FINDING AND ONLY FINDING LOCAL NASH EQUILIBRIA BY BOTH PRETENDING TO BE A FOLLOWER

Xuchan Bao, Guodong Zhang
University of Toronto, Vector Institute
{jennybao,gdzhang}@cs.toronto.edu

ABSTRACT

Finding local Nash equilibria in two-player differentiable games is a classical problem in game theory with important relevance in machine learning. We propose double Follow-the-Ridge (double-FTR), an algorithm that locally converges to and only to local Nash equilibria in general-sum two-player differentiable games. To our knowledge, double-FTR is the first algorithm with such guarantees for general-sum games. Furthermore, we show that by varying its preconditioner, double-FTR leads to a broader family of algorithms with the same properties. Double-FTR avoids oscillation near equilibria due to the real-eigenvalues of its Jacobian at critical points. Finally, we empirically verify the effectiveness of double-FTR in finding local Nash equilibria in two simple examples.

1 INTRODUCTION

Much of the recent success in deep learning can be attributed to the effectiveness of gradient-based optimization. It is well-known that for a minimization problem, with appropriate choice of learning rates, gradient descent has convergence guarantee to local minima (Lee et al., 2016; 2019). Based on this foundational result, an array of accelerated and higher-order methods have since been proposed and widely applied in training neural networks (Duchi et al., 2011; Kingma & Ba, 2014; Reddi et al., 2018; Zhang et al., 2019).

However, once we leave the realm of minimization problems and consider the multi-agent setting, the optimization landscape becomes much more complicated. Multi-agent optimization problems arise in diverse fields such as robotics, economics and machine learning (Foerster et al., 2016; Von Neumann & Morgenstern, 2007; Goodfellow et al., 2014; Ben-Tal & Nemirovski, 2002; Gemp et al., 2020; Anil et al., 2021).

A classical abstraction that is especially relevant for machine learning is two-player differentiable games, where the objective is to find local Nash equilibria. The equivalent of gradient descent in such a game would be that each agent applies gradient descent to minimize their own objective function. However, in stark contrast with minimization problems, this gradient-descent-style algorithm may converge to critical points that are not local Nash equilibria, and in the general-sum game case, local Nash equilibria might not even be stable critical points for this algorithm (Mazumdar et al., 2020)!

These negative results have driven a surge of recent interest in developing other gradient-based algorithms for finding Nash equilibria in differentiable games. Among them is Mazumdar et al. (2019), who proposed an update algorithm whose attracting critical points are only local Nash equilibria in the special case of zero-sum games. However, to the best of our knowledge, such guarantees have not been extended to general-sum games.

We propose double Follow-the-Ridge (double-FTR), a gradient-based algorithm for general-sum differentiable games whose attracting critical points are equivalent to differential Nash equilibria. Double-FTR is closely related to the Follow-the-Ridge (FTR) algorithm for Stackelberg games (Wang et al., 2019), which converge to and only to local Stackelberg equilibria (Fiez et al., 2019). Double-FTR can be viewed as its counterpart for simultaneous games, where each player adopts the "follower" strategy in FTR.

The rest of this paper is organized as follows. In Section 2, we give background on two-player differentiable games and equilibrium concepts. We also explain the issues with using gradient-descent-style algorithm on such games. In Section 3, we present the double-FTR algorithm and prove the equivalence of its attracting critical points and differential Nash equilibria. We also identify a more general class of algorithms that share the properties. We discuss recent work that is closely-connected to double-FTR in Section 4 and other related work in Section 5. In Section 6, we show empirical evidence of double-FTR's convergence to and only to local Nash equilibria.

2 BACKGROUND

2.1 TWO-PLAYER DIFFERENTIABLE GAMES AND EQUILIBRIUM CONCEPTS

In a general-sum two-player differentiable game, player 1 aims to minimize $f : \mathbb{R}^{n+m} \to \mathbb{R}$ with respect to $\boldsymbol{x} \in \mathbb{R}^n$, whereas player 2 aims to maximize $g : \mathbb{R}^{n+m} \to \mathbb{R}$ with respect to $\boldsymbol{y} \in \mathbb{R}^m$. Following the notation in Mazumdar et al. (2019), we denote such the game as $\{(f, -g), \mathbb{R}^{n+m}\}$. We also assume that f and g are twice-differentiable.

For two rational, non-cooperative players, their optimal outcome is to achieve a local Nash equilibrium. A point $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ is a local Nash equilibrium of $\{(f, -g), \mathbb{R}^{n+m}\}$ if there exists open sets $\mathcal{S}_{\boldsymbol{x}} \subset \mathbb{R}^n$, $\mathcal{S}_{\boldsymbol{y}} \subset \mathbb{R}^m$ such that $\boldsymbol{x}^* \in \mathcal{S}_{\boldsymbol{x}}$, $\boldsymbol{y}^* \in \mathcal{S}_{\boldsymbol{y}}$, and

$$\begin{aligned} f(\boldsymbol{x}^*, \boldsymbol{y}^*) &\leq f(\boldsymbol{x}, \boldsymbol{y}^*) \quad \forall \boldsymbol{x} \in \mathcal{S}_{\boldsymbol{x}} \setminus \{\boldsymbol{x}^*\}, \\ g(\boldsymbol{x}^*, \boldsymbol{y}^*) &\geq g(\boldsymbol{x}^*, \boldsymbol{y}\ ) \quad \forall \boldsymbol{y} \in \mathcal{S}_{\boldsymbol{y}} \setminus \{\boldsymbol{y}^*\}. \end{aligned}$$

As shown in Ratliff et al. (2013), when f and g are twice-differentiable, local Nash equilibria generally satisfy stronger conditions, and are differential Nash equilibria (DNE).

Definition 2.1 (Differential Nash equilibrium)**.** $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ is a differential Nash equilibrium of $\{(f, -g), \mathbb{R}^{n+m}\}$ if:

- $\nabla_{\boldsymbol{x}} f(\boldsymbol{x}^*, \boldsymbol{y}^*) = 0$ and $\nabla_{\boldsymbol{y}} g(\boldsymbol{x}^*, \boldsymbol{y}^*) = 0$.
- $\nabla^2_{\boldsymbol{xx}} f(\boldsymbol{x}^*, \boldsymbol{y}^*) \succ 0$ and $\nabla^2_{\boldsymbol{yy}} g(\boldsymbol{x}^*, \boldsymbol{y}^*) \prec 0$.

2.2 ISSUES WITH GRADIENT-BASED ALGORITHMS

A natural strategy for agents to search for local Nash equilibria in a differentiable game is to use gradient-based algorithms. The simplest gradient-based algorithm is the gradient descent-ascent (GDA) (Ryu & Boyd, 2016; Zhang et al., 2021b) (Algorithm 1) or its variants (Zhang et al., 2021a; Korpelevich, 1976; Mokhtari et al., 2020).

Algorithm 1 Gradient descent-ascent (GDA)

Require: Number of iterations T, learning rates $\gamma_1, \ldots, \gamma_T$
1: **for** $t = 1, \ldots, T$ **do**
2: $\quad \boldsymbol{x}_{t+1} = \boldsymbol{x}_t - \gamma_t \nabla_{\boldsymbol{x}} f(\boldsymbol{x}_t, \boldsymbol{y}_t)$
3: $\quad \boldsymbol{y}_{t+1} = \boldsymbol{y}_t + \gamma_t \nabla_{\boldsymbol{y}} g(\boldsymbol{x}_t, \boldsymbol{y}_t)$
4: **end for**

To help analyzing the dynamics, we study GDA in a continuous-time limit. Let $\boldsymbol{z} = (\boldsymbol{x}, \boldsymbol{y})$ and $\boldsymbol{\omega}_{\mathrm{GDA}}(\boldsymbol{z}) = \begin{bmatrix} \nabla_{\boldsymbol{x}} f(\boldsymbol{x}, \boldsymbol{y})^\top & -\nabla_{\boldsymbol{y}} g(\boldsymbol{x}, \boldsymbol{y})^\top \end{bmatrix}^\top$, the continuous-time flow of GDA is:

$$\dot{\boldsymbol{z}} = -\boldsymbol{\omega}_{\mathrm{GDA}}(\boldsymbol{z}) \tag{1}$$

where the Jacobian matrix ($\boldsymbol{J} := \frac{\partial \boldsymbol{\omega}(\boldsymbol{z})}{\partial \boldsymbol{z}}$) is:

$$\boldsymbol{J}_{\mathrm{GDA}} = \frac{\partial \boldsymbol{\omega}_{\mathrm{GDA}}(\boldsymbol{z})}{\partial \boldsymbol{z}} = \begin{bmatrix} \nabla^2_{\boldsymbol{xx}} f & \nabla^2_{\boldsymbol{xy}} f \\ -\nabla^2_{\boldsymbol{yx}} g & -\nabla^2_{\boldsymbol{yy}} g \end{bmatrix}$$

Using the Jacobian matrix, we then introduce some characterizations on the critical points of the update (as in Mazumdar et al. (2019)).

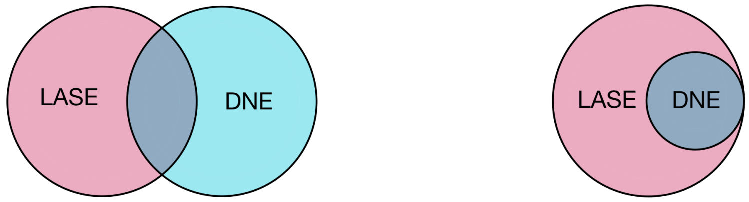

(a) General-sum (b) Zero-sum

Figure 1: Venn diagrams showing the relationship between the set of locally asymptotically stable equilibria (LASE) of the GDA flow and the set of differential Nash equilibria (DNE) in two-player differentiable games.

Definition 2.2 ((Hyperbolic) critical point)**.** $\boldsymbol{z}^* \in \mathbb{R}^{n+m}$ is a critical point of $\boldsymbol{\omega}$ if $\boldsymbol{\omega}(\boldsymbol{z}^*) = 0$. It is a hyperbolic critical point if $Re(\lambda) \neq 0$ for $\forall\lambda \in \mathrm{spec}(\boldsymbol{J}(\boldsymbol{z}^*))$.

Definition 2.3 (Locally asymptotically stable equilibrium (LASE))**.** $\boldsymbol{z}^*$ is a locally asymptotically stable equilibrium of the continuous-time dynamics $\dot{\boldsymbol{z}} = -\boldsymbol{\omega}(\boldsymbol{z})$ if

$$\boldsymbol{\omega}(\boldsymbol{z}^*) = 0 \quad \text{and} \quad \mathrm{Re}(\lambda) > 0 \text{ for } \forall\lambda \in \mathrm{spec}(\boldsymbol{J}(\boldsymbol{z}^*)).$$

LASE points are important for deriving local convergence results, as they are isolated and exponentially attracting under $\dot{\boldsymbol{z}} = -\boldsymbol{\omega}(\boldsymbol{z})$ (Sastry, 2013). This means that with proper discretization (i.e. with appropriate choice of learning rate η), when initialized within the neighbourhood of the LASE point, the discretized system $\boldsymbol{z}_{t+1} = \boldsymbol{z}_t - \eta\boldsymbol{\omega}(\boldsymbol{z}_t)$ has linear convergence.

We focus on games where all critical points $\boldsymbol{\omega}_{\mathrm{GDA}}$ are hyperbolic. This is a reasonable assumption, as hyperbolic critical points are generic for smooth dynamical systems (Sastry, 2013). Unfortunately, GDA is not guaranteed to converge to DNE, nor are DNE necessarily LASE of the GDA dynamics. Even in the special case of zero-sum games ($g = f$), GDA dynamics can still have LASE that are not DNE (Mazumdar et al., 2020). The relationship is shown in the Venn diagrams in Figure 1.

In Figure 2, we provide simple 2-D examples that demonstrate the failure modes of GDA in zero-sum games . In 2a, GDA converges to a spurious LASE which is not DNE. In 2b, GDA converges to a limit cycle instead of the unique DNE (Hsieh et al., 2020).

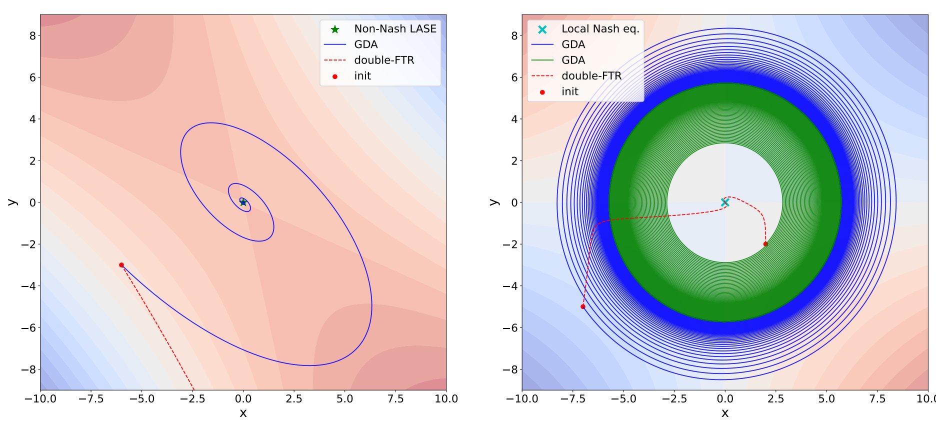

(a) $f(x, y) = g(x, y) = 2x^2 + 5xy + y^2$ (b) $f(x, y) = g(x, y) = xy + \epsilon(x^2 - y^4)$

Figure 2: Two examples of GDA failure modes in finding Nash equilibrium in zero-sum games. (a) GDA converges to the spurious LASE $(0, 0)$, which is not a Nash equilibrium. (b) GDA converges to a limit cycle (we use $\epsilon = 0.0001$, $\gamma_t = 0.01$ for $\forall t$) instead of the unique Nash equilibrium $(0, 0)$.

3 DOUBLE FOLLOW-THE-RIDGE

We propose double Follow-the-Ridge (double-FTR), an update rule for general-sum two-player continuous games whose attracting critical points are equivalent to differential Nash equilibria. The double-FTR update is shown in Algorithm 2 (the arguments $\boldsymbol{x}_t$, $\boldsymbol{y}_t$ of f and g are dropped to avoid notational clutter).

Algorithm 2 Double Follow-the-Ridge

Require: Learning rate $\eta_{\boldsymbol{x}}$ and $\eta_{\boldsymbol{y}}$; number of iterations T.
1: **for** $t = 1, \ldots, T$ **do**
2: $\quad \boldsymbol{x}_{t+1} \leftarrow \boldsymbol{x}_t - \eta_{\boldsymbol{x}} \nabla_{\boldsymbol{x}} f - \eta_{\boldsymbol{y}} (\nabla^2_{\boldsymbol{xx}} f)^{-1} \nabla^2_{\boldsymbol{xy}} g \nabla_{\boldsymbol{y}} g$
3: $\quad \boldsymbol{y}_{t+1} \leftarrow \boldsymbol{y}_t + \eta_{\boldsymbol{y}} \nabla_{\boldsymbol{y}} g + \eta_{\boldsymbol{x}} (\nabla^2_{\boldsymbol{yy}} g)^{-1} \nabla^2_{\boldsymbol{yx}} f \nabla_{\boldsymbol{x}} f$
4: **end for**

In the limit of infinitesimal learning rates, double-FTR behaves as a continuous flow (let $c = \frac{\eta_{\boldsymbol{y}}}{\eta_{\boldsymbol{x}}}$):

$$\begin{bmatrix} \dot{\boldsymbol{x}} \\ \dot{\boldsymbol{y}} \end{bmatrix} = - \begin{bmatrix} \boldsymbol{I} & -(\nabla^2_{\boldsymbol{xx}} f)^{-1} \nabla^2_{\boldsymbol{xy}} g \\ -(\nabla^2_{\boldsymbol{yy}} g)^{-1} \nabla^2_{\boldsymbol{yx}} f & \boldsymbol{I} \end{bmatrix} \begin{bmatrix} \nabla_{\boldsymbol{x}} f(\boldsymbol{x}, \boldsymbol{y}) \\ -c \nabla_{\boldsymbol{y}} g(\boldsymbol{x}, \boldsymbol{y}) \end{bmatrix}. \tag{2}$$

To characterize the stationary points of equation 2, we study its Jacobian in the neighbourhood of an equilibrium. Let $\boldsymbol{z} = (\boldsymbol{x}, \boldsymbol{y})$. Applying the approximation $\begin{bmatrix} \nabla_{\boldsymbol{x}} f \\ -\nabla_{\boldsymbol{y}} g \end{bmatrix} \approx \boldsymbol{J}_{\mathrm{GDA}} \boldsymbol{z}$ in the neighbourhood of an equilibrium, we get the following factorized flow for double-FTR:

$$\dot{\boldsymbol{z}} = -\boldsymbol{\omega}_{\mathrm{FTR}}(\boldsymbol{z}), \text{ where } \boldsymbol{\omega}_{\mathrm{FTR}}(\boldsymbol{z}) = \begin{bmatrix} (\nabla^2_{\boldsymbol{xx}} f)^{-1} & \\ & -(\nabla^2_{\boldsymbol{yy}} g)^{-1} \end{bmatrix} \boldsymbol{J}_{\mathrm{GDA}}^\top \begin{bmatrix} \boldsymbol{I} & \\ & c\boldsymbol{I} \end{bmatrix} \boldsymbol{J}_{\mathrm{GDA}} \boldsymbol{z}. \tag{3}$$

3.1 LASE OF DOUBLE-FTR $\iff$ DIFFERENTIAL NASH EQUILIBRIA

Our main theoretical result is stated below.

Theorem 1. *For a two-player game $\{(f, -g), \mathbb{R}^{n+m}\}$ where all critical points of the GDA flow (equation 1) are hyperbolic, the continuous-time dynamical system in equation 3 satisfies:*

- *$\boldsymbol{z}^*$ is a locally asymptotically stable equilibrium (LASE) of equation 3 $\iff$ $\boldsymbol{z}^*$ is a differential Nash equilibrium of the game $\{(f, -g), \mathbb{R}^{n+m}\}$.*
- *The Jacobian $\boldsymbol{J}_{\mathrm{FTR}} := \frac{\partial \boldsymbol{\omega}_{\mathrm{FTR}}}{\partial \boldsymbol{z}}$ at any critical point of equation 3 has real eigenvalues.*

In practice, we need to apply a discretization of the continuous flow in equation 3 with sufficiently small learning rates. We defer the proof of Theorem 1 to Appendix A.

Combining Theorem 1 with the local convergence property of LASE points, we conclude that equation 3 locally converges to and only to differential Nash equilibria. To the best of our knowledge, the double FTR is the first algorithm with such local convergence result for general-sum games.

3.2 GENERAL PRECONDITIONERS FOR DOUBLE-FTR

In the following remark, we show that double-FTR can be generalized to include family of algorithms.

Remark 1. Theorem 1 applies to a more general version of the double FTR algorithm. In particular, we can generalize equation 3 to allow a broader class of "preconditioners":

$$\dot{\boldsymbol{z}} = -\tilde{\boldsymbol{\omega}}_{\mathrm{FTR}}(\boldsymbol{z}), \;\; \tilde{\boldsymbol{\omega}}_{\mathrm{FTR}}(\boldsymbol{z}) = \begin{bmatrix} \boldsymbol{P}_{\boldsymbol{x}} & \\ & -\boldsymbol{P}_{\boldsymbol{y}} \end{bmatrix} \boldsymbol{J}_{\mathrm{GDA}}^\top(\boldsymbol{z}) \begin{bmatrix} \boldsymbol{I} & \\ & c\boldsymbol{I} \end{bmatrix} \boldsymbol{J}_{\mathrm{GDA}}(\boldsymbol{z}) \boldsymbol{z}, \tag{4}$$

where $\boldsymbol{P}_{\boldsymbol{x}}$, $\boldsymbol{P}_{\boldsymbol{y}}$ are functions of $\boldsymbol{x}, \boldsymbol{y}$ respectively, and satisfy $\boldsymbol{P}_{\boldsymbol{x}} \succ 0 \iff \nabla^2_{\boldsymbol{xx}} f \succ 0$, and $\boldsymbol{P}_{\boldsymbol{y}} \prec 0 \iff \nabla^2_{\boldsymbol{yy}} g \prec 0$.

Remark 1 also provides intuition on the convergence of double-FTR. Without the preconditioner $\boldsymbol{P_x}$ and $\boldsymbol{P_y}$, double-FTR reduces to Hamiltonian gradient descent (Mescheder et al., 2017; Balduzzi et al., 2018; Abernethy et al., 2021), which is not guaranteed to only converge to local Nash equilibria. It is the introduction of the preconditioner that enables LASE points to satisfy the second-order condition of DNE.

4 Connection with other algorithms

Mazumdar et al. (2019) proposed local symplectic surgery (LSS) – a gradient-based algorithm whose LASE are only local Nash equilibria in zero-sum games. LSS avoids oscillatory behaviour at local Nash equilibria (similar to double-FTR). Compared to LSS, double-FTR appears to have a simpler form and enables a broader family of algorithms with such LASE characterizations in general-sum games.

The Follow-the-Ridge (FTR) algorithm (Wang et al., 2019) is closely related to our proposed double-FTR. FTR was proposed for two-player sequential games and is guaranteed to converge to and only to local minimax. FTR applies a gradient correction term on the follower in a sequential game, so that the agents approximately follow a ridge in the landscape of the objective function. The double-FTR can be viewed as a counterpart of FTR for simultaneous games. The update rule of double-FTR resembles that of FTR, with the gradient modification term applied on both players.

Another related algorithm is the Hamiltonian gradient descent (HGD) (Abernethy et al., 2021; Loizou et al., 2020). HGD performs gradient-descent on the Hamiltonian, or the squared norm of the gradient. HGD is guaranteed to converge, as it is essentially a minimization problem. However, in general it is not guaranteed to converge only to local Nash equilibria. Interestingly, our double-FTR can be viewed as a preconditioned HGD.

5 Related work

Mazumdar et al. (2020) introduced a general framework for competitive gradient-based learning. They characterized local Nash equilibria in terms of the critical points of the gradient algorithms. They showed the lack of convergence of the gradient algorithm in games, which motivated the development of the double-FTR algorithm.

Much work has focused on improving the dynamics in finding stable fixed points, which is crucial in applications such as GANs, where oscillation caused by eigenvalues with zero real parts of large imaginary parts in the gradient Jacobian can lead to training instability. Mescheder et al. (2017) proposes Consensus Optimization, which encourages agreement between the two players by introducing a regularization term in the objectives of both players. The regularization term results in a more negative real-part for the eigenvalues of the gradient Jacobian, therefore reduces oscillation and allows larger learning rates. Balduzzi et al. (2018); Gemp & Mahadevan (2018) proposes Symplectic Gradient Adjustment (SGA), which decomposes the gradient Jacobian into symmetric (potential) and asymmetric (Hamiltonian) parts and adds a gradient adjustment term for rapid convergence to stable fixed points. Schäfer & Anandkumar (2019) proposes Competitive Gradient Descent (CGD), whose update is given by the Nash equilibrium of a regularized bilinear approximation of the original game. Compared to other methods, CGD has the advantage of not needing to adapt step size when the interaction strength changes between players. Many other methods have been proposed with different strategies for predicting other agents' moves, such as Learning with Opponent Learning Awareness (LOLA) (Foerster et al., 2016) and optimistic gradient descent-ascent (OGDA) (Popov, 1980; Rakhlin & Sridharan, 2013; Daskalakis et al., 2018; Mertikopoulos et al., 2018). However, none of these existing methods address the problem of spurious (i.e. non-Nash) stable fixed points.

6 Experiments

We conduct simple experiments to demonstrate the implications of our theoretical results. In Section 6.1, we show that the double-FTR algorithm empirically converges to and only to differential Nash equilibria, as predicted by Theorem 1. In Section 6.2, we demonstrate the practical implications of another property of double-FTR — eigenvalues of $\boldsymbol{J}_{\mathrm{FTR}}$ at critical points are real.

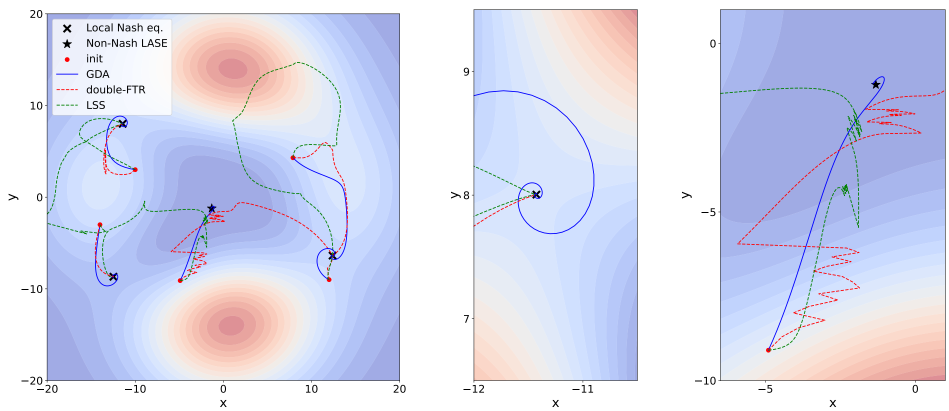

Figure 3: Let: evolution of GDA and double-FTR in the 2-D toy example from multiple initial points. Middle: zoom-in near a local Nash equilibrium point. Right: zoom-in near a non-Nash LASE for the GDA algorithm.

6.1 2-D TOY EXAMPLE

We consider the zero-sum game $\{f, -f\}, \mathbb{R}^2$ with the following 2-D function (same as in Mazumdar et al. (2019)):

$$f(x, y) = e^{-0.01(x^2+y^2)}\big((0.3x^2 + y)^2 + (0.5y^2 + x)^2\big).$$

This function has several LASE points for the GDA dynamics, among which some are DNE and some are not. As shown in Figure 3, while GDA may converge to critical points that are not local Nash equilibria, double-FTR avoids such spurious critical points. Also, in the neighbourhood of local Nash equilibria, GDA exhibits oscillatory behaviour due to complex eigenvalues of the Jacobian matrix. In contrast, the double-FTR does not have oscillatory behaviour near local Nash equilibria. For reference, we also show the trajectories of the Local Symplectic Surgery (LSS). In this experiment, LSS has similar convergence properties – it avoids spurious critical points and does not have oscillatory behaviour near local Nash equilibria.

6.2 PARAMETERIZED BILINEAR GAME

The fact that the Jacobian of the double-FTR at critical points have real eigenvalues is also beneficial for stochastic games, in which the GDA might not converge due to its oscillatory behaviour.

Consider the following stochastic parameterized bilinear game:

$$\min_{\mu_x, \sigma_x} r(x, y), \quad \min_{\mu_y, \sigma_y} -r(x, y)$$

$$\text{where } x \sim \mathcal{N}(\mu_x, \sigma_x^2), \;\; y \sim \mathcal{N}(\mu_y, \sigma_y^2), \;\; r(x, y) = xy.$$

The unique Nash equilibrium with respect to (x, y) is $(0, 0)$. However, the learnable parameters are the mean and the standard deviation of the distributions where x and y are drawn from. We denote the learnable parameters for x and y as $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ respectively. At each time step, we obtain an unbiased estimate of the gradient using REINFORCE, over a mini-batch of size B:

$$\tilde{\nabla}_{\boldsymbol{\theta}} r(\boldsymbol{\theta}, \boldsymbol{\phi}) = \frac{1}{B}\sum_{i=1}^{B} \nabla_{\boldsymbol{\theta}} \log \mathcal{N}(x_i; \boldsymbol{\theta}) r(x_i, y_i), \;\; \boldsymbol{\theta} = \begin{bmatrix} \mu_x \\ \sigma_x \end{bmatrix}$$
$$\tilde{\nabla}_{\boldsymbol{\phi}} r(\boldsymbol{\theta}, \boldsymbol{\phi}) = \frac{1}{B}\sum_{i=1}^{B} \nabla_{\boldsymbol{\phi}} \log \mathcal{N}(y_i; \boldsymbol{\phi}) r(x_i, y_i), \;\; \boldsymbol{\phi} = \begin{bmatrix} \mu_y \\ \sigma_y \end{bmatrix}.$$

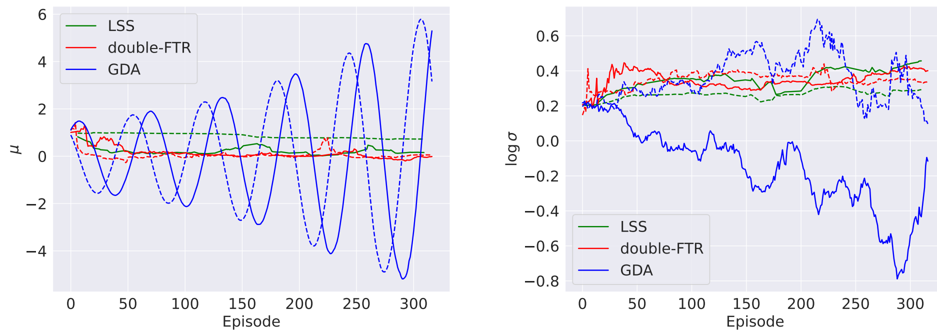

Figure 4: Weight evolution of the parameterized bilinear game under GDA and double-FTR. Solid line represent the weights for $\boldsymbol{\theta}$, and dashed lines represent the weights for $\boldsymbol{\phi}$.

The second-order derivatives are estimated as below ($\tilde{\nabla}^2_{\boldsymbol{\phi}\boldsymbol{\phi}} r(\boldsymbol{\theta}, \boldsymbol{\phi})$ and $\tilde{\nabla}^2_{\boldsymbol{\phi}\boldsymbol{\theta}} r(\boldsymbol{\theta}, \boldsymbol{\phi})$ are computed similarly):

$$\tilde{\nabla}^2_{\boldsymbol{\theta}\boldsymbol{\theta}} r(\boldsymbol{\theta}, \boldsymbol{\phi}) = \frac{1}{B} \sum_{i=1}^{B} \nabla_{\boldsymbol{\theta}} \log \mathcal{N}(x_i; \boldsymbol{\theta}) \big(\nabla_{\boldsymbol{\theta}} \log \mathcal{N}(x_i; \boldsymbol{\theta}) r(x_i, y_i)\big)^\top + \nabla^2_{\boldsymbol{\theta}\boldsymbol{\theta}} \log \mathcal{N}(x_i; \boldsymbol{\theta}) r(x_i, y_i)$$

$$\tilde{\nabla}^2_{\boldsymbol{\theta}\boldsymbol{\phi}} r(\boldsymbol{\theta}, \boldsymbol{\phi}) = \frac{1}{B} \sum_{i=1}^{B} \nabla_{\boldsymbol{\theta}} \log \mathcal{N}(x_i; \boldsymbol{\theta}) \big(\nabla_{\boldsymbol{\phi}} \log \mathcal{N}(y_i; \boldsymbol{\phi}) r(x_i, y_i)\big)^\top.$$

As is often the case, GDA has oscillatory behaviour due to the complex eigenvalues of its Jacobian at critical points. In this stochastic setting, the oscillation prevents convergence for GDA (Figure 4). In contrast, the double-FTR algorithm does not have rotational behaviour at critical points, and converges to the unique Nash equilibrium $(x, y) = (0, 0)$.

7 Conclusion

We propose double Follow-the-Ridge (double-FTR), an gradient-based algorithm for finding local Nash equilibria in differentiable games. We prove that when only hyperbolic critical points are present, double-FTR converges to and only to local Nash equilibria in the general-sum setting, and avoids oscillation in the neighbourhood of critical points. Furthermore, we remark that by varying the preconditioner, double-FTR leads to a broader family of algorithms that share the same convergence guarantee. Finally, we empirically verify the effectiveness of double-FTR in finding and only finding local Nash equilibria through two simple examples.

References

Jacob Abernethy, Kevin A Lai, and Andre Wibisono. Last-iterate convergence rates for min-max optimization: Convergence of hamiltonian gradient descent and consensus optimization. In *Algorithmic Learning Theory*, pp. 3–47. PMLR, 2021.

Cem Anil, Guodong Zhang, Yuhuai Wu, and Roger Grosse. Learning to give checkable answers with prover-verifier games. *arXiv preprint arXiv:2108.12099*, 2021.

David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In *International Conference on Machine Learning*, pp. 354–363. PMLR, 2018.

Aharon Ben-Tal and Arkadi Nemirovski. Robust optimization–methodology and applications. *Mathematical programming*, 92(3):453–480, 2002.

Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. In *International Conference on Learning Representations*, 2018.

John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.

Tanner Fiez, Benjamin Chasnov, and Lillian J Ratliff. Convergence of learning dynamics in stackelberg games. *arXiv preprint arXiv:1906.01217*, 2019.

Jakob Foerster, Ioannis Alexandros Assael, Nando De Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. *Advances in neural information processing systems*, 29, 2016.

Ian Gemp and Sridhar Mahadevan. Global convergence to the equilibrium of gans using variational inequalities. *arXiv preprint arXiv:1808.01531*, 2018.

Ian Gemp, Brian McWilliams, Claire Vernade, and Thore Graepel. Eigengame: Pca as a nash equilibrium. In *International Conference on Learning Representations*, 2020.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

Ya-Ping Hsieh, Panayotis Mertikopoulos, and Volkan Cevher. The limits of min-max optimization algorithms: convergence to spurious non-critical sets. *arXiv preprint arXiv:2006.09065*, 2020.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.

Jason D Lee, Max Simchowitz, Michael I Jordan, and Benjamin Recht. Gradient descent only converges to minimizers. In *Conference on learning theory*, pp. 1246–1257. PMLR, 2016.

Jason D Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I Jordan, and Benjamin Recht. First-order methods almost always avoid strict saddle points. *Mathematical programming*, 176(1):311–337, 2019.

Nicolas Loizou, Hugo Berard, Alexia Jolicoeur-Martineau, Pascal Vincent, Simon Lacoste-Julien, and Ioannis Mitliagkas. Stochastic hamiltonian gradient methods for smooth games. In *International Conference on Machine Learning*, pp. 6370–6381. PMLR, 2020.

Eric Mazumdar, Lillian J Ratliff, and S Shankar Sastry. On gradient-based learning in continuous games. *SIAM Journal on Mathematics of Data Science*, 2(1):103–131, 2020.

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 preprint arXiv:1901.00838*, 2019.

Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. In *International Conference on Learning Representations*, 2018.

Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. The numerics of gans. *Advances in neural information processing systems*, 30, 2017.

Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *International Conference on Artificial Intelligence and Statistics*, pp. 1497–1507. PMLR, 2020.

Leonid Denisovich Popov. A modification of the arrow-hurwicz method for search of saddle points. *Mathematical notes of the Academy of Sciences of the USSR*, 28(5):845–848, 1980.

Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Conference on Learning Theory*, pp. 993–1019. PMLR, 2013.

Lillian J Ratliff, Samuel A Burden, and S Shankar Sastry. Characterization and computation of local nash equilibria in continuous games. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 917–924. IEEE, 2013.

Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018.

Ernest K Ryu and Stephen Boyd. Primer on monotone operator methods. *Appl. Comput. Math*, 15 (1):3–43, 2016.

Shankar Sastry. *Nonlinear systems: analysis, stability, and control*, volume 10. Springer Science & Business Media, 2013.

Florian Schäfer and Anima Anandkumar. Competitive gradient descent. *Advances in Neural Information Processing Systems*, 32, 2019.

John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton university press, 2007.

Yuanhao Wang, Guodong Zhang, and Jimmy Ba. On solving minimax optimization locally: A follow-the-ridge approach. In *International Conference on Learning Representations*, 2019.

Guodong Zhang, Xuchan Bao, Laurent Lessard, and Roger Grosse. A unified analysis of first-order methods for smooth games via integral quadratic constraints. *Journal of Machine Learning Research*, 22:1–39, 2021a.

Guodong Zhang, Yuanhao Wang, Laurent Lessard, and Roger Grosse. Near-optimal local convergence of alternating gradient descent-ascent for minimax optimization. *arXiv preprint arXiv:2102.09468*, 2021b.

Michael Zhang, James Lucas, Jimmy Ba, and Geoffrey E Hinton. Lookahead optimizer: k steps forward, 1 step back. *Advances in Neural Information Processing Systems*, 32, 2019.

A Deferred proofs

Theorem 1. *For a two-player game $\{(f,-g),\mathbb{R}^{n+m}\}$ where all critical points of the GDA flow (equation 1) are hyperbolic, the continuous-time dynamical system in equation 3 satisfies:*

- *$\boldsymbol{z}^*$ is a locally asymptotically stable equilibrium (LASE) of equation 3 $\Longleftrightarrow$ $\boldsymbol{z}^*$ is a differential Nash equilibrium of the game $\{(f,-g),\mathbb{R}^{n+m}\}$.*
- *The Jacobian $\boldsymbol{J}_{\mathrm{FTR}} := \frac{\partial \boldsymbol{\omega}_{\mathrm{FTR}}}{\partial \boldsymbol{z}}$ at any critical point of equation 3 has real eigenvalues.*

Proof. To avoid notational clutter, we drop the subscripts for the Jacobian matrix of the GDA flow in this proof (i.e. we use $\boldsymbol{J}$ to refer to $\boldsymbol{J}_{\mathrm{GDA}}$). Define $\boldsymbol{A} := \begin{bmatrix} (\nabla^2_{\boldsymbol{xx}} f)^{-1} & \\ & -(\nabla^2_{\boldsymbol{yy}} g)^{-1} \end{bmatrix}$ and $\tilde{\boldsymbol{J}} := \begin{bmatrix} \boldsymbol{I} & \\ & \sqrt{c}\boldsymbol{I} \end{bmatrix} \boldsymbol{J}$, we re-write equation 2 as:

$$\begin{bmatrix} \dot{\boldsymbol{x}} \\ \dot{\boldsymbol{y}} \end{bmatrix} = -\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}} \begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{y} \end{bmatrix}. \tag{5}$$

Since we restrict the games to only have hyperbolic critical points, we have $Re(\lambda) \neq 0$ for all $\lambda \in \operatorname{spec}(\boldsymbol{J})$, hence $\boldsymbol{J}$, $\tilde{\boldsymbol{J}}$ have full rank, and $(\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}}$ exists. Then, notice that $\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}}$ is similar to $(\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}} \boldsymbol{A} (\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}}$, we conclude that all eigenvalues of $\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}}$ are real.

The rest of the proof consists of two directions: 1) we show that all locally asymptotically stable equilibria (LASE) of equation 5 are differential Nash equilibria (DNE); and 2) we show that for two-player zero games with only hyperbolic critical points, all DNE are LASE of equation 5.

(1) LASE $\implies$ DNE.

Let $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ be a LASE of equation 5, we have:

$$\dot{\boldsymbol{x}}^* = \boldsymbol{0}, \quad \dot{\boldsymbol{y}}^* = \boldsymbol{0}, \quad \text{and} \quad \operatorname{spec}(\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}}) \subset \mathbb{C}^+.$$

First-order condition for DNE Since all eigenvalues of $\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}}$ are in the strictly positive real part, we can conclude that the matrices $\boldsymbol{A}$, $\boldsymbol{J}$ and $\tilde{\boldsymbol{J}}$ are all of full-rank. Therefore, the nullspace of $\boldsymbol{A}\boldsymbol{J}^\top \begin{bmatrix} \boldsymbol{I} & \\ & c\boldsymbol{I} \end{bmatrix}$ has only the zero vector, and $\begin{bmatrix} \dot{\boldsymbol{x}}^* \\ \dot{\boldsymbol{y}}^* \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix}$ implies $\begin{bmatrix} \nabla_{\boldsymbol{x}} f(\boldsymbol{x}^*, \boldsymbol{y}^*) \\ \nabla_{\boldsymbol{y}} f(\boldsymbol{x}^*, \boldsymbol{y}^*) \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix}$. The first-order condition for DNE is satisfied.

Second-order condition for DNE

Since $\operatorname{spec}(\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}}) \subset \mathbb{C}^+$, we have that $(\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}} \boldsymbol{A} (\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}} \succ 0$ (which means that in fact, $\operatorname{spec}(\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}}) \subset \mathbb{R}^+$). Since positive-definiteness is preserved from congruence transformation, we have $\boldsymbol{A} \succ 0$.

Finally, $\boldsymbol{A} \succ 0 \implies (\nabla^2_{\boldsymbol{xx}} f)^{-1} \succ 0, (\nabla^2_{\boldsymbol{yy}} g)^{-1} \prec 0 \implies \nabla^2_{\boldsymbol{xx}} f \succ 0, \nabla^2_{\boldsymbol{yy}} g \prec 0$. The second-order condition for DNE is satisfied.

(2) DNE $\implies$ LASE.

Let $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ be a DNE of the game $\{(f,-f),\mathbb{R}^{n+m}\}$, we have:

$$\nabla_{\boldsymbol{x}} f(\boldsymbol{x}^*, \boldsymbol{y}^*) = \boldsymbol{0}, \quad \nabla_{\boldsymbol{y}} f(\boldsymbol{x}^*, \boldsymbol{y}^*) = \boldsymbol{0}, \quad \text{and} \quad \nabla^2_{\boldsymbol{xx}} f \succ 0, \quad \nabla^2_{\boldsymbol{yy}} g \prec 0.$$

First, from the update rule of simultaneous FTR, we immediately have $\begin{bmatrix} \nabla_{\boldsymbol{x}} f(\boldsymbol{x}^*, \boldsymbol{y}^*) \\ \nabla_{\boldsymbol{y}} f(\boldsymbol{x}^*, \boldsymbol{y}^*) \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix} \implies \begin{bmatrix} \dot{\boldsymbol{x}}^* \\ \dot{\boldsymbol{y}}^* \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix}$.

Then, since we restrict the games to only have hyperbolic critical points, we have $Re(\lambda) \neq 0$ for all $\lambda \in \mathrm{spec}(\boldsymbol{J})$, hence $\boldsymbol{J}$, $\tilde{\boldsymbol{J}}$ have full rank, and $(\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}}$ exists.

$\nabla^2_{\boldsymbol{xx}} f \succ 0,\ \nabla^2_{\boldsymbol{yy}} g \prec 0 \implies \boldsymbol{A} \succ 0 \implies (\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}} \boldsymbol{A} (\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}} \succ 0.$ Since $\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}}$ is similar to $(\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}} \boldsymbol{A} (\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}})^{\frac{1}{2}}$, we have that $\mathrm{spec}(\boldsymbol{A}\tilde{\boldsymbol{J}}^\top \tilde{\boldsymbol{J}}) \subset \mathbb{R}^+$.

$(\boldsymbol{x}^*, \boldsymbol{y}^*)$ is a LASE of equation 5.

$\square$