol{A}$ is denoted by $\mathrm{Sp}(\boldsymbol{A})$, and its spectral radius is defined as $\rho(A) := \max\{|\lambda| \text{ for } \lambda \in \mathrm{Sp}(\boldsymbol{A})\}$. We respectively note $\sigma_{\min}(\boldsymbol{A})$ and $\sigma_{\max}(\boldsymbol{A})$ the smallest and the largest positive singular values of $\boldsymbol{A}$. The identity matrix of $\mathbb{R}^{m\times m}$ is written $\boldsymbol{I}_m$. We use $\Re$ and $\Im$ to respectively denote the real and imaginary part of a complex number. $\mathcal{O}$, Ω and Θ stand for the standard asymptotic notations. Finally, all the omitted proofs can be found in §D.

Game theory formulation of GANs Generative adversarial networks consist of a discriminator $D_{\boldsymbol{\varphi}}$ and a generator $G_{\boldsymbol{\theta}}$. In this game, the discriminator's objective is to tell real from generated examples. The generator's goal is to produce examples that are sufficiently close to real examples to confuse the discriminator.

From a game theory point of view, GAN training is a differentiable two-player game: the discriminator $D_{\boldsymbol{\varphi}}$ aims at minimizing its cost function $\mathcal{L}_D$ and the generator $G_{\boldsymbol{\theta}}$ aims at minimizing its own cost function $\mathcal{L}_G$. Using the same formulation as the one in Mescheder et al. (2017) and Gidel et al. (2018), the GAN objective has the following form,

$$\begin{cases} \boldsymbol{\theta}^* \in \arg\min_{\boldsymbol{\theta}\in\Theta} \mathcal{L}_G(\boldsymbol{\theta}, \boldsymbol{\varphi}^*) \\ \boldsymbol{\varphi}^* \in \arg\min_{\boldsymbol{\varphi}\in\Phi} \mathcal{L}_D(\boldsymbol{\theta}^*, \boldsymbol{\varphi}) \,. \end{cases} \tag{1}$$

Given such a game setup, GAN training consists of finding a local Nash Equilibrium, which is a state $(\boldsymbol{\varphi}^*, \boldsymbol{\theta}^*)$ in which neither the discriminator nor the generator can improve their respective cost by a small change in their parameters. In order to analyze the dynamics of gradient-based methods near a Nash Equilibrium, we look at the *gradient vector field*,

$$\boldsymbol{v}(\boldsymbol{\varphi}, \boldsymbol{\theta}) := \begin{bmatrix} \nabla_{\boldsymbol{\varphi}}\mathcal{L}_D(\boldsymbol{\varphi}, \boldsymbol{\theta}) & \nabla_{\boldsymbol{\theta}}\mathcal{L}_G(\boldsymbol{\varphi}, \boldsymbol{\theta}) \end{bmatrix}^\top , \tag{2}$$

and its associated *Jacobian* $\nabla\boldsymbol{v}(\boldsymbol{\varphi}, \boldsymbol{\theta})$,

$$\begin{bmatrix} \nabla^2_{\boldsymbol{\varphi}}\mathcal{L}_D(\boldsymbol{\varphi}, \boldsymbol{\theta}) & \nabla_{\boldsymbol{\varphi}}\nabla_{\boldsymbol{\theta}}\mathcal{L}_D(\boldsymbol{\varphi}, \boldsymbol{\theta}) \\ \nabla_{\boldsymbol{\varphi}}\nabla_{\boldsymbol{\theta}}\mathcal{L}_G(\boldsymbol{\varphi}, \boldsymbol{\theta})^T & \nabla^2_{\boldsymbol{\theta}}\mathcal{L}_G(\boldsymbol{\varphi}, \boldsymbol{\theta}) \end{bmatrix} . \tag{3}$$

Games in which $\mathcal{L}_G = -\mathcal{L}_D$ are called *zero-sum games* and (1) can be reformulated as a min-max problem. This is the case for the original *min-max* GAN formulation, but not the case for the *non-saturating loss* (Goodfellow et al., 2014) which is commonly used in practice.

For a zero-sum game, we note $\mathcal{L}_G = -\mathcal{L}_D = \mathcal{L}$. When the matrices $\nabla^2_{\boldsymbol{\varphi}}\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\theta})$ and $\nabla^2_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\theta})$ are zero, the Jacobian is anti-symmetric and has pure imaginary eigenvalues. We call games with pure imaginary eigenvalues *purely adversarial games*. This is the case in a simple bilinear game $\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\theta}) := \boldsymbol{\varphi}^\top \boldsymbol{A}\boldsymbol{\theta}$. This game can be formulated as a GAN where the true distribution is a Dirac on 0, the generator is a Dirac on θ and the discriminator is linear. This setup was extensively studied in 2D by Gidel et al. (2018).

Conversely, when $\nabla_{\boldsymbol{\varphi}}\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\theta})$ is zero and the matrices $\nabla^2_{\boldsymbol{\varphi}}\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\theta})$ and $-\nabla^2_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\theta})$ are symmetric and definite positive, the Jacobian is symmetric and has real positive eigenvalues. We call games with real positive eigenvalues *purely cooperative games*. This is the case, for example, when the objective function $\mathcal{L}$ is separable such as $\mathcal{L}(\boldsymbol{\varphi}, \boldsymbol{\theta}) = f(\boldsymbol{\varphi}) - g(\boldsymbol{\theta})$ where f and g are two convex functions. Thus, the optimization can be reformulated as two separated minimization of f and g with respect to their respective parameters.

These notions of *adversarial* and *cooperative* games can be related to the notions of *potential* games (Monderer and Shapley, 1996) and *Hamiltonian* games recently introduced by Balduzzi et al. (2018): a game is a *potential game* (resp. *Hamiltonian game*) if its Jacobian is symmetric (resp. asymmetric). Our definition of *cooperative game* is a bit more general than the definition of *potential game* since some non-symmetric matrices may have positive eigenvalues. Similarly, the notion of

Method	β	Bounded	Converges	Bound on Δ_t
Simult. Thm. 5	>0	✗	✗	$\Omega\left(((1+\beta)^2+\eta^2\sigma^2_{\max}(A))^t\right)$
	0	✗	✗	$\Omega\left((1+\eta^2\sigma^2_{\max}(A))^t\right)$
	<0	✗	✗	$\Omega\left((1+\lvert\beta\rvert)^{2t}\right)$
Altern. Thm. 6	>0	✗	✗	**Conjecture:** $\Omega\left((1+\beta^2)^t\right)$
	0	✓	✗	$\Theta\left(\Delta_0\right)$
	<0	✓	✓	$\mathcal{O}\left(\Delta_0(1-\eta\sigma_{\min}(A)/4)^t\right)$

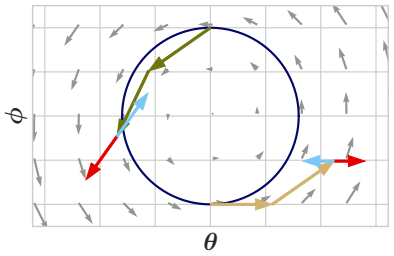

Figure 2: **Left:** Effect of gradient methods on an unconstrained bilinear example: $\min_{\boldsymbol{\theta}}\max_{\boldsymbol{\varphi}}\ \boldsymbol{\theta}^\top \boldsymbol{A}\boldsymbol{\varphi}$. The quantity Δ_t is the distance to the optimum (see formal definition in §5) and β is the momentum value. **Right:** Graphical intuition of the role of momentum in two steps of simultaneous updates (**tan**) or alternated updates (**olive**). Positive momentum (**red**) drives the iterates outwards whereas negative momentum (**blue**) pulls the iterates back towards the center, but it is only strong enough for alternated updates.

adversarial game generalizes the *Hamiltonian games* since some non-antisymmetric matrices may have pure imaginary eigenvalues, for instance,

$$\mathrm{Sp}\left(\begin{bmatrix}0 & -1\\ 2 & 3\end{bmatrix}\right)=\{1,2\}\,,\quad \mathrm{Sp}\left(\begin{bmatrix}-1 & 1\\ -2 & 1\end{bmatrix}\right)=\{\pm i\}\,.$$

In this work, we are interested in games *in between* purely adversarial games and purely cooperative ones, i.e., games which have eigenvalues with non-negative real part (cooperative component) and non-zero imaginary part (adversarial component). For $\boldsymbol{A}\in\mathbb{R}^{d\times p}$, a simple class of such games is parametrized by $\alpha\in[0,1]$,

$$\min_{\boldsymbol{\theta}\in\mathbb{R}^d}\max_{\boldsymbol{\varphi}\in\mathbb{R}^p}\ \alpha\|\boldsymbol{\theta}\|_2^2+(1-\alpha)\boldsymbol{\theta}^\top\boldsymbol{A}\boldsymbol{\varphi}-\alpha\|\boldsymbol{\varphi}\|_2^2\,, \tag{4}$$

Simultaneous Gradient Method. Let us consider the dynamics of the simultaneous gradient method. It is defined as the repeated application of the operator,

$$F_\eta(\boldsymbol{\varphi},\boldsymbol{\theta}):=\begin{bmatrix}\boldsymbol{\varphi} & \boldsymbol{\theta}\end{bmatrix}^\top-\eta\,\boldsymbol{v}(\boldsymbol{\varphi},\boldsymbol{\theta})\,,\quad(\boldsymbol{\varphi},\boldsymbol{\theta})\in\mathbb{R}^m\,, \tag{5}$$

where η is the learning rate. Now, for brevity we write the joint parameters $\boldsymbol{\omega}:=(\boldsymbol{\varphi},\boldsymbol{\theta})\in\mathbb{R}^m$. For $t\in\mathbb{N}$, let $\boldsymbol{\omega}_t=(\boldsymbol{\varphi}_t,\boldsymbol{\theta}_t)$ be the t^{th} point of the sequence computed by the gradient method,

$$\boldsymbol{\omega}_t=\underbrace{F_\eta\circ\ldots\circ F_\eta}_{t}(\boldsymbol{\omega}_0)=F_\eta^{(t)}(\boldsymbol{\omega}_0)\,. \tag{6}$$

Then, if the gradient method converges, and its limit point $\boldsymbol{\omega}^*=(\boldsymbol{\varphi}^*,\boldsymbol{\theta}^*)$ is a *fixed point* of F_η such that $\nabla v(\boldsymbol{\omega}^*)$ is positive-definite, then $\boldsymbol{\omega}^*$ is a local Nash equilibrium. Interestingly, some of the stable stationary points of gradient dynamics may not be Nash equilibrium (Adolphs et al., 2018). In this work, we focus on the local convergence properties near the stationary points of gradient . To the best of our knowledge, there is no first order method alleviating this issue. In the following, $\boldsymbol{\omega}^*$ is a stationary point of the gradient dynamics (i.e. a point such that $\boldsymbol{v}(\boldsymbol{\omega}^*)=0$).

3 TUNING THE STEP-SIZE

Under certain conditions on a fixed point operator, linear convergence is guaranteed in a neighborhood around a fixed point.

Theorem 1 (Prop. 4.4.1 Bertsekas (1999))**.** *If the spectral radius $\rho_{\max}:=\rho(\nabla F_\eta(\boldsymbol{\omega}^*))<1$, then, for $\boldsymbol{\omega}_0$ in a neighborhood of $\boldsymbol{\omega}^*$, the distance of $\boldsymbol{\omega}_t$ to the stationary point $\boldsymbol{\omega}^*$ converges at a linear rate of $\mathcal{O}\big((\rho_{\max}+\epsilon)^t\big)$, $\forall\epsilon>0$.*

From the definition in (5), we have:

$$\nabla F_\eta(\boldsymbol{\omega}^*)=\boldsymbol{I}_m-\eta\nabla\boldsymbol{v}(\boldsymbol{\omega}^*)\,, \tag{7}$$
$$\text{and}\quad \mathrm{Sp}(\nabla F_\eta(\boldsymbol{\omega}^*))=\left\{1-\eta\lambda\middle|\lambda\in\mathrm{Sp}(\nabla\boldsymbol{v}(\boldsymbol{\omega}^*))\right\}\,.$$

If the eigenvalues of $\nabla\boldsymbol{v}(\boldsymbol{\omega}^*)$ all have a positive real-part, then for small enough η, the eigenvalues of $\nabla F_\eta(\boldsymbol{\omega}^*)$ are inside a convergence circle of radius $\rho_{\max}<1$, as illustrated in Fig. 3. Thm. 1 guarantees the existence of an optimal step-size η_{best} which yields a non-trivial convergence rate $\rho_{\max}<1$. Thm. 2 gives analytic bounds on the optimal step-size η_{best}, and lower-bounds the best convergence rate $\rho_{\max}(\eta_{best})$ we can expect.

Theorem 2. *If the eigenvalues of $\nabla\boldsymbol{v}(\boldsymbol{\omega}^*)$ all have a positive real-part, then, the best step-size η_{best}, which minimizes the spectral radius $\rho_{\max}(\eta)$ of $\nabla F_\eta(\boldsymbol{\varphi}^*,\boldsymbol{\theta}^*)$, is the solution of a (convex) quadratic by parts problem, and satisfies,*

$$\max_{1\le k\le m}\sin(\psi_k)^2\le\rho_{\max}(\eta_{best})^2\le 1-\Re(1/\lambda_1)\delta\,, \tag{8}$$
$$\textit{with}\quad \delta:=\min_{1\le k\le m}|\lambda_k|^2(2\Re(1/\lambda_k)-\Re(1/\lambda_1)) \tag{9}$$
$$\textit{and}\quad \Re(1/\lambda_1)\le\eta_{best}\le 2\Re(1/\lambda_1) \tag{10}$$

where $(\lambda_k=r_ke^{i\psi_k})_{1\le k\le m}=\mathrm{Sp}(\nabla\boldsymbol{v}(\boldsymbol{\varphi}^*,\boldsymbol{\theta}^*))$ are sorted such that $0<\Re(1/\lambda_1)\le\cdots\le\Re(1/\lambda_m)$. Particularly, when $\eta_{best}=\Re(1/\lambda_1)$ we are in the case of the top plot of Fig.3 and $\rho_{\max}(\eta_{best})^2=\sin(\psi_1)^2$.

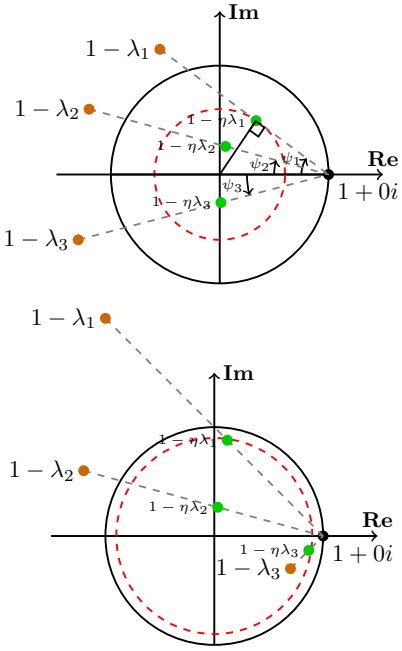

Figure 3: Eigenvalues λ_i of the Jacobian $\nabla \boldsymbol{v}(\boldsymbol{\phi}^*, \boldsymbol{\theta}^*)$ and their trajectories $1-\eta\lambda_i$ for growing step-sizes. The unit circle is drawn in **black**, and the **red** dashed circle has radius equal to the largest eigenvalue $\mu_{\max}$, which is directly related to the convergence rate. Therefore, smaller red circles mean better convergence rates. **Top:** The red circle is limited by the tangent trajectory line $1-\eta\lambda_1$, which means the best convergence rate is limited only by the eigenvalue which will pass furthest from the origin as η grows, i.e., $\lambda_i = \arg\min \Re(1/\lambda_i)$. **Bottom:** The red circle is cut (not tangent) by the trajectories at points $1-\eta\lambda_1$ and $1-\eta\lambda_3$. The η is optimal because any increase in η will push the eigenvalue λ_1 out of the red circle, while any decrease in step-size will retract the eigenvalue λ_3 out of the red circle, which will lower the convergence rate in any case. *Figure inspired by Mescheder et al. (2017).*

When $\nabla \boldsymbol{v}$ is positive-definite, the best η_{best} is attained either because of one or several limiting eigenvalues. We illustrate and interpret these two cases in Fig. 3. In multivariate convex optimization, the optimal step-size depends on the extreme eigenvalues and their ratio, the *condition number*. Unfortunately, the notion of the condition number does not trivially extend to games, but Thm. 2 seems to indicate that the real part of the inverse of the eigenvalues play an important role in the dynamics of smooth games. We think that a notion of condition number might be meaningful for such games and we propose an illustrative example to discuss this point in §B. Note that when the eigenvalues are pure positive real numbers belonging to $[\mu, L]$, (8) provides the standard bound $\rho_{\max} \leq 1 - \mu/L$ obtained with a step-size $\eta = 1/L$ (see §D.2 for details).

Note that, in (9), we have $\delta > 0$ because (λ_k) are sorted such that, $\Re(1/\lambda_k) \geq \Re(1/\lambda_1)\,, \forall 1 \leq k \leq m$. In (8), we can see that if the Jacobian of $\boldsymbol{v}$ has an almost purely imaginary eigenvalue $r_j e^{\psi_j}$ then $\sin(\psi_j)$ is close to 1 and thus, the convergence rate of the gradient method may be arbitrarily close to 1. Zhang and Mitliagkas (2017) provide an analysis of the momentum method for quadratics, showing that momentum can actually help to better condition the model. One interesting point from their work is that the best conditioning is achieved when the added momentum makes the Jacobian eigenvalues turn from positive reals into complex conjugate pairs. Our goal is to use momentum to wrangle game dynamics into convergence manipulating the eigenvalues of the Jacobian.

4 NEGATIVE MOMENTUM

As shown in (8), the presence of eigenvalues with large imaginary parts can restrict us to small step-sizes and lead to slow convergence rates. In order to improve convergence, we add a *negative* momentum term into the update rule. Informally, one can think of negative momentum as friction that can damp oscillations. The new momentum term leads to a modification of the *parameter update operator* $F_\eta(\boldsymbol{\omega})$ of (5). We use a similar state augmentation as Zhang and Mitliagkas (2017) to form a compound state $(\boldsymbol{\omega}_t, \boldsymbol{\omega}_{t-1}) := (\boldsymbol{\varphi}_t, \boldsymbol{\theta}_t, \boldsymbol{\varphi}_{t-1}, \boldsymbol{\theta}_{t-1}) \in \mathbb{R}^{2m}$. The update rule (5) turns into the following,

$$F_{\eta,\beta}(\boldsymbol{\omega}_t, \boldsymbol{\omega}_{t-1}) = (\boldsymbol{\omega}_{t+1}, \boldsymbol{\omega}_t) \tag{11}$$

$$\text{where} \quad \boldsymbol{\omega}_{t+1} := \boldsymbol{\omega}_t - \eta \boldsymbol{v}(\boldsymbol{\omega}_t) + \beta(\boldsymbol{\omega}_t - \boldsymbol{\omega}_{t-1})\,, \tag{12}$$

in which $\beta \in \mathbb{R}$ is the momentum parameter. Therefore, the Jacobian of $F_{\eta,\beta}$ has the following form,

$$\begin{bmatrix} \boldsymbol{I}_n & \boldsymbol{0}_n \\ \boldsymbol{I}_n & \boldsymbol{0}_n \end{bmatrix} - \eta \begin{bmatrix} \nabla \boldsymbol{v}(\boldsymbol{\omega}_t) & \boldsymbol{0}_n \\ \boldsymbol{0}_n & \boldsymbol{0}_n \end{bmatrix} + \beta \begin{bmatrix} \boldsymbol{I}_n & -\boldsymbol{I}_n \\ \boldsymbol{0}_n & \boldsymbol{0}_n \end{bmatrix} \tag{13}$$

Note that for $\beta = 0$, we recover the gradient method.

In some situations, if $\beta < 0$ is adjusted properly, negative momentum can improve the convergence rate to a local stationary point by pushing the eigenvalues of its Jacobian towards the origin. In the following theorem, we provide an explicit equation for the eigenvalues of the Jacobian of $F_{\eta,\beta}$.

Theorem 3. *The eigenvalues of* $\nabla F_{\eta,\beta}(\boldsymbol{\omega}^*)$ *are*

$$\mu_\pm(\beta, \eta, \lambda) := (1 - \eta\lambda + \beta)\frac{1 \pm \Delta^{\frac{1}{2}}}{2}, \tag{14}$$

where $\Delta := 1 - \frac{4\beta}{(1-\eta\lambda+\beta)^2}$, $\lambda \in \mathrm{Sp}(\nabla \boldsymbol{v}(\boldsymbol{\omega}^*))$ *and* $\Delta^{\frac{1}{2}}$ *is the* complex *square root of* Δ *with positive real part*[*]. *Moreover we have the following Taylor approximation,*

$$\mu_+(\beta, \eta, \lambda) = 1 - \eta\lambda - \beta\frac{\eta\lambda}{1-\eta\lambda} + O(\beta^2)\,, \tag{15}$$

$$\mu_-(\beta, \eta, \lambda) = \frac{\beta}{1-\eta\lambda} + O(\beta^2)\,. \tag{16}$$

[*] If Δ is a negative real number we set $\Delta^{\frac{1}{2}} := i\sqrt{-\Delta}$

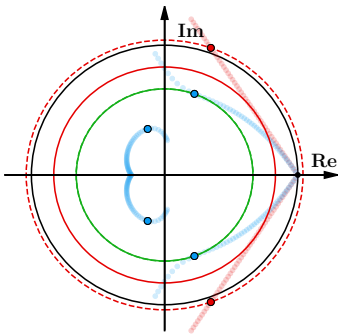

Figure 4: Transformation of the eigenvalues by a negative momentum method for a game introduced in (4) with $d = p = 1, A = 1, \alpha = .4, \eta = 1.55, \beta = -.25$. Convergence circles for gradient method are in **red**, negative momentum in **green**, and unit circle in **black**. Solid convergence circles are optimized over all step-sizes, while dashed circles are at a given step-size η. For a fixed η, original eigenvalues are in **red** and negative momentum eigenvalues are in **blue**. Their trajectories as η sweeps in $[0, 2]$ are in light colors. Negative momentum helps as the new convergence circle (green) is smaller, due to shifting the original eigenvalues (red dots) towards the origin (right blue dots), while the eigenvalues due to state augmentation (left blue dots) have smaller magnitude and do not influence the convergence rate. Negative momentum allows faster convergence (green circle inside the solid red circle) for a broad range of step-sizes.

When β is small enough, Δ is a complex number close to 1. Consequently, μ_+ is close to the original eigenvalue for gradient dynamics $1 - \eta\lambda$, and μ_-, the eigenvalue introduced by the state augmentation, is close to 0. We formalize this intuition by providing the first order approximation of both eigenvalues.

In Fig. 4, we illustrate the effects of negative momentum on a game described in (4). Negative momentum shifts the original eigenvalues (trajectories in light red) by pushing them to the left towards the origin (trajectories in light blue).

Since our goal is to minimize the largest magnitude of the eigenvalues of $F_{\eta,\beta}$ which are computed in Thm. 3, we want to understand the effect of β on these eigenvalues with potential large magnitude. Let $\lambda \in \mathrm{Sp}(\nabla \boldsymbol{v}(\boldsymbol{\omega}^*))$, we define the (squared) magnitude $\rho_{\lambda,\eta}(\beta)$ that we want to optimize,

$$\rho_{\lambda,\eta}(\beta) := \max\left\{|\mu_+(\beta,\eta,\lambda)|^2, |\mu_-(\beta,\eta,\lambda)|^2\right\}. \quad (17)$$

We study the local behavior of $\rho_{\lambda,\eta}$ for small β. The following theorem shows that a well suited β decreases $\rho_{\lambda,\eta}$, which corresponds to faster convergence.

Theorem 4. *For any $\lambda \in \mathrm{Sp}(\nabla \boldsymbol{v}(\boldsymbol{\omega}^*))$ s.t. $\Re(\lambda) > 0$,*

$$\rho'_{\lambda,\eta}(0) > 0 \Leftrightarrow \eta \in I(\lambda) := \left(\tfrac{|\lambda|-|\Im(\lambda)|}{|\lambda|\Re(\lambda)}, \tfrac{|\lambda|+|\Im(\lambda)|}{|\lambda|\Re(\lambda)}\right).$$

Particularly, we have $\rho'_{\lambda,\Re(1/\lambda)}(0) = 2\Re(\lambda)\Re(1/\lambda) > 0$ and $|Arg(\lambda)| \geq \frac{\pi}{4} \Rightarrow \left(\Re(1/\lambda), 2\Re(1/\lambda)\right) \subset I(\lambda)$.

As we have seen previously in Fig. 3 and Thm. 2, there are only few eigenvalues which slow down the convergence. Thm. 4 is a local result showing that a small negative momentum can improve the magnitude of the limiting eigenvalues in the following cases: when there is only one limiting eigenvalue λ_1 (since in that case the optimal step-size is $\eta_{best} = \Re(1/\lambda_1) \in I(\lambda_1)$) or when there are several limiting eigenvalues $\lambda_1, \ldots, \lambda_k$ and the intersection $I(\lambda_1) \cap \ldots \cap I(\lambda_k)$ is not empty. We point out that we do not provide any guarantees on whether this intersection is empty or not but note that if the absolute value of the argument of λ_1 is larger than $\pi/4$ then by (10), our theorem provides that the optimal step-size η_{best} belongs to $I(\lambda_1)$.

Since our result is local, it does not provide any guarantees on large negative values of β. Nevertheless, we numerically optimized (17) with respect to β and η and found that for any non-imaginary fixed eigenvalue λ, the optimal momentum is negative and the associated optimal step-size is larger than $\hat{\eta}(\lambda)$. Another interesting aspect of negative momentum is that it admits larger step-sizes (see Fig. 4 and 5).

For a game with purely imaginary eigenvalues, when $|\eta\lambda| \ll 1$, Thm. 3 shows that $\mu_+(\beta,\eta,\lambda) \approx 1-(1+\beta)\eta\lambda$. Therefore, at the first order, β only has an impact on the imaginary part of μ_+. Consequently μ_+ cannot be pushed into the unit circle, and the convergence guarantees of Thm. 1 do not apply. In other words, the analysis above provides convergence rates for games without any pure imaginary eigenvalues. It excludes the purely adversarial bilinear example ($\alpha = 0$ in Eq. 4) that is discussed in the next section.

5 BILINEAR SMOOTH GAMES

In this section we analyze the dynamics of a purely adversarial game described by,

$$\min_{\boldsymbol{\theta}\in\mathbb{R}^d} \max_{\boldsymbol{\varphi}\in\mathbb{R}^p} \boldsymbol{\theta}^\top \boldsymbol{A}\boldsymbol{\varphi} + \boldsymbol{\theta}^\top \boldsymbol{b} + \boldsymbol{c}^\top \boldsymbol{\varphi}, \quad \boldsymbol{A} \in \mathbb{R}^{d\times p}. \quad (18)$$

The first order stationary condition for this game characterizes the solutions $(\boldsymbol{\theta}^*, \boldsymbol{\varphi}^*)$ as

$$\boldsymbol{A}\boldsymbol{\varphi}^* = \boldsymbol{b} \quad \text{and} \quad \boldsymbol{A}^\top\boldsymbol{\theta}^* = \boldsymbol{c}. \quad (19)$$

If $\boldsymbol{b}$ (resp. $\boldsymbol{c}$) does not belong to the column space of $\boldsymbol{A}$ (resp. $\boldsymbol{A}^\top$), the game (18) admits no equilibrium. In the following, we assume that an equilibrium does exist for this game. Consequently, there exist $\boldsymbol{b}'$ and $\boldsymbol{c}'$ such that $\boldsymbol{b} = \boldsymbol{A}\boldsymbol{b}'$ and $\boldsymbol{c} = \boldsymbol{A}^\top\boldsymbol{c}'$. Using the translations $\boldsymbol{\theta} \to \boldsymbol{\theta} - \boldsymbol{c}'$ and $\boldsymbol{\varphi} \to \boldsymbol{\varphi} - \boldsymbol{b}'$, we can assume without loss of generality, that $p \geq d$, $\boldsymbol{b} = \boldsymbol{0}$ and $\boldsymbol{c} = \boldsymbol{0}$. We provide upper and lower bounds on the squared distance from the known equilibrium,

$$\Delta_t = \|\boldsymbol{\theta}_t - \boldsymbol{\theta}^*\|_2^2 + \|\boldsymbol{\varphi}_t - \boldsymbol{\varphi}^*\|_2^2 \quad (20)$$

where $(\boldsymbol{\theta}^*, \boldsymbol{\varphi}^*)$ is the projection of $(\boldsymbol{\theta}_t, \boldsymbol{\varphi}_t)$ onto the solution space. We show in §C, Lem. 2 that, for our methods of interest, this projection has a simple formulation that only depends on the initialization $(\boldsymbol{\theta}_0, \boldsymbol{\varphi}_0)$.

We aim to understand the difference between the dynamics of simultaneous steps and alternating steps. Practitioners have been widely using the latter instead of the former when optimizing GANs despite the rich optimization literature on simultaneous methods.

5.1 Simultaneous gradient descent

We define this class of methods with momentum using the following formulas,

$$F^{\text{sim}}_{\eta,\beta}(\boldsymbol{\theta}_t, \boldsymbol{\varphi}_t, \boldsymbol{\theta}_{t-1}, \boldsymbol{\varphi}_{t-1}) := (\boldsymbol{\theta}_{t+1}, \boldsymbol{\varphi}_{t+1}, \boldsymbol{\theta}_t, \boldsymbol{\varphi}_t) \quad (21)$$

$$\text{where } \begin{cases} \boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta_1 \boldsymbol{A}\boldsymbol{\varphi}_t + \beta_1(\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}) \\ \boldsymbol{\varphi}_{t+1} = \boldsymbol{\varphi}_t + \eta_2 \boldsymbol{A}^\top \boldsymbol{\theta}_t + \beta_2(\boldsymbol{\varphi}_t - \boldsymbol{\varphi}_{t-1}) \,. \end{cases}$$

In our simple setting, the operator $F^{\text{sim}}_{\eta,\beta}$ is linear. One way to study the asymptotic properties of the sequence $(\boldsymbol{\theta}_t, \boldsymbol{\varphi}_t)$ is to compute the eigenvalues of $\nabla F^{\text{sim}}_{\eta,\beta}$. The following proposition characterizes these eigenvalues.

Proposition 1. *The eigenvalues of $\nabla F^{sim}_{\eta,\beta}$ are the roots of the 4^{th} order polynomials:*

$$(x-1)^2(x-\beta_1)(x-\beta_2)+\eta_1\eta_2\lambda x^2,\ \lambda \in \text{Sp}(\boldsymbol{A}^\top \boldsymbol{A}) \quad (22)$$

Interestingly, these roots only depend on the product $\eta_1\eta_2$ meaning that any re-scaling $\eta_1 \to \gamma\eta_1\,,\ \eta_2 \to \frac{1}{\gamma}\eta_2$ does not change the eigenvalues of $\nabla F^{\text{sim}}_{\eta,\beta}$ and consequently the asymptotic dynamics of the iterates $(\boldsymbol{\theta}_t, \boldsymbol{\varphi}_t)$. The magnitude of the eigenvalues described in (22), characterizes the asymptotic properties for the iterates of the simultaneous method (21). We report the maximum magnitude of these roots for a given λ and for a grid of step-sizes and momentum values in Fig 7. We observe that they are always larger than 1, which transcribes a diverging behavior. The following theorem provides an analytical rate of divergence.

Theorem 5. *For any $\eta_1, \eta_2 \geq 0$ and $\beta_1 = \beta_2 = \beta$, the iterates of the simultaneous methods (21) diverge as,*

$$\Delta_t \in \begin{cases} \Omega\Big(\Delta_0[(1+\beta)^2 + \eta^2\sigma^2_{\max}(A)]^t\Big) & \text{if} \quad \beta \geq 0 \\ \Omega\Big(\Delta_0\,(1-\beta)^t\Big) & \text{if} \quad -\dfrac{1}{16} \leq \beta < 0\,. \end{cases}$$

This theorem states that the iterates of the simultaneous method (21) diverge geometrically for $\beta \geq -\frac{1}{16}$. Interestingly, this geometric divergence implies that *even* a uniform averaging of the iterates (standard in game optimization to ensure convergence (Freund et al., 1999)) cannot alleviate this divergence.

5.2 Alternating gradient descent

Alternating gradient methods take advantage of the fact that the iterates $\boldsymbol{\theta}_{t+1}$ and $\boldsymbol{\varphi}_{t+1}$ are computed sequentially, to plug the value of $\boldsymbol{\theta}_{t+1}$ (instead of $\boldsymbol{\theta}_t$ for simultaneous update rule) into the update of $\boldsymbol{\varphi}_{t+1}$,

$$F^{\text{alt}}_{\eta,\beta}(\boldsymbol{\theta}_t, \boldsymbol{\varphi}_t, \boldsymbol{\theta}_{t-1}, \boldsymbol{\varphi}_{t-1}) := (\boldsymbol{\theta}_{t+1}, \boldsymbol{\varphi}_{t+1}, \boldsymbol{\theta}_t, \boldsymbol{\varphi}_t) \quad (23)$$

$$\text{where } \begin{cases} \boldsymbol{\theta}_{t+1} = \boldsymbol{\theta}_t - \eta_1 \boldsymbol{A}\boldsymbol{\varphi}_t + \beta_1(\boldsymbol{\theta}_t - \boldsymbol{\theta}_{t-1}) \\ \boldsymbol{\varphi}_{t+1} = \boldsymbol{\varphi}_t + \eta_2 \boldsymbol{A}^\top \boldsymbol{\theta}_{t+1} + \beta_2(\boldsymbol{\varphi}_t - \boldsymbol{\varphi}_{t-1}) \,. \end{cases}$$

This slight change between (21) and (23) significantly shifts the eigenvalues of the Jacobian. We first characterize them with the following proposition.

Proposition 2. *The eigenvalues of $\nabla F^{alt}_{\eta,\beta}$ are the roots of the 4^{th} order polynomials:*

$$(x-1)^2(x-\beta_1)(x-\beta_2)+\eta_1\eta_2\lambda x^3\,,\ \lambda \in \text{Sp}(\boldsymbol{A}^\top \boldsymbol{A}) \quad (24)$$

The same way as in (22), these roots only depend on the product $\eta_1\eta_2$. The only difference is that the monomial with coefficient $\eta_1\eta_2\lambda$ is of degree 2 in (22) and of degree 3 in (24). This difference is major since, for well chosen values of negative momentum, the eigenvalues described in Prop. 2 lie in the unit disk (see Fig. 7). As a consequence, the iterates of the alternating method with no momentum are bounded and do converge if we add some well chosen negative momentum:

Theorem 6. *If we set $\eta \leq \frac{1}{\sigma_{\max}(A)}$, $\beta_1 = -\frac{1}{2}$ and $\beta_2 = 0$ then we have*

$$\Delta_{t+1} \in O\left(\max\{\tfrac{1}{2}, 1 - \tfrac{\eta\sigma_{\min}(A)}{4}\}^t \Delta_0\right) \quad (25)$$

If we set $\beta_1 = 0$ and $\beta_2 = 0$, then there exists $M > 1$ such that for any $\eta_1, \eta_2 \geq 0$, $\Delta_t = \Omega(\Delta_0)$.

Our results from this section, namely Thm. 5 and Thm. 6, are summarized in Fig. 2, and demonstrate how alternating steps can improve the convergence properties of the gradient method for bilinear smooth games. Moreover, combining them with negative momentum can surprisingly lead to a linearly convergent method. The conjecture provided in Fig. 2 (divergence of the alternating method with positive momentum) is backed-up by the results provided in Fig. 5 and §A.1.

6 EXPERIMENTS AND DISCUSSION

Min-Max Bilinear Game [Fig. 5] In our first experiments, we showcase the effect of negative momentum in a bilinear min-max optimization setup (4) where $\boldsymbol{\phi}, \boldsymbol{\theta} \in \mathbb{R}$ and $\boldsymbol{A} = 1$. We compare the effect of positive and negative momentum in both cases of alternating and simultaneous gradient steps.

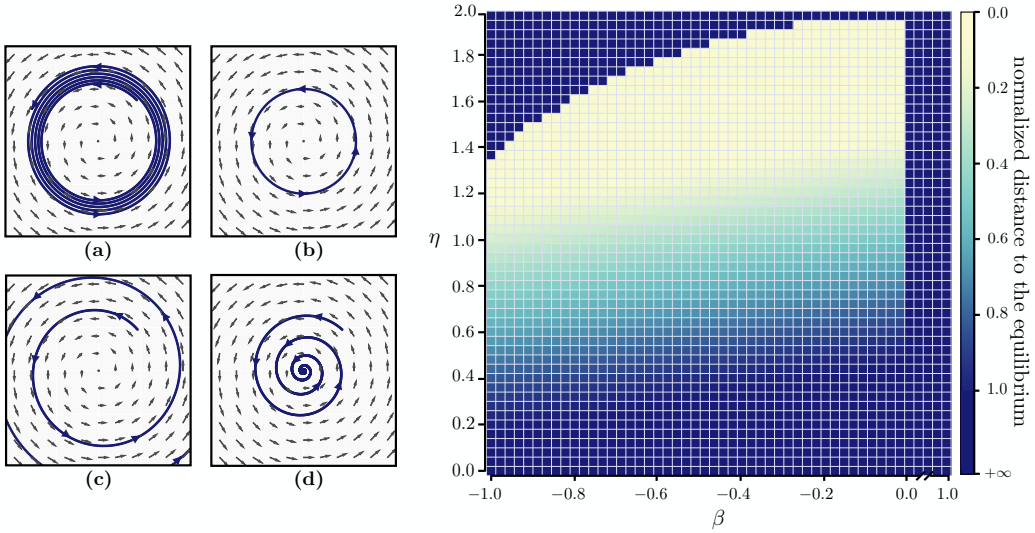

Figure 5: The effect of momentum in a simple min-max bilinear game where the equilibrium is at $(0, 0)$. **(left-a)** Simultaneous GD with no momentum **(left-b)** Alternating GD with no momentum. **(left-c)** Alternating GD with a momentum of $+0.1$. **(left-d)** Alternating GD with a momentum of -0.1. **(right)** A grid of experiments for alternating GD with different values of momentum (β) and step-sizes (η): While any positive momentum leads to divergence, small enough value of negative momentum allows for convergence with large step-sizes. The color in each cell indicates the normalized distance to the equilibrium after 500k iteration, such that 1.0 corresponds to the initial condition and values larger (smaller) than 1.0 correspond to divergence (convergence).

Fashion MNIST and CIFAR 10 [Fig. 6] In our third set of experiments, we use negative momentum in a GAN setup on CIFAR-10 (Krizhevsky and Hinton, 2009) and Fashion-MNIST (Xiao et al., 2017) with *saturating loss* and alternating steps. We use residual networks for both the generator and the discriminator with no batch-normalization. Following the same architecture as Gulrajani et al. (2017), each residual block is made of two 3×3 convolution layers with $ReLU$ activation function. Up-sampling and down-sampling layers are respectively used in the generator and discriminator. We experiment with different values of momentum on the discriminator and a constant value of 0.5 for the momentum of the generator. We observe that using a negative value can generally result in samples with higher quality and inception scores. Intuitively, using negative momentum only on the discriminator slows down the learning process of the discriminator and allows for better flow of the gradient to the generator. Note that we provide an additional experiment on mixture of Gaussians in § A.2.

7 RELATED WORK

Optimization From an optimization point of view, a lot of work has been done in the context of understanding momentum and its variants (Polyak, 1964; Qian, 1999; Nesterov, 2013; Sutskever et al., 2013). Some recent studies have emphasized the importance of momentum tuning in deep learning such as Sutskever et al. (2013), Kingma and Ba (2015), and Zhang and Mitliagkas (2017), however, none of them consider using negative momentum. Among recent work, using robust control theory, Lessard et al. (2016) study optimization procedures and cover a variety of algorithms including momentum methods. Their analysis is global and they establish worst-case bounds for smooth and strongly-convex functions. Mitliagkas et al. (2016) considered negative momentum in the context of asynchronous single-objective minimization. They show that asynchronous-parallel dynamics 'bleed' into optimization updates introducing momentum-like behavior into SGD. They argue that algorithmic momentum and asynchrony-induced momentum add up to create an effective 'total momentum' value. They conclude that to attain the optimal (positive) effective momentum in an asynchronous system, one would have to reduce algorithmic momentum to small or sometimes negative values. This differs from our work where we show that for games the optimal effective momentum may be negative. Ghadimi et al. (2015) analyze momentum and provide global convergence properties for functions with Lipschitz-continuous gradients. However, all the results mentioned above are restricted to minimization problems. The purpose of our work is to try to understand how momentum influences game dynamics which is intrinsically different from minimization dynamics.

GANs as games A lot of recent work has attempted to make GAN training easier with new optimization methods. Daskalakis et al. (2018) extrapolate the next value of the gradient using previous history and Gidel

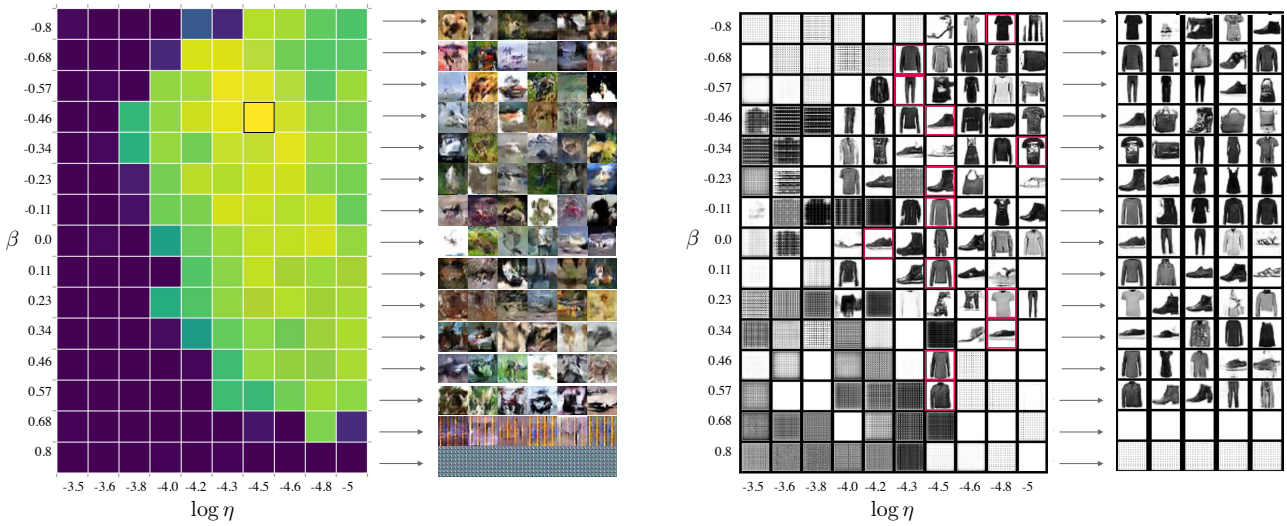

Figure 6: Comparison between negative and positive momentum on GANs with saturating loss on CIFAR-10 (left) and on Fashion MNIST (right) using a residual network. For each dataset, a grid of different values of momentum (β) and step-sizes (η) is provided which describes the discriminator's settings while a constant momentum of 0.5 and step-size of 10^{-4} is used for the generator. Each cell in CIFAR-10 (or Fashion MNIST) grid contains a single configuration in which its color (or its content) indicates the inception score (or a single sample) of the model. For CIFAR-10 experiments, yellow is higher while blue is the lower inception score. Along each row, the best configuration is chosen and more samples from that configuration are presented on the right side of each grid.

et al. (2018) explore averaging and introduce a variant of the extra-gradient algorithm.

Balduzzi et al. (2018) develop new methods to understand the dynamics of general games: they decompose second-order dynamics into two components using Helmholtz decomposition and use the fact that the optimization of Hamiltonian games is well understood. It differs from our work since we do not consider any decomposition of the Jacobian but focus on the manipulation of its eigenvalues. Recently, Liang and Stokes (2018) provide a unifying theory for smooth two-player games for non-asymptotic local convergence. They also provide theory for choosing the right step-size required for convergence.

From another perspective, Odena et al. (2018) show that in a GAN setup, the average conditioning of the Jacobian of the generator becomes ill-conditioned during training. They propose Jacobian clamping to improve the inception score and Frechet Inception Distance. Mescheder et al. (2017) provide discussion on how the eigenvalues of the Jacobian govern the local convergence properties of GANs. They argue that the presence of eigenvalues with zero real-part and large imaginary-part results in oscillatory behavior but do not provide results on the optimal step-size and on the impact of momentum. Nagarajan and Kolter (2017) also analyze the local stability of GANs as an approximated continuous dynamical system. They show that during training of a GAN, the eigenvalues of the Jacobian of the corresponding vector field are pushed away from one along the real axis.

8 CONCLUSION

In this paper, we show analytically and empirically that alternating updates with negative momentum is the only method within our study parameters (Fig.2) that converges in bilinear smooth games. We study the effects of using negative values of momentum in a GAN setup both theoretically and experimentally. We show that, for a large class of adversarial games, negative momentum may improve the convergence rate of gradient-based methods by shifting the eigenvalues of the Jacobian appropriately into a smaller convergence disk. We found that, in simple yet intuitive examples, using negative momentum makes convergence to the Nash Equilibrium easier. Our experiments support the use of negative momentum for saturating losses on mixtures of Gaussians, as well as on other tasks using CIFAR-10 and fashion MNIST. Altogether, fully stabilizing learning in GANs requires a deep understanding of the underlying highly non-linear dynamics. We believe our work is a step towards a better understanding of these dynamics. We encourage deep learning researchers and practitioners to include negative values of momentum in their hyper-parameter search.

We believe that our results explain a decreasing trend in momentum values used for training GANs in the past few years reported in Fig. 4. Some of the most successful papers use zero momentum (Arjovsky et al., 2017; Gulrajani et al., 2017) for architectures that would otherwise call for high momentum values in a non-adversarial setting.

Acknowledgments

This research was partially supported by the Canada CIFAR AI Chair Program, the FRQNT nouveaux chercheurs program, 2019-NC-257943, the Canada Excellence Research Chair in "Data Science for Real-time Decision-making", by the NSERC Discovery Grant RGPIN-2017-06936, a Google Focused Research Award and an IVADO grant. Authors would like to thank NVIDIA corporation for providing the NVIDIA DGX-1 used for this research. Authors are also grateful to Frédéric Bastien, Florian Bordes, Adam Beberg, Cam Moore and Nithya Natesan for their support.

References

L. Adolphs, H. Daneshmand, A. Lucchi, and T. Hofmann. Local saddle point optimization: A curvature exploitation approach. *arXiv preprint arXiv:1805.05751*, 2018.

M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.

D. Balduzzi, S. Racaniere, J. Martens, J. Foerster, K. Tuyls, and T. Graepel. The mechanics of n-player differentiable games. In *ICML*, 2018.

D. P. Bertsekas. *Nonlinear programming*. Athena scientific Belmont, 1999.

C. Daskalakis, A. Ilyas, V. Syrgkanis, and H. Zeng. Training GANs with optimism. In *ICLR*, 2018.

E. L. Denton, S. Chintala, R. Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pages 1486–1494, 2015.

Y. Freund, R. E. Schapire, et al. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 1999.

E. Ghadimi, H. R. Feyzmahdavian, and M. Johansson. Global convergence of the heavy-ball method for convex optimization. In *ECC*, 2015.

G. Gidel, H. Berard, P. Vincent, and S. Lacoste-Julien. A variational inequality perspective on generative adversarial nets. *arXiv preprint arXiv:1802.10551*, 2018.

I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014.

I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein GANs. In *NIPS*, 2017.

D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.

L. Lessard, B. Recht, and A. Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 2016.

T. Liang and J. Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. *arXiv preprint arXiv:1802.06132*, 2018.

L. Mescheder, S. Nowozin, and A. Geiger. The numerics of GANs. In *NIPS*, 2017.

M. Mirza and S. Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.

I. Mitliagkas, C. Zhang, S. Hadjis, and C. Ré. Asynchrony begets momentum, with an application to deep learning. In *2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2016.

T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.

D. Monderer and L. S. Shapley. Potential games. *Games and economic behavior*, 1996.

V. Nagarajan and J. Z. Kolter. Gradient descent GAN optimization is locally stable. In *NIPS*, 2017.

Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.

A. Odena, J. Buckman, C. Olsson, T. B. Brown, C. Olah, C. Raffel, and I. Goodfellow. Is generator conditioning causally related to gan performance? In *ICML*, 2018.

D. Pfau and O. Vinyals. Connecting generative adversarial networks and actor-critic methods. *arXiv preprint arXiv:1610.01945*, 2016.

B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 1964.

N. Qian. On the momentum term in gradient descent learning algorithms. *Neural networks*, 1999.

A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.

I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *ICML*, 2013.

H. Xiao, K. Rasul, and R. Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

F. Zhang. *The Schur complement and its applications*. Springer Science & Business Media, 2006.

J. Zhang and I. Mitliagkas. Yellowfin and the art of momentum tuning. *arXiv preprint arXiv:1706.03471*, 2017.

J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2223–2232, 2017.