SA-IGA: A Multiagent Reinforcement Learning Method Towards Socially Optimal Outcomes

Chengwei Zhang · **Xiaohong Li** · **Jianye Hao** ·
Siqi Chen · **Karl Tuyls** · **Wanli Xue**

Abstract In multiagent environments, the capability of learning is important for an agent to behave appropriately in face of unknown opponents and dynamic environment. From the system designer's perspective, it is desirable if the agents can learn to coordinate towards socially optimal outcomes, while also avoiding being exploited by selfish opponents. To this end, we propose a novel gradient ascent based algorithm (SA-IGA) which augments the basic gradient-ascent algorithm by incorporating social awareness into the policy update process. We theoretically analyze the learning dynamics of SA-IGA using dynamical system theory and SA-IGA is shown to have linear dynamics for a wide range of games including symmetric games. The learning dynamics of two representative games (the prisoner's dilemma game and the coordination game) are analyzed in details. Based on the idea of SA-IGA, we further propose a practical multiagent learning algorithm, called SA-PGA, based on Q-learning update rule. Simulation results show that SA-PGA agent can achieve higher social welfare than previous social-optimality oriented Conditional Joint Action Learner (CJAL) and also is robust against individually rational opponents by reaching Nash equilibrium solutions.

Chengwei Zhang, Xiaohong Li and Wanli Xue
School of Computer Science and Technology, Tianjin University
E-mail: chenvy,xiaohongli,xuewanli@tju.edu.cn

Jianye Hao, Corresponding author
School of Computer Software, Tianjin University
E-mail: jianye.hao@tju.edu.cn

Siqi Chen
School of Computer and Information Science, the Southwest University
E-mail: siqichen@swu.edu.cn

Karl Tuyls
University of Liverpool
E-mail: k.tuyls@liverpool.ac.uk

1 Introduction

In multiagent systems, the ability of learning is important for an agent to adaptively adjust its behaviors in response to coexisting agents and unknown environments in order to optimize its performance. Multiagent learning algorithms have received extensive investigation in the literature, and lots of learning strategies [6, 9, 18, 30] have been proposed to facilitate coordination among agents.

The multi-agent learning criteria proposed in [8] require that an agent should be able to converge to a stationary policy against some class of opponents (*convergence*) and the best-response policy against any stationary opponent (*rationality*). If both agents adopt a rational learning strategy in the context of repeated games and also their strategies converge, then they will converge to a Nash equilibrium of the stage game. Indeed, convergence to Nash equilibrium has been the most commonly accepted goal to pursue in multiagent learning literature. Until now, a number of gradient-ascent based multiagent learning algorithms [**?**, 1, 8, 23, 29] have been sequentially proposed towards converging to Nash equilibrium with improved convergence performance and more relaxed assumptions (less information is required). Under the same direction, another well-studied family of multiagent learning strategies is based on reinforcement learning (e.g., Q-learning [27]). Representative examples include distributed Q-learning in cooperative games [15], minimax Q-learning in zero-sum games [16], Nash Q-learning in general-sum games [14], and other extensions [9, 17], to name just a few.

1's payoff		Agent 2's actions	
2's payoff		C	D
Agent 1's actions	C	3/3	0/5
	D	5/0	1/1

Table 1 The Prisoner's Dilemma Game

All the aforementioned learning strategies pursue converging to Nash equilibrium under self-play, however, Nash equilibrium solution may be undesirable in many scenarios. One well-known example is the prisoner's dilemma (PD) game shown in Table 1. By converging to the Nash equilibrium (D, D), both agents obtain the payoff of 1, while they could have obtained a much higher payoff of 3 by coordinating on the non-equilibrium outcome (C, C). In situations like the PD game, converging to the socially optimal outcome, i.e., the maximal total reward of all players, under self-play would be more preferred. To address this issue, one natural modification for a gradient-ascent learner is to update its policy along the direction of maximizing the sum of all agents' expected payoff instead of its own. However, in an open environment, the agents are usually designed by different parties and may have not the incentive to follow the strategy we design. The above way of updating strategy would be easily exploited and taken advantage by (equilibrium-driven) self-interested agents. Thus it would be highly desirable if an agent can converge to socially optimal out-

comes under self-play and Nash equilibrium against self-interested agents to avoid being exploited.

In this paper, we propose a new gradient-ascent based algorithm (SA-IGA) which augments the basic gradient ascent algorithm by incorporating "social awareness" into the policy update process. Social awareness means that agents try to optimize social outcomes as well as its own outcome. A SA-IGA agent holds a social attitude to reflect its socially-aware degree, which can be adjusted adaptively based on the relative performance between its own and its opponent. A SA-IGA agent seeks to update its policy in the direction of increasing its overall payoff which is defined as the average of its individual and the social payoff weighted by its socially-aware degree. We theoretically show that for a wide range of games (e.g., symmetric games), the dynamics of SA-IGAs under self-play exhibits linear characteristics. For general-sum games, it may exhibit non-linear dynamics which can still be analyzed numerically. The learning dynamics of two representative games (the prisoner's dilemma game and the coordination game representing symmetric games and asymmetric games, respectively) are analyzed in details. Like previous theoretical multiagent learning algorithms, SA-IGA also requires additional assumption of knowing the opponent's policy and the game structure.

To relax the above assumption, we then propose a practical gradient ascent based multiagent learning strategy, called Socially-aware Policy Gradient Ascent (SA-PGA). SA-PGA relaxes the above assumptions by estimating the performance of its own and the opponent using Q-learning techniques. We empirically evaluate its performance in different types of benchmark games and simulation results show that SA-PGA agent outperforms previous learning strategies in terms of maximizing the social welfare and Nash product of the agents. Besides, SA-PGA is also shown to be robust against individually rational opponents and converges to Nash equilibrium solutions.

The remainder of the paper is organized as follows. Section 2 generally reviews some related works about Gradient Ascent Reinforcement Learning algorithms. Section 3 reviews normal-form game and the basic gradient ascent approach. Section 4 introduces the SA-IGA algorithm and analyzes its learning dynamics theoretically. Section 5 presents the practical multiagent learning algorithm SA-PGA in details. In Section 6, we extensively evaluate the performance of SA-PGA under various benchmark games. Lastly we conclude the paper and point out future directions in Section 7.

2 Related Works

The first gradient ascent multiagent reinforcement learning algorithm is Infinitesimal Gradient Ascent (IGA [23]), in which each learner updates its policy towards the gradient direction of its expected payoff. The purpose of IGA is to promote agents to converge to a particular Nash Equilibrium in a two-player two-action normal-form game. IGA has been proved that agents will converge to Nash equilibrium or if the strategies themselves do not converge, then their average payoffs will nevertheless converge to the average payoffs of Nash equilibrium. Soon after, M. Zinkevich et al.

[31] propose an algorithm called Generalized Infinitesimal Gradient Ascent(GIGA), which extends IGA to the game with an arbitrary number of actions.

Both IGA and GIGA can be combined with the Win or Learn Fast (WoLF) heuristic in order to improve performance in stochastic games (Wolf-IGA [8], Wolf-GIGA [7]). The intuition behind WoLF principle is that an agent should adapt quickly when it performs worse than expected, whereas it should maintain the current strategy when it receives payoff better than the expected one. By altering the learning rate according to the WoLF principle, a rational algorithm can be made convergent. The shortage of WoLF-IGA or WoLF-GIGA is that these two algorithms require a reference policy, i.e., they require the estimation of Nash equilibrium strategies and corresponding payoffs. To this end, Banerjee et al [4] propose an alternative criterion of WoLF-IGA, named Policy Dynamics based WoLF(PDWoLF), that can be accurately computed and guarantees convergence. The Weighted Policy Learner (WPL [1]) is another variation of IGA that also modulates the learning rate, meanwhile, it does not require a reference policy. Both of the WoLF and WPL are designed to guarantee convergence in stochastic repeated games.

Another direction for extending IGA is making improvements from the learning value functions. Zhang et al [29] propose a gradient-based learning algorithm by adjusting the expected payoff function of IGA, named Gradient Ascent with Policy Prediction Algorithm(IGA-PP). The algorithm is designed for games with two agents. The key idea behind this algorithm is that a player adjusts its strategy in response to forecasted strategies of the other player, instead of its current ones. It has been proved that, in two-player, two-action, general-sum matrix games, IGA-PP in self-play or against IGA would lead players' strategies to converge to a Nash equilibrium. Like other MARL algorithms, besides the common assumption, this algorithm also has additional requirements that a player knows the other players strategy and current strategy gradient (or payoff matrix) so that it can forecast the other players strategy.

All the aforementioned learning strategies pursue converging to Nash equilibriums. In contrast, in this work, we seek to incorporate the social awareness into GA-based strategy update and aim at improving the social welfare of the players under self-play rather than pursuing Nash equilibrium solutions. Meanwhile, individually rational behavior is employed when playing against a selfish agent. Similar idea of adaptively behaving differently against different opponents was also employed in previous algorithms [10,12,17,19]. However, all the existing works focus on maximizing an agent's individual payoff against different opponents in different types of games, but do not directly take into consideration the goal of maximizing social welfare (e.g., cooperate in the prisoner's dilemma game).

3 Background

In this section we introduce the necessary background for our contribution. First, we gave an overview of the relevant game theory definition. Then a brief review of gradient ascent based MARL (GA-MARL) algorithm is given.

3.1 Game theory

Game theory provides a framework for modeling agents' interaction, which was used by previous researchers in order to analyze the convergence properties of MARL algorithms [1, 8, 23, 29]. A game specifies, in a compact and simple manner, how the payoff of an agent depends on other agents actions. A (normal form) game is defined by the tuple $< N, A_1, ..., A_N, R_1, ..., R_N >$, where N is the number of players in the game, A_i is the set of actions available to agent i, and $R_i : A_1 \times ... \times A_N \rightarrow \mathbb{R}$ is the reward (payoff) of agent i which is defined as a function of the joint action executed by all agents. If the game has only two agents, then it is convenient to define their reward functions as a payoff matrix as follows,

$$R_i = \{r_i^{jk}\}_{|A_1| \times |A_2|}$$

where $i \in \{1, 2\}$, $j \in A_j$ and $k \in A_k$. Each element r_i^{jk} in the matrix represents the payoff received by agent i, if agent i plays action j and its opponent plays action k.

A *policy* (or a *strategy*) of an agent i is denoted by $\pi_i : A_i \rightarrow [0, 1]$, which maps its actions to a probability. The probability of choosing an action k according to policy π_i is $\pi_i(k)$. A policy is deterministic or pure if the probability of playing one action is 1 while the probability of playing other actions is 0, (i.e. $\exists \pi_i(k) = 1$ AND $\forall l \neq k, \pi_i(l) = 0$), otherwise the policy is stochastic or mixed. The joint policy of all agents is the collection of individual agents' policies, which is defined as $\pi = < \pi_1, ..., \pi_N >$. For continence, the joint policy is usually expressed as $\pi = < \pi_i, \pi_{-i} >$, where π_{-i} is the collection of all policies of agents other than agent i.

The *expected payoff* of an agent is defined as the reward averaged over the joint policy. Let $A_{-i} = \{< a_1, ..., a_N >: a_j \in A_j \wedge i \neq j\}$, if agents follow a joint policy π, then the *expected payoff* of agent i would be, $V_i(\pi) = \sum_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} \pi_i(a_i) \pi_{-i}(a_{-i}) R_i(a_i, a_{-i})$, where $R_i(a_i, a_{-i}) = r_i^{a_i a_{-i}}$.

The goal of each agent is to find such a policy that maximizes the players expected payoff. Ideally, we want all agents to reach the equilibrium that maximizes their individual payoffs. However, when agents do not communicate and/or agents are not cooperative, reaching a globally optimal equilibrium is not always attainable. An alternative goal is converging to the Nash Equilibrium (NE), which is by definition a local maximum across agents. A joint strategy is called a *Nash Equilibrium* (NE), if no player can get a better expected payoff by changing its current strategy unilaterally. Formally, $\pi^* = (\pi_i^*, \pi_{-i}^*)$ is a NE, iff $\forall i \forall \pi_i: V_i(\pi_i^*, \pi_{-i}^*) \geq V_i(\pi_i, \pi_{-i}^*)$. An NE is pure if all its constituting policies are pure. Otherwise the NE is called mixed or stochastic. Any game has at least one Nash equilibrium, but may not have any pure equilibrium.

Next subsection, we introduce the Gradient Ascent based MARL algorithm (GA-MARL), together with a brief review of the dynamic analysis of GA-MARL.

3.2 Gradient Ascent (GA) MARL Algorithms

Gradient ascent MARL algorithms (GA-MARL) learn a stochastic policy by directly following the expected reward gradient. The ability to learn a stochastic policy is

particularly important when the world is not fully observable or has a competitive nature. The basic GA-MARL algorithm whose dynamics were analyzed is the Infinitesimal Gradient Ascent(IGA [23]) . When a game is repeatedly played, an IGA player updates its strategy towards maximizing its expected payoffs. A player i employing GA-based algorithms updates its policy towards the direction of its expected reward gradient, as illustrated by the following equations,

$$\Delta \pi_i^{(t+1)} \leftarrow \alpha \frac{\partial V_i \left(\pi^{(t)} \right)}{\partial \pi_i} \tag{1}$$

$$\pi_i^{(t+1)} \leftarrow \Pi_{[0,1]} \left(\pi_i^{(t)} + \Delta \pi_i^{(t+1)} \right) \tag{2}$$

where parameter α is the gradient step size, and $\Pi_{[0,1]}$ is the projection function mapping the input value to the valid probability range of $[0,1]$, used to prevent the gradient moving the strategy out of the valid probability space. Formally, we have,

$$\Pi_{[0,1]} \left(x \right) = argmin_{z \in [0,1]} \left| x - z \right| \tag{3}$$

Singh, Kearns, and Mansour [23] examined the dynamics of using gradient ascent in two-player, two-action, iterated matrix games. We can represent this problem as two matrices,

$$R_i = \begin{bmatrix} r_i^{11} & r_i^{12} \\ r_i^{21} & r_i^{22} \end{bmatrix}, i \in \{1,2\}$$

We refer to the joint policy of the two players at time t by the probabilities of choosing the first action (p_1^t, p_2^t), where $\pi_i = (p_i^t, 1 - p_i^t)$, $i \in \{1,2\}$ is the policy of player i. The t notation will be omitted when it does not affect clarity (for example, when we are considering only one point in time). Then, for the two-player two-action case, the above way of GA-based updating in Equations 1 and 2 can be simplified as follows,

$$p_i^{(t+1)} \leftarrow \Pi_{[0,1]} \left(p_i^{(t)} + \alpha \left(u_i p_{-i}^{(t)} + c_i \right) \right) \tag{4}$$

where $u_i = r_i^{11} + r_i^{22} - r_i^{12} - r_i^{21}$, $c_i = r_i^{12} - r_i^{22}$.

In the case of infinitesimal gradient step size ($\eta \rightarrow 0$), the learning dynamics of the players can be modeled as a system of differential equations, i.e. $\dot{p}_i = u_i p_{-i} + c_i$, $i \in \{1,2\}$, which can be analyzed using dynamic system theory [11]. It is proved that the agents will converge to a Nash equilibrium, or if the strategies themselves do not converge, then their average payoffs will nevertheless converge to the average payoffs of a Nash equilibrium [23].

Combined with Q-learning [26], researchers propose a practical learning algorithm, i.e. the policy hill-climbing algorithm (PHC) [8], which is a simple extension of IGA and is shown in Table 1.

The algorithm performs hill-climbing in the space of mixed policies, which is similar to gradient ascent, but does not require as much knowledge. Q values are maintained just as in normal Q-learning. In addition the algorithm maintains the current mixed policy. The policy is improved by increasing the probability that it selects the highest valued action according to a learning rate $\alpha \in (0,1]$. After that, the policy is mapped back to the valid probability space. This technique, like Q-learning, is

---

Algorithm 1 PHC for player i

1: Lets $\alpha, \beta \in (0, 1)$ be learning rates.
2: Initialize,
  $Q_i(a) \leftarrow 0, \pi_i(a) \leftarrow \frac{1}{|A_i|}$.
3: **repeat**
4:   Select action $a \in A_i$ according to mixed strategy π_i with suitable exploration.
5:   Observing reward r. Update Q,
    $Q_i(a) \leftarrow (1 - \beta) Q_i(a) + \beta r$.
6:   Update π_i according to gradient ascent strategy,
    $\pi_i(a) \leftarrow \Pi_{[0,1]}[\pi_i(a) - \alpha]$, if $a \neq \underset{a' \in A}{argmax}\, Q(a')$,
    $\pi_i(a) \leftarrow 1 - \underset{a' \neq a}{\sum} \pi(a')$, if $a = \underset{a' \in A}{argmax}\, Q(a')$.
7: **until** the repeated game ends

---

rational and will converge to an optimal policy if other players are playing stationary strategies. The algorithm guarantees the Q values will converge to Q^* (the local optimal value of Q) with a suitable exploration policy. π will converge to a policy that is greedy according to Q, which is converging to Q^*, and therefore will converge to a best response. PHC is rational and has no limit on the number of agents and actions.

4 Socially-aware Infinitesimal Gradient Ascent (SA-IGA)

In our daily life, people usually do not always behave as a purely individually rational entity and seek to achieve Nash equilibrium solutions. For example, when two person subjects play a PD game, reaching mutual cooperation may be observed frequently. Similar phenomena have also been observed in extensive human-subject based experiments in games such as the Public Goods game [13] and Ultimatum game [2], in which human subjects are usually found to obtain much higher payoff by mutual cooperation rather than pursuing Nash equilibrium solutions. If the above phenomenon is transformed into computational models, it indicates that an agent may not only update its policy in the direction of maximizing its own payoff, but also take into consideration other's payoff. We call this type of agents as socially-aware agents.

In this paper, we incorporate the social awareness into the gradient-ascent based learning algorithm. In this way, apart from learning to maximizing its individual payoff, an agent is also equipped with the social awareness so that it can (1) reach mutually cooperative solutions faced with other socially-aware agents (self-play); (2) behave in a purely individually rational manner when others are purely rational.

Specifically, for each SA-IGA agent i, it distinguishes two types of expected payoffs, namely V_i^{idv} and V_i^{soc}. Payoffs $V_i^{\text{idv}}(\pi)$ and $V_i^{\text{soc}}(\pi)$ represent the individual and social payoff (the average payoff of all agents) that agent i perceives under the joint strategy π respectively. The payoff $V_i^{\text{idv}}(\pi)$ follows the same definition as IGA and the payoff $V_i^{\text{soc}}(\pi)$ is defined as the average of the individual payoffs of all agents.

$$V_i^{\text{soc}}(\pi) = \frac{1}{N} \sum_i V_i^{\text{idv}}(\pi), \tag{5}$$

Each agent i adopts a social attitude $w_i \in [0, 1]$ to reflect its socially-aware degree. The social attitude intuitively models an agent's social friendliness degree towards others. Specifically, it is used as the weighting factor to adjust the relative importance between V_i^{idv} and V_i^{soc}, and agent i's overall expected payoff is defined as follows,

$$V(\pi) = (1 - w_i) V_i^{\text{idv}}(\pi) + w_i V_i^{\text{soc}}(\pi) \tag{6}$$

Each agent i updates its strategy in the direction of maximizing the value of V_i. Formally we have,

$$\Delta \pi_i \leftarrow \alpha_\pi \frac{\partial V_i(\pi)}{\partial \pi_i}, \pi_i \leftarrow \Pi_{[0,1]}(\pi_i + \Delta \pi_i) \tag{7}$$

where parameter α_π is the gradient step size of π_i. If $w_i = 0$, it means that the agent seeks to maximize its individual payoff only, which is reduced to the case of traditional gradient-ascent updating; if $w_i = 1$, it means that the agent seeks to maximize the sum of the payoffs of both players.

Finally, each agent i's socially-aware degree is adaptively adjusted in response to the relative value of V_i^{idv} and V_i^{soc} as follows. During each round, if player i's own expected payoff V_i^{idv} exceeds the value of V_i^{soc}, then player i increases its social attitude w_i, (i.e., it becomes more social-friendly because it perceives itself to be earning more than the average). Conversely, if V_i^{idv} is less than V_i^{soc}, then the agent tends to care more about its own interest by decreasing the value of w_i. Formally,

$$w_i^{t+1} \leftarrow \Pi_{[0,1]}\left(w_i^t + \alpha_w\left(V_i^{idv} - V_i^{soc}\right)\right) \tag{8}$$

where parameter α_w is the learning rate of w_i.

4.1 Theoretical Modeling and Analysis of SA-IGA

An important aspect of understanding the behavior of a multiagent learning algorithm is theoretically modeling and analyzing its underlying dynamics [6, 20, 25]. In this section, we first show that the learning dynamics of SA-IGA under self-play can be modeled as a system of differential equations. To simplify analysis, we only considered two-player-two-action games.

Based on the adjustment rules in Equation (7) and (8), the learning dynamics of a SA-IGA agent can be modeled as a set of equations in (9). For ease of exposition, we concentrate on an unconstrained update equations by removing the policy projection function which does not affect our qualitative analytical results. Any trajectory with linear (non-linear) characteristic without constraints is still linear (non-linear) when a boundary is enforced.

$$\begin{aligned}
\Delta \pi_i^{(t+1)} &\leftarrow \alpha_\pi \frac{\partial V_i\left(\pi^{(t)}\right)}{\partial \pi_i} \\
\Delta w_i^{t+1} &\leftarrow \alpha_w(V_i^{\text{idv}} - V_i^{\text{soc}}) \\
\pi_i^{(t+1)} &\leftarrow \pi_i^{(t)} + \Delta \pi_i^{(t+1)} \\
w_i^{(t+1)} &\leftarrow w_i^{(t)} + \Delta w_i^{(t+1)}
\end{aligned} \tag{9}$$

Substituting V_i^{idv} and V_i^{soc} by their definitions (Equations 4 and 5), the learning dynamics of two SA-IGA agents can be expressed as follows,

$$\Delta p_i^{t+1} = \alpha_p \cdot \left[\left(u_i + \frac{u_{-i} - u_i}{2} w_i^t \right) p_{-i}^t + \frac{d_{-i} - c_i}{2} w_i^t + c_i \right] \tag{10}$$
$$\Delta w_i^{t+1} = \alpha_w \cdot \left[(u_i - u_{-i}) p_i^t p_{-i}^t + (c_i - d_{-i}) p_i^t + (c_{-i} - d_i) p_{-i}^t + e_i \right]$$

where $u_i = r_i^{11} + r_i^{22} - r_i^{12} - r_i^{21}, c_i = r_i^{12} - r_i^{22}, d_i = r_i^{21} - r_i^{22}$, and $e_i = r_i^{22} - r_{-i}^{22}$ with $i \in \{1, 2\}$.

As $\alpha_\pi \to 0$ and $\alpha_w \to 0$, it is straightforward to show that the above equations become differential. Thus the unconstrained dynamics of the strategy pair and social attitudes as a function of time is modeled by the following system of differential equations:

$$\dot{p}_i = \left(u_i + \frac{u_{-i} - u_i}{2} w_i \right) p_{-i} + \frac{d_{-i} - c_i}{2} w_i + c_i \tag{11}$$
$$\dot{w}_i = \varepsilon \cdot \left[(u_i - u_{-i}) p_i p_{-i} + (c_i - d_{-i}) p_i + (c_{-i} - d_i) p_{-i} + e_i \right]$$

where $\varepsilon = \frac{\alpha_w}{\alpha_p} > 0$.

Based on the above theoretical modeling, next we analyze the learning dynamics of SA-IGA qualitatively as follows.

Theorem 1 *SA-IGA has non-linear dynamics when $u_1 \neq u_2$.*

Proof : From differential equations in (11), it is straightforward to verify that the dynamics of SA-IGA learners are non-linear when $u_1 \neq u_2$ due to the existence of $w_1 p_2$, $w_2 p_1$ and $p_1 p_2$ in all equations.

Since SA-IGA's dynamics are non-linear when $u_1 \neq u_2$, in general we cannot obtain a closed-form solution, but we can still resort to solve the equations numerically to obtain useful insight of the system's dynamics. Moreover, a wide range of important games fall into the category of $u_1 = u_2$, in which the system of equations become linear. Therefore, it allows us to use dynamic system theory to systematically analyze the underlying dynamics of SA-IGA.

Theorem 2 *SA-IGA has linear dynamics when the game itself is symmetric.*

Proof : A two-player two-action symmetric game can be represented in Table 2 in general. It is obvious to check that it satisfies the constraint of $u_1 = u_2$, given that $u_i = r_i^{11} + r_i^{22} - r_i^{12} - r_i^{21}$, $i \in \{1, 2\}$. Thus the theorem holds.

4.2 Dynamics Analysis of SA-IGA

Previous section mainly analyzed the dynamics of SA-IGA in a qualitative manner. In this section, we move to provide detailed analysis of SA-IGA's learning dynamics. We first summarize a generalized conclusion for symmetric games, and then analysis symmetric circumstances in two representative games: the Prisoner's Dilemma game

1's payoff 2's payoff		Agent 2's actions	
		action 1	action 2
Agent 1's actions	action 1	a/a	b/c
	action 2	c/b	d/d

Table 2 The General Form of a Symmetric Game

and the Symmetric Coordination game. For asymmetric circumstances, because the complexity of nonlinear problem analysis, we only focus on the general coordination game (Table 3). Specifically we analyze the SA-IGA's learning dynamics of those games by identifying the existing equilibrium points, which provides useful insights into understanding of SA-IGA's dynamics.

For symmetric games, we have the following conclusion,

Theorem 3 *The dynamics of SA-IGA algorithm under self-play under a symmetric game have three types of equilibrium points:*

1. $\left\{ (0, 0, w_1^*, w_2^*) \left| \frac{c-b}{2} w_i^* + b - d < 0, w_i^* \in [0,1] \right. \right\}$;
   $\left\{ (1, 1, w_1^*, w_2^*) \left| \frac{c-b}{2} w_i^* + a - c > 0, w_i^* \in [0,1] \right. \right\}$;
2. $\left\{ (1, 0, 0, 1), (0, 1, 1, 0) \right\}$, *if* $c > b > d \wedge b + c > 2a$;
   $\left\{ (1, 0, 1, 0), (0, 1, 0, 1) \right\}$, *if* $b > c > a \wedge b + c > 2d$;
3. $\left\{ (p^*, p^*, w^*, w^*) \left| p^* = \frac{b-c}{2u} w^* + \frac{d-b}{u}, p^*, w^* \in [0,1] \right. \right\}$,

   where $u = a + d - b - c$. *The first and second types of equilibrium points are stable, while the last is not. We say an equilibrium point is stable if once the strategy starts "close enough" to the equilibrium (within a distance δ from it), it will remain "close enough" to the equilibrium point forever.*

Proof : Following the system of differential equations in Equations (11), we can express the dynamics of SA-IGA in Symmetric game as follows:

$$\dot{p}_i = u p_{-i} + \frac{c - b}{2} w_i + b - d$$
$$\dot{w}_i = \varepsilon (b - c) (p_i - p_{-i}) \tag{12}$$

where $\varepsilon = \frac{\eta_w}{\eta_p} > 0, u = a + d - b - c, i \in \{1, 2\}.$

We start with proving the last type of equilibrium points: If there exist an equilibrium $eq = (p_1^*, p_2^*, w_1^*, w_2^*)^T \in (0,1)^4$, *then we have* $\dot{p}_i(eq) = 0$ *and* $\dot{w}_i(eq) = 0$, $i \in \{1, 2\}$. *By solving the above equations, we have* $p_1^* = p_2^* = \frac{b-c}{2u} w^* + \frac{d-b}{u}$ *and* $w^* = w_1^* = w_2^*$. *Since* $p_1^*, p_2^* \in (0, 1)$, *then we have,*

$$0 < \frac{b - c}{2u} w^* + \frac{d - b}{u} < 1$$

Then $eq = (p_1^*, p_2^*, w_1^*, w_2^*)^T$ *is an equilibrium. The stability of* eq *can be verified using theories of non-linear dynamics [21]. By expressing the unconstrained update*

differential equations in the form of $\dot{x} = Ax + B$, we have

$$A = \begin{bmatrix} 0 & u & c-b & 0 \\ u & 0 & 0 & c-b \\ \varepsilon(b-c) & \varepsilon(c-b) & 0 & 0 \\ \varepsilon(c-b) & \varepsilon(b-c) & 0 & 0 \end{bmatrix}$$

After calculating matrix A's eigenvalue, then we have $\lambda_1 = 0$, $\lambda_2 = u$, $\lambda_3 = -\frac{u}{2} + k$ and $\lambda_4 = -\frac{u}{2} - k$, where k is a constant. Since there exist an eigenvalue $\lambda > 0$, the equilibrium eq is not stable.

Next we turn to consider cases that equilibriums are in the boundary. In these cases, we need to put the projection function back. If $p_i = 0, i \in \{1, 2\}$, according to the known conditions, we have $w\frac{c-b}{2}w_i^* < d - b$. Combined with the unconstrained update differential equations 12, we have $\lim_{t\to\infty} \dot{p}_i < 0$, then p_i remains unchanged. And because $p_1 = p_2 = 0$, then for $\forall w_i \in [0, 1]$, $\dot{w}_i = 0$, then $((0, 0, w_1^*, w_2^*))$ is an equilibrium.

Because $w\frac{c-b}{2}w_i^* < d-b$, there exist a $\delta > 0$, and a set $U(eq, \delta) = \left\{ x \in [0, 1]^4 \mid |x - eq| < \delta \right\}$, that for $\forall x \in U(eq, \delta)$, $\lim_{p_i} \dot{p}_i < 0$. Thus p will stabilize on the point of 0. Also, as $\lim_{t\to 0} \dot{w}_i = (b - c)\lim_{t\to 0}(p_1 - p_2) = (b - c)\lim_{t\to 0}(0 - 0) = 0$, w also stable, and thus the equilibrium eq is stable.

The case that $p_i = 1, i \in \{1, 2\}$ can be proved similarly, which is omitted here.

For the case $p_1 = 1 \wedge p_2 = 0$, if $((1, 0, w_1^*, w_2^*))$ is an equilibrium, combined with the unconstrained update differential equations 12, we have $\dot{w}_1 = -\dot{w}_2$, which means that w_i will keeps changing until $w_1 = 1 \wedge w_2 = 0$ or $w_1 = 0 \wedge w_2 = 1$. If $((1, 0, 0, 1))$ is an equilibrium, then $\dot{p}_1 > 0 \wedge \dot{p}_2 < 0$ and $\dot{w}_1 < 0 \wedge \dot{w}_2 > 0$. Take into Equations 12, we get $c > b > d \wedge b + c > 2a$. Other case are the same, thus we it omit here.

The stability of the second type of equilibriums can be proved by the way as the first type one, which is omitted here.

From Theorem 3, we know that there are three types of equilibriums if both players play SA-IGA policy, while only the first and second types of equilibrium points are stable. Besides, all equilibriums of the first two types are pure strategies, i.e., the probability p_i for selecting action 1 for agent $i \in \{1, 2\}$ equals to 1 or 0. Notably, the range of w (the social attitude) in these three types of equilibriums may be overlapped, resulting in that the final convergence of the algorithm also depends on the value of p. Next we concentrate on details of two representative symmetric games: the Prisoner's Dilemma (PD) game and the Symmetric Coordination game.

The Prisoner's Dilemma (PD) game is a symmetric game whose parameters meet the conditions: $c > a > d > b$. Combined with Theorem 3, we have the following conclusion,

Corollary 1 *The dynamics of SA-IGA algorithm under Prisoner's Dilemma (PD) game have two types of stable equilibrium points:*

1. $(0, 0, w_1^*, w_2^*)$, if $w_1^*, w_2^* < min\left\{\frac{2(c-a)}{c-b}, \frac{2(d-b)}{c-b}\right\}$;

2. $(1, 1, w_1^*, w_2^*)$, if $w_1^*, w_2^* > max\left\{\frac{2(c-a)}{c-b}, \frac{2(d-b)}{c-b}\right\}$;

Proof : Because the PD game is a symmetric game, we can use conclusions of Theorem 3 directly. From Theorem 3, we can see that the PD game have two types of stable equilibrium points:

1. $\left\{ (0,0,w_1^*,w_2^*) \left| \frac{c-b}{2}w_i^* + b - d < 0, w_i^* \in [0,1] \right. \right\}$;
   $\left\{ (1,1,w_1^*,w_2^*) \left| \frac{c-b}{2}w_i^* + a - c > 0, w_i^* \in [0,1] \right. \right\}$;
2. $\left\{ (1,0,0,1), (0,1,1,0) \right\}$, *if* $c > b > d \wedge b + c > 2a$;
   $\left\{ (1,0,1,0), (0,1,0,1) \right\}$, *if* $b > c > a \wedge b + c > 2d$;

For the first type of equilibrium, take $c > a > d > b$ into conditions in above formulas, we have: if $w_1^*, w_2^* < min \left\{ \frac{2(c-a)}{c-b}, \frac{2(d-b)}{c-b} \right\}$, then $(0,0,w_1^*,w_2^*)$ is an stable equilibrium; else if $w_1^*, w_2^* > max \left\{ \frac{2(c-a)}{c-b}, \frac{2(d-b)}{c-b} \right\}$, then $(0,0,w_1^*,w_2^*)$ is an stable equilibrium.

For the second type of equilibrium, take $c > a > d > b$ into consideration, we found that the conditions are in conflict with each other, which means there is no such type of equilibriums under Prisoner's Dilemma (PD) game.

Intuitively, for a PD game, from Corollary 1, we know that if both SA-IGA players are initially sufficiently social-friendly (the value of w is large than a certain threshold), then they will always converge to mutual cooperation of (C, C). In other words, given that the value of w exceeds certain threshold, the strategy point of $(1, 1)$ (or (C, C)) in the strategy space is asymptotically stable. If both players start with a low socially-aware degree (w is smaller than certain threshold), then they will always converge to mutual defection of (D, D) eventually. For the rest of cases, there exist infinite number of equilibrium points in-between the above two extreme cases, all of which are not stable, which means that the learning dynamic will never converge to those equilibrium points.

Next we turn to analyze the dynamics of SA-IGA playing the Symmetric Coordination game. The general form of a Coordination game is shown in Table 3. From the table, we can see that the Coordination game is asymmetric if any of the following conditions are met: $R \neq r$, $P \neq p$, $T \neq t$ or $S \neq s$. we analyze a simplified game first, i.e., the Symmetric Coordination game, the general circumstance of coordination game will be analyzed later. Similar to the analysis of Theorem 1, we have,

1's payoff 2's payoff		Agent 2's actions	
		1	2
Agent 1's actions	1	R/r	S/t
	2	T/s	P/p

Table 3 The General Form of a Coordination Game (where $R > T \wedge P > S$ and $r > t \wedge p > s$)

Corollary 2 *The dynamics of SA-IGA algorithm under a symmetric coordination game have two types of stable equilibrium points:*

1. $(1, 1, w_1^*, w_2^*)$, with $\frac{T-S}{2} w_1^* > T - R$;
2. $(0, 0, w_1^*, w_2^*)$, with $\frac{T-S}{2} w_i^* < P - S$;

where $u = R + P - S - T$.

Proof : The proof is the same with Theorem 1, thus we omit it here.

Intuitively, for a Symmetric Coordination game, from Corollary 2, there are two types of stable equilibrium if players playing SA-IGA policy, which means players will eventually converging to action $(1, 1)$ or $(0, 0)$, i.e., the Nash equilibriums of the Symmetric Coordination game. Besides, because the final convergence of the algorithm depends on the combined effect of p and w, we cannot give a theoretical conclusion about the condition under which the algorithm will converge to the social optimal for a symmetric Coordination game. In fact, experimental simulations in the following section show that the SA-IGA has a higher probability converging to social optimal.

Now we turn to consider the asymmetric case. As we mentioned before, SA-IGA under an asymmetric game may have nonlinear dynamics when $u_1 \neq u_2$, which has caused great difficulties for theoretical analysis. For this reason, we only analyze the general Coordination game which is a typical asymmetric game.

Theorem 4 *The dynamics of SA-IGA algorithm under a general coordination game have three types of equilibrium points:*

1. $(0, 0, w_1^*, w_2^*)$, with $w_1^* = 1 \wedge w_2^* = 0$ when $P > p > t$; $w_1^* = 0 \wedge w_2^* = 1$ when $T < P < p$; and $\left(\frac{t-S}{2} w_1^* < P - S\right) \wedge \left(\frac{T-s}{2} w_2^* < p - s\right)$ when $P = p$;
2. $(1, 1, w_1^*, w_2^*)$, with $w_1^* = 1 \wedge w_2^* = 0$ when $R > r > s$; $w_1^* = 0 \wedge w_2^* = 1$ when $T < R < r$; and $\left(\frac{T-s}{2} w_1^* < R - T\right) \wedge \left(\frac{S-t}{2} w_2^* < r - t\right)$ when $R = r$;
3. $(p_1^*, p_2^*, w_1^*, w_2^*)$, others.

The first and second types of equilibrium points are stable, while the last non-boundary equilibrium points is not.

Proof : Following the system of differential equations in Equations (11), we can express the dynamics of SA-IGA in coordination game as follows:

$$\dot{p}_1 = \left(u_1 + \frac{u_2 - u_1}{2} w_1\right) p_2 + \frac{d_2 - c_1}{2} w_1 + c_1$$

$$\dot{p}_2 = \left(u_2 + \frac{u_1 - u_2}{2} w_2\right) p_1 + \frac{d_1 - c_2}{2} w_2 + c_2 \qquad (13)$$

$$\dot{w}_1 = \varepsilon \cdot \left[(u_1 - u_2) p_1 p_1 + (c_1 - c_2) p_1 + (d_2 - d_1) p_2 + e_1\right]$$

$$\dot{w}_2 = -\dot{w}_1$$

where $\varepsilon = \frac{\eta_w}{\eta_p} > 0$, $u_1 = R + P - S - T > 0$, $u_2 = r + p - s - t > 0$, $c_1 = S - P$, $c_2 = s - p$, $d_1 = T - P$, $d_2 = t - p$, and $e_1 = P - p$. We can see that the dynamic of coordination game is nonlinear when $u_1 \neq u_2$. We start with proving the last type of equilibrium points first:

If there exit a equilibrium $eq = (p_1^*, p_2^*, w_1^*, w_2^*)^T \in (0, 1)^4$, then there have $\dot{p}_i(eq) = 0$ and $\dot{w}_i(eq) = 0$, $i \in \{1, 2\}$. By linearizing the unconstrained update

differential equations into the form of $\dot{x} = Ax + B$ in point $eq = (p_1^*, p_2^*, w_1^*, w_2^*)^T$, we have

$$A = \begin{bmatrix} 0 & u_1^* & a_{13} & 0 \\ u_2^* & 0 & 0 & a_{24} \\ -\varepsilon a_{13} & \varepsilon a_{24} & 0 & 0 \\ \varepsilon a_{13} & -\varepsilon a_{24} & 0 & 0 \end{bmatrix}$$

where $u_1^* = u_1 + \frac{u_2 - u_1}{2} w_1^*$ and $u_2^* = u_2 + \frac{u_1 - u_2}{2} w_2^*$, The parameters a_{ij} are represented as functions of p_1^*, p_2^*, w_1^* and w_2^*. Without loss of generality, we set $u_1 \geq u_2$. Because of $u_1 \geq u_2 > 0$, and $w_1^*, w_2^* \in [0,1]$, we have $u_1^* \in [\frac{u_1 + u_2}{2}, u_1]$ and $u_2^* \in [u_2, \frac{u_1 + u_2}{2}]$, which means $u_1^* > u_2^* > 0$.

After calculating matrix A's eigenvalue in Matlab, we have an eigenvalue $\lambda_1 = 0$, an eigenvalue λ_2 with its real part $Re(\lambda_2) > 0$, an eigenvalue λ_3 with $Re(\lambda_3) < 0$ and an eigenvalue λ_4 close to 0. Since there exists an eigenvalue $\lambda > 0$, the equilibrium eq is not stable [21].

Next we turn to prove the first type of equilibrium. In this case, we need to put the projection function back since we are dealing with boundary cases.

For the case $P > p > t$, we have $V_i^{\mathrm{idv}}(eq) > V_i^{\mathrm{soc}}(eq)$, thus $\dot{w}_r(eq) > 0$ and $\dot{w}_2(eq) < 0$, which means w_1 and w_2 will keep $w_1 = 1$ and $w_2 = 0$. Because $\dot{p}_1(eq) = \frac{t - p + S - P}{2} < 0$ and $\dot{p}_2(eq) = s - p < 0$, then p_r and p_c will keep $p_r = 0$ and $p_c = 0$. According to the continuity theorem of differential equations [11], $(0, 0, 1, 0)$ is a stable equilibrium. The case $p > P > T$ can be proved similarly, which is omitted here.

For the case $P = p$, we have $V_i^{\mathrm{idv}} = V_i^{\mathrm{soc}}$, then $\dot{w}_1(eq) = -\dot{w}_2(eq) = \varepsilon(V_1^{idv} - V_i^{soc}) = 0$. Because $(\frac{T - s}{2} w_2^* < p - s)$, we have $\dot{p}_1 = \frac{T - s}{2} w_2^* + s - p < 0$. Because $(\frac{t - S}{2} w_1^* < P - S)$, we have $\dot{p}_2 = \frac{t - S}{2} w_2^* + S - P < 0$. According to the continuity theorem of differential equations, $(0, 0, w_1^*, w_2^*)$ is a stable equilibrium. The stability of the second type of equilibrium points can be proved similarly, which is omitted here.

From Theorem 4, we find that conclusions of Corollary 2 is a special case of Theorem 4. Note that it can be verified by drawn the symmetry conditions into Theorem 4.

5 A Practical Algorithm

In SA-IGA, each agent needs to know the policy of others and the payoff function, which are usually not available before a repeated game starts. Based on the idea of SA-IGA, we relax the above assumptions and propose a practical multiagent learning algorithm called Socially-Aware Policy Gradient Ascent (SA-PGA). The overall flow of SA-PGA is shown in Algorithm 2. In SA-PGA, each agent only needs to observe the payoffs of both agents by the end of each round.

In SA-IGA, we know that agent i's policy (the probability of selecting each action) is updated based on the partial derivative of the expected value V_i, while the social attitude w is adjusted according to the relative value of V_i^{idv} and V_i^{soc}. Here

---

Algorithm 2 SA-PGA for player i

1: Let $\alpha_\pi, \alpha_w \in (0,1)$ and $\beta \in (0,1)$ be learning rates.
2: Initialize,
  $Q_i^{\mathrm{idv}}(a) \leftarrow 0, Q_i^{\mathrm{soc}}(a) \leftarrow 0, Q_i(a) \leftarrow 0,$
  $\pi_i(a) \leftarrow \frac{1}{|A_i|}, w_i \leftarrow w_0.$
3: **repeat**
4:     Same as PHC in Step 4 of Table 1.
5:     Observing reward r and the average of all agents' current rewards r_{all},
    $Q_i^{\mathrm{idv}}(a) \leftarrow (1-\beta) Q_i^{\mathrm{idv}}(a) + \beta r,$
    $Q_i^{\mathrm{soc}}(a) \leftarrow (1-\beta) Q_i^{\mathrm{soc}}(a) + \beta r_{\mathrm{all}},$
    $Q_i(a) \leftarrow (1-w_i) Q_i^{\mathrm{idv}}(a) + w_i Q_i^{\mathrm{soc}}(a).$
6:     Update π_i according to gradient ascent strategy, Same as PHC in Step 6 of Table 1.
7:     Update w_i,
    $V_i^{\mathrm{idv}} = \sum_{a \in A_i} \pi_i(a) Q_i^{\mathrm{idv}}(a).$
    $V_i^{\mathrm{soc}} = \sum_{a \in A_i} \pi_i(a) Q_i^{\mathrm{soc}}(a).$
    $w_i \leftarrow \Pi_{[0,1]} [w_i + \alpha_w (V_i^{\mathrm{idv}} - V_i^{\mathrm{soc}})].$
8: **until** the repeated game ends

---

in SA-PGA, we first estimate the value of V_i^{idv} and V_i^{soc} using Q-values, which are updated based on the immediate payoffs received during repeated interactions. Specifically, each agent i keeps a record of the Q-value of each action for both its own and the average of all agents (Q_i^{idv} and Q_i^{soc}) (Step 5). Both Q-values are updated following Q-learning update rules accordingly by the end of each round (Step 5). Then the overall Q-value of each agent is calculated as the weighted average of Q_i^{idv} and Q_i^{soc} weighted by its social attitude w (Step 5). The policy update strategy is the same as the Table 1 in Step 6. Finally, the social attitude of agent i is updated in Step 7. The value of V_i^{idv} and V_i^{soc} are estimated based on its current policy and Q-values. The updating direction of w_i is estimated as the difference between V_i^{idv} and V_i^{soc}. Note that a SA-PGA player in each interaction needs only to know its own reward and the average reward of all agents. Knowing the average reward of a group is a reasonable assumption in many realistic scenarios, such as elections and voting.

6 Experimental Evaluation

This section is divided into three parts. Subsection 6.1 compare SA-IGA and SA-PGA with simulation in different types of two-agents, two-actions, general-sum games. Subsection 6.2 presents the experimental results for the 2x2 benchmark games, specifically, performance of converging to the social optimal outcomes and against selfish agents. Subsection 6.3 presents the experimental results for games with multiple agents, i.e. public good game [3].

6.1 Simulation comparison of SA-IGA and SA-PGA

We start the performance evaluation with analyzing the learning performance of SA-PGA under two-player two-action repeated games. In general a two-player two-action game can be classified into three categories [24]:

1. $\exists i \in \{1, 2\}, (r_i^{11} - r_i^{21})(r_i^{12} - r_i^{22}) > 0$. In this case, each player has a dominant strategy and thus the game only has one pure strategy NE.
2. $\forall i \in \{1, 2\}, (r_i^{11} - r_i^{21})(r_i^{12} - r_i^{22}) < 0$ and $(r_1^{11} - r_1^{21})(r_2^{21} - r_2^{22}) > 0$. In this case, there are two pure strategy NEs and one mixed strategy NE.
3. $\forall i \in \{1, 2\}, (r_i^{11} - r_i^{21})(r_i^{12} - r_i^{22}) < 0$ and $(r_1^{11} - r_1^{21})(r_2^{21} - r_2^{22}) < 0$. In this case, there only exists one one mixed strategy NE.

where r_i^{jk} is the payoff of player i when player i takes action j while its opponent $-i$ takes action k. We select one representative game for each category for illustration.

6.1.1 Category 1

For category 1, we consider the PD game as shown in Table 1. In this game, both players have one dominant strategy D, and (D, D) is the only pure strategy NE, while there also exists one socially optimal outcome (C, C) under which both players can obtain higher payoffs.

(a) SA-PGA in PD game   (b) SA-IGA in PD game
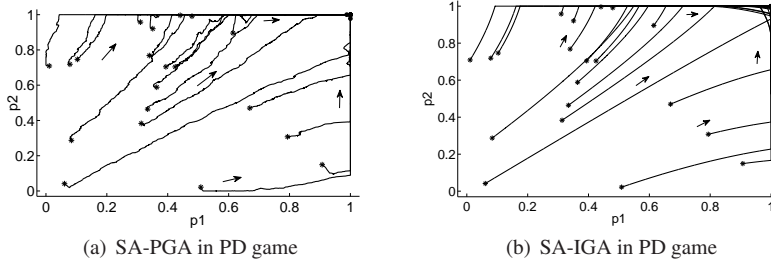
Fig. 1 The Learning Dynamics of SA-IGA and SA-PGA in PD game ($w_1(0) = w_2(0) = 0.85$, $\alpha_\pi = \alpha_w = 0.001$, $\beta = 0.8$)

Figure 1(a) show the learning dynamics of the practical SA-PGA algorithm playing the PD game. The x-axis $p1$ represents player 1's probability of playing action C and the y-axis $p2$ represents player 2's probability of playing action C. We randomly selected 20 initial policy points as the starting point for the SA-PGA agents. We can observe that the SA-PGA agents are able to converge to the mutual cooperation equilibrium point starting from different initial policies.
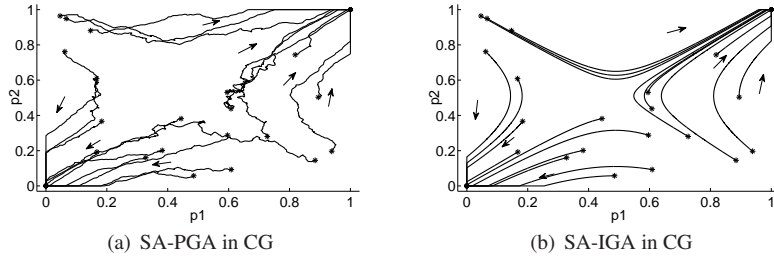
Figure 1(b) illustrates the learning dynamics predicted by the theoretical SA-IGA approach. Similar to the setting in Figure 1(a), the same set of initial policy points are selected and we plot all the learning curves accordingly. We can see that for each starting policy point, the learning dynamics predicted from the theoretical SA-IGA is well consistent with the learning curves from simulation. This indicates that we can better understand and predict the dynamics of SA-PGA algorithm using its corresponding theoretical SA-IGA model.

6.1.2 Category 2

For category 2, we consider the CG game as shown in Table 4. In this game, there exist two pure strategy Nash equilibria (C, C) and (D, D), and both of them are also socially optimal.

| 1's payoff | Agent 2's actions | |
2's payoff	C	D
Agent 1's actions C	3/4	0/0
D	0/0	4/3

Table 4 Coordination game (Category 2)

(a) SA-PGA in CG   (b) SA-IGA in CG

Fig. 2 The Learning Dynamics of SA-IGA and SA-PGA in coordination game ($w_1(0) = w_2(0) = 0.85$, $\alpha_\pi = \alpha_w = 0.001$, $\beta = 0.8$)

Figure 2(a) illustrates the learning dynamics of the practical SA-PGA algorithm playing a CG game. The x-axis $p1$ represents player 1's probability of playing action C and the y-axis $p2$ represents player 2's probability of playing action C. Similar to the case of PD game, 20 initial policy points are randomly selected as the starting points. We can see that the SA-PGA agents can converge to either of the aforementioned two equilibrium points depending on the initial policies they start with.

Figure 2(b) shows the learning dynamics predicted by the theoretical SA-IGA approach. Similar to the setting in Figure 2(a), we adopt the same set of 20 initial policy points for comparison purpose. All the learning curves starting from these 20 policy points are drawn accordingly. We can observe that for each starting policy point, the learning dynamics predicted from the theoretical SA-IGA is well consistent with the learning curves obtained from simulation. Therefore, the theoretical model can facilitate better understanding and predicting the dynamics of SA-PGA algorithm.

6.1.3 Category 3

The game we use in Category 3 is shown in Table 5. In this game, there only exists one mixed strategy Nash equilibrium, while the pure strategy outcome (C, D) is socially optimal.

1's payoff 2's payoff		Agent 2's actions	
		C	D
Agent 1's actions	C	3/2	4/4
	D	1/3	5/1

Table 5 An example game of Category 3

(a) SA-PGA for the game with one mix NE (b) SA-IGA for the game with one mix NE
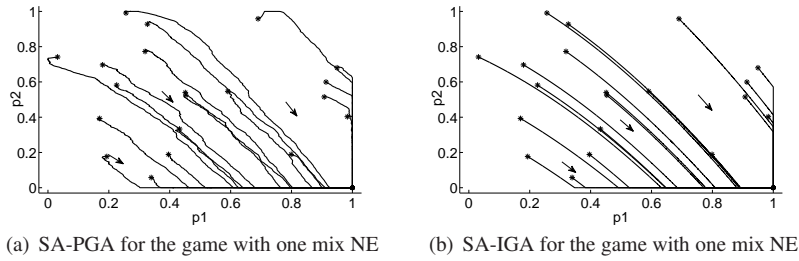
Fig. 3 The Learning Dynamics of SA-IGA and SA-PGA in game with one mix NE ($w_1(0) = w_2(0) = 0.85, \alpha_\pi = \alpha_w = 0.001, \beta = 0.8$)

Figure 3(a) illustrates the learning dynamics of the practical SA-PGA algorithm playing the game in Table 5. The x-axis $p1$ and y-axis $p2$ represent player 1's probability of playing action C and player 2's probability of playing action C respectively. Similar to the previous cases, 20 initial policy points are randomly selected as the starting points. From Figure 3(a), we can see that the SA-PGA agents can always converge to the socially optimal outcome (C, D) no matter where the initial policies start with.

Figure 3(b) presents the learning dynamics of agents predicted by the theoretical SA-IGA approach. Similar to the setting in Figure 3(a), we adopt the same set of 20 initial policy points for comparison purpose, and the corresponding learning curves are drawn accordingly. From Figure 3(b), we can observe that for each starting policy point, the theoretical SA-IGA model can well predict the simulation results of SA-PGA algorithm. Therefore, better understanding and insights of the dynamics of SA-PGA algorithm can be obtained through analyzing its corresponding theoretical model.

6.2 Performance in 2×2 General-sum Games

In this subsection we turn to evaluate the performance of SA-PGA in two-agents, two-actions, general-sum games. First we implement two previous representative learning algorithms for comparison: CJAL [5] and WoLF-PHC [8]. We compare their performance based on the following two criteria: utilitarian social welfare and Nash social welfare, which reflect the system-level efficiency of different learning strategies in terms of the total payoffs received for the agents. Then we evaluate the ability of SA-PGA against selfish opponents with the same three representative games used in previous sections.

6.2.1 Comparison of SA-PGA with CJAL and WoLF-PHC

we evaluate the performance of SA-PGA with CJAL [5] and WoLF-PHC [8] in two-player's repeated games under self-play. CJAL is selected since this algorithm is specifically designed to enable agents to achieve mutual cooperation (i.e., maximizing social welfare) instead of inefficient NE for games like prisoner's dilemma. WoLF-PHC is selected as one representative NE-oriented algorithm for baseline comparison purpose. For all previous strategies the same parameter settings used in their original papers are adopted.

|	Utilitarian Social Welfare	Nash Product
SA-PGA (our strategy) ($w_r(0) = w_c(0) = 0.85$)	7.241 ± 0.003	12.706 ± 0.015
CJAL [5]	6.504 ± 0.032	10.887 ± 0.114
WoLF-IGA [8]	6.536 ± 0.004	10.943 ± 0.145

Table 6 Performance comparison with CJAL and WoLF-PHC

We use all possible structurally distinct two-player, two-action conflict games as a testbed for SA-PGA. In each game, each player ranks the four possible outcomes from 1 to 4. We use the rank of an outcome as the payoff to that player for any outcome. We perform the evaluation under 100 randomly generated games with strict ordinal payoffs. We perform 10,000 interactions for each run and the results are averaged over 20 runs for each game.

We compare their performance based on the the following two criteria: utilitarian social welfare (USW) and Nash social welfare (NSW). Utilitarian social welfare is the sum of the payoffs obtained by the two players in their converged state, while Nash social welfare is the product of the payoffs obtained by two players in their converged state. Formally, USW $=V_1+V_2$ and $NSW = V_1V_2$, where V_1 and V_2 are payoffs obtained by the two players in their converged state, averaged over 100 randomly generated games. Both criteria reflect the system-level efficiency of different learning strategies in terms of the total payoffs received for the agents. Besides, Nash social welfare also partially reflects the fairness in terms of how equal the agents' payoffs are. The overall comparison results are summarized in Table 6. We can see that SA-

PGA outperforms the previous CJAL strategy and WoLF-PHC strategy under both criteria.

6.2.2 Performance Against Selfish Agents

Fig. 4 SA-PGA against a selfish agent for in PD game($w_r(0) = 1$, $p_r(0) = 0.2$ and $p_c(0) = 0.8$)

Fig. 5 SA-PGA against a selfish agent for in coordination game($w_r(0) = 1$, $p_r(0) = 0.2$ and $p_c(0) = 0.8$)

If a learning agent is facing selfish agents that attempt to exploit others, one reasonable choice for an effective algorithm is to learn a Nash equilibrium. In this section, we evaluate the ability of SA-PGA against selfish opponents. We adopt the same three representative games used in previous sections as the testbed and the results are given in Figure 4, 5 and 6 respectively. We can observe that for the PD and coordination games, the SA-PGA agent can successfully achieve the corresponding NE solution. This property is desirable since it prevents the SA-PGA agent from being taken advantage by selfish opponents. The results also show how the socially-aware degree w of SA-PGA agent changes, which varies depending on the game structure. For PD and coordination game, a SA-PGA agent eventually behaves as a purely individually rational entity and one pure strategy NE is eventually converged to. In

Fig. 6 SA-PGA against a selfish agent for the game with only one mix NE($w_r(0) = 1$, $p_r(0) = 0.2$ and $p_c(0) = 0.8$)

contrast, for the third type of game (Table 5), a SA-PGA agent behaves as a purely socially rational agent and cooperate with the selfish agent towards the socially optimal outcome (C, D) without fully exploiting the opponent. This indicates the cleverness of SA-PGA since higher individual payoff can be achieved under the outcome (C, D) than pursuing Nash equilibrium (C, C).

6.3 Performance in games with multiple agents

We use Public Goods Game (PGG) [3] to further evaluate the performance of SA-PGA in multiple agent cases. PGG is an extended version of the PD game in multiagent environment, which has attracted increasing attention to study cooperative behavior and, in particular, deviations from the rational equilibrium [22,28]. In a typical public goods experiment a group of players is endowed with one dollar each. The players then have the opportunity to invest their money into a common pool, knowing that the total amount will be doubled and split equally among all players, irrespective of their contributions. If everybody invests their money, they end up with two dollars. However, each player faces the temptation to free-ride on others' contributions by withholding the money because each invested dollar yields only a return of 50 cents. If everybody adopts this rational strategy, no one would increase the initial capital and forego the public good. The payoffs for cooperators R_C and defectors R_D in a group of N interacting individuals are then given by,

$$R_D = \frac{rN_Cc}{N}, R_C = R_D - c$$

where r denotes the multiplication factor of the public good, N_C the number of cooperators in the group and c the cost of the cooperative contributions, i.e. each agent's investment in the public good. From the definition, the defect action, i.e., the action of not contributing to the public, is the dominate strategy because $R_D > R_C$. The Nash equilibrium of all PGG players is that everyone chooses to defect, while the social optimal outcomes strategy of PGG is that everyone contributes the the public. We evaluate the performance of SA-PGA in PGG repeated games with three players under three circumstance: 1) games with three SA-PGA players, and

2) games with two SA-PGA player and one selfish opponent, and 3) games with one SA-PGA player and two selfish opponents. Without lose of generality, all players' initial policies $p(0)$ of each game are settled to 0.5. Other parameters such as r and c in the three experiments are exactly the same, $r = 2$, $c = 2$.
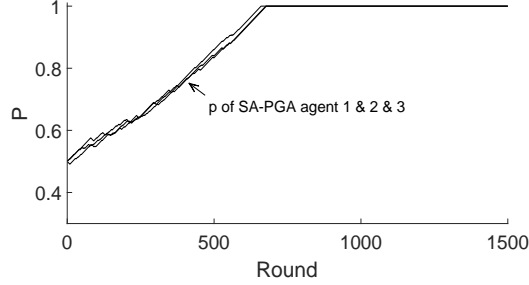
Fig. 7 The Learning Dynamics of SA-PGA in PGG with three SA-PGA players ($\alpha_\pi = \alpha_w = 0.001$, $\beta = 0.8$, $p_i(0) = 0.5$ and $w_i(0) = 0.85$)
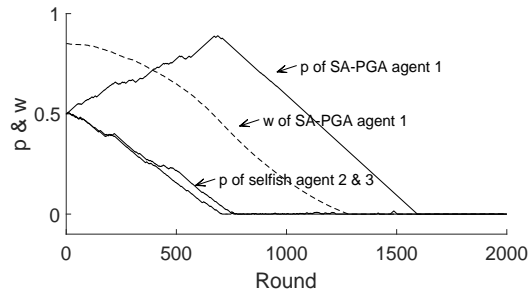

Fig. 8 The Learning Dynamics of SA-PGA in PGG with two SA-PGA players and one selfish opponents ($\alpha_\pi = \alpha_w = 0.001$, $\beta = 0.8$, $p_i(0) = 0.5$ and $w_1(0) = w_2(0) = 0.85$)

Figure 7 shows the learning dynamics of PGG games with three SA-PGA players. The y-axis p represents the probability of playing action C, i.e. the cooperate action, while the x-axis t is the timeline. Each line in Figure 7 shows the learning dynamic of one player's strategy. We can observe that the SA-PGA agents are able to converge to the mutual cooperation equilibrium point giving the initial value of $w(0)$ large enough (here we set $w = 0.85$).

Figure 8 shows the learning dynamics of PGG games with two SA-PGA players and one selfish opponent, while Figure 9 shows the learning dynamics of PGG games with one SA-PGA player and two selfish opponents. The y-axis $p\&w$ represents the probability of strategy p and the socially-aware degree w. The solid lines are learning dynamics of players' strategies, and dotted lines are learning dynamics of SA-PGA players' socially-aware degrees. From Figure 8 and 9, we can observe that agents

Fig. 9 The Learning Dynamics of SA-PGA in PGG with one SA-PGA players and two selfish opponents ($\alpha_\pi = \alpha_w = 0.001$, $\beta = 0.8$, $p_i(0) = 0.5$ and $w_1(0) = 0.85$)

initially tends to cooperate with others and later realizes that the other agents are not cooperating, thus converging to the pure strategy D eventually behaves as a purely individually rational entity. This property is desirable since it prevents the SA-PGA agent from being taken advantage by selfish opponents.

7 Conclusion

In this paper, we proposed a novel way of incorporating social awareness into traditional gradient-ascent algorithm to facilitate reaching mutually beneficial solutions (e.g., (C, C) in PD game). We first present a theoretical gradient-ascent based policy updating approach (SA-IGA) and analyzed its learning dynamics using dynamical system theory. For PD games, we show that mutual cooperation (C,C) is stable equilibrium point as long as both agents are strongly socially-aware. For Coordination games, either of the Nash equilibria (C,C) and (D,D) can be a stable equilibrium point depending on the agents' socially-aware degrees. Following that, we proposed a practical learning algorithm SA-PGA relaxing the impractical assumptions of SA-IGA. Experimental results show that a SA-PGA agent can achieve higher social welfare than previous algorithms under self-play and also is robust against individually rational opponents. As future work, more testbed scenarios (e.g., population of agents) will be applied to further evaluate the performance of SA-PGA. Another interesting direction is to investigate how to further improve the convergence rate of SA-PGA.

References

1. Abdallah, S., Lesser, V.: A multiagent reinforcement learning algorithm with non-linear dynamics. J. Artif. Int. Res. **33**(1), 521–549 (2008)
2. Alvard, M.S.: The ultimatum game, fairness, and cooperation among big game hunters. In: Foundations of Human Sociality, pp. 413–435 (2004)
3. Andreoni, J., Croson, R.: Partners versus Strangers: Random Rematching in Public Goods Experiments. Elsevier B.V. (1998)
4. Banerjee, B., Peng, J.: Adaptive policy gradient in multiagent learning. In: International Joint Conference on Autonomous Agents and Multiagent Systems, pp. 686–692 (2003)

5. Banerjee, D., Sen, S.: Reaching pareto optimality in prisoner's dilemma using conditional joint action learning. AAMAS'07 pp. 91–108 (2007)
6. Bloembergen, D., Tuyls, K., Hennes, D., Kaisers, M.: Evolutionary dynamics of multi-agent learning: a survey. Journal of Artificial Intelligence Research pp. 659–697 (2015)
7. Bowling, M.: Convergence and no-regret in multiagent learning. In: International Conference on Neural Information Processing Systems, pp. 209–216 (2004)
8. Bowling, M.H., Veloso, M.M.: Multiagent learning using a variable learning rate. Artificial Intelligence pp. 215–250 (2003)
9. Busoniu, L., Babuska, R., De Schutter, B.: A comprehensive survey of multiagent reinforcement learning. Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on **38**(2), 156–172 (2008)
10. Chakraborty, D., Stone, P.: Multiagent learning in the presence of memory-bounded agents. Autonomous agents and multi-agent systems **28**(2), 182–213 (2014)
11. Coddington E A, L.N.: Theory of ordinary differential equations. McGraw-Hill (1955)
12. Conitzer, V., Sandholm, T.: 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 (2007)
13. Hauert, C., Szab, G.: Prisoner's dilemma and public goods games in different geometries: Compulsory versus voluntary interactions. Complexity **8**(4), 31–38 (2003)
14. Hu, J., Wellman, M.P.: Nash q-learning for general-sum stochastic games. The Journal of Machine Learning Research **4**, 1039–1069 (2003)
15. Lauer, M., Rienmiller, M.: An algorithm for distributed reinforcement learning in cooperative multi-agent systems. In: ICML'00, pp. 535–542 (2000)
16. Littman, M.: Markov games as a framework for multi-agent reinforcement learning. In: Proceedings of the 11th international conference on machine learning, pp. 322–328 (1994)
17. Littman, M.L.: Friend-or-foe q-learning in general-sum games. In: ICML, vol. 1, pp. 322–328 (2001)
18. Matignon, L., Laurent, G.J., Le Fort-Piat, N.: Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems. The Knowledge Engineering Review **27**(01), 1–31 (2012)
19. Powers, R., Shoham, Y.: Learning against opponents with bounded memory. In: IJCAI, vol. 5, pp. 817–822 (2005)
20. Rodrigues Gomes, E., Kowalczyk, R.: Dynamic analysis of multiagent q-learning with ε-greedy exploration. In: Proceedings of the 26th Annual International Conference on Machine Learning, pp. 369–376. ACM (2009)
21. Shilnikov, L.P., Shilnikov, A.L., Turaev, D.V., Chua, L.O.: Methods of qualitative theory in nonlinear dynamics, vol. 5. World Scientific (2001)
22. Shivshankar, S., Jamalipour, A.: An evolutionary game theory-based approach to cooperation in vanets under different network conditions. IEEE Transactions on Vehicular Technology **64**(5), 2015–2022 (2015)
23. Singh, S., Kearns, M., Mansour, Y.: Nash convergence of gradient dynamics in general-sum games. In: Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence, pp. 541–548. Morgan Kaufmann Publishers Inc. (2000)
24. Tuyls, K., Hoen, P.J., Vanschoenwinkel, B.: An evolutionary dynamical analysis of multi-agent learning in iterated games. Autonomous Agents and Multi-Agent Systems **12**(1), 115–153 (2006)
25. Tuyls, K., Verbeeck, K., Lenaerts, T.: A selection-mutation model for q-learning in multi-agent systems. In: Proceedings of the second international joint conference on Autonomous agents and multi-agent systems, pp. 693–700. ACM (2003)
26. Watkins, C.J.C.H.: Learning from delayed rewards. Robotics & Autonomous Systems **15**(4), 233–235 (1989)
27. Watkins, C.J.C.H., Dayan, P.D.: Q-learning. Machine Learning pp. 279–292 (1992)
28. Wei, G., Zhu, P., Vasilakos, A.V., Mao, Y.: Cooperation dynamics on collaborative social networks of heterogeneous population. IEEE Journal on Selected Areas in Communications **31**(6), 1135–1146 (2013)
29. Zhang, C., Lesser, V.R.: Multi-agent learning with policy prediction. In: Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, pp. 927–934 (2010)
30. Zhang, Z., Zhao, D., Gao, J., Wang, D., Dai, Y.: Fmrq-a multiagent reinforcement learning algorithm for fully cooperative tasks. IEEE Transactions on Cybernetics **47**(6), 1367–1379 (2017)
31. Zinkevich, M.: Online convex programming and generalized infinitesimal gradient ascent. Icml pp. 928–936 (2003)