

Newton Optimization on Helmholtz Decomposition for Continuous Games

Giorgia Ramponi, Marcello Restelli

Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano
Piazza Leonardo da Vinci, 32, 20133, Milano, Italy
{giorgia.ramponi, marcello.restelli}@polimi.it

Abstract

Many learning problems involve multiple agents that optimize different interactive functions. In these problems, standard policy gradient algorithms fail due to the non-stationarity of the setting and the different interests of each agent. In fact, the learning algorithms must consider the complex dynamics of these systems to guarantee rapid convergence towards a (local) Nash equilibrium. In this paper, we propose NOHD (Newton Optimization on Helmholtz Decomposition), a Newton-like algorithm for multi-agent learning problems based on the decomposition of the system dynamics into its irrotational (Potential) and solenoidal (Hamiltonian) components. This method ensures quadratic convergence in purely irrotational systems and pure solenoidal systems. Furthermore, we show that NOHD is attracted to stable fixed points in general multi-agent systems and repelled by strict saddle ones. Finally, we empirically compare the NOHD's performance with state-of-the-art algorithms on some bimatrix games and in a continuous Gridworld environment.

1 Introduction

In recent years, Reinforcement Learning (RL) (Sutton, Barto et al. 1998) methods with multiple agents (Buşoniu, Babuška, and De Schutter 2010; Zhang, Yang, and Başar 2019) have made substantial progress in solving decision-making problems such as playing Go (Silver et al. 2016), robotic control problems (Lillicrap et al. 2015), playing card games (Brown and Sandholm 2019) and autonomous driving (Shalev-Shwartz, Shammah, and Shashua 2016). Furthermore, in other machine learning fields, powerful algorithms that optimize multiple losses have recently been proposed. Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) is an example, which achieves successful results in Computer Vision (Isola et al. 2017; Ledig et al. 2017) and Natural Language Generation (Nie, Narodytska, and Patel 2018; Yu et al. 2017). On the other hand, thanks to their ability to learn in the stochastic policy space and their effectiveness in solving high-dimensional, continuous state and action problems, policy-gradient algorithms (Peters and Schaal 2006) are natural candidates for use in multi-agent learning problems. Nonetheless, multiple

policy-gradient agents' interaction has proven unsuccessful in learning a set of policies that converges to a (local) Nash Equilibrium (Mertikopoulos, Papadimitriou, and Piliouras 2018; Papadimitriou and Piliouras 2016). More than one agent leads to the failure of standard optimization methods in most games due to the non-stationarity of the environment and the lack of cooperation between the agents.

How to optimize multiple policy-gradient agents is a problem of theoretical interest and practical importance. Over the past two years, a growing number of papers has addressed this problem by focusing on *continuous (differentiable) games*, i.e., games where the agent's objective functions are twice differentiable with respect to the policy parameters (Mazumdar, Ratliff, and Sastry 2020). Some of them only consider the competitive setting due to the success of GANs (Mescheder, Nowozin, and Geiger 2017). The general case was considered only recently, where the gradient descent update rule was combined with second-order terms (Balduzzi et al. 2018; Letcher et al. 2018; Foerster et al. 2018; Schäfer and Anandkumar 2019). Some of them guarantee linear convergence under specific assumptions.

In this paper, we study how to build a Newton-based algorithm (NOHD) for learning policies in multi-agent environments. First of all, we start by analyzing two specific game classes: Potential Games and Hamiltonian Games (Section 3). In Section 4, we propose a Newton-based update for these two classes of games, for which linear-rate algorithms that guarantee convergence are known, proving quadratic convergence rates. Then, we extend the algorithm to the general case, neither Hamiltonian nor Potential. We show that the proposed algorithm respects some desiderata, similar to those proposed in (Balduzzi et al. 2018): the algorithm has to guarantee convergence to (local) Nash Equilibria in (D1) Potential and (D2) Hamiltonian games; (D3) the algorithm has to be attracted by stable fixed points and (D4) repelled by unstable fixed points. Finally, in Section 6, we analyze the empirical performance of NOHD when agents optimize a Boltzmann policy in three bimatrix games: Prisoner's Dilemma, Matching Pennies, and Rock-Paper-Scissors. In the last experiment, we study the learning performance of NOHD in two continuous gridworld environments. In all experiments, NOHD achieves great results confirming the quadratic nature of the update. The proofs of all the results

presented in the paper are reported in Appendix ¹.

2 Related Works

The study of convergence in classic convex multiplayer games has been extensively studied and analyzed (Rosen 1965; Facchinei and Kanzow 2007; Singh, Kearns, and Mansour 2000). Unfortunately, the same algorithms cannot be used with neural networks due to the non-convexity of the objective functions. Various algorithms have been proposed that successfully guarantee convergence in specific classes of games: policy prediction in two-player two-action bimatrix games (Zhang and Lesser 2010; Song, Wang, and Zhang 2019); WoLF in two-player two-action games (Bowling and Veloso 2002); AWESOME in repeated games (Conitzer and Sandholm 2007); Optimistic Mirror Descent in two-player bilinear zero-sum games (Daskalakis et al. 2017); Consensus Optimization (Mescheder, Nowozin, and Geiger 2017), Competitive Gradient Descent (Schäfer and Anandkumar 2019) and (Mazumdar, Jordan, and Sastry 2019) in two-player zero-sum games. Other recent works propose methods to learn in smooth-markets (Balduzzi et al. 2020), sequential imperfect information games (Perolat et al. 2020), zero-sum linear-quadratic games (Zhang, Yang, and Basar 2019). Furthermore, the dynamics of Stackelberg games was studied by Fiez, Chasnov, and Ratliff (2019). Other insights on the convergence of gradient-based learning was proposed in (Chasnov et al. 2020). Many algorithms have also been proposed specifically for GANs (Heusel et al. 2017; Nagarajan and Kolter 2017; Gemp and Mahadevan 2018).

Some recent works have developed learning algorithms also for general games. The first example is the Iterated Gradient Ascent Policy Prediction (IGA-PP) algorithm, renamed LookAhead (Zhang and Lesser 2010). In (Letcher et al. 2018) it is proved that IGA-PP converges to local Nash Equilibria not only in two-player two-action bimatrix games but also in general games. Learning with opponent learning awareness (LOLA) (Foerster et al. 2018) is an attempt to use the other agents’ functions to account for the impact of one agent’s policy on the anticipated parameter update of the other agents. The empirical results show the effectiveness of LOLA, but no convergence guarantees are provided. Indeed, Letcher et al. (2018) have shown that LOLA may converge to non-fixed points and proposed Stable Opponent Shaping, an algorithm that maintains the theoretical convergence guarantees of IGA-PP, also exploiting the opponent dynamics like LOLA. In (Balduzzi et al. 2018; Letcher et al. 2019) the authors studied game dynamics by decomposing a game into its Potential and Hamiltonian components using the generalized Helmholtz decomposition. The authors propose Symplectic Gradient Adjustment (SGA), an algorithm for general games, which converges locally to stable fixed points, using the Hamiltonian part to adjust the gradient update. Instead, our algorithm uses information about the Potential component and the Hamiltonian component of the game to approximate the general game with one of these parts at each step.

¹A complete version of the paper, which includes the appendix, is available at <https://arxiv.org/pdf/2007.07804.pdf>.

3 Preliminaries

We cast the multi-agent learning problem as a *Continuous Stochastic Game*. We adapt the concept of Stochastic Game and Continuous (Differentiable) Game from (Foerster et al. 2018; Balduzzi et al. 2018; Ratliff, Burden, and Sastry 2016). In this section, after introducing Continuous Stochastic Games, we recall the game decomposition proposed by Balduzzi et al. (2018). Then, we describe the desired convergence points. In the end, we introduce Newton’s method.

Continuous Stochastic Games A continuous stochastic game (Foerster et al. 2018; Balduzzi et al. 2018) is a tuple $\mathcal{G} = (X, U_1, \dots, U_n, f, C_1, \dots, C_n, \gamma_1, \dots, \gamma_n)$ where n is the number of agents; X is the set of states; $U_i, 1 \leq i \leq n$, is the set of actions of agent i and $U = U_1 \times \dots \times U_n$ is the joint action set; $f : X \times U \rightarrow \Delta(X)$ is the state transition probability function (where $\Delta(\Omega)$ denotes the set of probability measures over a generic set Ω), $C_i : X \times U \rightarrow \mathbb{R}$ is the cost function of agent i , and $\gamma_i \in [0, 1)$ is its discount factor.² The agent’s behavior is described by means of a parametric twice differentiable policy $\pi_{\theta_i} : X \rightarrow \Delta(U_i)$, where $\theta_i \in \Theta \subseteq \mathbb{R}^d$ and $\pi_{\theta_i}(\cdot|x)$ specifies for each state x a distribution over the action space U_i .³ We denote by θ the vector of length nd obtained by stacking together the parameters of all the agents: $\theta = (\theta_1^T, \dots, \theta_n^T)^T$. We denote an infinite-horizon trajectory in game \mathcal{G} by $\tau = \{x_t, \mathbf{u}_t\}_{t=0}^\infty \in \mathbb{T}$, where $\mathbf{u}_t = (u_1(t), \dots, u_n(t))$ is the joint action at time t , $x_t \sim f(x_{t-1}, \mathbf{u}_{t-1})$ (for $t > 0$), and \mathbb{T} is the trajectory space. In stochastic games, all agents try to minimize their expected discounted cost separately, which is defined for the i -th agent as:

$$V_i(\theta) = \mathbb{E}_\tau [C_i(\tau)], \quad (1)$$

where $C_i(\tau) = \sum_{t=0}^\infty \gamma_i^t C_i(x_t, \mathbf{u}_t)$ and the expectation is with respect to the agents’ policies and the transition model. Note that the expectation depends on all the agents’ policies. We do not assume the convexity of the functions $V_i(\theta)$.

We define the *symultaneous gradient* (Letcher et al. 2018) as the concatenation of the gradient of each discounted return function respect to the parameters of each player:

$$\xi(\theta) = (\nabla_{\theta_1} V_1^T, \dots, \nabla_{\theta_n} V_n^T)^T. \quad (2)$$

The *Jacobian* of the game (Letcher et al. 2018; Ratliff, Burden, and Sastry 2013) J is an $nd \times nd$ matrix, where n is the number of agents and d the number of policy parameters for each agent. J is composed by the matrix of the derivatives of the simultaneous gradient, i.e., for each player i the i -th row of its hessian:

$$J = \nabla_\theta \xi = \begin{pmatrix} \nabla_{\theta_1}^2 V_1 & \nabla_{\theta_1, \theta_2} V_1 & \dots & \nabla_{\theta_1, \theta_n} V_1 \\ \nabla_{\theta_2, \theta_1} V_2 & \nabla_{\theta_2}^2 V_2 & \dots & \nabla_{\theta_2, \theta_n} V_2 \\ \vdots & \vdots & \ddots & \vdots \\ \nabla_{\theta_n, \theta_1} V_n & \nabla_{\theta_n, \theta_2} V_n & \dots & \nabla_{\theta_n}^2 V_n \end{pmatrix}.$$

Game dynamics J is a square matrix, not necessarily symmetric. The antisymmetric part of J is caused by each

²Given a vector \mathbf{v} and a matrix H , in the following we will denote by \mathbf{v}^T the transpose of \mathbf{v} and with $\|\mathbf{v}\|$ and $\|H\|$ the respective L2-norms.

³To ease the notation, we will drop θ (e.g., π_i instead of π_{θ_i}) when not necessary.

agent’s different cost functions and can cause cyclical behavior in the game (even in simple cases as bimatrix zero-sum games, see Figure 2). On the other hand, the symmetric part represents the “cooperative” part of the game. In (Balduzzi et al. 2018), the authors proposed how to decompose J in its symmetric and antisymmetric component using the Generalized Helmholtz decomposition (Wills 1958; Balduzzi et al. 2018)⁴.

Proposition 1. *The Jacobian of a game decomposes uniquely into two components $J = S + A$, where $S = \frac{1}{2}(J + J^T)$ and $A = \frac{1}{2}(J - J^T)$.*

Components S and A represent the irrotational (Potential), S , and the solenoidal (Hamiltonian), A , part of the game, respectively. The irrotational component is its curl-free component, and the solenoidal one is the divergence-free one.

Potential games are a class of games introduced by (Monderer and Shapley 1996). A game is a potential game if there exists a potential function $\phi : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$, such that: $\phi(\theta'_i, \theta_{-i}) - \phi(\theta''_i, \theta_{-i}) = \alpha(V_i(\theta'_i, \theta_{-i}) - V_i(\theta''_i, \theta_{-i}))$. A potential game is an *exact* potential game if $\alpha = 1$; exact potential games have $A = \mathbf{0}$. In these games, J is symmetric and it coincides with the hessian of the potential function. This class of games is widely studied because in these games gradient descent converges to a Nash Equilibrium (Rosenthal 1973; Lee et al. 2016). In the rest of the document, when we refer to potential games we refer to exact potential games.

Hamiltonian games, i.e., games with $S = \mathbf{0}$, were introduced in (Balduzzi et al. 2018). A Hamiltonian game is described by a Hamiltonian function, which specifies the conserved quantity of the game. Formally, a Hamiltonian system is fully described by a scalar function, $\mathcal{H} : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$. The state of a Hamiltonian system is represented by the generalized coordinates \mathbf{q} momentum and position \mathbf{p} , which are vectors of the same size. The evolution of the system is given by Hamilton’s equations: $\frac{d\mathbf{p}}{dt} = -\frac{\partial \mathcal{H}}{\partial \mathbf{q}}$, $\frac{d\mathbf{q}}{dt} = +\frac{\partial \mathcal{H}}{\partial \mathbf{p}}$. The gradient of \mathcal{H} corresponds to $(S + A^T)\xi$ (Balduzzi et al. 2018). In bimatrix games, Hamiltonian games coincide with zero-sum games, but this is not true in general games (Balduzzi et al. 2018).

Desired convergence points In classic game theory, the standard solution concept is the Nash Equilibrium (Nash et al. 1950). Since we focus on gradient-based methods and make no assumptions about the convexity of the return functions, we consider the concept of local Nash Equilibrium (Ratliff, Burden, and Sastry 2016).

Definition 1. *A point θ^* is a local Nash equilibrium if, $\forall i$, there is a neighborhood B_i of θ^*_i such that $V_i(\theta_i, \theta_{-i}^*) \geq V_i(\theta_i^*, \theta_{-i}^*)$ for any $\theta_i \in B_i$.*

Gradient-based methods can reliably find local (not global) optima even in single-agent non-convex problems (Lee et al. 2016, 2017), but they may fail to find local Nash equilibria in non-convex games.

⁴The Helmholtz decomposition applies to any vector field (Wills 1958).

Another desirable condition is that the algorithm converges into *stable fixed points* (Balduzzi et al. 2018).

Definition 2. *A fixed point θ^* with $\xi(\theta^*) = 0$ is stable if $S(\theta^*) \succeq 0$ and $S(\theta^*)$ is invertible, unstable if $S(\theta^*) \prec 0$ and a strict saddle if $S(\theta^*)$ has an eigenvalue with negative real part.*

Stable fixed points and local Nash equilibria are interesting solution concepts, the former from an optimization point of view and the latter from a game-theoretic perspective. It is important to notice that the stable fixed points are a subset of the all local Nash Equilibria of a game (see Lemma 2 in (Balduzzi et al. 2018)).

Newton method Newton’s method (Nocedal and Wright 2006) guarantees, under assumptions, a quadratic convergence rate to the root of a function, which, in optimization, is the derivative of the function to be optimized. This method is based on a second-order approximation of the twice differentiable function $g(\theta)$ that we are optimizing. Starting from an initial guess θ_0 , Newton’s method updates the parameters θ by setting the derivative of the second-order Taylor approximation of $g(\theta)$ to 0:

$$\theta_{t+1} = \theta_t - \nabla^2 g(\theta_t)^{-1} \nabla g(\theta_t). \quad (3)$$

For non-convex functions, the hessian $\nabla^2 g(\theta)$ is not necessarily positive semidefinite and all critical points are possible solutions for Newton’s method. Then, Newton’s update may converge to a local minimum, a saddle, or a local maximum. A possible solution to avoid this shortcoming is to use a modified version of the inverse of the hessian, called Positive Truncated inverse (PT-inverse) (Nocedal and Wright 2006; Paternain, Mokhtari, and Ribeiro 2019):

Definition 3 (PT-inverse). *Let $H \in \mathbb{R}^{n \times n}$ be a symmetric matrix, $Q \in \mathbb{R}^{n \times n}$ a basis of orthogonal eigenvectors of H , and $\Lambda \in \mathbb{R}^{n \times n}$ a diagonal matrix of corresponding eigenvalues. The $|\Lambda|_m \in \mathbb{R}^{n \times n}$ is the positive definite truncated eigenvalue matrix of Λ with parameter m :*

$$(|\Lambda|_m)_{ii} = \begin{cases} |\Lambda_{ii}| & \text{if } |\Lambda_{ii}| \geq m, \\ m & \text{otherwise.} \end{cases} \quad (4)$$

The PT-inverse of H with parameter m is the matrix $|H|_m^{-1} = Q|\Lambda|_m^{-1}Q^T$.

The PT-inverse flips the sign of negative eigenvalues and truncates small eigenvalues by replacing them with m . Then the usage of the PT-inverse, instead of the real one, guarantees convergence to a local minimum even in non-convex functions. These properties are necessary to obtain a convergent Newton-like method for non-convex functions.

4 Newton for Games

In this section, we describe how to apply Newton-based methods to Continuous Stochastic Games. We start by showing how Newton’s method can be applied to two game classes: Potential games and Hamiltonian Games. Then we describe an algorithm to extend Newton’s method in general games.




Figure 1: Phase of Potential games: from left phase dynamic of gradient descent, gradient on the potential function, and with NOHD.

Newton’s Method for Potential Games

Potential games are a class of games characterized by the existence of a potential function $\phi : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$ which describes the dynamics of the system. In these games, gradient descent on the simultaneous gradient or the potential function converges to a (local) Nash Equilibrium as shown in Figure 1. In these games, $J = S$, and the anti-symmetric component $A = \mathbf{0}$. Considering the existence of the potential function ϕ (see Section 3), it is sufficient to apply the Newton PT-inverse method (see Definition 3) on it to guarantee a quadratic convergence rate to a local Nash Equilibrium (see Figure 1). Newton’s update for Potential games is:

$$\theta_{t+1} = \theta_t - \eta S_m(\theta_t)^{-1} \xi(\theta_t), \quad (5)$$

where η is a learning rate and m the parameter of the PT-inverse (see Definition 3). So, in Potential games, we have the same convergence properties as in single-function optimization. The convergence into the local minima follows Newton’s convergence proofs for non-convex functions (Paternain, Mokhtari, and Ribeiro 2019).

Newton’s Method for Hamiltonian Games

Hamiltonian games are characterized by a Hamiltonian function $\mathcal{H} : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$. In these games, the gradient descent does not converge to a stable fixed point but causes cyclical behavior. Instead, the gradient descent on the Hamiltonian function converges to a Nash equilibrium. Figure 2 shows the dynamics of gradient descent w.r.t. ξ and $\nabla \mathcal{H}$ on a Hamiltonian game: the figure points out that a gradient descent on ξ cycles.

Example 4.1. Take a two-player bilinear game with agents with parameters θ_1 and θ_2 minimizing respectively $f(\theta) : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$ and $g(\theta) : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$ respectively. A point in this class of games is a Nash Equilibrium if $\xi(\theta) = 0$, i.e., $\nabla_{\theta_1} f = 0$ and $\nabla_{\theta_2} g = 0$, because $\nabla_{\theta_1}^2 f = 0$ and $\nabla_{\theta_2}^2 g = 0$. Considering this, the Nash Equilibrium can be calculated in closed form (if the inverse exists⁵) by setting the gradient equal to zero:

$$\begin{bmatrix} \nabla_{\theta_1} f \\ \nabla_{\theta_2} g \end{bmatrix} + \begin{bmatrix} 0 & \nabla_{\theta_1, \theta_2} f \\ \nabla_{\theta_2, \theta_1} g & 0 \end{bmatrix} \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = 0$$

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = - \begin{bmatrix} 0 & \nabla_{\theta_1, \theta_2} f \\ \nabla_{\theta_2, \theta_1} g & 0 \end{bmatrix}^{-1} \begin{bmatrix} \nabla_{\theta_1} f \\ \nabla_{\theta_2} g \end{bmatrix}.$$

⁵If the inverse does not exist and a Nash Equilibrium exists, the system is indeterminate. Using the Moore-Penrose inverse we find an approximate Nash Equilibrium (the one with the smallest Euclidean norm).




Figure 2: Phase of Hamiltonian games: from left phase dynamic of gradient descent, gradient on the Hamiltonian function, and with NOHD.

Algorithm 1 NOHD

input: discounted returns $\mathcal{V} = \{V_i\}_{i=1}^n$, PT inverse parameter m
output: update rule

```

Compute  $\xi, J, S, A, S_m^{-1}$ 
if  $\cos \nu_S \geq 0$  then
  if  $\cos \nu_S \geq \cos \nu_A$  then (5) else (6)
else
  if  $\cos \nu_S \leq \cos \nu_A$  then (5) else (6)
end if

```

The example above provides the intuition that the solution to quadratic Hamiltonian games is achieved by the following update rule (in Hamiltonian games $S = 0$):

$$\theta_{t+1} = \theta_t - A(\theta_t)^{-1} \xi(\theta_t). \quad (6)$$

In the following theorem, we state that even in this class of games the convergence to a local Nash Equilibrium (using the above update) is quadratic (see Figure 2). The proof is in Appendix A.

Theorem 1. Suppose that ξ and A are twice continuous differentiable and that A is invertible in the (local) Nash Equilibrium θ^* . Then, there exists $\epsilon > 0$ such that iterations starting from any point in the ball $\theta_0 \in B(\theta^*, \epsilon)$ with center θ^* and ray ϵ converge to θ^* . Furthermore, the convergence rate is quadratic.

Newton’s Method for General Games

In general games, it is not yet known whether and how the system’s dynamics can be reduced to a single function as for Potential and Hamiltonian games. Thus, finding a Newton-based update is more challenging: if we apply Newton’s Method with the Jacobian PT-transformation, we can alter the Hamiltonian dynamics of the game. Instead, applying the inverse of the Jacobian as in the Hamiltonian games can lead to local maxima. In this section, we show how to build a Newton-based learning rule that guarantees desiderata similar to those considered in (Balduzzi et al. 2018): the update rule has to be compatible (D1) with Potential dynamics if the game is a Potential game, with (D2) Hamiltonian dynamics if the game is a Hamiltonian game and has to be (D3) attracted by stable fixed points and (D4) repelled by unstable ones. By compatible we mean that given two vectors u, v then $u^T v > 0$.

The algorithm that we propose (see Algorithm 1) chooses the update to perform between the two updates in (5) and (6). The choice is based on the angles between the gradient of the Hamiltonian function \mathcal{H} and the two candidate updates⁶

directions. In particular, we compute

$$\cos \nu_S = \frac{(S_m^{-1}\xi)^T \nabla \mathcal{H}}{\|S_m^{-1}\xi\| \|\nabla \mathcal{H}\|}, \quad \cos \nu_A = \frac{(A^{-1}\xi)^T \nabla \mathcal{H}}{\|A^{-1}\xi\| \|\nabla \mathcal{H}\|}.$$

When the cosine is positive, the update rule follows a direction that reduces the value of the Hamiltonian function (i.e., reduces gradient norm), otherwise the update rule points in an increasing direction of the Hamiltonian function. Notably, there is a connection between the positive/negative definiteness of \mathcal{H} and the sign of $\cos \nu_S$.

Lemma 1. *Given the Jacobian $J = S + A$ and the simultaneous gradient ξ , if $S \succeq 0$ then $\cos \nu_S \geq 0$; instead if $S \prec 0$ then $\cos \nu_S < 0$.*

The idea of NOHD is to use the sign of $\cos \nu_S$ to decide whether to move in a direction that reduces the Hamiltonian function (aiming at converging to a stable fixed point) or not (aiming at getting away from unstable points). In case $\cos \nu_S$ is positive, the algorithm chooses the update rule with the largest cosine value (i.e., which minimizes the angle with $\nabla \mathcal{H}$), otherwise, NOHD tries to point in the opposite direction by taking the update rule that minimizes the cosine. In the following theorem, we show that the update performed by NOHD satisfies the desiderata described above.

Theorem 2. *The NOHD update rule satisfies requirements (D1), (D2), (D3), and (D4).*

Proof. (Sketch) The requirement (D1) is satisfied if $(S_m^{-1}\xi)^T \nabla \phi$ and $(A^{-1}\xi)^T \nabla \phi$ are nonnegative and \mathcal{G} is a Potential game; in this case the update rule is $S_m^{-1}\xi$ because $A = 0$. We have that $\nabla \phi = \xi$ and $\xi^T S_m^{-1}\xi \geq 0$ as said before. For requirement (D2) we can make similar considerations: in this case we have to show that $(S_m^{-1}\xi)^T \nabla \mathcal{H}$ and $(A^{-1}\xi)^T \nabla \mathcal{H}$ are nonnegative when \mathcal{G} is a Hamiltonian game, that it is equal to say: $(A^{-1}\xi)^T A^T \xi = \|\xi\|^2$. Finally, the fulfillment of desiderata (D3) and (D4) is a consequence of Lemma 1. \square

Given the results from Theorem 2 and Lemma 1, we can argue that if ξ points at a stable fixed point then NOHD points also to the stable fixed point otherwise if ξ points away from the fixed point also NOHD points away from it.

Then we prove that NOHD converges only to fixed points and, under some conditions, it converges locally to stable fixed points.

Lemma 2. *If NOHD converges to a θ^* then $\xi(\theta^*) = 0$.*

Theorem 3. *Suppose θ^* is a stable fixed point, and suppose A, S, J are bounded and Lipschitz continuous with modulus respectively M_A, M_S, M_J in the region of attraction of the stable fixed point $\xi(\theta^*)$. Furthermore assume that $\|A^{-1}\| \leq N_A$ and $\|S^{-1}\| \leq N_S$. Then there exists $\epsilon > 0$ such that the iterations starting from any point $\theta_0 \in B(\theta^*, \epsilon)$ converge to θ^* .*

More details about the proofs are given in Appendix A.

Remark 1. *In this paper, we focus our attention on convergence towards stable fixed points. However, some Nash Equilibria are not stable fixed points. On the other hand, stable fixed points are an interesting solution concept since in two-player zero-sum games (e.g., GANs (Goodfellow et al. 2014)), all local Nash Equilibria are also stable fixed points (Balduzzi et al. 2018).*

5 Learning via Policy Gradient in Stochastic Games

Usually, agents do not have access to the full gradient or the Jacobian, and we need to estimate them. We define a T -episodic trajectory as $\tau = (x_0, \mathbf{u}_0, C_0^1, \dots, C_0^n, \dots, x_T, \mathbf{u}_T, C_T^1, \dots, C_T^n)$. Following policy gradient derivation (Peters and Schaal 2008), $\nabla_{\theta_i} V_i$ can be estimated by:⁶

$$\widehat{\nabla}_{\theta_i}^M V_i = \frac{1}{M} \sum_{m=1}^M \sum_{t=0}^T \gamma^t C_{m,t}^i \sum_{t'=0}^t \nabla \log \pi_{\theta_i}(x_{m,t'}, u_{m,t'}^i).$$

Then, to estimate the Jacobian J , we have to compute the second-order gradient $\nabla_{\theta_k} \nabla_{\theta_j} V_i$, with $1 \leq k \leq n, 1 \leq j \leq n$. If $k \neq j$ we derive (Foerster et al. 2018) the second-order gradient, exploiting the independence of agents' policies:

$$\widehat{\nabla}_{\theta_k, \theta_j}^M V_i = \frac{1}{M} \sum_{m=1}^M \sum_{t=0}^T \gamma^t C_{m,t}^i \times \sum_{t'=0}^t (\nabla \log \pi_{\theta_k}(x_{m,t'}, u_{m,t'}^k)) (\nabla \log \pi_{\theta_j}(x_{m,t'}, u_{m,t'}^j))^T.$$

When $k = j$ we are evaluating the second-order gradient of π_{θ_j} . To evaluate this part we derive the second-order gradient as done for the single-agent case (Shen et al. 2019), with $g_m^i(\theta_k) = \sum_{t=0}^T C_t^i \sum_{t'=0}^t \nabla \log \pi_{\theta_k}(x_{m,t'}, u_{m,t'}^i)$:⁷

$$\widehat{\nabla}_{\theta_k}^{2,M} V_i = \frac{1}{M} \sum_{m=1}^M \nabla g_m^i(\theta_k) \nabla \log \pi_{\theta_k}(\tau_m)^T + \nabla^2 g_m^i(\theta_k).$$

6 Experiments

This section is devoted to the experimental evaluation of NOHD. The proposed algorithm is compared with Consensus Optimization (CO) (Mescheder, Nowozin, and Geiger 2017), Stable Opponent Shaping (SOS) (Letcher et al. 2018), Learning with Opponent-Learning Awareness (LOLA) (Foerster et al. 2018), Competitive Gradient Descent (CGD) (Schäfer and Anandkumar 2019), Iterated Gradient Ascent Policy Prediction (IGA-PP) (Zhang and Lesser 2010) and Symplectic Gradient Adjustment (SGA) (Balduzzi et al. 2018).

Matrix Games

We consider two matrix games: two-agent two-action Matching Pennies (MP) and two-agent three-action Rock

⁶With $\widehat{\nabla}_{\theta_i}^M$ we intend the estimator of ∇_{θ_i} over M samples.

⁷All derivations are reported in Appendix.

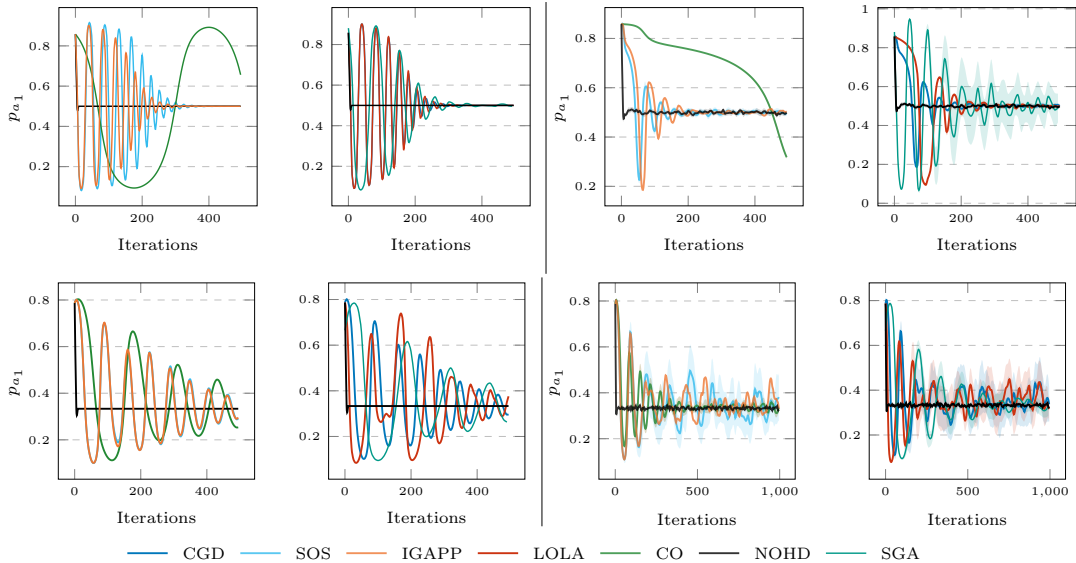


Figure 3: Matching Pennies and Rock Paper Scissors (on the rows): probabilities of agent 1 to perform first action with a Boltzmann parametrization of the policies. On the left: plots with exact gradients. On the right: plots with estimated gradients. Every algorithm with its best learning rate between 0.1, 0.5, 1.0. c.i. 95. %

Paper Scissors (RPS) (games’ rewards are reported in Appendix B). Considering a linear parameterization of the agents’ policies, MP and RPS are Hamiltonian games⁸. The Nash equilibria of MP, and RPS are respectively 0.5, and 0.333 regarding the probability of taking action 1. For the first experiment, we used a Boltzmann parametrization for the agents’ policies and exact computation of gradients and Jacobian. In this setting, games lose their Hamiltonian or Potential property, making the experiment more interesting and the behavior of NOHD not trivial. The results are shown in Figure 3 (left side). For each game, we perform experiments with learning rates 0.1, 0.5, 1.0. In the plots are reported only the best performance for each algorithm. In Matching Pennies we initialize probabilities to $[0.86, 0.14]$ for the first agent and to $[0.14, 0.86]$ for the second agent; instead in Rock Paper Scissors to $[0.66, 0.24, 0.1]$.⁹ The figure shows that each algorithm is able to converge to the Nash equilibrium (in MP, CO converges with a learning rate of 0.1 and takes more than 1000 iterations. The plot is reported in Appendix). In Table 1 we reported the ratio between the number of steps each algorithm takes to converge and the maximum number of steps in which the slowest algorithm converges. For this simulation, we sampled 50 random initializations of the parameters from a normal distribution with zero mean and standard deviation 0.5. Table 1 shows that NOHD significantly outperforms other algorithms even when starting from random initial probabilities.

In the second experiment, the gradients and the Jacobian are estimated from samples. The starting probabilities are the

⁸In Appendix we report results also for Dilemma game

⁹For readability, we show only a starting point. Appendix B contains the results for different starting points and more experiments.

same as in the previous experiment. We performed 20 runs for each setting. In each iteration, we sampled 300 trajectories of length 1. Figures 3 (right side) show that NOHD also in this experiment converges to the equilibrium in less than 100 iterations. Instead, the other algorithms exhibit oscillatory behaviors.

	NOHD	CGD	LOLA	IGA-PP	CO	SOS	SGA
MP	0.49	0.84	1.00	0.99	0.99	0.99	0.99
RPS	0.38	0.97	0.88	1.00	0.81	0.80	0.96

Table 1: Ratio between the mean convergence steps to Nash Equilibrium and the maximum mean convergence steps. 50 runs, sampling from a normal distribution $\mathcal{N}(0, 0.5^2)$.

$ \theta $	NOHD	IGAPP	LOLA	SOS	SGA	CGD	CO
4	0.7205	0.6979	0.6983	0.7186	0.7302	0.7265	0.7006
16	0.7898	0.7787	0.7735	0.7906	0.8358	0.8051	0.7758
36	1.1416	1.0625	1.0874	1.0705	1.1992	1.1444	1.1066
64	1.9486	1.6342	1.6162	1.6735	1.9555	1.8955	1.8850
100	3.4070	2.7734	2.6905	2.8169	3.4799	3.3126	3.2977
144	5.9191	4.6260	4.4351	4.8438	6.5164	5.8440	5.7876

Table 2: Computation time of one learning update of each algorithm with increasing parameter space size. 20 runs.

Continuous Gridworlds

This experiment aims at evaluating the performance of NOHD in two continuous Gridworld environments. The first gridworld is the continuous version of the second gridworld proposed in (Hu and Wellman 2003): the two agents are initialized in the two opposite lower corners and have to reach the same goal; when one of the two agents reaches

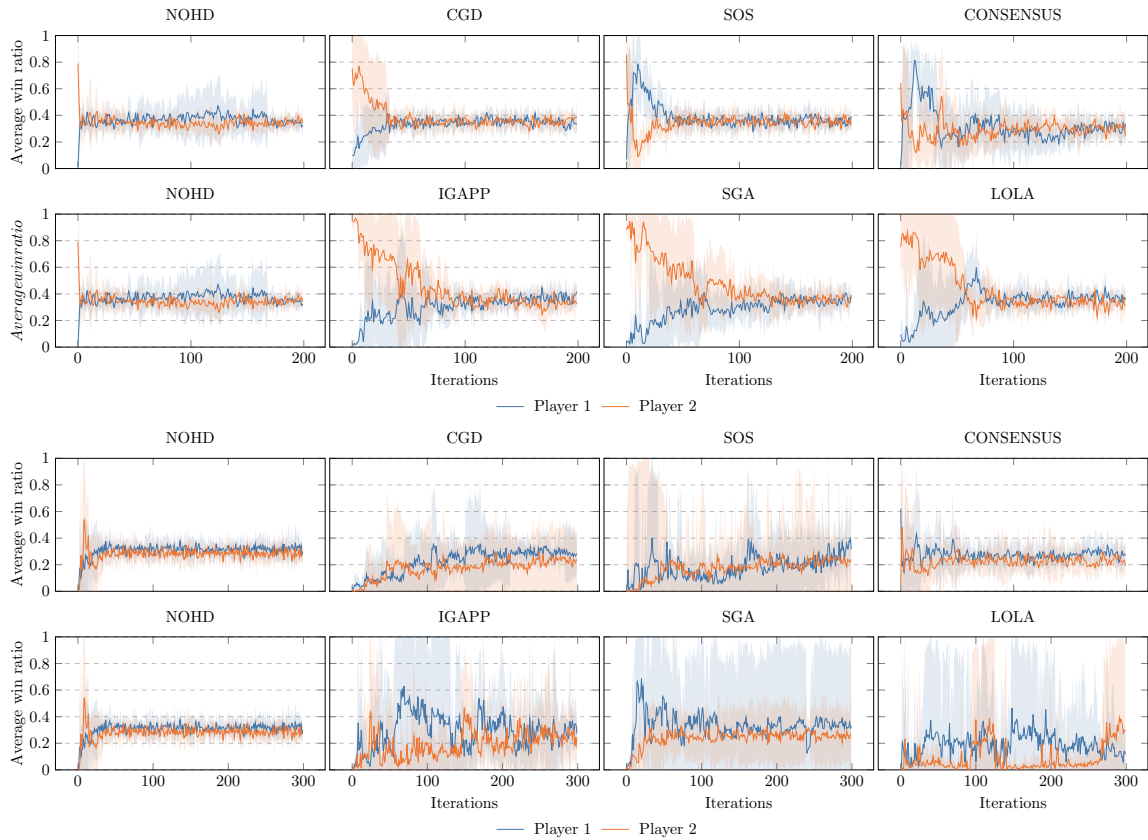


Figure 4: Average wins for player 1 and player 2 in two gridworld environments. 10 runs, c.i. 98%

the goal, the game ends, and this agent gets a positive reward. Each agent has to keep a distance of no less than 0.5 with the other agent and, if they decide to move to the same region, they cannot perform the action. Also, in the second gridworld the agents are initialized in the two lower corners and have to reach the same goal, but they have to reach the goal with a ball. An agent can take the opponent’s ball if their distance is less than 0.5; the ball is randomly given to one of the two agents at the beginning of each episode. The agents’ policies are Gaussian policies, linear in a set of respectively 72 and 68 radial basis functions, which generate the ν angle for the step’s direction.¹⁰ For each experiment, we perform 10 runs with random initialization. In Figure 4 we compare the performance of NOHD with CO, IGA-PP, LOLA, SOS, and CGD in the two gridworlds. The figure shows the mean average win of the two players at each learning iteration of the algorithms. In figure 4, at the top, we can see how NOHD outperforms the other algorithms in the first gridworld, converging in less than 30 steps to the equilibrium. At the bottom, results for the second gridworld are shown: NOHD converges quickly than the other algorithms, but Consensus has comparable performance.

¹⁰More information about the experiments are given in Appendix.

Computational Time

In Table 2 we report the computation time of an update of each algorithm with increasing policy parameter sizes $|\theta|$, from 4 to 144. As we can see, the computation time of NOHD is comparable to that of the other algorithms.

7 Conclusions

Although multi-agent reinforcement learning has achieved promising results in recent years, few algorithms consider the dynamics of the system. In this paper, we have shown how to apply Newton-based methods in the multi-agent setting. The paper’s first contribution is to propose a method to adapt Newton’s optimization to two simple game classes: Potential games and Hamiltonian games. Next, we propose a new algorithm NOHD that applies a Newton-based optimization to general games. The algorithm, such as SGA, SOS, and CGD, considers that agents can also act against their own interests to achieve a balance as quickly as possible. We also show that NOHD avoids unstable equilibria and is attracted to stable ones. We then show how the algorithm outperforms some baselines in matrix games with parametric Boltzmann policies. Furthermore, the algorithm manages to learn good policies in two continuous gridworld environments. In future work, we will try to extend the algorithm NOHD to settings where the agent cannot know the other agents’ policies and cost functions.

References

- Balduzzi, D.; Czarnecki, W. M.; Anthony, T. W.; Gemp, I. M.; Hughes, E.; Leibo, J. Z.; Piliouras, G.; and Graepel, T. 2020. Smooth markets: A basic mechanism for organizing gradient-based learners. *arXiv preprint arXiv:2001.04678* .
- Balduzzi, D.; Racaniere, S.; Martens, J.; Foerster, J.; Tuyls, K.; and Graepel, T. 2018. The mechanics of n-player differentiable games. *arXiv preprint arXiv:1802.05642* .
- Bowling, M.; and Veloso, M. 2002. Multiagent learning using a variable learning rate. *Artificial Intelligence* 136(2): 215–250.
- Brown, N.; and Sandholm, T. 2019. Superhuman AI for multiplayer poker. *Science* 365(6456): 885–890.
- Buşoniu, L.; Babuška, R.; and De Schutter, B. 2010. Multi-agent reinforcement learning: An overview. In *Innovations in multi-agent systems and applications-1*, 183–221. Springer.
- Chasnov, B.; Ratliff, L.; Mazumdar, E.; and Burden, S. 2020. Convergence analysis of gradient-based learning in continuous games. In *Uncertainty in Artificial Intelligence*, 935–944. PMLR.
- Chong, E. K.; and Zak, S. H. 2004. *An introduction to optimization*. John Wiley & Sons.
- Conitzer, V.; and Sandholm, T. 2007. AWESOME: A general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Machine Learning* 67(1-2): 23–43.
- Daskalakis, C.; Ilyas, A.; Syrgkanis, V.; and Zeng, H. 2017. Training gans with optimism. *arXiv preprint arXiv:1711.00141* .
- Facchinei, F.; and Kanzow, C. 2007. Generalized Nash equilibrium problems. *4or* 5(3): 173–210.
- Fiez, T.; Chasnov, B.; and Ratliff, L. J. 2019. Convergence of learning dynamics in stackelberg games. *arXiv preprint arXiv:1906.01217* .
- Foerster, J.; Chen, R. Y.; Al-Shedivat, M.; Whiteson, S.; Abbeel, P.; and Mordatch, I. 2018. Learning with opponent-learning awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 122–130. International Foundation for Autonomous Agents and Multiagent Systems.
- Gemp, I.; and Mahadevan, S. 2018. Global convergence to the equilibrium of gans using variational inequalities. *arXiv preprint arXiv:1808.01531* .
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, 6626–6637.
- Hu, J.; and Wellman, M. P. 2003. Nash Q-learning for general-sum stochastic games. *Journal of machine learning research* 4(Nov): 1039–1069.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.
- Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4681–4690.
- Lee, J. D.; Panageas, I.; Piliouras, G.; Simchowitz, M.; Jordan, M. I.; and Recht, B. 2017. First-order methods almost always avoid saddle points. *arXiv preprint arXiv:1710.07406* .
- Lee, J. D.; Simchowitz, M.; Jordan, M. I.; and Recht, B. 2016. Gradient descent converges to minimizers. *arXiv preprint arXiv:1602.04915* .
- Letcher, A.; Balduzzi, D.; Racaniere, S.; Martens, J.; Foerster, J. N.; Tuyls, K.; and Graepel, T. 2019. Differentiable Game Mechanics. *Journal of Machine Learning Research* 20(84): 1–40.
- Letcher, A.; Foerster, J.; Balduzzi, D.; Rocktäschel, T.; and Whiteson, S. 2018. Stable opponent shaping in differentiable games. *arXiv preprint arXiv:1811.08469* .
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971* .
- Mazumdar, E.; Ratliff, L. J.; and Sastry, S. S. 2020. On Gradient-Based Learning in Continuous Games. *SIAM Journal on Mathematics of Data Science* 2(1): 103–131.
- Mazumdar, E. V.; Jordan, M. I.; and Sastry, S. S. 2019. On finding local nash equilibria (and only local nash equilibria) in zero-sum games. *arXiv preprint arXiv:1901.00838* .
- Mertikopoulos, P.; Papadimitriou, C.; and Piliouras, G. 2018. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, 2703–2717. SIAM.
- Mescheder, L.; Nowozin, S.; and Geiger, A. 2017. The numerics of gans. In *Advances in Neural Information Processing Systems*, 1825–1835.
- Monderer, D.; and Shapley, L. S. 1996. Potential games. *Games and economic behavior* 14(1): 124–143.
- Nagarajan, V.; and Kolter, J. Z. 2017. Gradient descent GAN optimization is locally stable. In *Advances in neural information processing systems*, 5585–5595.
- Nash, J. F.; et al. 1950. Equilibrium points in n-person games. *Proceedings of the national academy of sciences* 36(1): 48–49.

- Nie, W.; Narodytska, N.; and Patel, A. 2018. Relgan: Relational generative adversarial networks for text generation .
- Nocedal, J.; and Wright, S. 2006. *Numerical optimization*. Springer Science & Business Media.
- Papadimitriou, C.; and Piliouras, G. 2016. From Nash equilibria to chain recurrent sets: Solution concepts and topology. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, 227–235.
- Paternain, S.; Mokhtari, A.; and Ribeiro, A. 2019. A Newton-based method for nonconvex optimization with fast evasion of saddle points. *SIAM Journal on Optimization* 29(1): 343–368.
- Perolat, J.; Munos, R.; Lespiau, J.-B.; Omidshafiei, S.; Rowland, M.; Ortega, P.; Burch, N.; Anthony, T.; Balduzzi, D.; De Vylder, B.; et al. 2020. From Poincaré Recurrence to Convergence in Imperfect Information Games: Finding Equilibrium via Regularization. *arXiv preprint arXiv:2002.08456* .
- Peters, J.; and Schaal, S. 2006. Policy gradient methods for robotics. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2219–2225. IEEE.
- Peters, J.; and Schaal, S. 2008. Reinforcement learning of motor skills with policy gradients. *Neural networks* 21(4): 682–697.
- Ratliff, L. J.; Burden, S. A.; and Sastry, S. S. 2013. Characterization and computation of local nash equilibria in continuous games. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 917–924. IEEE.
- Ratliff, L. J.; Burden, S. A.; and Sastry, S. S. 2016. On the characterization of local Nash equilibria in continuous games. *IEEE Transactions on Automatic Control* 61(8): 2301–2307.
- Rosen, J. B. 1965. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society* 520–534.
- Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2(1): 65–67.
- Schäfer, F.; and Anandkumar, A. 2019. Competitive gradient descent. In *Advances in Neural Information Processing Systems*, 7623–7633.
- Shalev-Shwartz, S.; Shammah, S.; and Shashua, A. 2016. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295* .
- Shen, Z.; Ribeiro, A.; Hassani, H.; Qian, H.; and Mi, C. 2019. Hessian Aided Policy Gradient. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 5729–5738. Long Beach, California, USA: PMLR. URL <http://proceedings.mlr.press/v97/shen19d.html>.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature* 529(7587): 484.
- Singh, S. P.; Kearns, M. J.; and Mansour, Y. 2000. Nash Convergence of Gradient Dynamics in General-Sum Games. In *UAI*, 541–548.
- Song, X.; Wang, T.; and Zhang, C. 2019. Convergence of multi-agent learning with a finite step size in general-sum games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 935–943. International Foundation for Autonomous Agents and Multiagent Systems.
- Sutton, R. S.; Barto, A. G.; et al. 1998. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge.
- Wills, A. P. 1958. Vector analysis with and introduction to tensor analysis. Technical report.
- Yu, L.; Zhang, W.; Wang, J.; and Yu, Y. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Zhang, C.; and Lesser, V. 2010. Multi-agent learning with policy prediction. In *Twenty-fourth AAAI conference on artificial intelligence*.
- Zhang, K.; Yang, Z.; and Başar, T. 2019. Multi-Agent Reinforcement Learning: A Selective Overview of Theories and Algorithms. *arXiv preprint arXiv:1911.10635* .
- Zhang, K.; Yang, Z.; and Basar, T. 2019. Policy optimization provably converges to Nash equilibria in zero-sum linear quadratic games. In *Advances in Neural Information Processing Systems*, 11602–11614.