

Multi-agent Inverse Reinforcement Learning for Certain General-Sum Stochastic Games

Xiaomin Lin

*Data Science, MassMutual Financial Group
470 Atlantic Ave., Boston, MA 02210 USA
University of Virginia
Charlottesville, VA 22904 USA*

XL5DB@VIRGINIA.EDU

Stephen C. Adams

Peter A. Beling

*University of Virginia
Charlottesville, VA 22904 USA*

SCA2C@VIRGINIA.EDU

PB3A@VIRGINIA.EDU

Abstract

This paper addresses a subset of problems of multi-agent inverse reinforcement learning (MIRL) in a two-player general-sum stochastic game framework. Specifically, five variants of MIRL are considered: *uCS-MIRL*, *advE-MIRL*, *cooE-MIRL*, *uCE-MIRL*, and *uNE-MIRL*, each distinguished by its solution concept. Problem *uCS-MIRL* is a cooperative game in which the agents employ cooperative strategies that aim to maximize the total game value. In problem *uCE-MIRL*, agents are assumed to follow strategies that constitute a correlated equilibrium while maximizing total game value. Problem *uNE-MIRL* is similar to *uCE-MIRL* in total game value maximization, but it is assumed that the agents are playing a Nash equilibrium. Problems *advE-MIRL* and *cooE-MIRL* assume agents are playing an adversarial equilibrium and a coordination equilibrium, respectively. We propose novel approaches to address these five problems under the assumption that the game observer either knows or is able to accurately estimate the policies and solution concepts for players. For *uCS-MIRL*, we first develop a characteristic set of solutions ensuring that the observed bi-policy is a *uCS* and then apply a Bayesian inverse learning method. For *uCE-MIRL*, we develop a linear programming problem subject to constraints that define necessary and sufficient conditions for the observed policies to be correlated equilibria. The objective is to choose a solution that not only minimizes the total game value difference between the observed bi-policy and a local *uCS*, but also maximizes the scale of the solution. We apply a similar treatment to the problem of *uNE-MIRL*. The remaining two problems can be solved efficiently by taking advantage of solution uniqueness and setting up a convex optimization problem. Results are validated on various benchmark grid-world games.

1. Introduction

This paper addresses the generalization of inverse reinforcement learning (IRL) to a multi-agent setting. IRL has been widely studied in recent years as part of a broad and growing interest in reinforcement learning (RL). The RL problem is commonly framed in terms of a single agent that aims to learn an optimal control policy for a Markov decision process (MDP) with reward function and state transition probabilities that are either known explicitly or that can be experienced through interaction with the environment. IRL, the inverse problem, has the objective of estimating the reward function given observations of a control

policy followed by an agent (Ng & Russell, 2000). IRL has been applied to a number of problems, most related to the problems of learning from demonstrations and apprenticeship learning (Abbeel & Ng, 2004; Baker, Saxe, & Tenenbaum, 2009; Ziebart, Maas, Bagnell, & Dey, 2008).

A key assumption in IRL is that a behavioral model for the agent is known. The most common behavior model is that the agent has acted optimally with respect to the MDP. Other behavioral assumptions can be adopted, however, such as a probabilistic selection of suboptimal behavior that attempts to model agents observed in the midst of their own learning process. IRL is inherently ill-defined, as more than one reward function can be consistent with an observed optimal policy (one way to see this is to note that any policy is optimal with respect to a reward function that is everywhere zero). The ill-defined nature of the problem typically is addressed by formulating an optimization objective that imposes additional structure or properties on the learned rewards. Two prominent examples of such problems are max-margin IRL (Abbeel & Ng, 2004), in which the rewards are chosen so as to maximize the value function of the observed policy relative to all other policies, and Bayesian IRL (Ramachandran & Amir, 2007), which aims to find maximum *a posteriori* estimate of rewards given priors and a likelihood model.

RL and IRL can be extended to multi-agent settings. Multi-agent reinforcement learning (MRL) is defined in terms of a stochastic game, rather than an MDP as in the single-agent case (see, e.g., Owen, 1968; Shapley, 1953). In a stochastic game, the players are presented with a sequence of games with the joint actions taken in each game determining the transition probabilities to the next game. Hence, games play the role of states in an MDP. The rewards experienced by each player are determined by the payoff matrices of the games and the joint actions (or *bi-policy*) of the players. To define MRL as a computational problem, it is necessary to specify a solution concept, such as a Nash equilibrium or a correlated equilibrium that the players attempt to achieve. MRL algorithms have been developed for a number of equilibria with varying success in terms of theoretical and empirical performance (Abdallah & Lesser, 2008; Cigler & Faltings, 2013; Conitzer & Sandholm, 2007; Greenwald & Hall, 2003; Hu & Wellman, 1998; Littman, 1994, 2001; Sodomka, Hilliard, Littman, & Greenwald, 2013).

The inverse learning problem for MRL, which we term multi-agent inverse reinforcement learning (MIRL), is to estimate the payoffs of a stochastic game given only observations of the actions taken by the players. Like its single-agent counterpart, the MIRL problem must be stated in terms of an assumed behavior model for the agents. Typically, the assumption is that the agents are playing equilibrium strategies. If one assumes that the agents reached the equilibrium by following an MRL algorithm, then it makes sense in MIRL to focus on developing solution methods for problems defined in terms of equilibria that have existing algorithms for the forward problem.

MIRL can be viewed as a generalization of IRL in the sense that the later treats other agents in the system as part of the environment, ignoring the difference between decision-making agents and the passive environment. A financial trading example can be used to illustrate the difference. Yang, Qiao, Beling, Scherer, and Kirilenko (2015) use the reward functions inferred from IRL as a feature space for the purpose of classifying high-frequency trading algorithms in the stock market. This model is reasonable for a typical stock trading market. Usually there are many traders involved and their activities give rise to cancellation

effects that make it reasonable for any one trader to model the collective actions of all the other traders as a stochastic system. However, if the market is dominated by a small number of traders—such as is the case currently with crypto-currency trading—the market should not be modeled as a passive system. Rather, each dominant trader should take the other’s possible strategies into account before making decisions. In such a case, a stochastic game framework (used in MIRL) would be more appropriate than a MDP framework (used in IRL).

MIRL has been studied far less than IRL, though there has been some relevant work on the topic in recent years. Natarajan et al. address multi-agent problems using an IRL model for multiple agents without dealing with interactions or interference among agents (Natarajan, Kunapuli, Judah, Tadepalli, Kersting, & Shavlik, 2010). Waugh, Ziebart, and Bagnell (2011) contribute to the inverse equilibrium problem in the context of simultaneous one-stage games, rather than the sequential stochastic games that are the subject of MIRL. Reddy, Gopikrishna, Zaruba, and Huber (2012) use the concept of a subgame perfect equilibrium (SPE) (Maskin & Tirole, 2001), a refinement of the Nash equilibrium used in dynamic games, to address MIRL for general-sum stochastic games that have the property that each player’s rewards do not depend on the actions of the others. Hadfield-Menell, Dragan, Abbeel, and Russell (2016) introduce a cooperative IRL problem, motivated from an autonomous system design problem, where the robot is required to align its value with those of the humans in its environment in such a way that its actions contribute to the maximization of values for the humans. Their problem is not modeled as a MIRL problem in a stochastic game context. Lin, Beling, and Cogill (2018) develop a Bayesian maximum *a posteriori* (MAP) estimation method. Their work is limited to two-person zero-sum stochastic games and is not applicable to arbitrary general-sum problems because the uniqueness property of the minimax equilibrium in zero-sum games does not carry over to solution concepts, such as the Nash equilibrium, that are important in general-sum games. In particular, the non-uniqueness of equilibria makes it unclear how to specify the likelihood function for a Bayesian inverse learning formulation. Wang and Klabjan (2018) propose an algorithm that takes sub-optimal demonstrations into account, rather than the optimal strategies assumed by Lin et al. (2018), and finds the reward function to minimize the margin between experts’ performance and Nash equilibria-directed results.

This paper studies five two-person general-sum MIRL problems, *uCS-MIRL*, *advE-MIRL*, *cooE-MIRL*, *uCE-MIRL*, and *uNE-MIRL*, each distinguished by a solution concept and corresponding class of equilibrium policies that the observed agents should be assumed to be playing. The first problem, *uCS-MIRL*, is a *cooperative game* in which the agents employ cooperative strategies (CSs) that aim to maximize the sum of their value functions, or the *total game value*. The second and third problems consider circumstances in which two special and unique Nash equilibria (NE) are employed: *advE* is in general a *win-or-lose* equilibrium, but not necessarily for a zero-sum game; *cooE* is such an equilibrium that players maximize their own payoffs by coordinating with others. In the fourth problem, *uCE-MIRL*, the agents are assumed to follow strategies that constitute a *utilitarian correlated equilibrium* (*uCE*), which achieves the maximum total game value among all CEs. In the last problem, *uNE-MIRL*, players are assumed to follow strategies that constitute a NE that maximizes total game value.

The equilibria that we study from an inverse perspective arise in a variety of application contexts. For example, *uCS-MIRL*, *uCE-MIRL*, and *uNE-MIRL* embed solution concepts that correspond to agents trying to achieve a socially efficient outcome (with or without certain constraints) that maximizes the sum of their value functions and is a Pareto optimum, meaning that it is not possible to make one player better off without making the other player worse off (Barr, 2012). These equilibria are of particular interest in welfare economics. For example, *uCS* applies to the situation where all players of a basketball team cooperate to maximize the total points rather than boosting any single player’s performance. *uCE* is motivated from the social governance problem where policy makers are required to design rules to achieve Pareto optimality in social welfare without harming individual interests. Note that *uCE* has been studied from an MRL perspective (Greenwald & Hall, 2003). Despite the fact the *advE* and *cooE-MIRL* equilibria are not guaranteed to exist in every game, they have been studied in MRL (Littman, 2001) and have potential applications. Consider an example in which two power suppliers compete with each other in a local market. Though to each supplier, this is obviously a win-or-lose game, it is usually not likely to evolve to a dominate-or-exit situation. Hence it might be more reasonable to assume the suppliers are playing an *advE* equilibrium rather than a minimax solution to a zero-sum game. As for *cooE*, the classic Stag Hunt game (see, e.g. Skyrms, 2004) is representative of a broad range of social cooperation games. In Stag Hunt there are two hunters, each can chose to hunt hare or stag, with symmetric payoffs. If they both hunt stag(hare), they both will get a payoff of 2(1); and if their targets are different, the one who hunts stag will fail to get anything and the other will get a payoff of 1. In this game, (stag, stag) is a *cooE*.

In addition to their importance in applications, the five solution concepts that we consider each has been studied from the MRL (i.e., forward) perspective and computational algorithms exist that players of these games could, if they chose, follow to reach an equilibrium solution. Equilibrium *uCS* is actually an extension of RL to a multi-agent context in which the RL optimality concept is still valid. Equilibria *cooE*, *advE* and *uCE* have been well studied as a forward problem and corresponding Q-learning algorithms that guarantee convergence to these equilibria have been developed (Greenwald & Hall, 2003; Hu & Wellman, 1998; Littman, 2001). Equilibrium *uNE* has similar properties as *uCE* and is also computationally achievable. Hence, each of the five MRL problems (*uCS-MIRL*, *advE-MIRL*, *cooE-MIRL*, *uCE-MIRL*, *uNE-MIRL*) that we study is based on an equilibrium concept that might be followed by a pair of decision-makers playing a stochastic game. Given that these problems cover cooperative, semi-cooperative and fully competitive games, we argue that they represent a reasonable starting point for the study of MRL in general-sum, stochastic games.

The principal contributions of this paper lie in framing and deriving solution methods for five MRL problems. For *uCS/advE/cooE-MIRL*, the key step is the development of a set of linear constraints on the reward function that are necessary and sufficient for the observed bi-policy to be a unique equilibrium of the assumed type. We then show that, by using these conditions in a Bayesian MAP-estimation formulation similar to that developed by Lin et al. (2018) for zero-sum games, *uCS/advE/cooE-MIRL* can be solved as a convex optimization problem. An alternate approach, for which we do not provide details, would be to use the same necessary and sufficient conditions as the basis for a max-margin formulation similar to those seen in IRL (Natarajan et al., 2010; Ng & Russell, 2000; Reddy et al., 2012).

Linear conditions specifying a unique equilibrium are not known for uCE/uNE. This circumstance forces us to abandon the MAP estimation formulation used in uCS/advE/cooE-MIRL and adopt an equilibrium selection approach. For uCE-MIRL, we derive linear conditions on reward that are necessary and sufficient for the observed bi-policy to be a CE, a class of equilibria that includes the unique uCE. These conditions are used in a linear program with a novel objective function that incorporates both the scale of the solution and a metric on the total game value difference between the observed bi-policy and a notion of a local equilibrium. We apply a similar treatment to the uNE-MIRL problem. From a high-level perspective, the idea behind our uCE/uNE-MIRL algorithm completely departs from previous ones from which many IRL/MIRL algorithms stem, such as margin-maximization, posterior maximization and entropy maximization. The ideas developed in this paper—while limited to five problems and to two-player settings—could be extended to other uniquely existing solution concepts and to n -player stochastic games.

The remainder of this paper is structured as follows. Section 2 introduces notation, terminology and definitions that will be used throughout this paper, as well as some basic game theory equilibrium concepts. Section 3 summarizes several conventional MIRL algorithms. Section 4 provides the main technical work, developing different approaches for different problems to learn rewards. Section 5 and Section 6 demonstrate our approaches through several benchmark experiments that include comparison with existing MIRL algorithms. Taking uNE-MIRL algorithm as an example, Section 7 explores the robustness of our algorithms given incomplete observations. Section 8 offers concluding remarks and a discussion of future work.

2. Preliminaries

This section serves to introduce concepts and notation for MRL that will be used throughout the paper. It also introduces relevant concepts and formalism for two-player general-sum games and the equilibria of interest in later sections.

2.1 Stochastic Game

A two-player general-sum discounted stochastic game is a tuple $\{\mathcal{S}, \mathcal{A}_i, R_i, P, \gamma\}$, where \mathcal{S} is the common state space for all players, \mathcal{A}_i and R_i are the action space and reward for player i , respectively. P is the probabilistic function controlling state transitions, conditioned on the past state and joint actions. The reward discount factor is $\gamma \in [0, 1)$. In this paper, we assume that both players share the same action space. The state and action spaces are both finite, i.e. $|\mathcal{S}| = N$ and $|\mathcal{A}_i| = M$. A stochastic game is a sequence of single-stage games, or *subgames*, induced in every state $s \in \mathcal{S}$, such that both players need to determine an *individual strategy* $\pi_i(s)$ (in a non-cooperative game) or coordinate a *bi-strategy* $\pi(s)$ (in a cooperative game) that guides their actions in every subgame. The collection of all bi-strategies is a *bi-policy* $\pi = (\pi_1, \pi_2) = (\pi_1(s), \pi_2(s)), \forall s \in \mathcal{S}$. Note that an individual strategy can be a mixed strategy, which is a probability distribution over all available actions. We define a *pure* bi-strategy $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ as a bi-strategy where both players select deterministic actions. Each player’s reward values are assumed dependent on state and possibly, bi-strategies, but are independent of each other.

2.2 Multi-agent Problems in RL/IRL Context

We now introduce some basic terms, notations and fundamental equations that will be used throughout this paper. The derivations, developed by Lin et al. (2018) are omitted here. $\tilde{r}_i^\pi(s)$ denotes the *expected reward value* received by agent i at state s under bi-policy π , specifically,

$$\begin{aligned}\tilde{r}_i^\pi(s) &= \sum_a \pi_1(a_1|s) \pi_2(a_2|s) R_i(s, a) \\ &= [\pi_1(s)]^T R_i(s) \pi_2(s), \forall s \in \mathcal{S},\end{aligned}\tag{1}$$

where a_i is one of player i 's available actions, and a is a pure bi-strategy. $\pi_i(s)$ is a $M \times 1$ vector denoting the probability distribution over actions in state s , each entry of which is the probability of taking action a_i at state s , denoting $\pi_i(a_i|s)$. $R_i(s)$ is a $M \times M$ matrix, each entry of which denotes a pure bi-strategy dependent reward value. Structuring all $R_i(s, a)$ into a column vector as r_i , we can simplify and represent (1) in matrix notation as

$$\tilde{r}_i^\pi = B_\pi r_i.\tag{2}$$

The linear transformation operator B_π is a $N \times NM^2$ matrix constructed from π , whose k th row is:

$$[\Phi_{1,1}^\pi(k), \Phi_{1,2}^\pi(k), \dots, \Phi_{M,M}^\pi(k)],$$

where

$$\Phi_{i,j}^\pi(k) = \left[\underbrace{0, \dots, 0}_{k-1}, \phi_{i,j}^\pi(k), \underbrace{0, \dots, 0}_{N-k} \right],$$

and

$$\phi_{i,j}^\pi(k) = \pi^1(i|k) \pi^2(j|k).$$

Player i 's *value function*, starting at state s and under π , is defined as

$$V_i^\pi(s) = \sum_{t=0}^{\infty} \gamma^t E(\tilde{r}_i^\pi(s_t) | s_0 = s),\tag{3}$$

and its Q -function, upon s and a , is

$$Q_i^\pi(s, a) = r_i(s, a) + \gamma P_{s,a} V_i^\pi,\tag{4}$$

where $P_{s,a}$ is a row vector, each entry of which represents the transition probability from s to all possible state given action a . Furthermore,

$$V_i^\pi(s) \in \text{solution concept}_i(Q_1^\pi(s), Q_2^\pi(s)), \forall s \in \mathcal{S}.\tag{5}$$

Presumably, players agree on a solution concept and choose individual strategies accordingly. In order to pose an inverse learning problem, the solution concept must be known to the game observers estimating the reward function because it cannot be inferred or observed when provided the actions of the players.

Let G_π denote a $N \times N$ transition matrix under bi-policy π , with elements

$$g_\pi(s'|s) = \sum_a \pi_1(a_1|s) \pi_2(a_2|s) p(s'|s, a). \quad (6)$$

Then

$$V_i^\pi = (I - \gamma G_\pi)^{-1} B_\pi r_i, \quad (7)$$

$$\vec{Q}_i^\pi = \left(I + \gamma P (I - \gamma G_\pi)^{-1} B_\pi \right) r_i. \quad (8)$$

V_i^π can be rewritten more compactly as

$$V_i^\pi = B_\pi \vec{Q}_i^\pi. \quad (9)$$

Lastly, we define the *total game value* of a two-player stochastic game starting at state s , under a bi-policy π , $V^\pi(s)$, as the sum of the value functions of both players, i.e., $V^\pi(s) = V_1^\pi(s) + V_2^\pi(s)$.

2.3 Non-cooperative Equilibrium

In non-cooperative game theory, *Nash equilibrium* (NE) and *correlated equilibrium* (CE) are two of the most important solution concepts. A NE is an equilibrium where no player will benefit from unilaterally deviating from their current strategy given the other players' strategies remain unchanged (Nash, 1950, 1951). In a two-player single-stage game (state s), $\pi(s)$ is a NE if and only if,

$$R_i(s, \pi(s)) \geq R_i(s, a_i, \pi_{-i}(s)), a_i \in \mathcal{A}_i, \quad (10)$$

where $\pi_{-i}(s)$ denotes the other player's strategy at state s . Unlike NE, in which all agents act *independently* on a *selfish* and *conservative* basis, a CE is a probability distribution over the joint space of actions, in which all agents optimize their payoff with respect to one another's probabilities, conditioned on their own probabilities (Greenwald & Hall, 2003). Let $\Delta(\mathcal{A})$ denote the set of probability distributions over \mathcal{A} , and X be a random variable taking values in $\mathcal{A} = \prod_{i \in \mathcal{I}} \mathcal{A}_i$ distributed according to $\pi \in \Delta(\mathcal{A})$. Then π is a correlated equilibrium if and only if (Ozdaglar, 2010)

$$\sum_{a_{-i} \in \mathcal{A}_{-i}} P(X = a | X_i = a_i) [R_i(s, a_i, a_{-i}) - R_i(s, \check{a}_i, a_{-i})] \geq 0,$$

for all $a_i \in \mathcal{A}_i$ such that $P(X_i = a_i) > 0$ and all $\check{a}_i \in \mathcal{A}_i \setminus a_i$.

It has been proved that every finite game has at least one NE (Nash, 1951; Owen, 1968), as well as at least one CE (Hart & Schmeidler, 1989). In fact, CE is a superset of NE (Aumann, 1974), and hence for any general sum game, the number of CEs is larger than or equal to the number of NEs. In regard to equilibrium search, finding a NE or determining the number of them is a NP-hard problem (Daskalakis, Goldberg, & Papadimitriou, 2009), while finding CEs can be done in polynomial time via linear programming (Papadimitriou & Roughgarden, 2008). Nevertheless, the non-uniqueness property causes the non-convergence issue and is a bottleneck for MRL/MIRL research based on these two equilibria.

2.4 Cooperative Strategy

Both NE and CE are equilibria of competitive games. In a cooperative game, by contrast, an agreement on a joint strategy for all players can be called a *cooperative strategy* (CS). A *characteristic function* v defines the type of cooperation between players (Ferguson, 2008), and for a two-player single-stage game (state s), can be defined as

$$v(s, a) = \text{Val}(R_1(s, a), R_2(s, a)), a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2. \quad (11)$$

$\text{Val}(\cdot)$ is self-defined, based on the type of cooperation.

3. Conventional MIRL Approaches

Before introducing our approaches, we review several existing approaches to MIRL and related problems. The first of these is a *decentralized* MIRL (d-MIRL) approach developed by Reddy et al. (2012). This approach is decentralized in the sense that it infers agents' rewards one by one, rather than all at once. All agents are assumed to follow a Nash equilibrium at every single game. The key idea is to find a reward that maximize the difference between the Q value of the observed policy and those of pure strategies, which is analogous to the classical approach to single-agent IRL proposed by Ng and Russell (2000). Though in the original version of their approach reward is assumed dependent only on the state, we can extend it to treat action dependency as well. Using our notation and player 1 as an example, the d-MIRL approach to a two-person general-sum MIRL problem is to solve the following linear program:

$$\begin{aligned} \text{maximize:} \quad & \sum_{s=1}^N \min_{a_1} (B_\pi - B_{\pi|a_1}) D_\pi r_1 - \lambda \|r_1\|_1 \\ \text{subject to:} \quad & (B_{\pi|a_1} - B_\pi) D_\pi r_1 \leq 0, \end{aligned}$$

where D_π is defined in (19) in the following section, λ is an adjustable penalty coefficient for having too many non-zero values in the reward vector.

The key idea of the second approach is to model a two-person general-sum MIRL as an IRL problem. This approach requires us to select one player (e.g. player 1) and treat the other player as part of the passive environment. A Bayesian IRL (BIRL) algorithm developed by Lin et al. (2018) extends Qiao and Beling's (2011) work by considering action-dependent reward cases in addition to state-dependent reward cases. Note that the reward of player 1 to be recovered is $R_1(s, a_1)$ instead of $R_1(s, a_1, a_2)$, as player 2 is not considered adaptive. That is to say, $R_1(s, a_1, a_2) = R_1(s, a_1)$ for all $a_2 \in \mathcal{A}_2$. Using our notation, the approach to recover player 1's reward is:

$$\begin{aligned} \text{minimize:} \quad & \frac{1}{2} (r_1 - \mu_{r_1})^T \Sigma_{r_1}^{-1} (r_1 - \mu_{r_1}) \\ \text{subject to:} \quad & (F_{a_1}^{\pi_1} - C_{a_1}) r_1 \geq 0, \end{aligned} \quad (12)$$

for all $a_1 \in \mathcal{A}_1$, where

$$F_{a_1}^{\pi_1} = \left[\gamma (G_\pi - G_{\pi_1|a_1}) (I - \gamma G_\pi)^{-1} + I \right] C_{\pi_1},$$

and C_{π_1} is a $N \times NM$ sparse matrix constructed from π_1 , whose i th row is,

$$\left[\underbrace{0, \dots, \pi^1(i, 1), \dots, 0}_N, \underbrace{\dots}_{(M-2)N}, \underbrace{0, \dots, \pi^1(i, M), \dots, 0}_N \right],$$

and C_{a_1} is conceptually similar to C_{π_1} , except for being constructed from a pure strategy a_1 for all states.

Strictly speaking, the BIRL approach is not a dedicated algorithm for MIRL problems but rather a way of shoehorning the multi-agent problem into a single-agent IRL setting. BIRL will provide a useful point of comparison to quantify the benefits of explicitly modeling the decisions of all players.

The third approach is not applicable to a general-sum MIRL problem but a restricted family: zero-sum games. It is

$$\begin{aligned} \text{minimize:} & \quad \frac{1}{2} (r - \mu_r)^T \Sigma_r^{-1} (r - \mu_r) \\ \text{subject to:} & \quad (B_{\pi|a_1} - B_{\pi}) D_{\pi} r \leq 0 \\ & \quad (B_{\pi|a_2} - B_{\pi}) D_{\pi} r \geq 0, \end{aligned} \tag{13}$$

for all $a_1 \in \mathcal{A}_1$ and $a_2 \in \mathcal{A}_2$. Lin et al. (2018) provide more details.

These three approaches will be revisited as benchmarks in later sections.

4. MIRL Model Development

This section proposes five two-player general-sum MIRL problems and corresponding approaches to them. We first informally define an MIRL problem. Given a bipolicy π being played in a two-player, general-sum game with states, actions, dynamics, and discount $\{\mathcal{S}, \mathcal{A}_i, P, \gamma\}$, the MIRL problem is to find rewards r_1, r_2 that best explain the observed policy. Though we will not do so in this paper, MIRL may defined in terms of an set of observed state-action values \mathcal{O} rather a bipolicy.

The MRL literature suggests that an agreement over a specific solution concept may be needed to solve a MRL problem. Similarly, in our approaches to MIRL, one basic assumption is required: both players agree on a specific strategy/equilibrium to play and this information is available to the coordinator in posing the MIRL problem. We limit attention to the following five solution concepts:

1. **utilitarian Cooperative Strategy (uCS)**. In (11), consider $\text{Val}(\cdot) = \sum(\cdot)$. A single-stage game in state s and taking action a is a *utilitarian* cooperative strategy (uCS) if and only if

$$\sum_i R_i(s, a) \geq \sum_i R_i(s, a'), a' \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \setminus a. \tag{14}$$

2. **Adversarial Equilibrium (advE)** An advE is a type of NE. It has another feature that no player is hurt by any change of others (Hu & Wellman, 1998; Littman, 2001;

Piotrowski, Ślaskowski, & Szczypińska, 2010). That is to say, in a two-player single-stage game (state s), $\pi(s)$ is an advE if and only if, in addition to (10),

$$R_i(s, \pi_i(s), \pi_{-i}(s)) \leq R_i(s, \pi_i(s), a_{-i}), a_{-i} \in \mathcal{A}_{-i}, \quad (15)$$

3. **Coordination Equilibrium (cooE)**. A cooE is also a type of NE. It has another feature that all players' maximum expected payoffs are achieved (Hu & Wellman, 1998; Littman, 2001; Piotrowski et al., 2010). Mathematically, in a two-player single-stage game (state s), $\pi(s)$ is a cooE if and only if, in addition to (10),

$$R_i(s, \pi(s)) \geq R_i(s, a), a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2. \quad (16)$$

4. **utilitarian Correlated Equilibrium (uCE)**. We borrow the concept of *utilitarian* correlated equilibrium (uCE) defined by Greenwald and Hall (2003) and state that in a two-player single-stage game (state s), $\pi(s)$ is a uCE if and only if,

$$\sum_i R_i(s, \pi(s)) \geq \sum_i R_i(s, \tilde{\pi}(s)), \tilde{\pi}(s) \in \pi_{\text{CE}} \setminus \pi(s). \quad (17)$$

5. **utilitarian Nash Equilibrium (uNE)**. Similar to uCE, in a two-player single-stage game (state s), a NE $\pi(s)$ is a *utilitarian* Nash Equilibrium (uNE) if and only if

$$\sum_i R_i(s, \pi(s)) \geq \sum_i R_i(s, \tilde{\pi}(s)), \tilde{\pi}(s) \in \pi_{\text{NE}} \setminus \pi(s). \quad (18)$$

Among the above five equilibria, it is easy to show that uCS, uCE and uNE always exist and are unique in any games (uCS for cooperative while uCE and uNE for noncooperative). Both advE and cooE have been shown to be unique in a noncooperative game, though neither of them is guaranteed to exist (Hu & Wellman, 1998; Littman, 2001) in any games.

The distinctions between cooE and uCS are worth noting. Intuitively, cooE is a non-cooperative game equilibrium, which means that agents are essentially selfish. They are *forced* to cooperate in order to maximize their individual benefits. When following a uCS, by contrast, agents cooperate with each other *actively* and may even sacrifice their own benefits to achieve a better overall outcome. Section 5 will help illustrate the differences.

4.1 Extension to Stochastic Games

Filar and Vrieze (1996) show how the Q function links a stochastic game to a single stage game. Q functions at one particular state with different bi-strategy are treated as payoffs for that particular single stage game (note the terms “game” and “state” can be used interchangeably), and the stochastic game is said to be *in an equilibrium* if and only if all single games (over all states) are in equilibrium. We now extend our definitions of the five strategies/equilibria from a single game to a two-player stochastic game, as follows,

Definition 4.1. *A bi-policy π is a uCS/advE/cooE/uNE/uCE of a two-player stochastic game \mathcal{G} if only if $\pi(s)$ is a uCS/advE/CooE/uNE/uCE of its sub-game $\mathcal{G}(s)$, for all $s \in \mathcal{S}$.*

Correspondingly, we define that a uCS/advE/cooE/uNE/uCE-MIRL problem is an MIRL problem in which the players are assumed to employ a uCS/advE/CooE/uNE/uCE.

4.2 uCS-MIRL

A main result characterizing the set of solutions to a two-player uCS-MIRL problem is the following:

Theorem 4.1. *Given a two-player stochastic game $\{\mathcal{S}, \mathcal{A}_i, r_i, P, \gamma\}$, an observed bi-policy π is a uCS if and only if*

$$(B_\pi - B_a) D_\pi (r_1 + r_2) \geq 0, a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \quad (19)$$

where $D_\pi = I + \gamma P (I - \gamma G_\pi)^{-1} B_\pi$. B_a is obtained from such a bi-policy that players employ the bi-strategy a in all states.

Proof. According to the definition of uCS, π is a uCS if and only if, for any state s and pure bi-strategy $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$, we have

$$\begin{aligned} & \pi(s) \in \arg \max_{a \in \mathcal{A}} \sum_i Q_i^\pi(s, a) \\ \Leftrightarrow & \sum_i Q_i^\pi(s, \pi(s)) \geq \sum_i Q_i^\pi(s, a) \\ \Leftrightarrow & r_1(s, \pi(s)) + r_2(s, \pi(s)) + \gamma P_{s, \pi(s)} (V_1^\pi + V_2^\pi) \\ & \geq r_1(s, a) + r_2(s, a) + \gamma P_{s, a} (V_1^\pi + V_2^\pi) \quad (20) \\ \Leftrightarrow & B_\pi (r_1 + r_2) + \gamma B_\pi P (I - \gamma G_\pi)^{-1} B_\pi (r_1 + r_2) \\ & \geq B_a (r_1 + r_2) + \gamma B_a P (I - \gamma G_\pi)^{-1} B_\pi (r_1 + r_2) \\ \Leftrightarrow & (B_\pi - B_a) \left(I + \gamma P (I - \gamma G_\pi)^{-1} B_\pi \right) (r_1 + r_2) \geq 0 \\ \Leftrightarrow & (B_\pi - B_a) D_\pi (r_1 + r_2) \geq 0 \end{aligned}$$

□

Since any solution that is consistent with (19) ensures a unique uCS, we can borrow the idea introduced by Lin et al. (2018) and propose a Bayesian approach. The general idea is to maximize the posterior probability of the inferred rewards, $p(r_1, r_2 | \pi)$, which can be expressed as

$$p(r_1, r_2 | \pi) \propto f(r_1, r_2) p(\pi | r_1, r_2), \quad (21)$$

where $p(\pi | r_1, r_2)$ is the likelihood of observing π given r_1 and r_2 and $f(r_1, r_2)$ is a joint prior of r_1 and r_2 that we need to specify. Recall the assumption that

$$f(r_1, r_2) = f(r_1) f(r_2), \quad (22)$$

which allows specification of the prior over r_1 and r_2 independently. We adopt a Gaussian prior on both rewards; that is, $r_i \sim \mathcal{N}(\mu_{r_i}, \Sigma_{r_i})$, where μ_{r_i} is the mean of r_i and Σ_{r_i} is the covariance. Then the probability density function of r_i is

$$f(r_i) = \frac{1}{(2\pi)^{N/2} |\Sigma_{r_i}|^{1/2}} \exp\left(-\frac{1}{2} (r_i - \mu_{r_i})^T \Sigma_{r_i}^{-1} (r_i - \mu_{r_i})\right), i = 1, 2. \quad (23)$$

To model the likelihood function $p(\pi|r_1, r_2)$, assume that the bi-policy which the two agents follow is a unique uCS given r_1, r_2 . The likelihood is then a probability mass function given by

$$p(\pi|r_1, r_2) = \begin{cases} 1, & \text{if } \pi \text{ is uCS for } r_1, r_2 \\ 0, & \text{otherwise.} \end{cases} \quad (24)$$

Thus, the optimization problem for uCS-MIRL can be formulated as

$$\begin{aligned} & \text{maximize: } f(r_1, r_2) \\ & \text{subject to: } p(\pi|r_1, r_2) = 1. \end{aligned} \quad (25)$$

Equivalently,

$$\begin{aligned} & \text{minimize: } \frac{1}{2} \sum_i (r_i - \mu_{r_i})^T \Sigma_{r_i}^{-1} (r_i - \mu_{r_i}) \\ & \text{subject to: } (B_\pi - B_{\pi|a}) D_\pi (r_1 + r_2) \geq 0, a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2. \end{aligned} \quad (26)$$

4.3 advE-MIRL

The main result characterizing the set of solutions to a two-player advE-MIRL problem is the following:

Theorem 4.2. *Given a two-player stochastic game $\{\mathcal{S}, \mathcal{A}_i, r_i, P, \gamma\}$, the observed bi-policy π is an advE if and only if*

$$\begin{aligned} & (B_{\pi|a_1} - B_\pi) D_\pi r_1 \leq 0, \forall a_1 \in \mathcal{A}_1 \\ & (B_{\pi|a_2} - B_\pi) D_\pi r_2 \leq 0, \forall a_2 \in \mathcal{A}_2 \\ & (B_{\pi|a_1} - B_\pi) D_\pi r_2 \geq 0, \forall a_1 \in \mathcal{A}_1 \\ & (B_{\pi|a_2} - B_\pi) D_\pi r_1 \geq 0, \forall a_2 \in \mathcal{A}_2, \end{aligned} \quad (27)$$

where $B_{\pi|a_1}$ is obtained from such a bi-policy that player 2 employs their original policy while player 1 always chooses action a_1 in any state (game).

Proof. Eqs (27) contain four inequalities. In this proof, we will first show that the first and second inequalities constitute a necessary and sufficient condition for π being a NE. Recall that a bi-policy π is a minimax equilibrium for a two-player zero-sum game if and only if (Lin et al., 2018)

$$\begin{aligned} & [Q^\pi(s)]^T \pi_1(s) \geq V^\pi(s) \mathbf{1}_M \\ & Q^\pi(s) \pi_2(s) \leq V^\pi(s) \mathbf{1}_M. \end{aligned} \quad (28)$$

Similarly, π is a NE if and only if

$$\begin{aligned} & [Q_2^\pi(s)]^T \pi_1(s) \leq V_2^\pi(s) \mathbf{1}_M \\ & Q_1^\pi(s) \pi_2(s) \leq V_1^\pi(s) \mathbf{1}_M. \end{aligned} \quad (29)$$

Combining (9) and (29) leads to

$$\begin{aligned} & B_{\pi|a_2} \vec{Q}_2^\pi \leq B_\pi \vec{Q}_2^\pi, \forall a_2 \in \mathcal{A}_2 \\ & B_{\pi|a_1} \vec{Q}_1^\pi \leq B_\pi \vec{Q}_1^\pi, \forall a_1 \in \mathcal{A}_1. \end{aligned} \quad (30)$$

Substituting (8) into (30) and rearranging the two sides of the inequalities yields

$$\begin{aligned} (B_{\pi|a_1} - B_{\pi}) D_{\pi} r_1 &\leq 0, \forall a_1 \in \mathcal{A}_1 \\ (B_{\pi|a_2} - B_{\pi}) D_{\pi} r_2 &\leq 0, \forall a_2 \in \mathcal{A}_2. \end{aligned} \quad (31)$$

We now turn to the additional feature of advE. Recall (15), it is easy to derive that a bi-policy π for a two-player general-sum game is an advE, if and only if, in addition to (29)

$$\begin{aligned} [Q_1^{\pi}(s)]^T \pi_1(s) &\geq V_1^{\pi}(s) \mathbf{1}_M \\ Q_2^{\pi}(s) \pi_2(s) &\geq V_2^{\pi}(s) \mathbf{1}_M. \end{aligned} \quad (32)$$

Following similar steps to those used to derive (30), the additional constraints (32) can be reduced to

$$\begin{aligned} (B_{\pi|a_1} - B_{\pi}) D_{\pi} r_2 &\geq 0, \forall a_1 \in \mathcal{A}_1 \\ (B_{\pi|a_2} - B_{\pi}) D_{\pi} r_1 &\geq 0, \forall a_2 \in \mathcal{A}_2. \end{aligned} \quad (33)$$

□

Since it has been proved that, in a one-stage game, if an advE exists it must be unique (Littman, 2001), an advE for a stochastic game, must also be unique, if it exists. Therefore, we can still use a Bayesian approach to solve advE-MIRL problems. The prior (23) is also valid here but the likelihood is modified as follows

$$p(\pi|r_1, r_2) = \begin{cases} 1, & \text{if } \pi \text{ is an AdvE for } r_1, r_2 \\ 0, & \text{otherwise.} \end{cases} \quad (34)$$

And the optimization problem for advE-MIRL is

$$\begin{aligned} \text{minimize: } & \frac{1}{2} \sum_i (r_i - \mu_{r_i})^T \Sigma_{r_i}^{-1} (r_i - \mu_{r_i}) \\ \text{subject to: } & (B_{\pi|a_1} - B_{\pi}) D_{\pi} r_1 \leq 0, \forall a_1 \in \mathcal{A}_1 \\ & (B_{\pi|a_2} - B_{\pi}) D_{\pi} r_2 \leq 0, \forall a_2 \in \mathcal{A}_2 \\ & (B_{\pi|a_1} - B_{\pi}) D_{\pi} r_2 \geq 0, \forall a_1 \in \mathcal{A}_1 \\ & (B_{\pi|a_2} - B_{\pi}) D_{\pi} r_1 \geq 0, \forall a_2 \in \mathcal{A}_2. \end{aligned} \quad (35)$$

In fact, there is a direct link between the minimax equilibrium of a competitive zero-sum game and an advE for a special zero-sum case, as the following proposition,

Proposition 4.1. *The minimax equilibrium of a single competitive zero-sum game is an advE, and vice versa.*

Proof. Let $r_1 = r = -r_2$. Then (27) reduces to

$$\begin{aligned} (B_{\pi|a_1} - B_{\pi}) D_{\pi} r &\leq 0, \forall a_1 \in \mathcal{A}_1 \\ (B_{\pi|a_2} - B_{\pi}) D_{\pi} r &\geq 0, \forall a_2 \in \mathcal{A}_2, \end{aligned} \quad (36)$$

which are exactly the constraints of (13), the necessary and sufficient conditions for π being a minimax equilibrium for a zero-sum game (see Lin et al., 2018). □

From Proposition 4.1 we can see that a advE is a more general concept for general-sum games, whereas the minimax equilibrium is specific to zero-sum games.

4.4 cooE-MIRL

The main result characterizing the set of solutions to a two-player cooE-MIRL problem is the following:

Theorem 4.3. *Given a two-player stochastic game $\{\mathcal{S}, \mathcal{A}_i, r_i, P, \gamma\}$, the observed bi-policy π is an CooE if and only if*

$$\begin{aligned}
 (B_{\pi|a_1} - B_{\pi}) D_{\pi} r_1 &\leq 0, \forall a_1 \in \mathcal{A}_1 \\
 (B_{\pi|a_2} - B_{\pi}) D_{\pi} r_2 &\leq 0, \forall a_2 \in \mathcal{A}_2 \\
 (B_{\pi} - B_a) D_{\pi} r_1 &\geq 0, \forall a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \\
 (B_{\pi} - B_a) D_{\pi} r_2 &\geq 0, \forall a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2.
 \end{aligned} \tag{37}$$

In (37), the first two inequalities, which guarantee π is a NE, have been established in Section 4.3. The latter two inequalities warrant the unique property of cooE, the proof of which is outlined below.

Proof. According to the definition of cooE, π is a cooE if and only if, for any state s and pure bi-strategy $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$,

$$\begin{aligned}
 \pi(s) &\in \arg \max_{a \in \mathcal{A}} Q_i^{\pi}(s, a) \\
 \Leftrightarrow Q_i^{\pi}(s, \pi(s)) &\geq Q_i^{\pi}(s, a) \\
 \Leftrightarrow r_i(s, \pi(s)) + \gamma P_{s, \pi(s)} V_i^{\pi} &\geq r_i(s, a) + \gamma P_{s, a} V_i^{\pi} \\
 \Leftrightarrow B_{\pi} r_i + \gamma B_{\pi} P (I - \gamma G_{\pi})^{-1} B_{\pi} r_i & \\
 &\geq B_a r_i + \gamma B_a P (I - \gamma G_{\pi})^{-1} B_{\pi} r_i \\
 \Leftrightarrow (B_{\pi} - B_a) \left(I + \gamma P (I - \gamma G_{\pi})^{-1} B_{\pi} \right) r_i &\geq 0 \\
 \Leftrightarrow (B_{\pi} - B_a) D_{\pi} r_i &\geq 0.
 \end{aligned} \tag{38}$$

□

Using the same reasoning as in the case of advE, it is easy to show that a cooE for a stochastic game is unique, if it exists. As a result, the Bayesian approach is also valid here, with the same prior (23) but a different likelihood as follows

$$p(\pi|r_1, r_2) = \begin{cases} 1, & \text{if } \pi \text{ is an cooE for } r_1, r_2 \\ 0, & \text{otherwise.} \end{cases} \tag{39}$$

Hence the optimization problem for cooE-MIRL is

$$\begin{aligned}
 \text{minimize: } & \frac{1}{2} \sum_i (r_i - \mu_{r_i})^T \Sigma_{r_i}^{-1} (r_i - \mu_{r_i}) \\
 \text{subject to: } & (B_{\pi|a_1} - B_{\pi}) D_{\pi} r_1 \leq 0, \forall a_1 \in \mathcal{A}_1 \\
 & (B_{\pi|a_2} - B_{\pi}) D_{\pi} r_2 \leq 0, \forall a_2 \in \mathcal{A}_2 \\
 & (B_{\pi} - B_a) D_{\pi} r_1 \geq 0, \forall a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \\
 & (B_{\pi} - B_a) D_{\pi} r_2 \geq 0, \forall a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2.
 \end{aligned} \tag{40}$$

4.5 uCE-MIRL

The result that characterizes the set of solutions to a two-player CE-MIRL problem is as follows:

Theorem 4.4. *Given a two-player stochastic game $\{\mathcal{S}, \mathcal{A}_i, r_i, P, \gamma\}$, the observed bi-policy π is a CE if and only if*

$$\bar{\pi}^T H(s, a_i)^T [H(s, a_i) - H(s, \check{a}_i)] D_\pi r_i \geq 0, i = 1, 2, \forall a_i \in \mathcal{A}_i, \check{a}_i \in \mathcal{A}_i \setminus a_i, \quad (41)$$

where $\bar{\pi}$ is restructured from π to be a column vector of length NM^2 . $H(s, a_i)$ is a sparse matrix of size $M \times NM^2$. Like B_π defined in (2), $H(s, a_i)$ is also a linear transformation operator. Specifically, $[R_1(s, a_1, \cdot)]^T = H(s, a_1) r_1$ and $R_2(s, \cdot, a_2) = H(s, a_2) r_2$. Recall that here $R_1(s, a_1, \cdot)$ is a $1 \times M$ row vector and $R_2(s, \cdot, a_2)$ is a $M \times 1$ column vector. Both r_1 and r_2 are $NM^2 \times 1$ column vectors.

Proof. By definition of CE, for a two-player general-sum stochastic game \mathcal{G} , a bi-policy π is a CE if and only if

$$\begin{aligned} \sum_{a_2} \pi(a_1, a_2 | s) Q_1^\pi(s, a_1, a_2) &\geq \sum_{a_2} \pi(a_1, a_2 | s) Q_1^\pi(s, \check{a}_1, a_2), \forall a_1 \in \mathcal{A}_1, \check{a}_1 \in \mathcal{A}_1 \setminus a_1 \\ \sum_{a_1} \pi(a_1, a_2 | s) Q_2^\pi(s, a_1, a_2) &\geq \sum_{a_1} \pi(a_1, a_2 | s) Q_2^\pi(s, a_1, \check{a}_2), \forall a_2 \in \mathcal{A}_2, \check{a}_2 \in \mathcal{A}_2 \setminus a_2, \end{aligned} \quad (42)$$

for all $s \in \mathcal{S}$. Rearranging (42) yields

$$\begin{aligned} \pi(a_1, \cdot | s) \left([Q_1^\pi(s, a_1, \cdot)]^T - [Q_1^\pi(s, \check{a}_1, \cdot)]^T \right) &\geq 0 \\ [\pi(\cdot, a_2 | s)]^T (Q_2^\pi(s, \cdot, a_2) - Q_2^\pi(s, \cdot, \check{a}_2)) &\geq 0, \end{aligned} \quad (43)$$

where $\pi(a_1, \cdot | s)$ is a row vector of $1 \times M$, spanning over all $a_2 \in \mathcal{A}_2$, and $\pi(\cdot, a_2 | s)$ is a column vector of $M \times 1$, spanning over all $a_1 \in \mathcal{A}_1$. Recall

$$Q_i^\pi(s, a) = R_i(s, a) + \gamma \sum_{s'} p(s' | s, a) V_i^\pi(s'). \quad (44)$$

So

$$\begin{aligned} [Q_1^\pi(s, a_1, \cdot)]^T &= [R_1(s, a_1, \cdot)]^T + \gamma p(\cdot | s, a_1, \cdot) V_1^\pi \\ Q_2^\pi(s, \cdot, a_2) &= R_2(s, \cdot, a_2) + \gamma p(\cdot | s, \cdot, a_2) V_2^\pi. \end{aligned} \quad (45)$$

Substituting (45) into (43) leads to

$$\begin{aligned} \pi(a_1, \cdot | s) \left\{ [R_1(s, a_1, \cdot)]^T - [r_1(s, \check{a}_1, \cdot)]^T + \gamma [p(\cdot | s, a_1, \cdot) - p(\cdot | s, \check{a}_1, \cdot)] V_1^\pi \right\} &\geq 0 \\ [\pi(\cdot, a_2 | s)]^T \left\{ R_2(s, \cdot, a_2) - r_2(s, \cdot, \check{a}_2) + \gamma [p(\cdot | s, \cdot, a_2) - p(\cdot | s, \cdot, \check{a}_2)] V_2^\pi \right\} &\geq 0. \end{aligned} \quad (46)$$

The above inequality can be further simplified. It is also easy to see $p(\cdot | s, a_1, \cdot) = H(s, a_1) P$ and $p(\cdot | s, \cdot, a_2) = H(s, a_2) P$. In addition, we can also have $\pi(a_1, \cdot | s) = [H(s, a_1) \bar{\pi}]^T = \bar{\pi}^T H(s, a_1)^T$, and $\pi(\cdot, a_2 | s) = H(s, a_2) \bar{\pi}$. Substituting (7) into (46) and rearranging it, we can get

$$\bar{\pi}^T H(s, a_i)^T [H(s, a_i) - H(s, \check{a}_i)] \left(I + \gamma P (I - \gamma G_\pi)^{-1} B_\pi \right) r_i \geq 0, i = 1, 2, \quad (47)$$

Recall

$$D_\pi = I + \gamma P (I - \gamma G_\pi)^{-1} B_\pi, \quad (48)$$

we can express (47) compactly as

$$\bar{\pi}^T H(s, a_i)^T [H(s, a_i) - H(s, \check{a}_i)] D_\pi r_i \geq 0, i = 1, 2, \forall a_i \in \mathcal{A}_i, \check{a}_i \in \mathcal{A}_i \setminus a_i, \quad (49)$$

□

Any sensible point that is consistent with (49) constitutes a CE for the stochastic game. Many points in the convex hull of CE, however, are less meaningful because only the uCE is of interest. Hence, we desire to find some way to choose between solutions satisfying (49). A first idea is to maximize $\sum_s V^\pi(s)$. Finding a uCS is a much easier problem, by contrast. This fact gives rise to another idea. Before going into details, we introduce four concepts: *cooperation gap*, *local uCS*, *local improvement* and *local reduced gap*.

Definition 4.2. *The cooperation gap $I_{cg}^\pi(s)$, corresponding to a starting state s and a bi-policy π in a two-player general-sum stochastic game, is the total game value difference between π and π^* , where π^* is any uCS; specifically,*

$$I_{cg}^\pi(s) = V^{\pi^*}(s) - V^\pi(s), s \in \mathcal{S}.$$

Definition 4.3. *The local uCS, corresponding to a starting state s and a bi-policy π in a two-player general-sum stochastic game, is a bi-policy for which the two players employ a uCS bi-policy π^* at s and then employ π afterwards.*

Definition 4.4. *The local improvement $I_{imp}^\pi(s)$, corresponding to a starting state s and a bi-policy π in a two-player general-sum stochastic game, is the total game value gain by employing the local uCS.*

Definition 4.5. *The local reduced gap $I_{rg}^\pi(s)$, corresponding to a starting state s and a bi-policy π in a two-player general-sum stochastic game, is the total game value difference between a uCS and a local uCS; specifically,*

$$I_{rg}^\pi(s) = V^{\pi^*}(s) - Q^\pi(s, \pi^*(s)), s \in \mathcal{S},$$

and the total local improvement for π is

$$I_{rg}^\pi = \sum_s I_{rg}^\pi(s) = \sum_s V^{\pi^*}(s) - Q^\pi(s, \pi^*(s)), s \in \mathcal{S}. \quad (50)$$

An implication from the above definitions is that for a starting state s , $I_{rg}^\pi(s) = I_{cg}^\pi(s) - I_{imp}^\pi(s)$, shown in Figure 1.

It is obvious that for a two-player general sum stochastic game, among all its CEs, the uCE is closest to its uCS in terms of the total game value, as illustrated in Figure 1. In a uCE-MIRL problem, however, all CEs except uCE are unobservable. Therefore, we need to find a way to infer a set of rewards $\{r_1, r_2\}$ such that the observed π is most likely the uCE of the game.

By definition, a local uCS improves $V^\pi(s)$ by employing a uCS strategy only at current state s , resulting in a *local improvement* with respect to π and a local reduced gap with

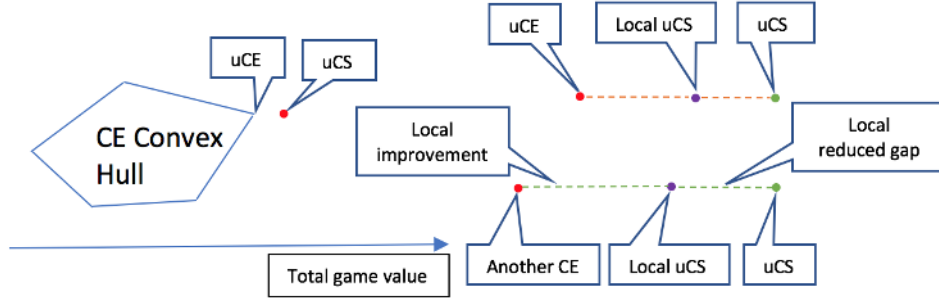


Figure 1: The game value distance relationship between uCE, uCS and other CEs

respect to a uCS. Adding up all those local reduced gaps over all states gives a measure of how close a bi-policy π is to a uCS, in terms of the total game value. We now propose an important theorem that captures the relationship between the local reduced gap and uCE, as follows:

Theorem 4.5. *Consider a two-person general-sum stochastic game Γ with a collection of CEs, Π_{CE} and a bi-policy π^* that is a uCS. Then $\pi_{CE}^* \in \Pi_{CE}$ is a uCE if and only if its total local reduced gap $I_{rg}^{\pi_{CE}^*}$ is no greater than that of any other CE, specifically,*

$$I_{rg}^{\pi_{CE}^*} \leq I_{rg}^{\pi_{CE}}, \forall \pi_{CE} \in \Pi_{CE}. \quad (51)$$

Proof. We first show necessity. From (50) and the properties of value function, we have

$$\begin{aligned} & I_{rg}^{\pi_{CE}^*} - I_{rg}^{\pi_{CE}} \\ &= \left[\sum_s V^{\pi^*}(s) - Q^{\pi_{CE}^*}(s, \pi^*(s)) \right] - \left[\sum_s V^{\pi^*}(s) - Q^{\pi_{CE}}(s, \pi^*(s)) \right] \\ &= \sum_s Q^{\pi_{CE}}(s, \pi^*(s)) - Q^{\pi_{CE}^*}(s, \pi^*(s)) \\ &= \sum_s \left\{ \left[\tilde{r}_1(s, \pi^*(s)) + \tilde{r}_2(s, \pi^*(s)) + \gamma P_{s, \pi^*(s)} V^{\pi_{CE}} \right] \right. \\ &\quad \left. - \left[\tilde{r}_1(s, \pi^*(s)) + \tilde{r}_2(s, \pi^*(s)) + \gamma P_{s, \pi^*(s)} V^{\pi_{CE}^*} \right] \right\} \\ &= \gamma \sum_s P_{s, \pi^*(s)} \left(V^{\pi_{CE}} - V^{\pi_{CE}^*} \right). \end{aligned} \quad (52)$$

Since the definition of π_{CE}^* implies that $V^{\pi_{CE}}(s) \leq V^{\pi_{CE}^*}(s)$ for all s , the column vector $V^{\pi_{CE}} - V^{\pi_{CE}^*}$ is non-positive. Also, $P_{s, \pi^*(s)}$ is a non-negative row vector as all its entries are probabilities. Therefore, $P_{s, \pi^*(s)} (V^{\pi_{CE}} - V^{\pi_{CE}^*}) \leq 0$ for all $s \in \mathcal{S}$, and consequently, $I_{rg}^{\pi_{CE}^*} \leq I_{rg}^{\pi_{CE}}$.

Next, we show sufficiency by assuming $I_{rg}^{\pi_{CE}^*} \leq I_{rg}^{\pi_{CE}}$, for all $\pi_{CE}^* \in \pi_{CE}$ and that π_{CE}^* is not a uCE. Since π_{CE}^* is not a uCE, there must exist a uCE, π_{uCE} , such that $V^{\pi_{uCE}}(s) > V^{\pi_{CE}^*}(s)$, for all $s \in \mathcal{S}$. Then from (52) we can conclude

$$I_{rg}^{\pi_{CE}^*} - I_{rg}^{\pi_{uCE}} = \gamma \sum_s P_{s, \pi^*(s)} \left(V^{\pi_{uCE}} - V^{\pi_{CE}^*} \right) > 0,$$

which contradicts our assumption that $I^{\pi_{\text{CE}}^*} \leq I^{\pi_{\text{CE}}}$, for all $\pi_{\text{CE}} \in \Pi_{\text{CE}}$. \square

The intuition behind Theorem 4.5 is: comparing to any other CE, a uCE is closer to the uCS. Its corresponding local uCS is even closer to uCS and as a result, there is less room to further shrink the local reduced gap. Hence the smaller the local reduced gap is, the more likely a CE π is a uCE. Thus, given a CE π , a desired pair of r_1 and r_2 satisfies

$$\begin{aligned} & \text{minimize: } \sum_s y(s) - V^\pi(s) \\ & \text{subject to: } Q_1^\pi(s, a) + Q_2^\pi(s, a) \leq y(s) \\ & \quad V^\pi(s) \leq y(s), \end{aligned} \tag{53}$$

for all $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$.

However, $r_1 = r_2 = \mathbf{0}$ is the optimal solution to (53). The reason is that $V^\pi(s)$ needs to be enlarged so that picking a uCE is achievable with higher probability. Putting all the above together, we propose the following linear programming problem to find the desired r_1 and r_2 ,

$$\begin{aligned} & \text{maximize: } \sum_s V^\pi(s) - \lambda(y(s) - V^\pi(s)) \\ & \quad + \text{regularization terms} \\ & \text{subject to: } Q_1^\pi(s, a) + Q_2^\pi(s, a) \leq y(s) \\ & \quad V^\pi(s) \leq y(s) \\ & \quad \text{Constraint (49),} \end{aligned} \tag{54}$$

where λ is a regularization parameter. Expressing V_i^π and $Q_i^\pi(s, a)$ as functions of r_i and reformulating those inequalities more compactly in matrix notation leads to

$$\begin{aligned} & \text{maximize: } \mathbf{1}_{1 \times N} \times \left[(1 + \lambda) (I - \gamma G_\pi)^{-1} B_\pi (r_1 + r_2) - \lambda y \right] \\ & \quad + \text{regularization terms} \\ & \text{subject to: } \pi^T H(s, a_i)^T [H(s, a_i) - H(s, \check{a}_i)] D_\pi r_i \geq 0, \\ & \quad i = 1, 2, \forall a_i \in \mathcal{A}_i, \check{a}_i \in \mathcal{A}_i \setminus a_i \\ & \quad D_\pi (r_1 + r_2) \leq y \cdot \mathbf{1}_{NM^2} \\ & \quad (I - \gamma G_\pi)^{-1} B_\pi (r_1 + r_2) \leq y \cdot \mathbf{1}_N. \end{aligned} \tag{55}$$

We now discuss the regularization terms in (55). One challenging issue for MIRL is that there often exists many solutions that are equally sensible so that it is more likely than IRL to recover rewards which are far from actual ones. For example, Lin et al. (2018) emphasize the importance of the structure of rewards. Therefore, some prior knowledge or assumption of the game, as well as the structure of the unknown rewards, is very helpful. For example, it is often assumed that, all other things being equal, an unknown reward vector is sparse (Ng & Russell, 2000). One easy way to incorporate this assumption is to add a penalty term to the objective function to regularize non-sparsity, which is $-\beta (\|r_A\|_1 + \|r_B\|_1)$, where $\beta > 0$ and $\|\cdot\|_1$ denotes the L_1 norm. There might be other problem-specific knowledge/assumption regarding to reward available and taking advantage of it by incorporating it in the regularization terms will help infer higher-quality rewards.

4.6 uNE-MIRL

Recall that the necessary and sufficient condition for an observed bi-policy π being a NE for a two-player general-sum stochastic game is given by

$$\begin{aligned} (B_{\pi|a_1} - B_\pi) D_\pi r_1 &\leq 0, \forall a_1 \in \mathcal{A}_1 \\ (B_{\pi|a_2} - B_\pi) D_\pi r_2 &\leq 0, \forall a_2 \in \mathcal{A}_2. \end{aligned} \tag{56}$$

Since NE is a subset of CE, we can borrow the idea proposed in Section 4.5 and solve a uNE-MIRL problem by solving the following LP problem

$$\begin{aligned} \text{maximize: } & \mathbf{1}_{1 \times N} \times \left[(1 + \lambda) (I - \gamma G_\pi)^{-1} B_\pi (r_1 + r_2) - \lambda y \right] \\ & + \text{other problem-specific regularized terms} \\ \text{subject to: } & (B_{\pi|a_1} - B_\pi) D_\pi r_1 \leq 0, \forall a_1 \in \mathcal{A}_1 \\ & (B_{\pi|a_2} - B_\pi) D_\pi r_2 \leq 0, \forall a_2 \in \mathcal{A}_2 \\ & D_\pi (r_1 + r_2) \leq y \cdot \mathbf{1}_{M \times M} \\ & (I - \gamma G_\pi)^{-1} B_\pi (r_1 + r_2) \leq y. \end{aligned} \tag{57}$$

5. Numerical Examples I: GridWorld

This section describes the behaviour of our approaches (except advE-MIRL) using two grid games (GGs), shown in Figure 2, namely GG1 for the left and GG2 for the right. These games have been used extensively in many theory-oriented MRL works (Hu & Wellman, 1998; Littman, 2001; Greenwald & Hall, 2003). In both GGs, there are two agents, A and B, and two goals (or homes). The two agents act simultaneously and can move only one step in any of the four compass directions. When adjacent to a wall, choosing a direction into a wall results in a no-op, where the agent remains in the current position. If both agents attempt to move into the same cell, a collision occurs and they are pushed back to their original positions immediately, except for cells in the bottom row. Each agent is rewarded upon reaching its goal. However, since the reward is discounted with time, the earlier to reach the goal, the better. GG1 and GG2 are similar in basic game rules but different in board setup in two aspects. First, in GG1, the two players' goals are separate while their goals coincide in GG2. Second, in GG2, there are two barriers and if any agent attempts to move downward through the barrier from the top, then with 1/2 probability this move fails and results in a no-op. We let agents A and B play the *go-back-home* games together according to uCS, uCE, uNE or cooE. Our task is to recover their rewards given the equilibrium, the bi-policy, and the state transition dynamics. The *basic* rewarding rule is: either player receives reward 1 (discounted with time) once reaching home and the game stops immediately, and 0 otherwise. When employing cooE, however, neither player receives reward unless they reach home *simultaneously*.

Our experiments are conducted as follows. First, we apply the cooE-MRL approach proposed by Hu and Wellman (2003), and the uCE-MRL approach proposed by Greenwald and Hall (2003) to obtain cooE and uCE bi-policies, respectively. Then we develop similar Q-learning based iterative algorithms for uCS-MRL and uNE-MRL. The general procedure, namely multi-Q-learning algorithm, is the same for all these four MRL approaches and

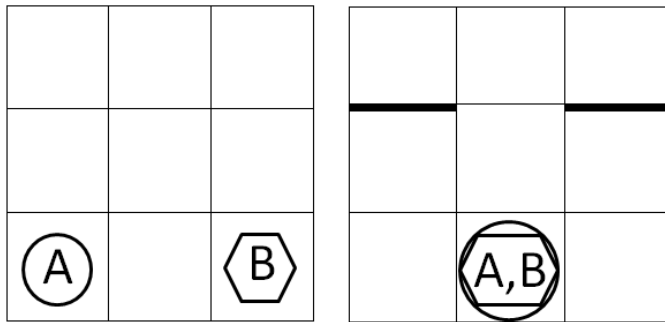


Figure 2: Grid games. The circle indicates A’s goal and the hexagon indicates B’s goal.

described in Algorithm 1. It is worth emphasizing that the multi-Q-learning algorithm can be applied to many variants of Q-learning problems as long as the equilibria exists and is unique (Hu & Wellman, 1998; Littman, 2001; Greenwald & Hall, 2003). It is easy to show that uCS, uCE, uNE and cooE all meet this requirement.

The second step is to apply our uCS-MIRL, cooE-MIRL uCE-MIRL and uNE-MIRL approaches accordingly, incorporating basic knowledge and some reasonable assumptions into Gaussian priors for uCS and cooE. For example, one assumption is that both players’ reward vectors are sparse, only depending on reaching home or not. In addition, one agent’s position might have a small effect on the other agent’s reward, or possibly no affect.

For each experiment, we compare recovered rewards of both players r_A^{rec} and r_B^{rec} , with the true values r_A and r_B numerically. We use a *normalized root mean squared error* (NRMSE) metric, where we first normalize a recovered reward vector r^{rec} on $[0, 1]$, as follows:

$$r^{\text{nrec}} = \frac{r^{\text{rec}} - \min(r^{\text{rec}})}{\max(r^{\text{rec}}) - \min(r^{\text{rec}})},$$

and then compute

$$\text{NRMSE} = \frac{1}{2} \left(\sqrt{\frac{\|r_A^{\text{nrec}} - r_A\|_2}{\dim(r_A)}} + \sqrt{\frac{\|r_B^{\text{nrec}} - r_B\|_2}{\dim(r_B)}} \right)$$

We compare our MIRL approach with IRL and d-MIRL approaches. First, we use an IRL approach to solve the uCS-, cooE-, uCE- and uNE-MIRL problems. Specifically, we focus on B, and try to infer its reward. Obviously, the inferred IRL reward is a function of the state and B’s own action. The IRL approach we use is BIRL, proposed by Qiao and Beling (2011). Note that the reward vector recovered from IRL can be extended to a MIRL reward vector by letting $R(s, a_1, a_2) = R(s, a_2)$ for all $a_1 \in \mathcal{A}_1$. Second, we use the d-MIRL approach to solve the above four problems. Recall that d-MIRL simply assumes agents employing a Nash equilibrium.

To further evaluate the quality of uCS-MIRL reward, let agents B_1 , B_2 and B_3 learn uCS-MIRL, IRL and d-MIRL rewards, respectively, and figure out their own policies π_{B_1} , π_{B_2} and π_{B_3} . Then let these three agents play with A by using their policies while A still employs π_A as if it plays with B. We compute their total game value over all states and

Algorithm 1 General Multi-Q-learning algorithm**Require:** f : uCS, cooE, uCE or uNE; α : learning rate

```

1: procedure MULTI-Q( $f, T, r_1, r_2, \alpha$ )
2:   Initialize:  $s, a, Q_1, Q_2, t = 0$ 
3:   while  $t < T$  do
4:     agents choose bi-strategy  $a$  in state  $s$ 
5:     observe rewards and next state  $s'$ 
6:
7:     for  $i = 1 \rightarrow N$  do
8:        $V_i(s') = f_i(Q_1(s'), Q_2(s'))$ 
9:        $Q_i(s, a) = (1 - \alpha)Q_i(s, a) + \alpha[(1 - \gamma)r_i(s, a) + \gamma V_i(s')]$ 
10:
11:    agents choose bi-strategy  $\vec{a}'$ 
12:     $s = s', a = a'$ 
13:    decay  $\alpha$ 
14:     $t = t + 1$ 

```

compare with the true total game values, which are the maximum. For any reward, a larger the total game value indicates a better reward.

Numerical results are summarized in Table 1 and Table 2. Table 2 summarizes how close the total game value is to the true total game value, using the recovered reward from each method, by computing $\frac{1}{N} \sum_{s=1}^N (V_{\text{true}}(s) - V_{\text{recovered}}(s))^2$, where N is the number of states, V_{true} and $V_{\text{recovered}}$ are true total game value vector and the one using recovered reward, respectively.

We can easily conclude that our MIRL approaches generate satisfactory results and performs much better than IRL and d-MIRL approaches for all the four problems.

	Grid Game #1	Grid Game #2
uCS-MIRL	1.50×10^{-3}	2.27×10^{-4}
IRL	0.061	0.060
dMIRL	0.104	0.106
cooE-MIRL	0.026	0.026
IRL	0.409	0.319
dMIRL	0.085	0.083
uCE-MIRL	1.30×10^{-3}	1.39×10^{-10}
IRL	0.287	0.311
dMIRL	0.089	0.083
uNE-MIRL	0	0
IRL	0.271	0.283
dMIRL	0.104	0.103

Table 1: NRMSE results for reward values comparison

	Grid Game #1	Grid Game #2
uCS-MIRL	0.099	2.50×10^{-4}
IRL	0.223	0.179
dMIRL	0.197	0.202

Table 2: total game value comparison for uCS

6. Numerical Examples II: Abstract Soccer Game

This section presents a demonstration of the advE-MIRL approach on a stylized soccer game. Two-player soccer games of many forms are popular among MRL researchers for algorithm demonstration and comparison purposes (Greenwald & Hall, 2003; Lin et al., 2018; Littman, 1994). Lin et al. (2018) propose a zero-sum MIRL approach is proposed and demonstrate the performance on a game that is similar to the one used here. However, their approach requires that the game be zero-sum. In this section, we relax the zero-sum assumption and require only that the two players be foes, which enables us to rely on a weaker assumption that they employ an advE.

The soccer game (see Figure 3) is depicted as follows. Players A and B compete with each other, aiming to score by either bringing or kicking the ball (represented by a circle) into their opponents’ goals (A’s goal are 6 and 11, and B’s goal are 10 and 15). Both players can move simultaneously either in four compass directions ending in a neighbouring cell or remain in their current cell. A ball exchange may occur with some probability in case of a collision in the same cell. A *kick* action is also available to players. Each one has a perception of how likely they are to score on a shot, or the *probability of a successful shot* (PSS), if kicking the ball at a given position. For simplicity, PSS is assumed to be independent of the opponent’s position. The position based PSS distribution is shown in Table 3.

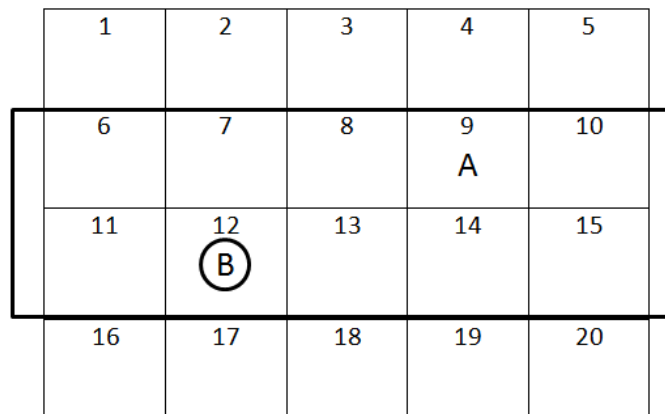


Figure 3: Soccer game: initial board

Note that a player’s PSS at a particular spot is their *perceived* likelihood of a scoring shot, rather than the *actual* probability of a successful shot. So statistically calculating the

	PSS = 0.7	PSS = 0.5	PSS = 0.3	PSS = 0.1	PSS = 0
A	1, 7, 12, 16	2, 8, 13, 17	3, 9, 14, 18	4, 10, 15, 19	5, 20
B	5, 9, 14, 20	4, 8, 13, 19	3, 7, 12, 18	2, 6, 11, 17	1, 16

Table 3: Original PSS distribution of each player

successful shot rate from observation data would not help reflect the player’s own beliefs about their shooting skills, which is the reward.

The two players play against each other, employing a minimax equilibrium in a zero-sum game. But this information is *not* available to us when solving the MIRL problem. Instead, we are given: (1) the bi-policy of the two players over all states, and; (2) the state transition dynamics (including the ball exchange rate $\beta = 0.6$). In fact, this information can be statistically calculated or estimated with sufficient observations. We simply skip this data pre-processing stage as it is not the emphasis of this paper. We then assume that the two players follow an advE and try to infer their rewards on this basis.

6.1 Prior Specification

The numerical experiments are performed using six different prior distributions. Multivariate Gaussian distributions are used as the prior on the reward. We test three settings for the mean of the prior distribution and two settings for the covariance. The specifics for each parameter setting are outlined below.

- *Weak Mean (WM)*: The weak mean parameter setting assumes the least amount of prior knowledge about the game. The mean parameter of the Gaussian prior is 0.5 in each state where A has possession of the ball and -0.5 in each state where A does not have the ball. The prior distribution for B is the same, i.e. the mean parameter is 0.5 in each state where B has possession of the ball and -0.5 in each state where B does not have the ball.
- *Medium Mean (MM)*: The medium mean parameter setting has more prior information about the game than the weak mean parameter setting. Specifically, if A possesses the ball and is in the row in front of the goal (1, 6, 11, and 16), the mean of the Gaussian prior is 1 for A and -1 for B. Symmetrically, the mean parameter is 1 for B and -1 for A if B possesses the ball and is in the row in front of the goal (5, 10, 15, and 20). In addition, whenever a player takes a shot, the mean of the prior is 0.5 in all board positions. Whenever the opposing player takes a shot, the mean of the prior is -0.5 in all board positions. The mean of the prior is 0 for all other states and actions.
- *Strong Mean (SM)*: The strong mean parameter setting assumes the most prior information about the game. The strong mean parameter setting is similar to the medium mean but the parameter is 1 only when the player has has the ball and is directly in front of the goal (6 and 11 for A and 10 and 15 for B).

- *Weak Covariance (WC)*: The weak covariance parameter setting assumes an identity matrix for all states and actions.
- *Strong Covariance (SC)*: The strong covariance parameter assumes the most prior information about the game. This covariance is built based upon the following two assumptions about action and player position:
 1. when one player has the ball and takes a shot, its PSS depends only on their current position in the field. For example, let s_1 and s_2 be states where Player A is in the same position. The correlation for these states and Player A’s action of shooting the ball is $\rho(r_A(s_1, a_A = \text{shoot}, a_B), r_A(s_2, a'_A = \text{shoot}, a'_B)) = 1$.
 2. when one player has the ball and takes any action other than a shot (a movement or choosing to stay in the same position), their reward is state dependent, i.e. $\rho(r_A(s_1, a_A \neq \text{shoot}, a_B), r_A(s_1, a'_A \neq \text{shoot}, a'_B)) = 1$.

6.2 Monte Carlo Simulation using Recovered Rewards

This section evaluates the advE-MIRL approach. By solving an advE-MIRL problem (35), we recover A and B’s reward vectors, over all states and all actions. There are 6 advE-MIRL reward vectors recovered corresponding to the 6 pairs of means and covariance matrices outlined in 6.1.

The zero-sum MIRL approach by Lin et al. (2018) is validated using the abstract soccer game and measure the quality of the 6 types of zero-sum MIRL reward in two ways: numerical (Section VI, Fig.2-7) and Monte Carlo simulation, both against the true reward. They also conduct a thorough sensitivity analysis with respect to the prior (Section VII, Fig.10). Their conclusion is that though there is a trend that the more the reward numerically deviates from the true value the worse it performs in Monte Carlo simulation, measuring the deviation from the true reward is not an accurate representation of the quality of the estimated reward. Hence, simply measuring the numerical difference from the true value may lead to misleading conclusions. In addition, given the same games settings and the same prior, advE-MIRL reward is, in theory, supposed to perform no better than zero-sum MIRL due to the stronger zero-sum MIRL restriction. Taking all the above into account, we adopt a Monte Carlo simulation approach, as follows:

1. Create agents:
 - C , which uses advE-MIRL reward;
 - D , which uses true reward;
 - E , which uses zero-sum MIRL reward;
 - F , which uses dMIRL reward;
 - G , which uses BIRL reward.
2. Simulate competitive games:
 - C against D ;
 - C against E ;

- C against F;
- C against G:

Note that agents E, F and G use rewards recovered from three conventional MIRL approaches covered in Section 3. Here we let C plays the role of A and others take the place of B (due to symmetry, two parties can switch roles as well). All those games are simulated in three different environmental settings, where the the ball exchange rates β are 0, 0.4 and 1. 5000 round games are simulated per case.

The simulation results are presented in Tables 4–7, where *WM*, *MM*, *SM*, *WC* and *SC* stand for *weak mean*, *medium mean*, *strong mean*, *weak covariance matrix* and *strong covariance matrix*, respectively. To interpret the result, take the 2nd row of Table 5 as an example: C uses WM and SC as the prior and recovers A’s advE-MIRL reward, while E uses the same prior and learns a zero-sum MIRL reward vector of B. They come up with their own minimax policies according to their learned rewards and environmental settings and play against each other. 32.28/36.40 means C beats E with probability 32.28%, loses with probability 36.40%, and end in a tie with probability 31.32%, when the ball exchange rate is 1. Note that 0/0 shown in these tables indicates that both parties learned rewards such that neither is able to score a single point even if its opponent is also poorly skilled.

advE-MIRL Rewards	W/L% ($\beta = 0.4$)	W/L% ($\beta = 1$)	W/L% ($\beta = 0$)
WM & WC	0/32.44	0/58.00	0/49.98
WM & SC	20.40/25.46	20.50/38.24	42.88/50.16
MM & WC	4.60/30.12	9.36/44.00	10.44/49.88
MM & SC	24.86/24.94	25.10/24.80	49.97/50.02
SM & WC	14.90/30.52	6.80/42.50	15.42/50.08
SM & SC	25.26/24.80	25.00/24.80	50.14/49.86

Table 4: C vs D

advE-MIRL Rewards	W/L% ($\beta = 0.4$)	W/L% ($\beta = 1$)	W/L% ($\beta = 0$)
WM & WC	0/2.40	0/0	0/0
WM & SC	22.76/28.94	32.28/36.40	43.14/50.14
MM & WC	0/0	9.20/5.60	4.12/16.86
MM & SC	24.86/25.12	25.04/24.96	49.54/50.44
SM & WC	11.24/10.60	8.80/9.18	16.10/24.46
SM & SC	25.28/25.06	24.94/25.12	50.13/49.86

Table 5: C vs E

The results in Tables 4–7 support the following conclusions:

1. The advE-MIRL approach performs, if not better, comparatively with zero-sum MIRL approach given same priors. Considering that the zero-sum MIRL has a stronger constraint, our advE-MIRL approach’s performance has reached its upper limit.
2. The advE-MIRL approach performs notably better than d-MIRL and BIRL approaches.

advE-MIRL Rewards	W/L% ($\beta = 0.4$)	W/L% ($\beta = 1$)	W/L% ($\beta = 0$)
WM & WC	0/0	0/0	0/0
WM & SC	27.10/0	25.42/0	50.04/0
MM & WC	6.04/0	8.64/0	18.02/0
MM & SC	25.28/0	26.06/0	49.86/0
SM & WC	13.98/0	9.00/0	49.26/0
SM & SC	24.90/0	26.08/0	49.90/0

Table 6: C vs F

advE-MIRL Rewards	W/L% ($\beta = 0.4$)	W/L% ($\beta = 1$)	W/L% ($\beta = 0$)
WM & WC	0/0	0/0	0/0
WM & SC	25.10/0	24.84/0	50.12/0
MM & WC	5.52/0	8.76/0	16.20/0
MM & SC	28.50/10.12	25.12/12.00	49.26/20.46
SM & WC	14.20/0	8.64/0	44.25/0
SM & SC	25.80/0	25.28/0	50.12/0

Table 7: C vs G

- As for the sensitivity with respect to prior, overall, the stronger the covariance matrix or the mean is when selecting a prior, the better the solution is. In particular, the covariance matrix has a greater influence than mean in recovering a reasonable reward.

7. Numerical Examples III: Incomplete Policy

All the approaches developed in this paper rely on a strong assumption that a complete bi-policy over all states is available, which is hardly possible in practice. Therefore, it is interesting to assess how much is lost in recovered reward when using our approaches if one player’s (or more) policy is not observed on all states. One way is to conduct imputation to the policy missing states. In this section, we will do a simple computational experiment in GG1 on uNE-MIRL. It is worth emphasizing that this section does not attempt to address incomplete observations issue.

Note that in GG1, there are 72 states in total. The experiment is conducted as follows. First, we randomly pick k states as if the bi-strategies of those states are unavailable. Second, in the complete bi-policy, replace the bi-strategies of those k states with some pre-defined bi-strategies for imputation purpose. Finally, use the imputed bi-policy as an input to (56) and recover the rewards. The imputation scheme we use for those “unavailable” states is uniform mix-strategy - each player picks an action randomly with equal probabilities over all actions. NRMSE is used as the evaluation metric as in Section 5.

The results are summarized in Table 8. We can see that the more missing states there are, the less accurate the result is, which is in line with our expectation.

Addressing incomplete/partial/noisy observations is a vital topic in IRL/MIRL as it helps bridge the gap between theory and practice. Some representative work on that subject

# of missing states	Result of imputed bi-policy
0	0
1	0.0521
5	0.0614
10	0.1155
15	0.1168

Table 8: NRMSE results for uNE-MIRL reward in GG1 with imputed bi-policy

has been conducted by Choi and Kim (2011) and Shahryari and Doshi (2017). Further investigation into this topic with respect to MIRL is beyond the scope of this paper.

8. Conclusions and Future Work

This paper has introduced solution approaches for five general-sum MIRL problems, each distinguished by its assumption about the equilibrium being played by the observed agents. Each solution approach was demonstrated on benchmark grid-world examples.

The advE-MIRL problem formulation requires weaker assumptions but performs comparably with zero-sum MIRL. Additionally, advE-MIRL outperforms both d-MIRL and BIRL. Both the uCS- and cooE-MIRL approaches generate good results if the estimated reward can be scaled correctly. The uCE- and uNE-MIRL approaches perform remarkably well in two benchmark grid-world examples, accurately estimating the value of the true reward. We offer three possible explanations for the high quality of the empirical results. First, these two small GGs are well-defined in the sense that there is no chance of moving in another direction by accident once a certain direction is selected (no noise in action). Second, the bi-policy π we use is exactly the equilibrium of interest because it is generated from a corresponding MRL-Q-learning algorithm. Third, we have incorporated strong prior information about the game, and a good solution can be achieved by tuning the regularized coefficients.

Although this paper is restricted to the two-person case, an extension of the methods to n -player cases would be straightforward because the equilibria we study are unique in n -player games, if they exist. In this sense, advE-MIRL has advantages over zero-sum MIRL as how to extend zero-sum MIRL from two-player to n -player is not yet clear.

Our work is limited in at least two aspects. First, we consider only the case where both state and action spaces are discrete and limited. Second, though it is not explicitly emphasized, we use a strong assumption that a full stationary bi-policy over all states is available. In practice, it might not be possible to observe the play of the game long enough to obtain an accurate estimate of the true bi-policy over all states. Two potential directions for future work are worth pursuing. One is to derive continuous versions of the proposed approaches. The other is to treat the condition when only a partial bi-policy is available. Ideally, future methods would have the characteristic that the more complete the bi-policy observations, the more robust the recovered rewards.

Our solution approaches center on the formulation of optimization problems that can be solved in polynomial time in the size of the state and action spaces. Like MDPs, however,

the sizes of stochastic games tend to scale very poorly as one moves away from toy examples, and the development of a large-scale MRL method remains an open problem.

References

- Abbeel, P., & Ng, A. (2004). Apprenticeship learning via inverse reinforcement learning. In *The 21st International Conference on Machine Learning, ICML'04*, pp. 1–8.
- Abdallah, S., & Lesser, V. (2008). A multiagent reinforcement learning algorithm with non-linear dynamics. *Journal of Artificial Intelligence Research*, *33*, 521–549.
- Aumann, R. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, *1*, 67–96.
- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3), 329–349.
- Barr, N. (2012). *Economics of the Welfare State* (5 edition). Oxford University Press.
- Choi, J., & Kim, K. (2011). Map inference for bayesian inverse reinforcement learning. In *Proceedings of the 15th Neural Information Processing Systems, NIPS01*, pp. 1989–1997.
- Cigler, L., & Faltings, B. (2013). Decentralized anti-coordination through multi-agent learning. *Journal of Artificial Intelligence Research*, *47*, 441–473.
- Conitzer, V., & Sandholm, T. (2007). Awesome: A general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Mach. Learning*, *67*(1–2), 23–24.
- Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, *39*(1), 195–259.
- Ferguson, T. S. (2008). *Game Theory*. UCLA.
- Filar, J., & Vrieze, K. (1996). *Competitive Markov Decision Processes* (1st edition). Springer-Verlag, New York, NY.
- Greenwald, A., & Hall, K. (2003). Correlated q-learning. In *Proceedings of the 20th International Conference on Machine Learning, ICML'03*, pp. 242–249.
- Hadfield-Menell, D., Dragan, A., Abbeel, P., & Russell, S. (2016). Cooperative inverse reinforcement learning. In *Proceedings of the 30th Neural Information Processing Systems (NIPS'16)*, pp. 3916–3924.
- Hart, S., & Schmeidler, D. (1989). Existence of correlated equilibria. *Mathematics of Operations Research*, *14*(1), 18–25.
- Hu, J., & Wellman, M. P. (1998). Multiagent reinforcement learning: Theoretical framework and an algorithm. In *Proc. Intl. Conf. on Mach. Learning (ICML'98)*, pp. 242–250.
- Hu, J., & Wellman, M. P. (2003). Nash q-learning for general-sum stochastic games. *The Journal of Machine Learning Research*, *4*, 1039–1069.
- Lin, X., Beling, P. A., & Cogill, R. (2018). Multi-agent inverse reinforcement learning for two-person zero-sum games. *IEEE Transactions on Games*, *10*(1), 56–68.

- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the 11th International Conference on Machine Learning, ICML'94*, pp. 157–163.
- Littman, M. L. (2001). Friend-or-foe q-learning in general-sum games. In *Proceedings of the 18th International Conference on Machine Learning, ICML'01*, pp. 322–328.
- Maskin, E., & Tirole, J. A. (2001). Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory*, 100(2), 191–219.
- Nash, J. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1), 48–49.
- Nash, J. (1951). Non-cooperative games. *The Annals of Mathematics*, 54(2), 286–295.
- Natarajan, S., Kunapuli, G., Judah, K., Tadepalli, P., Kersting, K., & Shavlik, J. W. (2010). Multi-agent inverse reinforcement learning. In *Proceedings of the 9th International Conference on Machine Learning and Applications, ICMLA'10*, pp. 395–400.
- Ng, A. Y., & Russell, S. (2000). Algorithms for inverse reinforcement learning. In *Proceedings of the 17th International Conference on Machine Learning, ICML'00*, pp. 663–670.
- Owen, G. (1968). *Game Theory* (1st edition). W. B. Saunders Company, Philadelphia, PA.
- Ozdaglar, A. (2010). Lecture 4: Strategic form games: solution concepts correlated rationalizability. In *6.254 Game Theory with Engineering Applications*. MIT OpenCourseWare.
- Papadimitriou, C. H., & Roughgarden, T. (2008). Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3), 1–29.
- Piotrowski, E. W., Śładkowski, J., & Szczypińska, A. (2010). Reinforced learning in market games. In *Econophysics and Economics of Games, Social Choices and Quantitative Techniques*, pp. 17–23. Springer.
- Qiao, Q., & Beling, P. A. (2011). Inverse reinforcement learning with gaussian process. In *Proc. American Control Conf. (ACC'11)*, pp. 113–118.
- Ramachandran, D., & Amir, E. (2007). Bayesian inverse reinforcement learning.. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, Vol. 7, pp. 2586–2591.
- Reddy, T. S., Gopikrishna, V., Zaruba, G., & Huber, M. (2012). Inverse reinforcement learning for decentralized non-cooperative multiagent systems. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, SMC'12*.
- Shahryari, S., & Doshi, P. (2017). Inverse reinforcement learning under noisy observations (extended abstract). In *The 16th International Conference on Autonomous Agents and Multiagent Systems, AAMAS'17*, pp. 1733–1735.
- Shapley, L. S. (1953). Stochastic games. *Proc. Nat. Academy Sci., Math.*, 39, 1095–1100.
- Skyrms, B. (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, Cambridge, UK.

- Sodomka, E., Hilliard, E., Littman, M., & Greenwald, A. (2013). Coco-q: Learning in stochastic games with side payments. In *In proceedings of the 30th International Conference on Machine Learning (ICML'13)*, pp. 1471–1479.
- Wang, X., & Klabjan, D. (2018). Competitive multi-agent inverse reinforcement learning with sub-optimal demonstrations. In *The 25th International Conference on Machine Learning, ICML'18*.
- Waugh, K., Ziebart, B., & Bagnell, J. (2011). Computational rationalization: The inverse equilibrium problem. In *Proceedings of the 28th International Conference on Machine Learning, ICML'11*, pp. 1169–1176.
- Yang, S. Y., Qiao, Q., Beling, P. A., Scherer, W. T., & Kirilenko, A. A. (2015). Gaussian process-based algorithmic trading strategy identification. *Quantitative Finance*, 15(10), 1683–1703.
- Ziebart, B. D., Maas, A. L., Bagnell, J. A., & Dey, A. K. (2008). Maximum entropy inverse reinforcement learning. In *Proceedings of the 23rd National Conference on Artificial Intelligence, AAAI'08*, Vol. 3, pp. 1433–1438.