

Differentiable Arbitrating in Zero-sum Markov Games

Jing Wang^{*†}
New York University
New York, United States
jw5665@nyu.edu

Meichen Song^{*†}
Stony Brook University
New York, United States
meichen.song@stonybrook.edu

Feng Gao[‡]
IIS, Tsinghua University;
Shanghai Qi Zhi Institute
feng.gao220@gmail.com

Boyi Liu
Northwestern University
Illinois, United States
boyiliu2018@u.northwestern.edu

Zhaoran Wang
Northwestern University
Illinois, United States
zhaoran.wang@northwestern.edu

Yi Wu
IIS, Tsinghua University;
Shanghai Qi Zhi Institute
jxwuyi@gmail.com

ABSTRACT

We initiate the study of how to perturb the reward in a zero-sum Markov game with two players to induce a desirable Nash equilibrium, namely arbitrating. Such a problem admits a bi-level optimization formulation. The lower level requires solving the Nash equilibrium under a given reward function, which makes the overall problem challenging to optimize in an end-to-end way. We propose a backpropagation scheme that differentiates through the Nash equilibrium, which provides the gradient feedback for the upper level. In particular, our method only requires a black-box solver for the (regularized) Nash equilibrium (NE). We develop the convergence analysis for the proposed framework with proper black-box NE solvers and demonstrate the empirical successes in two multi-agent reinforcement learning (MARL) environments. Supplementary for all the proofs in this paper could be found in: <https://arxiv.org/abs/2302.10058>.

KEYWORDS

Zero-sum game; Equilibrium refinement; Bi-level optimization

ACM Reference Format:

Jing Wang^{*†}, Meichen Song^{*†}, Feng Gao[‡], Boyi Liu, Zhaoran Wang, and Yi Wu. 2023. Differentiable Arbitrating in Zero-sum Markov Games. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 10 pages.

1 INTRODUCTION

Arbitrating the conflict between self-interest and collective interests permeates the development of human societies [32]. We study Markov games [53] as abstractions of human societies. Nash equilibrium (NE) [8], where none of the players could benefit from unilaterally deviating its strategy, is an essential concept in Markov games. While players modeled by Markov games act rationally to maximize their own rewards, the lack of collective consideration may lead to an undesirable NE, which undermines the overall welfare from a system perspective [23].

^{*}These authors contributed equally to this work.

[†]Work finished during the internship in Shanghai Qi Zhi Institute.

[‡]The author contributed the entire experiment part of this work.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Incentive design [73], which was originally developed for bandit settings, aims to arbitrate such a conflict by perturbing the rewards to refine NE so that the self-interested players reach a desirable one, e.g., the higher social welfare or the larger exploration rate. More specifically, the incentive design problem is naturally formulated as a bi-level structure [19, 58]. The upper-level designer aims to determine the optimal incentives and the lower-level players perform under incentive-perturbed rewards. To this end, it needs to anticipate how the lower-level players react to the prescribed incentives when they play a Markov game. However, taking a derivative through the bi-level structure is difficult [27], especially when the lower level has multiple agents. In particular, the lower-level equilibrium appears to be a black box for the upper-level designer, which makes it challenging for the upper-level designer to assess how the prescribed incentives will influence the lower-level equilibrium. As a result, the upper-level designer lacks the gradient feedback for effectively updating its policy for prescribing the incentives [51].

Some attempts have been proposed to work around the difficulty of derivatives in solving bi-level optimization. One direction is to get rid of the bi-level problem structure by converting the problem into a multi-objective problem: i.e., the low-level players simultaneously optimize the game reward and the designer’s objective by exploring multiple NEs [15, 22, 38, 49, 71, 81]. The simplified problem can be directly solved by an NE solver. However, there is no theoretical guarantee for finding the optimal NE subject to the designer’s objective [21]. The other direction is to keep the bi-level problem structure and avoid the derivative issue by applying a gradient-free optimizer to the upper level [64]. However, this suffers from a high computation cost since a zeroth-order method typically requires a large number of queries to the lower-level solver to derive the desired NE. Therefore, a first-order method that can be applied to complicated environments under the bi-level structure with a convergence guarantee is urgently needed to overcome the sample efficiency issue.

We extend the incentive design problem to the MARL settings and construct a provably *differentiable arbitrating* (DA) framework. We derive the derivatives through the NE for the DA framework, which enables the upper level to utilize the policy gradient feedback from the lower level to obtain the gradient of the designer’s objective. Theoretical convergence proof for DA is presented in the setting of the two-player zero-sum game by choosing proper NE solvers which have convergence guarantee, e.g. Policy Extragradients Method (PEM) [14] and Entropy-regularized OMWU [13]. In practice, we implement the DA framework using a multi-agent

variant of the soft actor-critic algorithm (MASAC) [30], leading to a practical algorithm, DASAC. We evaluate our algorithm in two two-player zero-sum MARL environments. Empirical results show that the incentive proposed by the designer arbitrates the conflict well: i.e., the upper-level designer’s loss function reduces while lower-level self-interested players still attain an NE. The emergent arbitrated behaviors are consistent with human intuitions in each setting. Besides, the sample efficiency for the upper level is significantly improved compared with a zeroth-order method.

Contributions. **i)** We consider the problem of incentive design in MARL settings and tackle the challenge of deriving the gradient of the upper-level objective through the bi-level structure. **ii)** We develop the *first* differentiable *first-order* framework, *Differentiable Arbitrating* (DA), which can arbitrate the conflict to obtain a desirable NE in Markov games. **iii)** We theoretically prove the convergence of DA with proper NE solvers for two-player zero-sum Markov games. **iv)** We empirically show that a desirable NE with interpretable behaviors can be found by DA more efficiently than zeroth-order methods in two MARL environments.

2 RELATED WORKS

Nash equilibrium for Markov Games. There is a large amount of work on finding the Nash equilibrium in multi-agent Markov games. Claus and Boutilier propose *fictitious play* [18] (FS), i.e., all players presuming that the opponents are stationary, for equilibrium selection with guaranteed convergence to a Nash equilibrium. FS was originally developed for normal form games and therefore was not widely applied to complex applications [34, 45, 61] until the *fictitious self-play* (FSP) [33] was proposed. [35] extends SP from extensive-form games to imperfect-information games. [55] proposes the smooth FSP algorithm, which expands the PPO algorithm [74] from the single-agent to the multi-agent zero-sum Markov game and provides with convergence guarantee. In the multi-agent reinforcement learning setting, MADDPG in [57] and MAAC in [39] adapt the actor-critic [43] and soft-actor-critic [30] algorithms from the single-agent to complex multi-agent setting and empirically show better performances fully than decentralized multi-agent RL on finding NEs. MADDPG and MAAC do not have theoretical convergence guarantee even though they achieve empirical successes. On the other hand, many works have developed convergence guarantee for value-based methods [13, 14, 16] and policy-based methods [3, 94] for two-player zero-sum Markov games. Cen [13, 14] introduce the entropy regularization term into the two-player zero-sum Markov game and propose the Policy Extragradients Methods and Entropy-regularized OMWU Methods, which have convergence guarantee without the NE uniqueness assumption on the two-player zero-sum Markov game.

Equilibrium Selection/Refinement. Since even team Markov Games (where the players have common interests) could have multiple NEs, there exists a large volume of work in game theory on NE selection [66, 75], e.g. admissibility [7], subgame perfection [75], Pareto efficiency [10] or stability against opponent’s deviation from best response [24]. [73] raises emphasis on *adaptive incentive design*, which modifies intrinsic rewards via additional incentive to balance the individuals’ self-interests and system’s social welfare for desirable NE selection. The incentive design was first proposed in economics [70] and attracts substantial attentions with many follow-up

works [11, 44, 88]. During the past decades, people explored the economic incentive to various domains, including energy [12, 65], transportation [42, 62], healthcare [5, 59], and education [60]. More recently, incentive design has also been extended to the area of game theory and reinforcement learning. [54] utilize the incentive design in the multi-bandit problem and prove that the proposed algorithm converges to the global optimum at a sub-linear rate for a broad class of games. [64] provide an incentive-design mechanism for an uncooperative multi-agent system and optimize the upper-level incentive objective with Bayesian optimization, a sample-efficient optimization algorithm, instead of the gradient-based methods because the lower-level MARL problem is a black box. [90] propose a decentralized incentive mechanism that allows each individual to directly give rewards to others and learn its own incentive function, respectively. Also, [78] proposes a bi-level incentive mechanism for a single lower-level player, which can be viewed as a special case of our setting. Stackelberg game [48, 52] is a strategic game in economics in which the leader moves first and the followers behave sequentially. The incentive design could be reformulated as the Stackelberg game by treating the incentive design objective as the leading player. [86] proposed a gradient-descent algorithm to find NE for the Stackelberg bandit problem. [95] propose a value iteration method for solving the Stackelberg Markov game with convergence guarantee. [85] features differentiation through policy gradient on Sequential Decision Problems without global convergence guarantee.

Bi-level optimization. In machine learning, a large amount of tasks can be formulated as bi-level optimization problems, e.g. adversarial learning [40], meta-learning [26], hyperparameter optimization [9], and end-to-end learning [4]. Many gradient-based algorithms are used for solving these problems. Most of them are first-order gradient-based methods [9, 26, 28, 67] thanks to both computational efficiency of first-order gradients and fast empirical convergence. By contrast, zeroth-order methods [79, 87] suffer from a heavy workload of data sampling, which requires tremendous computing resources. [69] develop a second-order method. However, the proposed algorithm requires computation of the inverse of Hessian, which restricts the algorithm to simple tabular settings. The most challenging part for gradient-based methods to solve the bi-level optimization problem is to derive the derivative of lower-level solution with respect to the incentive variables in the higher level. Feng [25] and Yang [91] expand the lower-level optimization process as a sequence of gradient updates and directly compute gradients throughout the entire optimization process. Despite of the simplicity, such a method suffers from tremendous computation troubles due to memory explosion and numerical issues. Therefore, approximation techniques such as truncated gradient with a sliding window would be necessary. By contrast, our DA framework gives an accurate first-order formula for exact gradient computation for the lower-level optima. Wang [85] constructed a bi-level problem to infer the missing parameters in MDP by learning a predictive model. This paper is under the single-agent setting where the objective in the lower-level is to minimize a clear performance measure while our setting is more complex by requiring the computation of NE.

3 PROBLEM FORMULATION

This section organizes as follows. Section 3.1 begins with a typical definition for two-player zero-sum Markov games \mathcal{G} . To illustrate how the designer induces the behaviors of players, Section 3.2 introduces the incentive parameter θ into the reward of the two-player zero-sum Markov game and formulate the incentivized Markov game \mathcal{G}_θ . Section 3.3 designs an entropy-regularized Markov game \mathcal{G}'_θ by adding entropy regularization terms into the reward of \mathcal{G}_θ and demonstrates \mathcal{G}'_θ owns the unique NE. Section 3.4 formulates the arbitrating system as a bi-level optimization scheme based on the regularized Markov game with the incentive-perturbed reward.

3.1 Two-player Zero-sum Markov Game

Consider a two-player zero-sum Markov game $\mathcal{G} = (\mathcal{S}, \{\mathcal{A}^i\}_{i \in \{1,2\}}, \mathcal{P}, \{r^i\}_{i \in \{1,2\}}, \gamma)$, where \mathcal{S} is the state space observed by all players, \mathcal{A}^i is the action space of player i and $\mathcal{A} := \mathcal{A}^1 \times \mathcal{A}^2$ is the joint action space of two players, then $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ denotes the transition probability from state $s \in \mathcal{S}$ to state $s' \in \mathcal{S}$ for taking joint action $a \in \mathcal{A}$. $r^i : \mathcal{S} \times \mathcal{A} \rightarrow [-1, 1]$ is the immediate reward function of player i , which implies $r^1 = -r^2$ in zero-sum game. $\gamma \in [0, 1)$ is the discounted factor. We remark that we consider the two-player zero-sum setting for notation simplicity. Extensions to more general settings will be discussed in Appendix C.

3.2 Markov Game with the Incentive-perturbed Reward

Let $\mathcal{G}_\theta = (\mathcal{S}, \{\mathcal{A}^i\}_{i \in \{1,2\}}, \mathcal{P}, \{r^i(\cdot; \theta)\}_{i \in \{1,2\}}, \gamma)$ be a two-player zero-sum Markov game with an incentive-perturbed reward function $r^i(\cdot; \theta)$ explicitly parameterized by θ , where $\theta \in \mathbb{R}^m$ represents the incentives added by the designer. We assume that $r^i(\cdot; \theta) : \mathcal{S} \times \mathcal{A} \rightarrow [-1, 1]$ is uniformly bound in θ and remains zero-sum. For a given trajectory $\tau = (s_0, \{a_t^i\}_{i \in \{1,2\}}, s_1, \dots, \{a_{T-1}^i\}_{i \in \{1,2\}}, s_T)$, we denote the total discounted reward for player $i \in \{1, 2\}$ as

$$R^i(\tau; \theta) = \sum_{t=0}^{T-1} \gamma^t r^i(s_t, a_t; \theta).$$

Define $\pi^i : \mathcal{S} \times \mathcal{A}^i \rightarrow [0, 1]$ as the probability distribution of policy of player i and $\pi = (\pi^1, \pi^2)$ as that of the joint policy.

Let $D^\pi(\tau)$ denote the probability distribution of the trajectory τ based on the policy pair π . That is,

$$D^\pi(\tau) := \rho_0(s_0) \prod_{t=0}^{T-1} P(s_{t+1}|s_t, a_t) \left(\prod_{i=1}^2 \pi^i(a_t^i|s_t) \right),$$

where $\rho_0(s_0)$ is the distribution for the initial state. We denote $\mathbb{E}_\pi[\cdot]$ as the expectation over the trajectory $\tau \sim D^\pi(\tau)$. Then the performance of the policy pair π for the game \mathcal{G}_θ is evaluated by the state-value function $V_\pi^i : \mathcal{S} \rightarrow \mathbb{R}$, which is defined as

$$\begin{aligned} V_\pi^i(s; \theta) &= \mathbb{E}_\pi \left[\sum_{t=0}^{T-1} \gamma^t \cdot r^i(s_t, a_t^i, a_t^{-i}; \theta) \mid s_0 = s \right] \\ &= \mathbb{E}_{\tau \sim D^\pi} [R^i(\tau; \theta) \mid s_0 = s]. \end{aligned} \quad (3.1)$$

The Nash equilibrium (NE) is an essential concept where no agent could benefit by individually changing its policy. The NE for the incentivized two-player zero-sum Markov game \mathcal{G}_θ is the solution for the following min-max game:

$$\min_{\pi^2} \max_{\pi^1} V_\pi^1(s; \theta), \quad s \in \mathcal{S}. \quad (3.2)$$

Shapley [76] proves that there exists a NE pair for the min-max game 3.2 for all state $s \in \mathcal{S}$ and the min-max value is unique [84].

3.3 Entropy-regularized Markov Game

To ensure the NE to be unique, we follow [14, 20, 93] and introduce the entropy-regularized counterpart $\mathcal{G}'_\theta = (\mathcal{S}, \{\mathcal{A}^i\}_{i \in \{1,2\}}, \mathcal{P}, \{r^i(\cdot; \theta)\}_{i \in \{1,2\}}, \gamma, \lambda)$, where $\lambda \geq 0$ is the regularization parameter. The entropy regularization technique is commonly utilized in reinforcement learning to encourage exploration to avoid being trapped at sub-optimal solutions [2, 29, 46] and keep the policies of different agents away from being heavily affected by the opponents' strategy [56, 89].

Specifically, the reward function for player i is replaced by its entropy-regularized reward function $r_\pi^{(i)}(\cdot; \theta) : \mathcal{S} \times \mathcal{A}^i \times \mathcal{A}^{-i} \rightarrow \mathbb{R}$, which is defined as,

$$\begin{aligned} r_\pi^{(i)}(s, a^i, a^{-i}; \theta) &= r^i(s, a^i, a^{-i}; \theta) - \lambda \log(\pi^i(a^i|s)) \\ &\quad + \lambda \log(\pi^{-i}(a^{-i}|s)). \end{aligned} \quad (3.3)$$

Here we clarify what the entropy we use with a light abuse of notation. The state-reward function and the entropy-regularized state-reward function associated with policy pair π are defined as

$$\begin{aligned} r_\pi^i(s; \theta) &= \mathbb{E}_{(a^i, a^{-i}) \sim \pi} [r^i(s, a^i, a^{-i}; \theta)], \\ r_\pi^{(i)}(s; \theta) &= \mathbb{E}_{(a^i, a^{-i}) \sim \pi} [r_\pi^{(i)}(s, a^i, a^{-i}; \theta)] \\ &= r_\pi^i(s; \theta) + \lambda H(\pi^i(\cdot|s)) - \lambda H(\pi^{-i}(\cdot|s)), \end{aligned} \quad (3.4)$$

where $H(\pi^i(\cdot|s)) = -\sum_{a^i \in \mathcal{A}^i} \pi^i(a^i|s) \log(\pi^i(a^i|s))$ is the Shannon entropy. Correspondingly, the total discounted reward $R_\pi^{(i)}$ for a given trajectory τ of the player i in game \mathcal{G}'_θ and the entropy-regularized state-value function with incentive parameter $V_\pi^{(i)} : \mathcal{S} \rightarrow \mathbb{R}$ are defined as

$$\begin{aligned} R_\pi^{(i)}(\tau; \theta) &= \sum_{t=0}^{T-1} \gamma^t r_\pi^{(i)}(s_t, a_t^i, a_t^{-i}; \theta), \\ V_\pi^{(i)}(s; \theta) &= \mathbb{E}_\pi \left[\sum_{t=0}^{T-1} \gamma^t \cdot r_\pi^{(i)}(s_t, a_t^i, a_t^{-i}; \theta) \mid s_0 = s \right] \\ &= \mathbb{E}_{\tau \sim D^\pi} [R_\pi^{(i)}(\tau; \theta) \mid s_0 = s]. \end{aligned} \quad (3.5)$$

Suppose the NE of the regularized Markov game with incentive-perturbed reward \mathcal{G}'_θ is $\pi_*(\theta) = (\pi_*^1(\theta), \pi_*^2(\theta))$. We write Player $-i$ as Player i 's opponent then the joint policy could also be rewritten as the policy pair $\pi = (\pi^i, \pi^{-i})$, where $i \in \{1, 2\}$. Then the NE can be defined as follows, for $\forall i \in \{1, 2\}$, the NE π_* satisfies

$$V_{\pi_*}^{(i)}(s; \theta) = \max_{\pi^i} V_{(\pi^i, \pi_*^{-i}(\theta))}^{(i)}(s; \theta), \quad \forall s \in \mathcal{S}. \quad (3.6)$$

The NE for the regularized Markov game \mathcal{G}'_θ is equivalent to the solution of the following min-max optimization problem

$$\min_{\pi^2} \max_{\pi^1} V_\pi^{(1)}(s; \theta), \quad \forall s \in \mathcal{S}. \quad (3.7)$$

As is known from [13, 14, 63], there exists a unique policy pair π for the min-max optimization problem (3.7). Therefore the entropy-regularized Markov game has a unique NE. Moreover, [14] proves that an $\epsilon/2$ -optimal NE for the **regularized** Markov game \mathcal{G}'_θ is an ϵ -optimal NE for the original when the regularization parameter $\lambda \leq \frac{(1-\gamma)\epsilon}{2(\log(|\mathcal{A}_1|) + \log(|\mathcal{A}_2|))}$, which indicates the difference between

NE for the regularized Markov game and NE for the original can be controlled by the regularization parameter. The definition for ϵ -optimal NE is clearly defined in [14], which means the optimal gap of the policy is smaller than ϵ .

3.4 Bi-level Optimization Scheme

Our final goal is to arbitrate conflicts between self-interested players and the collective-interested designer by navigating all individuals to the desirable refined NE with higher system welfare. Section 3.2 introduces the incentive parameter θ to modify the intrinsic reward $r^i(\cdot; \theta)$ of Markov games, which is utilized as a tool to design the system objective and refine NE for arbitration. In this section, we formulate a bi-level optimization scheme with equilibrium constraints induced by such an arbitrating system. Among the bi-level optimization scheme, for the lower-level players, we aim at finding the NE for the given incentivized Markov game \mathcal{G}'_θ . Regarding the upper-level designer, we define a system arbitrating objective targeting the optimal θ that could channel all individuals towards a desirable NE.

For the purpose of gradient computation in Section 4 and experiments in Section 7, we parameterize policy $\pi_{\phi^i}^i$ with $\phi^i \in \mathbb{R}^{d_i}$ for the player i ($i \in \{1, 2\}$) and denote the policy pair as $\pi_\phi = (\pi_{\phi^1}^1, \pi_{\phi^2}^2)$, where $\phi = (\phi^1, \phi^2) \in \mathbb{R}^d$, $d = d_1 + d_2$. To avoid the abuse of subscription, we omit i of ϕ^i in $\pi_{\phi^i}^i$ (w.r.t $\pi_{\phi^i}^i$). The simplification will not cause confusion because the policy π^i of a single player i is only determined by ϕ^i . Then the joint policy could also be rewritten as the policy pair $\pi_\phi = (\pi_{\phi^1}^1, \pi_{\phi^2}^2)$, where $i \in \{1, 2\}$. The NE of the regularized Markov game with incentive-perturbed reward \mathcal{G}'_θ is denoted as $\pi_{\phi_*}(\theta) = (\pi_{\phi_*^1}^1, \pi_{\phi_*^2}^2)$, where $\phi_*(\theta) = (\phi_*^1(\theta), \phi_*^2(\theta))$ is the parameter of NE under incentive θ . The definition of NE for regularized Markov game in (3.6) can be rewritten as,

$$V_{\pi_{\phi_*}(\theta)}^{(i)}(s; \theta) = \max_{\phi^i \in \mathbb{R}^{d_i}} V_{(\pi_{\phi^i}^i, \pi_{\phi_*}^{-i}(\theta))}^{(i)}(s; \theta), \quad \forall s \in \mathcal{S}.$$

Arbitrating system. Suppose the arbitrating objective for game system \mathcal{G}'_θ is only determined by the incentive parameter θ and the joint policy π_ϕ , written as mapping $f(\theta, \phi) : \Theta \times \mathbb{R}^d \rightarrow \mathbb{R}$. The system aims to minimize the arbitrating objective loss at the NE of \mathcal{G}'_θ , $f(\theta, \phi_*(\theta))$. Therefore, the arbitrating system could be formulated as a bi-level optimization problem:

$$\begin{aligned} \min_{\theta \in \Theta} \quad & f_*(\theta) = f(\theta, \phi_*(\theta)) \\ \text{s.t.} \quad & \mathbb{E}_{V^*} \left[V_{\pi_{\phi_*}(\theta)}^{(i)}(s; \theta) \right] \\ & = \max_{\phi^i \in \mathbb{R}^{d_i}} \mathbb{E}_{V^*} \left[V_{(\pi_{\phi^i}^i, \pi_{\phi_*}^{-i}(\theta))}^{(i)}(s; \theta) \right], \quad \forall i \in \{1, 2\} \end{aligned} \quad (3.8)$$

Remark: Eq. 3.8 is our major objective. We primarily focus on efficiently learning the upper-level incentive parameter while treating the lower level as a black-box NE solver. In the following sections, we focus on deriving the first-order gradient of the upper-level loss function $f_*(\theta)$, then establish our *Differentiable Arbitrating* (DA) framework based on it.

4 METHOD

The bi-level arbitrating system (Eq. 3.8) with NE constraints guides players indirectly via the incentive parameter θ to safeguard the arbitrating objective f . For the lower level, we attempt to determine the policy parameter of NE under a certain incentivized environment \mathcal{G}'_θ , remarked as $\phi^*(\theta)$. For the upper-level designer, we aim to derive the first-order algorithm to find the optimal incentive parameter θ such that the objective $f_*(\theta) = f(\theta, \phi_*(\theta))$ can be optimized while players attain their NE. In such a bi-level problem, it is challenging to obtain the gradient of f_* , especially when the solver for regularized NE in the lower level is a black box, since the incentive parameter θ affects objective f_* implicitly via refining the regularized NE for players. Our goal is to derive the gradient of f_* w.r.t θ with NE constraints and then develop a backpropagation scheme that utilize the feedback of policy gradients from the lower level to establish the gradient of the incentive objective function at a regularized NE, $\nabla_\theta f(\theta, \phi^*(\theta))$. To be more specific, the gradient $\nabla_\theta f(\theta, \phi^*(\theta))$ can be derived as,

$$\nabla_\theta f_* = \left[\nabla_\theta f(\theta, \phi) + \nabla_\theta \phi^*(\theta)^\top \nabla_\phi f(\theta, \phi) \right] \Big|_{\phi=\phi^*(\theta)}. \quad (4.1)$$

where $\nabla_\theta f(\theta, \phi)$ and $\nabla_\phi f(\theta, \phi)$ are known while the NE gradient $\nabla_\theta \phi_*(\theta)$ need to be derived based on the policy gradient feedback from the lower level. The computing of the NE gradient $\nabla_\theta \phi^*(\theta)$ is in detail discussed in Section 4.1 and then a general backpropagation framework that for the regularized NE is proposed in Section 4.2.

4.1 Gradient for the Incentive

In this section, we will derive the gradient at NE $\nabla_\theta f(\theta, \phi^*(\theta))$, in which the key point is to derive NE gradient $\nabla_\theta \phi^*(\theta)$. For notation simplification, we denote $u_\theta^{(i)}(\phi; s)$ as the gradient of $V_{\pi_\phi}^{(i)}(s; \theta)$ w.r.t ϕ^i then we define

$$u_\theta(\phi) := \mathbb{E}_{V^{(*)}} [u_\theta(\phi; s)], \quad (4.2)$$

$$u_\theta(\phi; s) := (u_\theta^{(1)}(\phi; s), u_\theta^{(2)}(\phi; s)) = (\nabla_{\phi^1} V_{\pi_\phi}^{(1)}(s; \theta), \nabla_{\phi^2} V_{\pi_\phi}^{(2)}(s; \theta)).$$

Then, the gradient of $u_\theta(\phi)$ w.r.t. to θ and ϕ are straightforward as follows,

$$\nabla_\theta u_\theta(\phi) = \begin{bmatrix} \mathbb{E}_{V^{(*)}} \nabla_{\theta \phi^1}^2 V_{\pi_\phi}^{(1)}(s; \theta) \\ \mathbb{E}_{V^{(*)}} \nabla_{\theta \phi^2}^2 V_{\pi_\phi}^{(2)}(s; \theta) \end{bmatrix} \in \mathbb{R}^{d \times m}, \quad (4.3)$$

$$\nabla_\phi u_\theta(\phi) = \begin{bmatrix} \mathbb{E}_{V^{(*)}} \nabla_{\phi \phi^1}^2 V_{\pi_\phi}^{(1)}(s; \theta) \\ \mathbb{E}_{V^{(*)}} \nabla_{\phi \phi^2}^2 V_{\pi_\phi}^{(2)}(s; \theta) \end{bmatrix} \in \mathbb{R}^{d \times d}. \quad (4.4)$$

The following Lemma 4.1 implies the computation of $\nabla_{\theta \phi^i}^2 V_{\pi_\phi}^{(i)}(s; \theta)$ and $\nabla_{\phi \phi^i}^2 V_{\pi_\phi}^{(i)}(s; \theta)$ can be explicitly derived based on the policy gradient, which are backpropagated from NE solver in the lower level. The computation details for the gradients and Hessians in Lemma 4.1 are presented throughout Lemma A.1. in appendix due to limitation of space. Then in light of (4.3-4.4) and Lemma 4.1, $\nabla_\theta \phi^*(\theta)$ can be expressed in policy gradient as derived in the following Lemma 4.2, inducing the final expression of objective gradient at NE $\nabla_\theta f_*(\theta)$.

Lemma 4.1. *In an incentive regularized Markov game \mathcal{G}'_θ , we denote $\pi_{\phi^i}^i := \pi_{\phi^i}^i(a_t^i | s_t)$, $\pi_{\phi^{-i}}^{-i} := \pi_{\phi^{-i}}^{-i}(a_t^{-i} | s_t)$ and $\pi_\phi := \pi_\phi(a_t | s_t)$ for*

notation simplicity, then we have

$$\begin{aligned}
\nabla_{\theta} V_{\pi_{\phi}}^{(i)}(s; \theta) &= \nabla_{\theta} V_{\pi_{\phi}}^i(s; \theta); \quad \nabla_{\theta \phi^i}^2 V_{\pi_{\phi}}^{(i)}(s; \theta) = \nabla_{\theta \phi^i}^2 V_{\pi_{\phi}}^i(s; \theta); \\
\nabla_{\phi^i} V_{\pi_{\phi}}^{(i)}(s; \theta) &= \nabla_{\phi^i} V_{\pi_{\phi}}^i(s; \theta) - \lambda \mathbb{E}_{D^{\pi_{\phi}}} \left[\sum_{t=0}^{T-1} \gamma^t \nabla_{\phi^i} \log \pi_{\phi^i}^i \right. \\
&\quad \left. + \left(\sum_{t=0}^{T-1} \gamma^t (\log \pi_{\phi^i}^i - \log \pi_{\phi^i}^{-i}) \right) \left(\sum_{t=0}^{T-1} \nabla_{\phi^i} \log \pi_{\phi^i}^i \right) \right]; \\
\nabla_{\phi \phi^i}^2 V_{\pi_{\phi}}^{(i)}(s; \theta) &= \nabla_{\phi \phi^i}^2 V_{\pi_{\phi}}^i(s; \theta) \\
&\quad - \lambda \mathbb{E}_{D^{\pi_{\phi}}} \left[\left(\sum_{t=0}^{T-1} \nabla_{\phi^i} \log \pi_{\phi^i}^i \right) \left(\sum_{t=0}^{T-1} \gamma^t (\nabla_{\phi} \log \pi_{\phi^i}^i - \nabla_{\phi} \log \pi_{\phi^i}^{-i}) \right)^{\top} \right. \\
&\quad \left. + \left(\sum_{t=0}^{T-1} \gamma^t (\log \pi_{\phi^i}^i - \log \pi_{\phi^i}^{-i}) \right) \left(\sum_{t=0}^{T-1} \nabla_{\phi^i} \log \pi_{\phi^i}^i \right) \left(\sum_{t=0}^{T-1} \nabla_{\phi} \log \pi_{\phi} \right)^{\top} \right. \\
&\quad \left. + \left(\sum_{t=0}^{T-1} \gamma^t (\log \pi_{\phi^i}^i - \log \pi_{\phi^i}^{-i}) \right) \left(\sum_{t=0}^{T-1} \nabla_{\phi \phi^i}^2 \log \pi_{\phi} \right)^{\top} \right. \\
&\quad \left. + \left(\sum_{t=0}^{T-1} \gamma^t \nabla_{\phi^i} \log \pi_{\phi^i}^i \right) \left(\sum_{t=0}^{T-1} \nabla_{\phi} \log \pi_{\phi} \right)^{\top} + \left(\sum_{t=0}^{T-1} \gamma^t \nabla_{\phi \phi^i}^2 \log \pi_{\phi} \right) \right],
\end{aligned}$$

where $\mathbb{E}_{D^{\pi_{\phi}}}[\cdot]$ denotes the expectation over the trajectory $\tau \sim D^{\pi_{\phi}}(\tau)$ and $\nabla_{\theta} V_{\pi_{\phi}}^i(s; \theta)$, $\nabla_{\phi^i} V_{\pi_{\phi}}^i(s; \theta)$, $\nabla_{\theta \phi^i}^2 V_{\pi_{\phi}}^i(s; \theta)$ and $\nabla_{\phi \phi^i}^2 V_{\pi_{\phi}}^i(s; \theta)$ are computed from Lemma A.1. in appendix.

Lemma 4.1 utilizes the policy gradient information, which is easily accessible in a NE solver, to describe how self-interested players react to the adjustment in the incentive parameter. It is straightforward to prove but is essential to provide the lower-level information embedded in the policy gradient to guide the direction of updates for the incentive parameter against a high collective loss.

In light of Lemma 4.1, Lemma 4.2 presents the gradient of NE w.r.t. the incentive parameter, which is crucial to leverage the back-propagated information from the lower level to update the incentive parameter. It describes how the equilibrium among players is affected by the incentive from the designer and shows that such information is related to the Hessian of $V_{\pi_{\phi}}^{(i)}(s; \theta)$ at NE.

Lemma 4.2. *The following formula holds for the policy parameter $\phi^{*}(\theta)$ of NE $\pi_{\phi^{*}(\theta)}$ for incentive regularized Markov game \mathcal{G}_{θ} :*

$$\nabla_{\theta} \phi^{*}(\theta) = - \left[\nabla_{\phi} u_{\theta}(\phi) \right]^{-1} \nabla_{\theta} u_{\theta}(\phi) \Big|_{\phi=\phi^{*}(\theta)}. \quad (4.5)$$

Then substituting Eq. (4.5) into Eq. (4.1) leads to the final explicit expression of $\nabla_{\theta} f_{*}(\theta)$:

$$\begin{aligned}
&\nabla_{\theta} f_{*}(\theta, \phi^{*}(\theta)) \\
&= \left[\nabla_{\theta} f(\theta, \phi) - \nabla_{\theta} u_{\theta}(\phi) \right]^{\top} \left[\nabla_{\phi} u_{\theta}(\phi) \right]^{-1} \nabla_{\phi} f(\theta, \phi) \Big|_{\phi=\phi^{*}(\theta)}. \quad (4.6)
\end{aligned}$$

4.2 Differentiable Arbitrating

Based on the gradient computation in Section 4.1, we propose the *Differentiable Arbitrating* (DA) framework to solve the arbitrating system (Eq. 3.8). DA differentiates the upper-level designer's objective function through NE and backpropagates the gradient information of policies from the lower level to derive the upper-level

Framework 1 DA: Differentiable Arbitrating in MARL.

Input: β_k : learning rate for upper-level iteration

Output: θ : incent. param.; ϕ : param. of policy π_{ϕ}

Initialize the incentive parameter $\theta = \theta_0$.

for $k=0,1,\dots$ **do**

Initialize param. $\phi = \phi_0$ of policy π_{ϕ_0} for game \mathcal{G}'_{θ_k} .

for $t=0,1,\dots$ **do**

Update $\phi = \phi_t$ with NE solver until corresponding policy π_{ϕ_t} converge to the Nash equilibrium of game \mathcal{G}'_{θ_k} , w.r.t $\phi^{*}(\theta_k)$.

end for

Update incentive parameter

$$\theta = \theta_{k+1} \leftarrow \theta_k - \beta_k \nabla f_{*}(\theta_k),$$

where $\nabla f_{*}(\theta_k)$ is defined in (4.6).

end for

gradient. In this scheme, the black-box solver for the regularized NE in the lower level could be any methods containing gradient information of the policies, for instance, SAC [30], MADDPG [57], MAAC [39], Policy Extragradient Method (PEM) [14] and Entropy-regularized OMWU [13].

The workflow of DA is shown in Framework 1. In this work, we specifically develop the convergence guarantee for the DA framework with lower-level NE solvers which have convergence guarantee and implement our DA framework using a multi-agent variant of SAC, which will be described in detail later in Section 6.

DA framework is a first-order method. This is because the update of upper-level incentive θ only relies on the first-order gradient $\nabla f_{*}(\theta)$ without any hessian information of θ . The second-order gradients defined in Lemma 4.1 are hessian matrices of lower-level value function V and are used for the computation of first-order gradient $\nabla f_{*}(\theta)$.

Extension to general settings. Although DA framework is described under two-player zero-sum fully observable Markov games, it could be extended to more general settings: *N-players, general-sum* and *partially observable Markov decision process (POMDP)*. For N-player general-sum Markov games, the game \mathcal{G}_{θ} is required to be monotone to ensure the entropy-regularized game \mathcal{G}'_{θ} to be strongly-monotone and therefore has unique NE. Then the DA framework could leverage a proper lower-level black-box NE solver for the (regularized) NE. For POMDPs, some minor adjustments are required, replacing the state-based policy $\pi^i(a^i|s)$ with the observation-based policy $\pi^i(a^i|o^i)$, where $o^i = O^i(s)$ is the observation of the player i in state s . Such adjustments will not change the gradient computation, which is essential for the DA framework. The detailed discussions of problem formulation and the gradient derivation for the general settings are deferred to Appendix C.

5 CONVERGENCE ANALYSIS

In this section, we develop the convergence analysis for the DA framework with proper lower-level NE solvers in the two-player zero-sum setting. Section 5.1 discuss lower-level NE solvers with convergence guarantee in literature that could be applied to our DA framework. Section 5.2 guarantees the convergence of incentive parameter θ in the upper level for the bi-level DA framework under convergent lower-level NE solvers.

5.1 Convergence of the Lower Level

The convergence of the bi-level DA framework requires the lower-level NE solver for having the convergence guarantee to the NE π_* of the entropy-regularized Markov game \mathcal{G}'_θ . There are diverse NE solvers that have good empirical performance in practical games [30, 39, 57]. Some proposed NE solvers have the theoretical convergence proofs [37, 77, 80]. For the illustration of the DA framework convergence, we take two NE solvers proposed under entropy-regularization that are directly related with our setting. It is known that the Policy Extragradient Method (PEM) [14] and Entropy-regularized OMWU [13] all have the convergence guarantee under our setting. This section restates the main convergence results for PEM and Entropy-regularized OMWU for the purpose of completeness of convergence guarantee for DA framework.

Proposition 5.1 (Theorem 3 in [14]). *Assume $|\mathcal{S}_1| \geq |\mathcal{S}_2|$ and $\lambda \leq 1$. Setting learning rate $\eta_t = \eta = \frac{1-\gamma}{2(1+\lambda(\log|\mathcal{S}_1|+1-\gamma))}$, the Policy Extragradient Method (Algorithm 4 in supplementary) takes no more than $\tilde{O}\left(\frac{1}{\lambda(1-\gamma)^2} \log^2\left(\frac{1}{\epsilon}\right)\right)$ iterations to achieve the policy pair $\pi = (\pi_1, \pi_2)$ that satisfies:*

$$\max_{\pi'_1, \pi'_2} \mathbb{E}_{s \sim \rho} \left[\left(V_{\pi'_1, \pi_2}^{(1)}(s; \theta) - V_{\pi_1, \pi'_2}^{(1)}(s; \theta) \right) \right] \leq \epsilon,$$

where ρ is an arbitrary distribution over the state space \mathcal{S} .

Proposition 5.2 (Theorem 1 in [13]). *Setting $0 < \eta \leq \frac{(1-\gamma)^3}{32000|\mathcal{S}|}$ and $\alpha_t = \eta\lambda$, the Entropy-regularized OMWU (Algorithm 5 in supplementary) takes no more than $\tilde{O}\left(\frac{1}{(1-\gamma)\eta\lambda} \log \frac{1}{\epsilon}\right)$ iterations to achieve the policy pair $\pi = (\pi_1, \pi_2)$ that satisfies:*

$$\max_{s \in \mathcal{S}, \pi'_1, \pi'_2} \left(V_{\pi'_1, \pi_2}^{(1)}(s; \theta) - V_{\pi_1, \pi'_2}^{(1)}(s; \theta) \right) \leq \epsilon.$$

Proposition 5.1 and Proposition 5.2 demonstrate that the PEM and Entropy-regularized OMWU algorithms converges to the NE for the lower-level incentivized entropy-regularized Markov games \mathcal{G}'_θ in linear rate.

5.2 Convergence of the Upper Level

Assume that $\theta^* = \arg \min_\theta f_*(\theta)$ represents the optimal point of incentive parameter θ in the arbitrating system (3.8). In this section, we are going to show that the framework 1 guarantees that the incentive parameter θ converges to θ^* under certain assumptions.

Assumption 5.3. We assume that $\nabla_\theta f_*(\theta)$ is L-Lipschitz continuous (L-smooth) w.r.t $\theta \in \Theta$, which means

$$\|\nabla_\theta f_*(\theta) - \nabla_\theta f_*(\theta')\| \leq L \|\theta - \theta'\|, \quad \forall \theta, \theta' \in \Theta$$

We also assume that $f_*(\theta)$ is bounded w.r.t $\theta \in \Theta$, which means there exists $M > 0$ such that

$$|f_*(\theta)| \leq M, \quad \forall \theta \in \Theta$$

L-smooth assumption is a common assumption in convergence proofs for different algorithms [1, 41, 50, 72, 92]. We also illustrate a concrete example to justify Assumption 5.3 is rational in practical setting in Appendix E. Assumption 5.3 does not require the composite function $f_*(\theta) = f(\theta, \phi^*(\theta))$ to be convex and the convergence guarantee for the upper level even suitable for non-convex cases.

Theorem 5.4. *Suppose that Assumption 5.3 holds and the lower-level NE solver has convergence guarantee. In Framework 1, let the update rule $\theta_{k+1} = \theta_k - \beta_k \nabla_\theta f_*(\theta_k)$ for incentive parameter θ_k run for T iterations with learning rates $\beta_k = \frac{1}{T}$, then we have*

$$\min_{k=0, \dots, T} \|\nabla_\theta f_*(\theta_k)\|^2 \leq \frac{2L(f_*(\theta_0) - f_*(\theta^*))}{T+1} \leq \frac{4LM}{T+1}.$$

Theorem 5.4 implies that there exist a subsequence in the sequence of gradient norms $\{\|\nabla_\theta f_*(\theta_k)\|\}_{k=0}^\infty$ which converges to zero. Therefore, Theorem 5.4 guarantees that the incentive parameter θ could always converge to a sub-optimal point with sublinear convergence rate in the upper level. It is easy to know if the composite function $f^*(\theta)$ is convex, it converges to the global optimal point. The detailed proof is shown in Appendix F.

6 IMPLEMENTATIONS

The direct implementation of PEM and Entropy-regularized OMWU are not practical/efficient although they have the theoretical convergence guarantee. PEM is a double-loop algorithm that invokes solving the NE of an entropy-regularized Matrix game for each inner-loop. Although Entropy-regularized OMWU is a single-loop algorithm that adapts a two-timescale iteration, its update rules are for the tabular setting, in which the policy are per state updated. Therefore, we modify the vanilla algorithm of Entropy-regularized OMWU with some techniques inspired by the multi-agent soft-actor-critic (MASAC) and propose the DA-SAC algorithm (Algo. 6 in appendix). We organize the modifications as the following:

- i) In stead of the per-state policy update in Entropy-regularized OMWU, we optimize a separate policy model via minimizing its KL divergence from the exponential energy function, i.e., we establish an actor-critic structure as widely accepted in SAC.
- ii) We use one-step gradient update with great practical success accepted by most of MARL practical algorithms, e.g. SAC [30], MADDPG [57] and MAAC [39], instead of the extragradient technique introduced in Entropy-regularized OMWU.
- iii) We follow SAC to set the energy function as a multiple of the Q function, where the scale factor is the reciprocal of the auto-tuned temperature coefficient [31].

Updated pseudo-code and all remaining details are included in Appendix G.

7 EXPERIMENTS

We evaluate the performance of the DA framework in sense of sample efficiency for the upper level in two zero-sum Markov games with two players and their incentivized variant, running with scissors (RWS) [83] which contains cyclic reinforcement learning challenges [6, 47, 68] and a standard predator-prey, both of which are implemented on a grid-world environment [17].

7.1 Evaluation Environments

Running with scissors (RWS). In 5×5 grid-world RWS, resources (rock, scissors, or paper) are tailed in resources pools (Figure 2). Three of them are deterministic pools spaced with fixed resources and six others are nondeterministic filled with random resources.

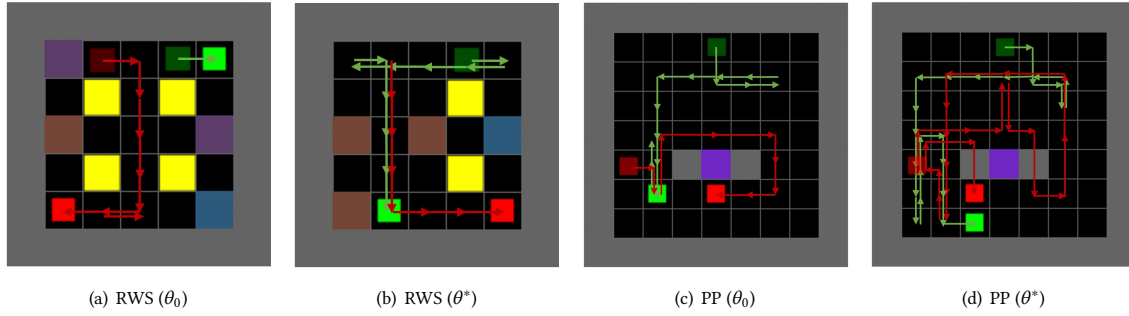


Figure 1: Trace for incentivized RWS and incentivized PP with different incentive parameters θ_0 and θ^* (left to right). θ_0 is the initial incentive parameter and θ^* is the optimized parameter with DA framework. The NEs after arbitration (Figure (b) & (d)) ensure higher exploration rate in the grid-world games compared with the initial NEs (Figure (a) & (c)) without arbitration.




Figure 2: Configuration of the Rock-with-Scissors (RWS). Two players are randomly spawned among the free grids and collect resources for confrontation. The bonus coins are the incentive added by the designer to encourage exploration.

Two players randomly spawned among the free grids and fully observe the environment and resources owned by their opponent. Each chooses to move one grid in four directions at one step. When the player i steps on the grid with the resource, it collects the resource to its inventory v^i , and the resource is removed from the grid. Four rocks, papers and scissors are randomly distributed to two players at the initial state and each player is assigned at least one for each resource. After 25 steps, the confrontation occurs and the payoffs for each player, r^0 and r^1 , are calculated on the basis of the standard antisymmetric matrix M [36] as

$$r^0 = \frac{v^0}{\|v^0\|} M \left(\frac{v^1}{\|v^1\|} \right)^T = -r^1, \quad \text{where } M = \begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix}.$$

Besides, the rules of the game are assumed to be a blank to the players and they must explore to discover them.

RWS is an extension of the classic matrix game rock-paper-scissors (RPS) with increasing complexity, and the game-theoretic structures of the RPS are satisfied in RWS. Since NE is to keep the same number of each resource in their inventory [82], players tend to stay around a grid after they achieve NE instead of continuing to explore, while the upper-level designer hopes more grids can be explored. Given such a conflict, the designer intervenes in the rewards by scattering gold coins as an incentive to channel players exploring more grids.

There are four grids that generate coins. Two of them are filled with fixed bonuses 0.5, and others are filled with adjustable bonuses marked as the incentive parameter $\theta = (\theta_1, \theta_2) \in \mathbb{R}^2$. Once the

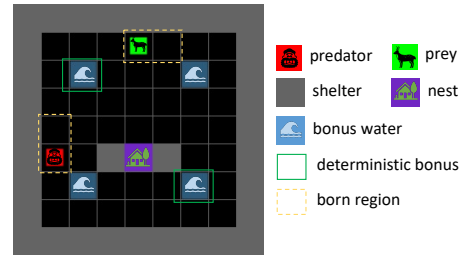


Figure 3: Configuration of the Predator-Prey (PP). The predator aims to catch the prey and the prey aims to return to the nest. Both of them are randomly spawned in born region. The water pools are set as the incentive by the designer.

player i collects coins, the corresponding bonuses are added to its total reward. Conversely, the bonuses collected by player i are deducted from its opponent $-i$. The total reward for the player i is

$$r_{total}^i = r^i + \sum_{j=1,2} \theta_j (I_j^i - I_j^{-i}) + 0.5(N_{fixed}^i - N_{fixed}^{-i}),$$

where $I_j^i = 1$ if the coin with incentive bonus θ_j is collected by the player i , otherwise $I_j^i = 0$ and N_{fixed}^i is the number of fixed coins collected by the player i . Then the exploration rate (ER) can be defined as

$$ER(\theta) = \frac{\sum_i G_e^i(\theta)}{G_{total}}, \quad (7.1)$$

where the $G_e^i(\theta)$ denotes the number of grids explored by player i and G_{total} denotes the total number of grids. Therefore, the designer's loss is defined as $f_*(\theta) = 1 - ER(\theta)$.

Predator-prey (PP). The environment for the predator-prey is set in a 7×7 grid-world as Figure 3, containing the prey's nest, two shelters, and water pools. Both the predator and prey cannot stop at the shelter. At each step, the prey can move one grid and the predator can move one or two grids in four directions. The game is terminated under three conditions: i) the episode achieves maximum length 25; ii) the prey returns the nest, then the reward for the prey, r_{prey} , is +1 and the reward for the predator, r_{pred} , is -1; iii) the predator catches the prey (they arrive on the same grid at the same time), then the reward for prey, r_{prey} , is -1 and the reward for predator, r_{pred} , is +1. From a whole-ecosystem perspective, like the dissemination of plant seeds, the designer stimulates

them to explore more places by setting four pools with different volumes of water in the four fixed grids. Two of them are filled with fixed water volume with an additional 0.1 bonuses and others are filled with adjustable water volume, whose bonuses are denoted as the incentive parameter $\theta = (\theta_1, \theta_2) \in \mathbb{R}^2$. The player gets the corresponding bonus if it finds and drinks up a pool of water while its opponent gets the corresponding penalty. Then the total pay-off for the predator is

$$r_{\text{pred}} = I_{\text{catch}} - I_{\text{nest}} + \sum_{j=1,2} \theta_j (I_j^{\text{pred}} - I_j^{\text{prey}}) + 0.5(N_{\text{fixed}}^{\text{pred}} - N_{\text{fixed}}^{\text{prey}}),$$

where I_{catch} and I_{nest} are two indicators for that the predator catch the prey or the prey back to the nest, I_j^{pred} (or I_j^{prey}) is the indication functions of whether the pool with θ_j amount of water is found by the predator (or prey) and $N_{\text{fixed}}^{\text{pred}}$ (or $N_{\text{fixed}}^{\text{prey}}$) represents the number of fixed pools discovered by the predator (or prey). In addition, we define the reward for prey as $r_{\text{prey}} = -r_{\text{pred}}$. The arbitrating objective for PP is the same as what is defined for RWS.

7.2 Results and Analysis

We use *GridSearch*(M) to denote the zeroth-order method, which spreads M number of points into the feasible region for incentive parameter θ , as a competitor to DASAC. Additionally, we adopt Bayesian optimization to improve the efficiency of the zeroth-order method, denoted as *BayesOpt*, which is also used in [64]. The experimental details of *GridSearch*(M) and *BayesOpt* are included in Appendix G.2. Figure 4 shows the trend of the arbitrating objective f_* after applying DASAC onto the environments. The green dotted line in Figure 4 is the best objective score of *GridSearch*(100). On the one hand, Figure 4 shows the differentiable first-order framework we propose is capable to admit a better incentive parameter θ with a lower objective loss. On the other hand, Figure 4 indicates that *GridSearch*(M) requires 100 samples, in another word, solving the lower-level Markov game 100 times, to achieve an equivalent performance of DASAC. DASAC can achieve better efficiency (5 outer-loop iterations) than these zeroth-order methods even though *BayesOpt* (12 outer-loop iterations, >2 times slower than ours) has improved the efficiency of *GridSearch* (100 outer-loop iterations). After considering the gradient computation, DASAC requires 3.5h for one iteration and the zeroth-order methods (*Gridsearch* and *BayesOpt*) requires 3h for one iteration under the same GPU setting. Our method is 2x faster than *BayesOpt* and 17x faster than *Gridsearch*(100). It means our first-order framework requires fewer evaluations for the upper level to obtain a desirable NE policy, leading to a higher sample efficiency.

Furthermore, we will present our DASAC tends to reach a desirable NE via analyzing the NE behavior at the initial incentive θ_0 and the optimal incentive θ^* (marked in Figure 4). Figure 1 illustrates the trace of players when taking the NE policies in RWS and PP environments with initial incentive θ_0 and the optimal incentive θ^* respectively. It demonstrates that the optimal θ^* navigates the system to a better NE that players are more exploration-minded. To be more specific, with the initial incentive parameter θ_0 , although players start to take account of the exploration rate and think outside the habitual behaviors in an unincentivized system, e.g., standstill but not collecting resources after a few steps in RWS or stalemate in a relative diagonal position in PP, they still tend

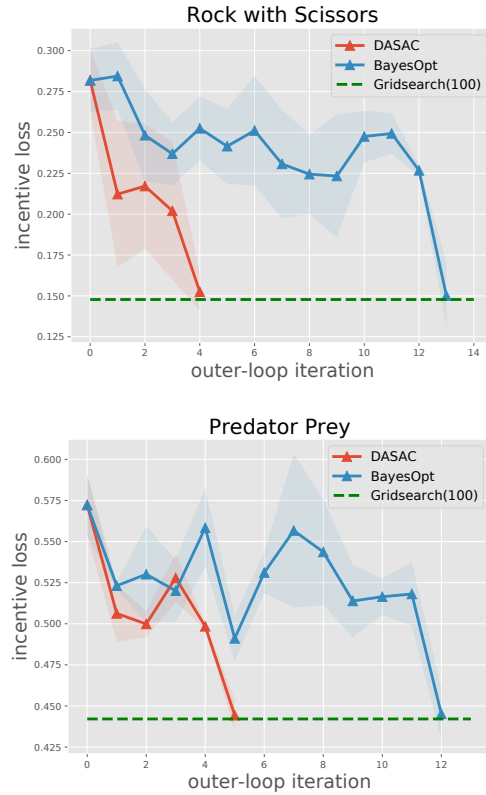


Figure 4: Arbitrating objective loss f_* for DASAC, BayesOpt, and GridSearch(100) in RWS (top) and PP (bottom). The green dotted line indicates the best objective score of GridSearch(100). Our DASAC can outperform two zeroth-order baselines, GridSearch(100) and BayesOpt, by a large margin, which shows that our first-order framework requires fewer evaluations for the upper level to obtain a desirable NE policy, leading to a higher sample efficiency.

to maintain the habitual behavior after a few explorations. As the training for the incentive parameter goes on, the habitual behaviors in an unincentivized system are gradually abandoned and they develop new patterns of behavior that lead to a higher exploration rate and ensure their own goals at the same time. For example, in PP, the predator still tries their best to catch the prey, and the prey tries to escape as soon as possible. The difference is that the prey tends to circle around, exploring the water source to replenish its energy instead of staying in one place and avoiding risky moves.

8 CONCLUSION

Our work initiates a provably differentiable framework in context with MARL to solve a bi-level arbitrating problem. We provide the convergence proof and empirically validate the effectiveness of the DASAC on arbitrating in two Markov games. Our work can be extended to multiple-player and general-sum settings with proper NE solvers whose convergence properties have been empirically shown. Therefore, we will empirically test the performance of the performance of the DA framework on more general cases in our future work.

REFERENCES

- [1] Hadi Abbaszadehpavasti, Etienne de Klerk, and Moslem Zamani. 2022. The exact worst-case convergence rate of the gradient method with fixed step lengths for L-smooth functions. *Optimization Letters* 16, 6 (2022), 1649–1661.
- [2] Zafarali Ahmed, Nicolas Le Roux, Mohammad Norouzi, and Dale Schuurmans. 2019. Understanding the impact of entropy on policy optimization. In *International Conference on Machine Learning*. PMLR, 151–160.
- [3] Ahmet Alacaoglu, Luca Viano, Niao He, and Volkan Cevher. 2022. A natural actor-critic framework for zero-sum Markov games. In *International Conference on Machine Learning*. PMLR, 307–366.
- [4] Brandon Amos and J Zico Kolter. 2017. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*. PMLR, 136–145.
- [5] Haris Aziz and Florian Brandl. 2021. Efficient, fair, and incentive-compatible healthcare rationing. *arXiv preprint arXiv:2102.04384* (2021).
- [6] David Balduzzi, Karl Tuyls, Julien Perolat, and Thore Graepel. 2018. Re-evaluating evaluation. *arXiv preprint arXiv:1806.02643* (2018).
- [7] Jeffrey S Banks and Joel Sobel. 1987. Equilibrium selection in signaling games. *Econometrica: Journal of the Econometric Society* (1987), 647–661.
- [8] Tamer Başar and Geert Jan Olsder. 1998. *Dynamic noncooperative game theory*. SIAM.
- [9] Atilim Gunes Baydin, Robert Cornish, David Martinez Rubio, Mark Schmidt, and Frank Wood. 2018. Online Learning Rate Adaptation with Hypergradient Descent. In *International Conference on Learning Representations*.
- [10] B Douglas Bernheim and Michael D Whinston. 1987. Coalition-proof Nash equilibria II. applications. *Journal of Economic Theory* 42, 1 (1987), 13–29.
- [11] Patrick Bolton and Mathias Dewatripont. 2004. *Contract theory*. MIT press.
- [12] Carlo Cambini and Laura Rondi. 2010. Incentive regulation and investment: evidence from European energy utilities. *Journal of regulatory economics* 38, 1 (2010), 1–26.
- [13] Shicong Cen, Yuejie Chi, Simon S Du, and Lin Xiao. 2022. Faster Last-iterate Convergence of Policy Optimization in Zero-Sum Markov Games. *arXiv preprint arXiv:2210.01050* (2022).
- [14] Shicong Cen, Yuting Wei, and Yuejie Chi. 2021. Fast policy extragradient methods for competitive games with entropy regularization. *Advances in Neural Information Processing Systems* 34 (2021), 27952–27964.
- [15] Georgios Chalkiadakis and Craig Boutilier. 2003. Coordination in Multiagent Reinforcement Learning: A Bayesian Approach. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems* (Melbourne, Australia) (AAMAS '03). Association for Computing Machinery, New York, NY, USA, 709–716. <https://doi.org/10.1145/860575.860689>
- [16] Ziyi Chen, Shaocong Ma, and Yi Zhou. 2021. Sample Efficient Stochastic Policy Extragradient Algorithm for Zero-Sum Markov Game. In *International Conference on Learning Representations*.
- [17] Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. 2018. Minimalistic Gridworld Environment for OpenAI Gym. <https://github.com/maximecb/gym-minigrid>.
- [18] Caroline Claus and Craig Boutilier. 1998. The dynamics of reinforcement learning in cooperative multiagent systems. *AAAI/IAAI* 1998, 746–752 (1998), 2.
- [19] Benoit Colson, Patrice Marcotte, and Gilles Savard. 2007. An overview of bilevel optimization. *Annals of operations research* 153, 1 (2007), 235–256.
- [20] Bo Dai, Albert Shaw, Lihong Li, Lin Xiao, Niao He, Zhen Liu, Jianshu Chen, and Le Song. 2018. SBEDD: Convergent Reinforcement Learning with Nonlinear Function Approximation. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 1125–1134. <https://proceedings.mlr.press/v80/dai18c.html>
- [21] Sam Devlin and Daniel Kudenko. 2011. Theoretical considerations of potential-based reward shaping for multi-agent systems. In *The 10th International Conference on Autonomous Agents and Multiagent Systems*. ACM, 225–232.
- [22] Yali Du, Lei Han, Meng Fang, Ji Liu, Tianhong Dai, and Dacheng Tao. 2019. Liir: Learning individual intrinsic reward in multi-agent reinforcement learning. (2019).
- [23] Pradeep Dubey. 1986. Inefficiency of Nash Equilibria. *Mathematics of Operations Research* 11, 1 (1986), 1–8. <http://www.jstor.org/stable/3690047>
- [24] Fei Fang, Albert Xin Jiang, and Milind Tambe. 2013. Protecting moving targets with multiple mobile resources. *Journal of Artificial Intelligence Research* 48 (2013), 583–634.
- [25] Xidong Feng, Oliver Slumbers, Ziyu Wan, Bo Liu, Stephen McAleer, Ying Wen, Jun Wang, and Yaodong Yang. 2021. Neural auto-curricula in two-player zero-sum games. *Advances in Neural Information Processing Systems* 34 (2021), 3504–3517.
- [26] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*. PMLR, 1126–1135.
- [27] Terry L Friesz, Roger L Tobin, Hsun-Jung Cho, and Nihal J Mehta. 1990. Sensitivity analysis based heuristic algorithms for mathematical programs with variational inequality constraints. *Mathematical Programming* 48, 1 (1990), 265–284.
- [28] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [29] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine. 2017. Reinforcement learning with deep energy-based policies. In *International Conference on Machine Learning*. PMLR, 1352–1361.
- [30] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*. PMLR, 1861–1870.
- [31] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
- [32] Yuval Noah Harari. 2014. *Sapiens: A brief history of humankind*. Random House.
- [33] Johannes Heinrich, Marc Lanctot, and David Silver. 2015. Fictitious Self-Play in Extensive-Form Games. In *Proceedings of the 32nd International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 37)*, Francis Bach and David Blei (Eds.). PMLR, Lille, France, 805–813. <https://proceedings.mlr.press/v37/heinrich15.html>
- [34] Johannes Heinrich and David Silver. 2015. Smooth uct search in computer poker. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- [35] Johannes Heinrich and David Silver. 2016. Deep reinforcement learning from self-play in imperfect-information games. *arXiv preprint arXiv:1603.01121* (2016).
- [36] Josef Hofbauer and Karl Sigmund. 2003. Evolutionary game dynamics. *Bulletin of the American mathematical society* 40, 4 (2003), 479–519.
- [37] Junling Hu and Michael P Wellman. 2003. Nash Q-learning for general-sum stochastic games. *Journal of machine learning research* 4, Nov (2003), 1039–1069.
- [38] Aly Ibrahim, Anirudha Jitani, Daoud Piracha, and Doina Precup. [n.d.]. Reward Redistribution Mechanisms in Multi-agent Reinforcement Learning. ([n.d.]).
- [39] Shariq Iqbal and Fei Sha. 2019. Actor-attention-critic for multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 2961–2970.
- [40] Haoming Jiang, Zhehui Chen, Yuyang Shi, Bo Dai, and Tuo Zhao. 2021. Learning to Defend by Learning to Attack. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 577–585.
- [41] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [42] Zafeiris Kokkinogonis, Nuno Monteiro, Rosaldo JF Rossetti, Ana LC Bazzan, and Pedro Campos. 2014. Policy and incentive designs evaluation: A social-oriented framework for Artificial Transportation Systems. In *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 151–156.
- [43] Vijay R Konda and John N Tsitsiklis. 2000. Actor-critic algorithms. In *Advances in neural information processing systems*. 1008–1014.
- [44] Jean-Jacques Laffont and David Martimort. 2009. *The theory of incentives*. Princeton university press.
- [45] Theodore J Lambert III, Marina A Epelman, and Robert L Smith. 2005. A fictitious play approach to large-scale optimization. *Operations Research* 53, 3 (2005), 477–489.
- [46] Kyungjae Lee, Sungjoon Choi, and Songhwa Oh. 2018. Sparse markov decision processes with causal sparsification regularization for reinforcement learning. *IEEE Robotics and Automation Letters* 3, 3 (2018), 1466–1473.
- [47] Joel Z Leibo, Edward Hughes, Marc Lanctot, and Thore Graepel. 2019. Autocurricula and the emergence of innovation from social interaction: A manifesto for multi-agent intelligence research. *arXiv preprint arXiv:1903.00742* (2019).
- [48] George Leitmann. 1978. On generalized Stackelberg strategies. *Journal of optimization theory and applications* 26, 4 (1978), 637–643.
- [49] Chenghao Li, Chengjie Wu, Tonghan Wang, Jun Yang, Qianchuan Zhao, and Chongjie Zhang. 2021. Celebrating Diversity in Shared Multi-Agent Reinforcement Learning. *arXiv preprint arXiv:2106.02195* (2021).
- [50] Huan Li, Cong Fang, Wotao Yin, and Zhouchen Lin. 2018. A sharp convergence rate analysis for distributed accelerated gradient methods. *arXiv preprint arXiv:1810.01053* (2018).
- [51] Jiayang Li, Jing Yu, Yu Nie, and Zhaoran Wang. 2020. End-to-end learning and intervention in games. *Advances in Neural Information Processing Systems* 33 (2020).
- [52] Tao Li and Suresh P Sethi. 2017. A review of dynamic Stackelberg game models. *Discrete & Continuous Dynamical Systems-B* 22, 1 (2017), 125.
- [53] Michael L Littman. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*. Elsevier, 157–163.
- [54] Boyi Liu, Jiayang Li, Zhuoran Yang, Hoi-To Wai, Mingyi Hong, Yu Marco Nie, and Zhaoran Wang. 2021. Inducing Equilibria via Incentives: Simultaneous Design-and-Play Finds Global Optima. *arXiv preprint arXiv:2110.01212* (2021).
- [55] Boyi Liu, Zhuoran Yang, and Zhaoran Wang. 2020. Policy Optimization in Zero-Sum Markov Games: Fictitious Self-Play Provably Attains Nash Equilibria. (2020).
- [56] Jingbin Liu, Xinyang Gu, and Shuai Liu. 2019. Policy optimization reinforcement learning with entropy regularization. *arXiv preprint arXiv:1912.01557* (2019).

- [57] Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 6382–6393.
- [58] Zhi-Quan Luo, Jong-Shi Pang, and Daniel Ralph. 1996. *Mathematical programs with equilibrium constraints*. Cambridge University Press.
- [59] Xiaochen Ma, Hong Wang, Li Yang, Leiyu Shi, and Xiaoyun Liu. 2019. Realigning the incentive system for China’s primary healthcare providers. *Bmj* 365 (2019).
- [60] Hugh Macartney, Robert McMillan, and Uros Petronijevic. 2016. *Incentive design in education: An empirical analysis*. Technical Report. National Bureau of Economic Research.
- [61] H Brendan McMahan and Geoffrey J Gordon. 2007. A fast bundle-based anytime algorithm for poker and other convex games. In *Artificial Intelligence and Statistics*. PMLR, 323–330.
- [62] Joy Melnikow, Mary Paliescheskey, and Gary K Stewart. 1997. Effect of a transportation incentive on compliance with the first prenatal appointment: a randomized trial. *Obstetrics & Gynecology* 89, 6 (1997), 1023–1027.
- [63] Panayotis Mertikopoulos and William H Sandholm. 2016. Learning in games via reinforcement and regularization. *Mathematics of Operations Research* 41, 4 (2016), 1297–1324.
- [64] David Mguni, Joel Jennings, Emilio Sison, Sergio Valcarcel Macua, Sofia Ceppi, and Enrique Munoz de Cote. 2019. Coordinating the Crowd: Inducing Desirable Equilibria in Non-Cooperative Systems. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 386–394.
- [65] Amir-Hamed Mohsenian-Rad, Vincent WS Wong, Juri Jatskevich, and Robert Schober. 2010. Optimal and autonomous incentive-based energy consumption scheduling algorithm for smart grid. In *2010 Innovative Smart Grid Technologies (ISGT)*. IEEE, 1–6.
- [66] Roger B Myerson. 1978. Refinements of the Nash equilibrium concept. *International journal of game theory* 7, 2 (1978), 73–80.
- [67] Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999* (2018).
- [68] Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Remi Munos. 2019. α -rank: Multi-agent evaluation by evolution. *Scientific reports* 9, 1 (2019), 1–29.
- [69] Jack Parker-Holder, Luke Metz, Cinjon Resnick, Hengyuan Hu, Adam Lerer, Alistair Letcher, Alex Peysakhovich, Aldo Pacchiano, and Jakob Foerster. 2020. Ridge Rider: Finding Diverse Solutions by Following Eigenvectors of the Hessian. *arXiv preprint arXiv:2011.06505* (2020).
- [70] A.C. Pigou and N. Aslanbeigui. 1920. *The Economics of Welfare (1st ed.)*. Routledge.
- [71] Aida Rahmattalabi, Jen Jen Chung, Mitchell Colby, and Kagan Tumer. 2016. D++: Structural credit assignment in tightly coupled multiagent domains. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 4424–4429.
- [72] Ali Ramezani-Kebrya, Ashish Khisti, and Ben Liang. 2018. On the generalization of stochastic gradient descent with momentum. *arXiv preprint arXiv:1809.04564* (2018).
- [73] Lillian J Ratliff, Roy Dong, Shreyas Sekar, and Tanner Fiez. 2019. A perspective on incentive design: Challenges and opportunities. *Annual Review of Control, Robotics, and Autonomous Systems* 2 (2019), 305–338.
- [74] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [75] Reinhard Selten. 1965. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit: Teil i: Bestimmung des dynamischen preisgleichgewichts. *Zeitschrift für die gesamte Staatswissenschaft/Journal of Institutional and Theoretical Economics* H. 2 (1965), 301–324.
- [76] Lloyd S Shapley. 1953. Stochastic games. *Proceedings of the national academy of sciences* 39, 10 (1953), 1095–1100.
- [77] Satinder Singh, Tommi Jaakkola, Michael L Littman, and Csaba Szepesvári. 2000. Convergence results for single-step on-policy reinforcement-learning algorithms. *Machine learning* 38, 3 (2000), 287–308.
- [78] Chen Siyu, Yang Donglin, Li Jiayang, Wang Senmiao, Yang Zhuoran, and Wang Zhaoran. 2022. *Adaptive Model Design for Markov Decision Process*. Technical Report.
- [79] Xingyou Song, Wenbo Gao, Yuxiang Yang, Krzysztof Choromanski, Aldo Pacchiano, and Yunhao Tang. 2020. ES-MAML: Simple Hessian-Free Meta Learning. In *ICLR*.
- [80] Csaba Szepesvári and Michael L Littman. 1999. A unified analysis of value-function-based reinforcement-learning algorithms. *Neural computation* 11, 8 (1999), 2017–2060.
- [81] Zhenggang Tang, Chao Yu, Boyuan Chen, Huazhe Xu, Xiaolong Wang, Fei Fang, Simon Shaolei Du, Yu Wang, and Yi Wu. 2021. Discovering Diverse Multi-Agent Strategic Behavior via Reward Randomization. In *International Conference on Learning Representations*. https://openreview.net/forum?id=lvRTC669EY_
- [82] Anne van den Nouweland. 2007. Rock-paper-scissors; a new and elegant proof. (2007).
- [83] Alexander Vezhnevets, Yuhuai Wu, Maria Eckstein, Rémi Leblond, and Joel Z Leibo. 2020. OPTIONS as REsponses: Grounding behavioural hierarchies in multi-agent reinforcement learning. In *International Conference on Machine Learning*. PMLR, 9733–9742.
- [84] John Von Neumann and Oskar Morgenstern. 2007. *Theory of games and economic behavior*. In *Theory of games and economic behavior*. Princeton university press.
- [85] Kai Wang, Sanket Shah, Haipeng Chen, Andrew Perrault, Finale Doshi-Velez, and Milind Tambe. 2021. Learning MDPs from Features: Predict-Then-Optimize for Sequential Decision Making by Reinforcement Learning. *Advances in Neural Information Processing Systems* 34 (2021), 8795–8806.
- [86] Kai Wang, Lily Xu, Andrew Perrault, Michael K Reiter, and Milind Tambe. 2022. Coordinating followers to reach better equilibria: End-to-end gradient descent for stackelberg games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 5219–5227.
- [87] Xiaoxing Wang, Wenxuan Guo, Junchi Yan, Jianlin Su, and Xiaokang Yang. 2021. ZARTS: On Zero-order Optimization for Neural Architecture Search. *arXiv preprint arXiv:2110.04743* (2021).
- [88] Thomas A Weber and AV Kryazhimskiy. 2011. *Optimal control theory with applications in economics*. Vol. 10. MIT press Cambridge, MA.
- [89] Changnan Xiao, Haosen Shi, Jiajun Fan, and Shihong Deng. 2021. An Entropy Regularization Free Mechanism for Policy-based Reinforcement Learning. *arXiv preprint arXiv:2106.00707* (2021).
- [90] Jiachen Yang, Ang Li, Mehrdad Farajtabar, Peter Sunehag, Edward Hughes, and Hongyuan Zha. 2020. Learning to Incentivize Other Learning Agents. *Advances in Neural Information Processing Systems* 33 (2020).
- [91] Jiachen Yang, Ethan Wang, Rakshit Trivedi, Tuo Zhao, and Hongyuan Zha. 2022. Adaptive Incentive Design with Multi-Agent Meta-Gradient Reinforcement Learning. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*. 1436–1445.
- [92] Tianbao Yang, Qihang Lin, and Zhe Li. 2016. Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization. *arXiv preprint arXiv:1604.03257* (2016).
- [93] Rui Zhao, Xudong Sun, and Volker Tresp. 2019. Maximum Entropy-Regularized Multi-Goal Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 7553–7562. <https://proceedings.mlr.press/v97/zhao19d.html>
- [94] Yulai Zhao, Yuandong Tian, Jason Lee, and Simon Du. 2022. Provably Efficient Policy Optimization for Two-Player Zero-Sum Markov Games. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2736–2761.
- [95] Han Zhong, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. 2021. Can Reinforcement Learning Find Stackelberg-Nash Equilibria in General-Sum Markov Games with Myopic Followers? *arXiv preprint arXiv:2112.13521* (2021).