Maximum-Entropy Multi-Agent Dynamic Games: Forward and Inverse Solutions

Negar Mehr[1], Mingyu Wang[2], and Mac Schwager[3]

Abstract—In this paper, we study the problem of multiple stochastic agents interacting in a dynamic game scenario with continuous state and action spaces. We define a new notion of stochastic Nash equilibrium for boundedly rational agents, which we call the Entropic Cost Equilibrium (ECE). We show that ECE is a natural extension to multiple agents of Maximum Entropy optimality for single agents. We solve both the "forward" and "inverse" problems for the multi-agent ECE game. For the forward problem, we provide a Riccati algorithm to compute closed-form ECE feedback policies for the agents, which are exact in the Linear-Quadratic-Gaussian case. We give an iterative variant to find locally ECE feedback policies for the nonlinear case. For the inverse problem, we present an algorithm to infer the cost functions of the multiple interacting agents given noisy, boundedly rational input and state trajectory examples from agents acting in an ECE. The effectiveness of our algorithms is demonstrated in a simulated multi-agent collision avoidance scenario, and with data from the INTERACTION traffic dataset. In both cases, we show that, by taking into account the agents' game theoretic interactions using our algorithm, a more accurate model of agents' costs can be learned, compared with standard inverse optimal control methods.

I. Introduction

In this article, we seek to learn the cost functions of a group of interacting dynamic agents from a set of trajectory demonstrations of those interactions. We call this problem an Inverse Dynamic Game (IDG), analogously to Inverse Reinforcement Learning (IRL) and Inverse Optimal Control (IOC) in the single-agent setting. To solve the inverse dynamic game, we first formulate a new notion of stochastic Nash equilibrium to describe the boundedly rational equilibrium condition found in natural human demonstrations. We call this equilibrium an Entropic Cost Equilibrium (ECE). We then present an algorithm to find feedback policies for agents engaged in an ECE game, that is, we solve the "forward" problem. We then use this forward solution within an algorithm to infer the cost functions of the multiple agents given trajectory demonstrations from those agents, i.e., we solve the "inverse" problem.

For many robotic applications, it is not trivial to design a cost function that mimics an expert's behavior, such as a human driver's actions. In such applications, a miss-specified

cost function may lead to undesired behaviors, and designing the correct cost function is notoriously challenging. A common practice is to infer the cost function from experts' demonstrations through the framework of Inverse Reinforcement Learning (IRL), also sometimes called Inverse Optimal Control (IOC). An IRL algorithm infers the cost function by observing an agent's behavior assuming that the agent behaves approximately optimally. However, real-world robotic applications, such as autonomous driving, usually involve multiple interactive agents whose behaviors are coupled through the feedback interactions between the agents. Consequently, when learning from interactive agents' behaviors, we cannot treat them as acting in isolation. Instead, we need to take into account the game theoretic coupling between agent's behaviors. In this paper, we develop inverse methods for such interactive multi-agent settings, where we learn each agent's cost function while taking into account their feedback interactions. We call this an inverse dynamic game (IDG).

One of the main challenges in solving IDGs is that agents no longer optimize their own cost functions in isolation. Rather, they reach a notion of game theoretic equilibrium. Hence, agents' cost functions must be learned such that the learned cost functions rationalize the set of demonstrations as game theoretic equilibrium strategies, rather than optimal strategies. Moreover, when learning from experts such as humans, we need to account for the humans' noisy behavior and bounded rationality. Reaching exact equilibria requires perfect rationality of agents; however, in decision-making settings, humans' rationality is normally bounded due to the limited information they have, their cognitive limitations and the finite amount of time available for decision making. Thus, we need to account for the noise in humans' decision making in a multi-agent setting.

In this paper, we address these challenges by defining a notion of "noisy" equilibrium for capturing the outcome of interactions between multiple boundedly rational agents. We call this equilibrium an *Entropic Cost Equilibrium* (ECE) and show its connections to the maximum entropy framework common in IRL [1]. We prove that the ECE concept is indeed an extension of maximum-entropy optimality to multi-agent settings. Once we formalize the notion of Entropic Cost Equilibrium, to assess the quality of a set of learned agents' cost functions, one must find the ECE policies under the learned costs and compare them to the set of agents' demonstrations. To enable such comparisons, we develop an algorithm to find ECE policies for a given set of agents' cost functions. We prove how ECE policies can be obtained in closed-form for a special class of games, namely linear-quadratic-

Gaussian games, using a Riccati solution reminiscent of the well-known Linear-Quadratic Regulator (LQR) and Linear-Quadratic Game (LQGame) solutions. Leveraging this result, we provide an iterative algorithm for approximating ECE policies in general multi-agent games with general nonlinear dynamics and costs.

Knowing how to approximate ECE policies for a set of given costs, we propose a multi-agent inverse dynamic game algorithm for learning agents' costs. Similar to the common practice in IRL [1], [2], [3], [4], we assume that each agent's cost function is parameterized as a linear combination of a set of known features. We then propose an iterative algorithm for learning the weights of the features in each agent's cost function such that the feature expectations under the learned costs match the empirical feature average from the demonstrations.

To demonstrate the effectiveness of our algorithm, we verify our algorithm using both synthetic datasets and real-world traffic data. First, we consider a goal-reaching and collision-avoidance scenario involving two and three agents. We show that by taking into account the agents' feedback interactions, a more accurate model of agents' costs can be learned. We then validate the performance of our algorithm on the INTERACION dataset [5] which involves highly interactive multi-agent driving scenarios collected from different countries. We again demonstrate that by taking into account the agents' interactions through our inverse dynamic game framework, more accurate predictions of agents' behavior can be made. We show that the prediction accuracy of our framework is very close to the intelligent driver's model (IDM) [6], which is a highly accurate human-designed driving model used for modeling human drivers' leader-following behavior.

We summarize our contributions as follows:

- We define and formalize the notion of Entropic Cost Equilibrium (ECE) for capturing the interaction of noisy agents.
- We develop an algorithm for computing approximate ECE policies for general nonlinear multi-agent systems (the forward problem).
- We propose an iterative algorithm for learning the agents' costs from a set of interactive demonstrations (the inverse problem).
- We validate our proposed inverse dynamic game algorithm in both synthetic and real-world driving scenarios.

The organization of this paper is as follows. Section II provides an overview of the related work. In Section III, we introduce our notation and discuss the preliminaries. In IV, we define the notion of Entropic Cost Equilibrium and discuss its connection to the maximum-entropy framework. We discuss finding ECE policies for general nonlinear multi-agent systems in Section V. Section VI provides the description of our inverse dynamic game algorithm. Simulations and experiments with the analysis of the resulting performance are incorporated in Section VII. Finally, we conclude the paper and discuss future directions in Section VIII.

II. RELATED WORK

A. Single-Agent Inverse Reinforcement Learning

Inferring and learning cost functions from system trajectories has been widely studied for single-agent systems. The problem of inferring an agent's cost function was first studied by Kalman in the context of inverse optimal control for linear quadratic systems with a linear control law [7]. Learning an agent's cost was later studied in [8] and [9] where the assumption was that the demonstrations satisfy optimality conditions. This assumption was relaxed to take into account the bounded rationality and human's noisy demonstrations in the framework of maximum entropy inverse reinforcement learning in [1].

Despite the success of these IRL methods, they were mostly developed for discrete state and action spaces. In [3], maximum-entropy cost inference was studied for systems with continuous state and action spaces. The common assumption in these works is that the agent's cost function is paramterized as a weighted sum of features. This assumption was further relaxed in [10], where the underlying cost function was learned as a neural network in a maximum-entropy framework. In [11], a framework was proposed for directly extracting a policy from an agent's demonstrations, as if it were obtained by reinforcement learning following maximum entropy inverse reinforcement learning. Inferring an agent's cost function was also studied as a bilevel optimization where, in the outer loop, the cost parameters were found such that the estimation error of system trajectories or the likelihood of demonstrations was optimized [12], [13]. Recent works have proposed learning cost parameters that minimize the residual of Karush–Kuhn–Tucker (KKT) conditions at the demonstrations [14], [15], [16]. Our work draws inspiration from the Maximum Entropy IRL framework where the noisy behavior of the agent is captured.

B. Multi-Agent Inverse Reinforcement Learning

Extension of single-agent IRL algorithms to the multi-agent settings have been studied in the past. Multi-agent IRL was studied for discrete state and action spaces in [17]. The inverse equilibrium problem was considered in [18] where the maximum entropy principle and the notion of regret were employed for solving the inverse equilibrium problem. No systems dynamics were considered in [18], and the inverse equilibrium problem was considered in matrix games. Inverse reinforcement learning was considered in [19] for two-person zero-sum stochastic games with discrete state and action spaces where the cost learning problem was formulated as an optimization problem. This framework was later extended to general-sum games with discrete state and action spaces in [20]. In [21], a new framework for multi-agent inverse reinforcement learning utilizing adversarial machine learning was proposed for high-dimensional state and action spaces with unknown dynamics. In [22], generative adversarial imitation learning in the single agent case [11] was extended to the multi-agent setting. The current work is distinct in that it focuses on multi-agent IRL in general-sum games with *known* system dynamics and continuous state and action spaces.

Multi-agent IRL has been studied as an estimation problem too. In [23], a particle filtering algorithm was utilized for online estimation of human behavior parameters where the critical role of accurate human motion prediction was demonstrated. In [24], a filtering technique based on an unscented Kalman filter was developed for online estimation of cost parameters in multi-agent settings. In [25], inverse reinforcement learning was considered for the class of linear quadratic games where the equilibrium strategies of all but one agent were known. This assumption simplified the problem and reduced it to effectively an instance of the single-agent cost inference problem. [26] extended this by proposing to minimize the residuals of the first-order necessary conditions for open-loop Nash equilibria. In [27], residual errors of optimality conditions for open-loop Nash equilibria were minimized in a maximum-entropy framework. In [28], state and input constraints are also considered in a maximum-entropy residual-minimization framework. In [29], agents' cost functions were estimated under partial observability. In this work, we focus on cost learning for general *nonlinear* games, and find cost parameters which rationalize interactions under *feedback* information structure. We further capture the *noisy* behavior of humans with our Entropic Cost Equilibrium concept. In [27], maximum-entropy MA-IRL was considered as a maximum likelihood problem. However, when computing the probability distribution over agents' trajectories, the coupling between the agents was not considered. Moreover, in this work, it was assumed that agents' feedback policies were known a-priori which is a restrictive assumption. In this paper, we show how the maximum-entropy framework can be formalized in a multi-agent game theoretic setting to account for agents' feedback interactions. We further provide an algorithm which does not require knowledge of the agents' policies and verify it on a real-world data set.

III. PRELIMINARIES

Consider $N \geq 1$ agents interacting in an environment. We use $[N] = \{1, \cdots, N\}$ to refer to the set of all agents' indices. Let $s_t \in \mathcal{S}$ denote the vector of joint states of all agents at each time t, where the set $\mathcal{S} \subseteq \mathbb{R}^n$ is the joint state space observed by all agents, and n is the dimension of the state space. Each agent $i \in [N]$ decides on its action $a_t^i \in \mathcal{A}^i$ at time step t, where $\mathcal{A}^i \subseteq \mathbb{R}^{m^i}$ is the action space of agent i, and m^i is the dimension of the action space of agent i. Throughout this paper, we use bold letters to refer to the concatenation of variables for all agents. For a given time step t, we use $\boldsymbol{a}_t = (a_t^1, \cdots, a_t^N)$ to denote the
vector of all agents' actions at time t. Following the conventional notation used in the game theory literature, we utilize the superscript $-i$ to indicate all agents expect agent i. For example, $\boldsymbol{a}_t^{-i}$ represents the vector of all agents' actions excluding the action of agent i at time t. We define $\mathcal{A} = \{\mathcal{A}^i\}_{i=1}^N$ to denote the collection of the action spaces of all the agents.

We assume that agents choose their actions through a stochastic Markovian policy. For each agent i, we use $\pi_t^i(.|s_t)$ to denote such a policy for agent i where $\pi_t^i(a_t^i|s_t)$ encodes the probability of agent i taking the action a_t^i at time t given that the system is in state s_t. Note that we are assuming that each agent selects its action independently i.e., $\pi_t^i(a_t^i|s_t, \boldsymbol{a}_t^{-i}) = \pi_t^i(a_t^i|s_t)$. For a finite horizon T, we use $\pi^i = \{\pi_t^i\}_{t=1}^T$ to refer to the agent i's policy for the entire horizon T. Moreover, we use $\boldsymbol{\pi} = \{\pi^i\}_{i=1}^N$ to refer to the set of all agents' time-dependent policies. The discrete-time dynamics of the system are represented by state updates of the following form

$$
\begin{aligned}
s_{t+1} &= f(t, s_t, \boldsymbol{a}_t) + g(s_t)w_t, \\
s_1 &\sim p_1(s_1), \quad w_t \sim p_w.
\end{aligned}
\tag{1}
$$

where p_1 and p_w are the distributions of the system initial state and system noise respectively. We assume that each agent i has a bounded per-stage cost function $c^i : \mathcal{S} \times \mathcal{A}^1 \cdots \times \mathcal{A}^N \to \mathbb{R}$. We further let $\boldsymbol{c} = \{c^i\}_{i=1}^N$ represent the vector of all agents' per-stage costs. We assume that each agent i is seeking to minimize its own expected cumulative cost $\mathbb{E}_{\boldsymbol{\pi}} \sum_{t=1}^T c^i(s_t, \boldsymbol{a}_t)$.

We model the agents' interaction as a dynamic game. We use the notation $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$ to refer to a game with a time horizon of length T between N agents whose action spaces and costs are defined via $\mathcal{A}$ and $\boldsymbol{c}$ on the state space $\mathcal{S}$ with the prespecified dynamics (1).

It is well known that the outcome of interaction between *perfectly* rational agents is best represented via Nash equilibria of the underlying game. Given a game $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$, a set of (mixed strategy) Nash equilibrium policies are defined through the following.

Definition 1. Given a game $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$, a set of agents' policies $\boldsymbol{\pi}^*$ is a (mixed-strategy) Nash equilibrium if and only if for each agent $i \in [N]$, we have

$$
\begin{aligned}
\mathbb{E}_{\pi^{i*}, \boldsymbol{\pi}^{-i*}} \sum_{k=1}^T c^i(s_k, a_k^i, \boldsymbol{a}_k^{-i}) \leq \\
\mathbb{E}_{\pi^i, \boldsymbol{\pi}^{-i*}} \sum_{k=1}^T c^i(s_k, a_k^i, \boldsymbol{a}_k^{-i}), \ \forall \pi^i.
\end{aligned}
\tag{2}
$$

Definition 1 implies that, at a Nash equilibrium, no agent will reduce its accumulated cost by unilaterally changing its policy from the equilibrium policy π^{i*} to another policy π^i.

Although Nash equilibrium is a powerful concept for modeling the interaction of agents, achieving Nash euilibria requires perfect rationality of agents. Moreover, it is well known that computing Nash equilibria is in general intractable even in normal-form games [30]. Thus, when learning from a set of demonstrations that are collected from experts such as humans, assuming that humans have achieved a Nash equilibrium may be unreasonable. Not only are humans computationally bounded, but they also may act under noisy information, or produce actions that are different than what they intend, making them appear to act irrationally to some degree. This concept is known in game theory as *bounded rationality*. Instead of making the "best" choices, humans often make choices that are "satisfactory on average" [31], [32]. In the next section, we generalize the notion of Nash equilibrium to capture the bounded rationality of experts such as humans during their interactions.

IV. ENTROPIC COST EQUILIBRIUM

In the seminal work of [33], [34], it was shown that to take into account the bounded-rationality of humans, their interaction can be modeled via the notion of quantal response equilibrium in normal form and extended games. It was demonstrated that the quantal response equilibrium can successfully model humans' choices in a set of lab experiments, while the predictions made by the Nash equilibrium deviated largely from the lab experiments. At quantal response equilibrium, every agent maintains a probability distribution over its actions. In fact, in quantal response equilibrium, the noisy behavior of humans is captured, where the probability of an action taken by a human is related to the cost associated to that action for the human. In [35], a continuous version of the notion of quantal response equilibrium, called logit equilibrium, was developed for repeated continuous games from the perspective of evolutionary game theory. In this setup, at logit equilibrium, every agent computes its expected cost with respect to the probability distribution over actions of all agents. Each agent takes actions that are exponentially proportional to the negative of this expectation. In the following, we extend this notion of logit equilibrium to dynamic games with continuous state and action spaces. We use the multi-agent extension of Q functions and prove certain properties of this notion of equilibrium.

Given a game $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$, at every time step $t < T$, for each agent $i \in [N]$, we define the quality of a state s_t and a vector of agents' actions $\boldsymbol{a}_t$ under a given set of agents' policies $\boldsymbol{\pi}$ via the following

$$
\begin{aligned}
Q_{t,\boldsymbol{\pi}}^i(s_t, \boldsymbol{a}_t) = c^i(s_t, \boldsymbol{a}_t) + \\
\mathbb{E}_{s_{t+1:T}, \boldsymbol{\pi}} \sum_{k=t+1}^{T} \left[c^i(s_k, \boldsymbol{a}_k) | s_t, \boldsymbol{a}_t \right].
\end{aligned}
\tag{3}
$$

For the final time step T, the cost associated with a state s_T and a vector of actions $\boldsymbol{a}_T$ for an agent i is

$$
Q_{T,\boldsymbol{\pi}}^i(s_T, \boldsymbol{a}_T) = c^i(s_T, \boldsymbol{a}_T).
\tag{4}
$$

Equations (3) and (4) are the extensions of the definition of the Q function to the multi-agent game theoretic setting. Following [35], for each agent $i \in [N]$ and every time step $t \leq T$, we define

$$
\bar{Q}_{t,\boldsymbol{\pi}}^i(s_t, a_t^i) = \mathbb{E}_{\boldsymbol{a}^{-i} \sim \boldsymbol{\pi}^{-i}} Q_{t,\boldsymbol{\pi}}^i(s_t, \boldsymbol{a}_t).
\tag{5}
$$

In (5), the expectation of the Q function is computed with respect to the actions of all the other agents $\boldsymbol{a}^{-i}$.

Note that $\bar{Q}_{\boldsymbol{\pi}}^i(s_t, a_t^i)$ depends only on the action of agent i and the state, not the actions of the other agents. In fact, $\bar{Q}_{\boldsymbol{\pi}}^i(s_t, a_t^i)$ determines the quality of a pair of the system state and an agent's action given the set of other agents' policies $\boldsymbol{\pi}^{-i}$. Ideally, if agents were perfectly rational, given knowledge of the other agents' policies, at every time step, each agent would have taken actions that minimize $\bar{Q}_{\boldsymbol{\pi}}^i(s_t, a_t^i)$ at equilibrium. However, when agents are boundedly rational, we propose that agents' noisy behavior can be modeled through the following notion of equilibrium.

Definition 2. For a given game $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$, a set of agents' policies $\boldsymbol{\pi}^* = \{\pi^{i^*}\}_{i=1}^N$ is an entropic cost equilibrium

(ECE) if and only if for every agent $i \in [N]$ and every action a_t^i, the following holds at every time step $t \leq T$:

$$
\pi_t^{i^*}(a_t^i | s_t) = \frac{e^{-\bar{Q}_{t,\boldsymbol{\pi}^*}^i(s_t, a_t^i)}}{\int e^{-\bar{Q}_{t,\boldsymbol{\pi}^*}^i(s_t, \tilde{a}_t^i)} d\tilde{a}_t^i}.
\tag{6}
$$

Note that (6) must hold for all agents. An equilibrium policy $\boldsymbol{\pi}^*$ is indeed the fixed point of (6), and $\bar{Q}_{\boldsymbol{\pi}^*}^i(s_t, a_t^i)$ depends on the policy of all agents. It is important to note that in entropic cost equilibrium (ECE), at every time step t, the agents' actions a_t^i are independent. No agent needs to know the action choice of any other agent to choose its own action. However, the probability of taking an action is implicitly dependent on the *policies* of other agents $\boldsymbol{\pi}^{-i^*}$ through the expectation with respect to agents' policies in $\bar{Q}_{\boldsymbol{\pi}^*}^i(s_t, a_t^i)$. Therefore, although the agents' instantanous actions are independent, their policies, i.e. the probability distribution over their actions, are related in ECE through (6), and the ECE policies $\boldsymbol{\pi}^*$ are indeed the fixed points of (6) for all agents. This is in contrast to the definition of equilibrium in [21] where a stochastic version of correlated equilibrium was developed, and the actions were assumed to be correlated.

We would like to highlight the connection between (6) and the maximum entropy framework. If there exists only one single agent in the environment, (6) reverts to the well-known maximum entropy formulation, which is widely used in the development of both reinforcement learning and inverse reinforcement learning algorithms [3], [36], [37], [38], [39]. If there is only one agent, the probability of taking an action a_t given a state s_t is proportional to the exponential of the negative accumulated cost from that state (s_t, a_t).

Now that we have defined ECE as our equilibrium notion for capturing humans' noisy interactions, we will prove a property of ECE which demonstrates its applicability and relevance to modeling the interaction of multiple agents with bounded rationality.

Theorem 1. *For a game $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$, a set of agents' policies $\boldsymbol{\pi}^* = \{\pi^{i^*}\}_{i=1}^N$ is an ECE if and only if $\boldsymbol{\pi}^*$ is a (mixed-strategy) feedback Nash equilibrium for the maximum entropy game $\tilde{G} = (\mathcal{S}, \mathcal{A}, f, g, \tilde{\boldsymbol{c}}, T)$ where $\tilde{\boldsymbol{c}} = \{\tilde{c}^i\}_{i=1}^N$ is defined as*

$$
\tilde{c}^i(s_t, \boldsymbol{a}_t) = c^i(s_t, \boldsymbol{a}_t) - \mathcal{H}(\pi^i(\cdot | s_t)),
\tag{7}
$$

and $\mathcal{H}(\pi_t^i(\cdot | s_t))$ is the entropy of policy $\pi^i(\cdot | s_t)$.

Proof. First, we show that if $\boldsymbol{\pi}^*$ is a feedback Nash equilibrium for the game $\tilde{G}$, then, $\boldsymbol{\pi}^*$ is an ECE for the game G. Let $\boldsymbol{\pi}^*$ be a feedback Nash equilibrium of game $\tilde{G}$. Fix an agent $i \in [N]$ and the policy of all the other agents $\boldsymbol{\pi}^{-i^*}$. Then, at Nash equilibrium of game $\tilde{G}$, the agent's policy π^{i^*} optimizes the following

$$
\min_{\pi^i} \mathbb{E} \sum_{k=1}^{T} \left(c^i(s_k, a_k^i, \boldsymbol{a}_k^{-i^*}) - \mathcal{H}(\pi_k^i(\cdot | s_k)) \right),
\tag{8}
$$

where the expectation is with respect to $\boldsymbol{a}_k^{-i^*} \sim \boldsymbol{\pi}_k^{-i^*}$, and $a_k^i \sim \pi_k^i$. We can solve the above for finding Nash equilibrim

policies π^{i^*} using dynamic programming. Starting from the final time step T, for each state s_T, we have

$$\min_{\pi_T^i} \mathbb{E}\left(c^i(s_T, a_T^i, \boldsymbol{a}_T^{-i^*}) - \mathcal{H}(\pi_T^i(\cdot|s_T))\right), \qquad (9)$$

where the expectation is with respect to $\boldsymbol{a}_T^{-i^*} \sim \boldsymbol{\pi}_T^{-i^*}$, and $a_T^i \sim \pi_T^i$. (9) can be rewritten as

$$\min_{\pi_T^i} \mathbb{E}_{\pi_T^i}\left(\mathbb{E}_{\boldsymbol{\pi}_T^{-i^*}}\left[c^i(s_T, a_T^i, \boldsymbol{a}_T^{-i^*})\right] - \mathcal{H}(\pi_T^i(\cdot|s_T))\right). \quad (10)$$

Define the following

$$\bar{c}^i(s_T, a_T^i) = \mathbb{E}_{\boldsymbol{a}_T^{-i^*} \sim \boldsymbol{\pi}_T^{-i^*}}\left(c^i(s_T, a_T^i, \boldsymbol{a}_T^{-i^*})\right). \qquad (11)$$

Then, (10) can be rewritten as minimizing the following:

$$\mathbb{E}_{\pi_T^i}\left[\bar{c}^i(s_T, a_T^i) - \mathcal{H}(\pi_T^i(\cdot|s_T))\right] = \qquad (12)$$

$$\mathbb{E}_{\pi_T^i}\left[\bar{c}^i(s_T, a_T^i) + \log(\pi_T^i(\cdot|s_T))\right] \qquad (13)$$

Now, following the results in [40], (13) can be rewritten as

$$D_{\mathrm{KL}}\left(\pi(.|s_t)\|\frac{1}{\exp\left(V_T^i(s_T)\right)}\exp\left(-\bar{c}^i(s_T, a_T^i)\right)\right) + V_T^i(s_T), \tag{14}$$

where D_{KL} denotes the KL divergence, and $V_T^i(s_T)$ is

$$V_T^i(s_T) = \log \int e^{-\bar{c}^i(s_T, \tilde{a}_T^i)}d\tilde{a}_T^i. \qquad (15)$$

Now, to minimize (14), note that $\exp\left(V_T^i(s_T)\right)$ is a constant for a given state s_T. Since the KL divergence is minimum when the two arguements are the same, the policy which minimizes (14) is

$$\pi_T^{i^*}(a_T^i|s_T) = \frac{e^{-\bar{c}^i(s_T, a_T^i)}}{e^{V_T^i(s_T)}}. \qquad (16)$$

For every state s_T, $V^i(s_T)$ is indeed the cost-to-go (or value function) of agent i at time T. We can use using dynamic programming to propagate this cost-to-go backwards in time and find the equilibrium policy for agent i at time $T-1$. At time $t-1$, we need to solve for

$$\min_{\pi_{T-1}^i} \mathbb{E}\big[c^i(s_{T-1}, a_{T-1}^i, \boldsymbol{a}_{T-1}^{-i^*}) \\ - \mathcal{H}(\pi_{T-1}^i(\cdot|s_{T-1})) + \mathbb{E}_{s_T}(V_T^i(s_T))\big], \tag{17}$$

where the first expectation is with respect to $a_{T-1}^i \sim \pi_{T-1}^i$ and $\boldsymbol{a}_{T-1}^{-i^*} \sim \boldsymbol{\pi}_{T-1}^{-i^*}$. Again, we can separate the expectations in (17) as

$$\min_{\pi_{T-1}^i} \mathbb{E}_{\pi_{T-1}^i}\Big(\mathbb{E}_{\boldsymbol{\pi}_{T-1}^{-i^*}}\left[c^i(s_{T-1}, a_{T-1}^i, \boldsymbol{a}_{T-1}^{-i^*}) + \mathbb{E}_{s_T}V_T^i(s_T)\right] \\ -\mathcal{H}(\pi_{T-1}^i(\cdot|s_{T-1}))\Big). \tag{18}$$

Then, using (5) we see that (18) can be rewritten as

$$\min_{\pi_{T-1}^i} \mathbb{E}_{\pi_{T-1}^i}\left[\bar{Q}_{T-1}^i(s_{T-1}, a_{T-1}^i) - \mathcal{H}(\pi_{T-1}^i(\cdot|s_{T-1}))\right]. \tag{19}$$

Thus, similar to minimizing (13), to minimize (19), the equilibrium policy for agent i at time step $T-1$ is

$$\pi_{T-1}^{i^*}(a_{T-1}^i|s_{T-1}) = \frac{e^{-\bar{Q}_{T-1}^i(s_{T-1}, a_{T-1}^i)}}{e^{V_{T-1}^i(s_{T-1})}}, \qquad (20)$$

where the cost-to-go $V_{T-1}^i(s_{T-1})$ is defined as

$$V_{T-1}^i(s_{T-1}) = \log \int e^{-\bar{Q}_{T-1}^i(s_{T-1}, \tilde{a}_{T-1}^i)}d\tilde{a}_{T-1}^i.$$

Repeating this procedure, we can solve for the mixed-strategy Nash equilibrium policies backwards in time. Thus, we show via induction that mixed-strategy Nash equilibrium policies of the game $\tilde{G} = (\mathcal{S}, \mathcal{A}, f, g, \tilde{\boldsymbol{c}}, T)$ are in fact the ECE policies of the original game $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$. To prove the reverse, i.e., every ECE of the game G is a mixed strategy Nash equilibrium for $\tilde{G} = (\mathcal{S}, \mathcal{A}, f, g, \tilde{\boldsymbol{c}}, T)$, we follow the same reasoning. Starting from the final time step, one can show that ECE policies optimize the cost-to-go in the game $\tilde{G}$, and use induction to prove that ECE policies of G are indeed Nash equilibria of $\tilde{G}$. $\qquad\square$

Theorem 1 connects ECE to the Nash equilibria of a game between agents who aim to maximize the entropy of their policy while minimizing their accumulated cost. In other words, ECE is in fact an extension of the maximum entropy-framework in the single-agent setting to the setting of multiple interactive agents. It is well known that when it comes to learning from demonstrations in the single-agent scenarios, to capture the bounded rationality and noisiness of the demonstrator, the demonstrator is best modeled via the maximum-entropy framework [1], [2]. Theorem 1 extends this notion to multi-agent games. We will use Theorem 1 in the remainder of this paper for computing the ECE policies in general dynamic games.

Remark 1. The notion of ECE can incorporate a temperature parameter γ^i for each agent such that $p(a^i|s_t) \propto \exp\left(\frac{-\bar{Q}_{\boldsymbol{\pi}}^i(s_t, a_t^i)}{\gamma^i}\right)$. This in turn will lead to an additional weight γ^i on the entropy of each agent's policy in (7). The temperature weight γ^i captures each agent's rationality. The higher γ^i is, the noisier the agent acts. If for each agent $i \in [N]$, $\gamma^i \to \infty$, then in the limit, all agents act completely randomly with a uniform probability distribution over actions. On the other hand, when $\gamma^i \to 0$ for all agents, the Nash equilibrium policies are recovered. Therefore, the temperature coefficient γ^i reflects the rationality of each agent. For simplicity, in this paper, we assume that $\gamma^i = 1$ for all agents.

V. SOLVING FOR ECE POLICIES

So far, we have defined ECE for capturing the interaction of boundedly rational and noisy agents. In this section, we give an algorithm for computing ECE policies for a given game $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$ using Theorem 1. In the next section we use this "forward" solution to learn cost function parameters given trajectory demonstrations, i.e., we solve the inverse dynamic game. To solve the maximum entropy dynamic game, first, we prove that ECE policies can be

computed in closed-form for the class of linear quadratic Gaussian games using a Riccati solution similar to the classic LQR [7] and LQGames [41] solutions. Then, we extend this algorithm to general nonlinear dynamic games through iterative linear-quadratic approximations, similar to Differential Dynamic Programming [42], iLQR [43], and iLQGames [44].

A. Maximum Entropy Linear Quadratic Gaussian Games

Consider a class of games $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$ where the system dynamics are linear, and the system noise is normally distributed, i.e., the state update (1) is of the following form

$$
\begin{aligned}
s_{t+1} &= As_t + \sum_{j \in [N]} B^j a_t^j + w_t, \\
w_t &\sim \mathcal{N}(0, I),
\end{aligned} \tag{21}
$$

where $A \in \mathbb{R}^{n \times n}$, and $B^j \in \mathbb{R}^{m^j \times m^j}$ are known time-invariant matrices of appropriate dimensions, and w_t is a zero-mean normally distributed random variable with covariance matrix being identity. Note that in general, the dynamics matrices can be time-variant, and the process noise can be any normal distribution. Here, for the ease of description, we assume system dynamics are linear time-invariant subject to process noise with covariance matrix being identity. Moreover, consider the class of cost functions where every agent $i \in [N]$ minimizes a convex quadratic cost function for every time step $t \leq T$:

$$
\begin{aligned}
c^i(s_t, a_t^1, \cdots, a_t^N) = \frac{1}{2} \big(s_t^\mathsf{T} Q^i s_t + {l^i}^\mathsf{T} s_t + \\
\sum_{j \in [N]} {a_t^j}^\mathsf{T} R^{ij} a_t^j \big),
\end{aligned} \tag{22}
$$

where $Q^i \in \mathbb{R}^{n \times n}$ is a positive semi-definite matrix, and l^i is a vector capturing the affine penalty of agent i for the system state. Moreover for every agent $j \in [N], j \neq i$, $R^{ij} \in \mathbb{R}^{m^j \times m^j}$ is a positive semi-definite matrix while R^{ii} is a positive definite matrix. Equations (21) and (22) define the class of linear quadratic games where the dynamics are linear and the stage costs are quadratic in the states and actions.

Definition 3. A given game $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$ is a linear quadratic Gaussian game if its dynamics are of the form (21), and further, for each agent $i \in [N]$, the cost function c^i is of the form (22).

Note that for a linear quadratic Gaussian game G, in the resulting maximum entropy game $\tilde{G}$, every agent $i \in [N]$ minimizes the following accumulated cost when we fix the policy of all the other agents to be $\boldsymbol{\pi}^{-i}$,

$$
\begin{aligned}
\min_{\pi^i} \mathbb{E}_{\pi^i} \mathbb{E}_{\boldsymbol{\pi}^{-i}} \sum_{k=1}^{T} \frac{1}{2} \bigg(s_k^\mathsf{T} Q^i s_k + {l^i}^\mathsf{T} s_k + \\
\sum_{j \in [N]} {a_k^j}^\mathsf{T} R^{ij} a_k^j \bigg) - \sum_{k=1}^{T} \mathcal{H}(\pi^i(.|s_k)).
\end{aligned} \tag{23}
$$

Now that we have defined the class of linear-quadratic games, using (23) and Theorem 1, we will prove how ECE policies can be obtained in closed form for this class of games.

Theorem 2. *Consider a linear quadratic Gaussian game* $G = (\mathcal{S}, \mathcal{A}, f, g, \boldsymbol{c}, T)$. *For every agent* $i \in [N]$, *at every time step* $t < T$, *the ECE policy* $\pi_t^{i^*}$ *is a normal distribution* $\pi_t^{i^*} \sim \mathcal{N}(\mu_t^i, \Sigma_t^i)$ *where the mean* μ_t^i *and the covariance* Σ_t^i *are computed by*

$$
\mu_t^i = -P_t^i s_t - \alpha_t^i, \tag{24}
$$

$$
\Sigma_t^i = (R^{ii} + {B^i}^\mathsf{T} Z_{t+1}^i B^i)^{-1}. \tag{25}
$$

The matrices P_t^i *and the vectors* α_t^i *satisfy the following sets of linear equations*

$$
\Big[R^{ii} + {B^i}^\mathsf{T} Z_{t+1}^i B^i \Big] P_t^i + {B^i}^\mathsf{T} Z_{t+1}^i \sum_{j \in [N], j \neq i} B^j P_t^j = \\
{B^i}^\mathsf{T} Z_{t+1}^i A, \tag{26}
$$

$$
\Big[R^{ii} + {B^i}^\mathsf{T} Z_{t+1}^i B^i \Big] \alpha_t^i + {B^i}^\mathsf{T} Z_{t+1}^i \sum_{j \in [N], j \neq i} B^j \alpha_t^j = \\
{B^i}^\mathsf{T} \xi_{t+1}^i, \tag{27}
$$

where Z_t^i, ξ_t^i *are recursively computed from the following*

$$
Z_t^i = F_t^\mathsf{T} Z_{t+1}^i F_t + \sum_{j \in [N]} {P_t^j}^\mathsf{T} R^{ij} P_t^j + Q^i \tag{28}
$$

$$
\xi_t^i = F_t^\mathsf{T} (\xi_{t+1}^i + Z_{t+1}^i \beta_t) + \sum_{j \in [N]} {P_t^j}^\mathsf{T} R^{ij} \alpha_t^j, \tag{29}
$$

where

$$
F_t = A_t - \sum_{j \in [N]} B^j P_t^j \tag{30}
$$

$$
\beta_t = - \sum_{j \in [N]} B^j \alpha_t^j. \tag{31}
$$

The terminal conditions for Equations (28) *and* (29) *are*

$$
\xi_T^i = l^i, \quad Z_T = Q^i. \tag{32}
$$

Proof. The proof can be found in Appendix (A). $\square$

Note that Theorem 2 is in fact an extension of Maximum Entropy LQR derived in [2] to the multi-agent setting. Equations (26), (27), (28), (29) are the extensions of the Riccati backward recursion to the multi-agent setting. Indeed, these equations are similar to the backward recursions for deterministic LQ games derived in [41]. Theorem 2 suggests that for Maximum Entropy linear quadratic games, the optimal policy is obtained by a normal distribution whose mean is found by solving LQ games. Then, the variance of the normal policies at every time step is found via (25).

B. General Nonlinear Games

Theorem 2 provides a closed-loop expression for computing ECE policies of linear quadratic Gaussian games over a finite horizon of time T. Inspired by [44], [45], we will leverage this result for approximating the ECE trajectories in general games with nonlinear dynamics and non-quadratic cost functions. We

assume that each agent i can minimize a possibly nonlinear cost function of the form

$$c^i(s_t, \boldsymbol{a}_t) = v^i_{s,t}(s_t) + \sum_{j \in [N]} a_t^{j\top} R^{ij} a_t^j, \quad (33)$$

where $v^i_{s,t}(s_t)$ is a nonlinear state cost. Note that in general, we could have non-quadratic nonlinear costs on the actions as well; however, since (33) captures a wide range of applications, we consider cost functions of the form (33).

To approximate ECE policies, we propose an iterative algorithm where we start by a nominal trajectory. Then, we linearize the dynamics, find a quadratic approximation of the cost function for every agent, and solve the resulting Maximum Entropy linear quadratic Gaussian game. We then update our reference trajectory and repeat this process until convergence (see Algorithm 1). Consequently, our algorithm shares a similar structure with DDP [42], [43] and iLQR methods [46]. More precisely, we start with a nominal sequence of the gain matrices and offset vectors $\{P^i_t, \alpha^i_t\}$ for every time step $t \in T$, and every agent $i \in [N]$. If such a nominal sequence of control matrices is not available, a trivial initialization is initializing all matrices to zeros. We choose the mean value found by (24) as our reference controller. Then, at every iteration, a sequence of states and actions $\eta = \{\bar{s}, \bar{\boldsymbol{a}}^1 = (\mu^1_1, \cdots, \mu^1_T), \cdots, \bar{\boldsymbol{a}}^N = (\mu^N_1, \cdots, \mu^N_T)\}$ is found by forward simulating the system dynamics using the nominal control inputs (which are chosen to be μ^i_t for each agent i at time t). We then linearize the system dynamics (1) around this nominal trajectory to obtain

$$\delta s_{t+1} \approx A_t \delta s_t + \sum_{j \in [N]} B^j_t \delta a^j_t, \quad (34)$$

where $\delta s_t = s_t - \bar{s}_t$, $\delta a^i_t = a^i_t - \bar{a}^i_t$, and $A_t = D_s f(\cdot)$ and $B^j_t = D_{a^j_t} f(\cdot)$ are the Jacobians of the system dynamics (1) with respect to the state s_t, and actions a^j_t, respectively. We further acquire a quadratic approximation of the cost function for each agent. For each agent $i \in [N]$, we let

$$c^i(\bar{s}_t + \delta s_t, \bar{\boldsymbol{a}} + \delta \boldsymbol{a}) \approx v^i_{s,t}(\bar{s}_t) + \frac{1}{2} \delta s_t^\top H^i_t \delta s_t + l_t^{i\top} \delta s_t, \quad (35)$$

denote such an approximation where $H^i_t = D_{s_t s_t} v^i_{s,t}(\cdot)$ and $l^i_t = D_{s_t} v^i_{s,t}(\cdot)$ are the Hessian and the gradient of the cost function $v^i_{s,t}(\cdot)$ with respect to s_t. Note that our formulation only considers nonlinear costs on state variables, and the dependence on agents' actions is quadratic. For a more general case, where the cost function is nonlinear in control actions too, a similar approximation could be used to derive the quadratic terms and linear terms in a^i_t. Note that all the approximations A_t, B^j_t, H^i_t, l^j_t are evaluated at η.

For the linearized system dynamics (34) and quadratized cost (35), we reach a new approximated linear quadratic Gaussian game with new variable sequences $\delta s, \delta \boldsymbol{a}^1, \cdots, \delta \boldsymbol{a}^N$. These approximations result in a new game that can be solved using Theorem 2. Once the approximated game is solved, we obtain a new sequence of mean control actions

$$\{\bar{a}^i_t + \delta a^{i*}_t, t = 0, \cdots, T-1\}, \quad (36)$$

where δa^{i*}_t is the mean value of the action distribution found by solving for the ECE policies of the approximated linear quadratic Gaussian game. A new $\bar{s}_t$ is attained from the forward simulation of the original system dynamics (1) using the newly obtained nominal control actions. We repeat the above process until convergence, i.e., the deviation of the new state trajectory from the state trajectory in the previous iteration lies within a desired tolerance. Note that in our iterative process, we always forward simulate the mean value of the distribution of agents' action. Once the mean trajectories converge, we sample ECE policies by sampling control actions from (24) and (25) computed around the converged mean trajectory.

Remark 2. Applying a^i_t directly from (36) may lead to non-convergence since the resulting trajectory could deviate too much from the original non-linear system which we approximated around η_t. As in other iterative linear quadratic methods [47], [48], [44], we only take a small step in the proposed direction. At each iteration, rather than (36), we apply the following control input

$$\bar{a}^i_t - P^i_t \delta s_t - \epsilon \alpha^i_t, \quad (37)$$

where ϵ is the step size for improving our control strategy. Initially, we set $\epsilon = 1$. Drawing inspiration from line search method in optimization problems, we decrease ϵ by half until the new trajectory's deviation from the nominal trajectory is within a threshold.

Algorithm 1 Approximating ECE Trajectories

1: **Inputs**
2: system dynamics (1), agents' cost functions (33)
3: **Initialization**
4: initialize the control policy using $P^i_t = 0$, and $\alpha^i_t = 0$, $\forall i \in [N]$
5: forward simulation and obtain $(\bar{s}_0, \bar{s}_1, \cdots, \bar{s}_T)$, $\bar{\boldsymbol{a}}^i$, $\cdots, \bar{\boldsymbol{a}}^N$
6: **while** not converged **do**
7: 　linear approximation of (1)
8: 　quadratic approximation of (33)
9: 　solve the backward recursion with (24-32)
10: 　forward simulation using μ^j's and obtain the new trajectories
11: **end while**
12: **return** P^i_t, α^i_t, and Z^i_{t+1}

VI. LEARNING AGENTS' COSTS

So far, we have discussed the concept of noisy equilibrium, ECE, that we adopt and how to solve for approximate equilibrium policies under this model. Now, we are ready to discuss how to learn agents' cost functions c^i from a set of interaction demonstrations from humans or other experts. The main intuition behind our algorithm is similar to the single-agent Maximum Entropy IRL [1]. The common assumption in single-agent IRL is that the agent's cost function is parameterized as a linear combination of a set of features, and the weight

of features is learned such that the expectation of the features under the learned cost function matches the empirical feature means under the demonstrations. This is achieved through an iterative algorithm where at every iteration, the difference between the feature expectations under the demonstrations and under the learned cost functions is utilized for updating the cost parameters.

We extend such an iterative cost learning framework to the multi-agent game setting. We assume that agents' dynamics are known and each agent's cost function $c^i(\cdot)$ is parameterized as a weighted sum of features, i.e., for each agent $i \in [N]$, we have

$$c^i(s_t, \boldsymbol{a}_t) = {w^i}^\mathsf{T} \phi^i(s_t, \boldsymbol{a}_t), \tag{38}$$

where w^i is the vector of weight parameters for agent i and $\phi^i(s_t, \boldsymbol{a}_t)$ is the vector of features for agent i. Note that for each agent i, the vector of features $\phi^i(s_t, \boldsymbol{a}_t)$ depends on the entire state vector and the actions of all agents. We define $\boldsymbol{w} = (w^1, \cdots, w^N)$ to be the collection of cost weights for all agents.

Our goal is to find the weight parameters $\boldsymbol{w}$ such that the feature expectation under the learned cost weights matches the empirical feature mean under the demonstrations. Let $\mathcal{D}_{\boldsymbol{w}}$ be the probability distribution over equilibrium trajectories induced by the cost parameters $\boldsymbol{w}$. We further let $\bar{\mathcal{D}}$ be the emirical distrubution of equilibrium trajectories in the interaction demonstrations. Moreover, we define a trajectory $(s, \boldsymbol{a}) = (s_t, \boldsymbol{a}_t)_{t=1}^T$ to be the vector of agents' states and actions over a trajectorty of length T. We assume that agents maintain an entropic cost equilibrium in their demonstrations.

With a slight abuse of notation, for each agent i, we denote the empirical mean of the features under the interaction demonstrations by $\mathbb{E}_{(s,\boldsymbol{a})\sim\bar{\mathcal{D}}} \phi^i(s, \boldsymbol{a})$. Likewise, we let $\mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a})$ be the expected value of the features under the equilibrium trajectories induced by the weight vector $\boldsymbol{w}$.

We want to find weight parameters $\boldsymbol{w}$ to induce a probability distribution $\mathcal{D}$ over equilibrium trajectories such that feature expectation under the learned cost weights matches the empirical feature mean under demonstrations, i.e., $\mathbb{E}_{(s,\boldsymbol{a})\sim\bar{\mathcal{D}}} \phi^i(s, \boldsymbol{a}) = \mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a})$ for all agents $i \in [N]$. To this end, we propose an iterative algorithm. We initialize the cost weights $\boldsymbol{w}$ for all agents. At each iteration of the algorithm, we iterate over all the agents. For agent i, we compute the difference between the feature expectations under the demonstrations and the current model $\mathbb{E}_{(s,\boldsymbol{a})\sim\bar{\mathcal{D}}} \phi^i(s, \boldsymbol{a}) - \mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a})$ and use this difference for updating the weight parameters

$$w^i \leftarrow w^i - \gamma \left(\mathbb{E}_{(s,\boldsymbol{a})\sim\bar{\mathcal{D}}} \phi^i(s, \boldsymbol{a}) - \mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a}) \right), \tag{39}$$

where γ is the learning rate. Note that once w^i is updated, the entire vector of cost weights $\boldsymbol{w}$ is in fact updated. In the single-agent setting where an agent is nosily optimizing a cost function, the update rule (39) is the gradient of the likelihood of trajectories under weight parameters $\boldsymbol{w}$. In other words, in the single-agent setting, (39) solves a maximum likelihood problem for finding weight parameters $\boldsymbol{w}$ that maximize the likelihood of demonstrations. In the multi-agent setting, if we

fix the policy and cost parameters of all agents except agent i, the equilibrium policy of agent i is the optimal policy which minimizes agent i's cost, i.e., the problem reduces to a single-agent setting where agent i noisly minimizes its cost. Thus, if we fix all agents except agent i, similar to the single-agent setting, we can interpret (39) as a gradient-descent update rule. Once the cost parameters of agent i are updated, we compute the feature expectations under the new set of cost parameters, and update the cost parameters for the next agent. We iterate over all i, and repeat this process until convergence. This can be interpreted as a block coordinate descent method for solving the coupled parameter estimation problems for the multi-agent game.

Algorithm 2 summarizes the algorithm steps. Note that in a multi-agent setting, once the cost parameters of one agent are updated, we need to recompute the feature expectation for updating the next agent's cost, since agents' cost parameters and feature expectations are interdependent, i.e., a change in one agent's cost parameters affects the expectation of features for other agents. This is due to the coupling between agents at entropic cost equilibrium.

Algorithm 2 Multi-Agent Inverse Reinforcemenr Leaning

1: **Inputs**
2: system dynamics (1), agents' cost features $\phi^i(\cdot)$, and the set of interaction demonstrations.
3: **Initialization**
4: for each agent $i \in [N]$, compute the empirical feature mean under demonstrations $\mathbb{E}_{(s,\boldsymbol{a})\sim\bar{\mathcal{D}}} \phi^i(s, \boldsymbol{a})$
5: initialize the cost weights $\boldsymbol{w}$
6: **while** not converged **do**
7:     **for** each agent $i \in [N]$ **do**
8:         Compute feature expectations $\mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a})$
9:         Compute the difference

$$\mathbb{E}_{(s,\boldsymbol{a})\sim\bar{\mathcal{D}}} \phi^i(s, \boldsymbol{a}) - \mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a})$$

10:         Update cost weight for agent i:

$$w^i \leftarrow w^i - \gamma \left(\mathbb{E}_{(s,\boldsymbol{a})\sim\bar{\mathcal{D}}} \phi^i(s, \boldsymbol{a}) - \mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a}) \right)$$

11:     **end for**
12: **end while**
13: **return** cost parameters w^i for all agents $i \in [N]$

Each iteration of Algorithm 2 at line 8 requires computing agents' feature expectation $\mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a})$ under ECE policies. However, computing $\mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a})$ for general nonlinear games is not tractable. We propose to approximate this by sampling ECE policies through Algorithm 1 under the current cost parameter estimates. In each iteration, we sample p ECE equilibrium trajectories $(s_1, \boldsymbol{a}_1), \cdots, (s_p, \boldsymbol{a}_p)$ using Algorithm 1, and approximate these through the following

$$\mathbb{E}_{(s,\boldsymbol{a})\sim\mathcal{D}_{\boldsymbol{w}}} \phi^i(s, \boldsymbol{a}) \approx \frac{1}{p} \sum_{j=1}^p \phi^i(s_j, \boldsymbol{a}_j). \tag{40}$$

Thus, with known system dynamics, Algorithm 1 lets us compare equilibrium trajectories under the learned cost parameters

with the demonstration trajectories and update the cost parameters if needed. Algorithm 3 summarizes the computation of feature expectations.

Algorithm 3 Approximating Feature Expectations

1: **Inputs**
2: system dynamics (1), agents' cost features $\phi^i(\cdot)$, agents' cost wights w^i, number of samples p
3: **for** $1 \leq j \leq p$ **do**
4:     Sample an ECE trajectory $(s_j, \boldsymbol{a}_j)$
5: **end for**
6: **return** $\frac{1}{p} \sum_j \phi^i(s_j, \boldsymbol{a}_j)$


VII. EXPERIMENTS

In this section, we evaluate the performance of the MA-IRL algorithm in two different scenarios. In the first scenario, we consider a motion planning task with multiple agents that navigate to goal locations while avoiding collisions with each other. In the second scenario, we use the INTERACTION dataset [5] which contains trajectories in interactive traffic scenes, such as intersection, roundabouts, and highway merging, collected from different locations in different countries.

We compare our algorithm with a baseline method, continuous inverse optimal control (CIOC) [3], [49], which is an IOC algorithm for large, continuous domains and does not assume any feedback interaction among the observed agents. Additionally, for the INTERACTION dataset, we also compare with the intelligent driver model (IDM) [6], [50] as our reference model. IDM is a widely used expert-designed model to simulate traffic flow that is known to yield accurate predictions of drivers' trajectories.

A. Motion Planning with Collision Avoidance

We first consider a synthetic environment where 2 or 3 agents move to their goal locations and try to avoid collisions with each other. Demonstrations for all scenarios are generated by solving an approximate entropic cost equilibrium as described in Sec. V. Fig. 1 illustrates the demonstration trajectories for 2 and 3 agents scenarios. Each agent wants to minimize a cost function which is linear in the set of features: tracking a reference trajectory, penalizing close proximity to other agents, and control effort. Plots of different features are shown in Fig. 2. The reference trajectory is a straight line between initial location and goal location. Each agent has different weightings on the features, which leads to different interactive behaviors as shown in Fig. 1.

The demonstration dataset contains 200 trajectories. We use MA-IRL to learn the cost function coefficients for all agents jointly. To apply CIOC method, we learn the cost coefficients of each agent individually and treat all other agents as obstacles. In our implementation of the CIOC method, we use a generic optimization solver. The computation time scales cubically with the planning horizon. For computational tractability, we partition the demonstration trajectories into sections of shorter trajectories.

(a) Two-player case   (b) Three-player case

Fig. 1: Demonstrations for two-player and three-player collision avoidance motion planning. Trajectories of different players are shown in green, yellow, and blue, respectively. Players have different cost functions. Due to the stochastic policy, we observe different modes of interactions.

(a) Collision avoidance feature   (b) Reach goal location feature.

Fig. 2: Cost features for motion planning with collision avoidance task. Costs are high close to other agents and away from the goal location.

We evaluate the performance of both algorithms in test scenarios where the initial configurations are generated randomly and are different from the demonstration data set. To evaluate the performance, we first solve the ECE solutions with the learned cost functions using both our MA-IRL approach and the CIOC approach, as well as with the true cost functions in test scenarios. Then, we simulate the stochastic equilibrium policies for 200 trials in each case and compute the feature expectations. The cost of each simulation for all agents are then computed with the true cost function. If the feature distribution from the learned cost functions is similar to the feature distribution in the demonstration data set, it means our learned costs are close to the true costs. The simulation results are shown in Table I and II.

KL DIV		agent 1			agent 2		
		tracking	control	obstacle	tracking	control	obstacle
task 1	MA-IRL (ours)	**0.054**	**0.024**	**0.044**	**0.134**	**0.059**	**0.044**
	CIOC	1.701	0.165	1.766	0.836	0.065	1.766
task 2	MA-IRL (ours)	**0.025**	**0.080**	**0.061**	**0.034**	**0.053**	**0.061**
	CIOC	1.904	0.146	1.086	0.152	0.144	1.086

TABLE I: Two-agent collision avoidance simulations. KL-divergence of feature distribution between demonstrations and simulations from learned costs. Smaller values are better.

We compute the Kullback–Leibler divergence, which captures the difference between two distributions. Lower values means the distributions are closer. In the two different test

KL DIV		agent 1			agent 2			agent 3		
		tracking	control	obstacle	tracking	control	obstacle	tracking	control	obstacle
task 1	MA-IRL (ours)	**0.031**	0.063	**0.167**	**0.074**	0.144	**0.052**	**0.050**	0.066	**0.127**
	CIOC	2.583	**0.054**	0.939	1.083	**0.039**	0.489	0.082	**0.022**	0.629
task 2	MA-IRL (ours)	**0.096**	**0.019**	**0.116**	**0.054**	**0.019**	**0.148**	**0.030**	0.108	**0.080**
	CIOC	1.104	0.073	0.701	1.685	0.032	0.780	2.198	0.171	0.487

TABLE II: Three-agent collision avoidance simulations. KL-divergence of feature distribution between demonstrations and simulations from learned costs. Smaller values are better.

avg. dist. to goal		true	MA-IRL	CIOC
Task 1	agent 1	0.205±0.103	0.229±0.117	0.144±0.083
	agent 2	0.319±0.153	0.323±0.160	0.248±0.101
Task 2	agent 1	0.720±0.702	0.864±0.737	0.381±0.428
	agent 2	2.079±1.880	1.797±1.953	1.347±1.624

TABLE III: Task-relevant statistics for two player collision avoidance case. Distances to goal location at the end of simulations are shown. The statistics of the MA-IRL approach are closer to the true cost samples.

tasks, MA-IRL algorithm achieves closer similarity to true feature distribution, which means the MA-IRL learned cost functions generalize to different tasks better than CIOC.

In addition, we follow prior work and also compute task-relevant statistics to evaluate how much the recovered cost and corresponding stochastic policy resembles the demonstrations. We compute the distance to goal location at the end of simulation (6 seconds) and the results are shown in Table. III. We can see that MA-IRL better resembles the true cost functions.

B. INTERACTION dataset

We further apply the MA-IRL algorithm on a real world traffic dataset, the INTERACTION dataset [5], which has highly interactive driving scenarios collected from different countries. Scenarios such as a roundabout, highway merging, and intersections are included.

1) Data Processing: We pick a highway merging scenario as it involves the vehicles negotiating their merging actions, leading to challenging game-theoretic planning for all vehicles. Among all the merging trajectories, each data point consists of three vehicles: a merging vehicle whose trajectory starts from the onramp and ends on a main highway lane, a leader vehicle that is ahead of the merging vehicle, and a follower vehicle that is behind the merging vehicle in the highway. Leader vehicles are treated as obstacles which do not interact with the other two vehicles, and we assume that the merging vehicle and the follower vehicle follow an ECE solution. We discard the trajectories that are too long or too short (longer than 9 seconds or shorter than 5 seconds) and end up with 87 trajectories in total. All 87 data points are recorded in the same highway onramp in China from different recordings. The average time to finish merging is 6.3 seconds in the data set. Fig. 3 shows an example snapshot, where car 8 is merging onto highway after car 13.

2) Features: For this challenging and complex driving task, we propose features that capture efficiency (driving forward and achieving the lane merge for the merging vehicle), comfort and energy efficiency (penalizing large accelerations), and safety (penalizing relative speed and a Gaussian-cost on close proximity with other cars). In total, we learn the cost coefficients for 9 features for both merging and follower vehicles.

To evaluate the learning results on the INTERACTION data set, since we do not know the true cost function, we cannot compute the corresponding cost distribution. Instead, we compute the root mean squared error on trajectories that are generated from the learned cost functions. For each trajectory tuple in the data set, we run 10 trials using the ECE policy with the learned cost function from both MA-IRL and CIOC. Then, we compute the average deviation from true trajectories over 5 seconds. For position errors, MA-IRL achieves significantly better prediction accuracy than the CIOC model due to the interactive nature of this scenario. IDM's performance is roughly on par with MA-IRL. We note that IDM is a popular expert designed, and empirically validated driving model for traffic simulation, and is only applicable for car following scenarios. Hence, IDM acts as aproxy for the unknown ground truth policy in this scenario. Moreover, we learn the parameters of IDM to specifically fit this data set, thereby directly approximating the control policy of the vehicles in the data set.

In addition, we follow previous work and evaluate the learned policy by computing task-related statistics from demonstrations and generated samples. We measure average speed of both merging and following cars, the average Euclidean distance between leader and merging cars, and average Euclidean distance between merging and following cars. The results are summarized in Table IV. In Fig. 4, the position rmse of MA-IRL is similar to that of IDM and better than CIOC. From Table IV, we can see that MA-IRL achieves good approximation results in average speed of merging vehicle and follower vehicle, which is on par with IDM's performance and better than CIOC. For average distance between merging vehicle and other vehicles, the performance of all three algorithms are similar and have larger deviation than speed statistics. We think the reason is that average distance is subject to more uncertainty.

VIII. CONCLUSION

In this paper, we presented an algorithm for Maximum Entropy Inverse Reinforcement Learning in multi-agent set-

Fig. 3: A snapshot of highway merging scenario in the INTERACTION data set. Vehicle 8 is a merging vehicle between car 13 and 6.


Fig. 4: Root Mean Square Error

task statistics	true	MA-IRL	IDM	CIOC
M avg. speed [m/s]	3.272	3.372	2.876	2.949
F avg. speed [m/s]	2.893	2.592	2.574	2.410
ML avg. dist [m]	7.770	7.438	7.951	8.041
MF avg. dist [m]	7.447	8.716	6.844	7.009

TABLE IV: Task statistics for generated simulations using learned cost function/policy compared to human demonstrations from the same initial configuration.

tings. To enable learning from boundedly rational agents such as humans, we defined a notion of noisy equilibrium called Entropic Cost Equilibrium (ECE). We proved that this noisy equilibrium is indeed an extension of the maximum entropy principle to the multi-agent setting. We then proved that ECE policies can be obtained in closed form for the special class of linear quadratic Gaussian games. We further presented an algorithm for approximating ECE policies for general nonlinear games. Finally, knowing how to find maximum-entropy multi-agent policies, we provided an iterative algorithm for learning agents' costs from a set of demonstrations when cost functions are represented as a linear combination of a set of features. We verified and validated our algorithm using both synthetic and real-world data set demonstrating that our MA-IRL algorithm can successfully capture the interactions between agents and

learn accurate models of agents' costs.

Limitations and Future Work: While our simulations and experiments show the effectiveness of our MA-IRL algorithm, a key future direction is to enable learning general nonlinear costs. In this paper, we assumed that agents' costs can be represented as weighted sum of a set of features. An important future direction is to relax this assumption. Moreover, to avoid handpicking the set of relevant features a-priori, we would like to study the extensions of this work to enable learning cost functions that are represented in the form of function approximators such as neural networks. Finally, the current work can be extended to the case of unknown system dynamics as well. We assumed that the system dynamics were known and fixed. We believe that our method can be further extended to systems with unknown dynamics.

APPENDIX A
PROOF OF THEOREM 2

We provide the proof for Theorem 2. We know from Theorem 1 that to find the ECE policies of a linear quadratic game $G = (\mathcal{S}, \mathcal{A}, f, g, p_0, c)$, we can equivalently find the Nash equilibrium policies of the auxiliary game $\tilde{G} = (\mathcal{S}, \mathcal{A}, f, g, p_0, \tilde{c}, T)$. The key idea is that at Nash equilibria of the game $\tilde{G}$, the cost-to-go or the value function for each agent at any given state can be represented in closed form via a quadratic function of the state. Using dynamic programming, the closed-form solution of the cost-to-go function can be propagated backwards in time to find both the Nash equilibrium policy and the cost-to-go of agent i.

More precisely, let $\boldsymbol{\pi}^*$ denote a set of Nash equilibrium policies of the game $\tilde{G}$. Consider an agent $i \in [N]$. Fix the policies of all the other agents $\boldsymbol{\pi}^{-i^*}$. Then, using Definition 1, for the game $\tilde{G}$, the (Nash) equilibrium policy of the agent π^{i^*} optimizes

$$\min_{\pi^i} \mathbb{E}_{\pi^i, \boldsymbol{\pi}^{-i^*}} \sum_{k=1}^{T} \frac{1}{2} \left(s_k^\mathsf{T} Q^i s_k + l_i^\mathsf{T} s_k + \sum_{j \in [N]} a_k^{j\mathsf{T}} R^{ij} a_k^j \right) - \sum_{k=1}^{T} \mathcal{H}(\pi_k^i(.|s_t)). \quad (41)$$

First, note that for any policy $\pi^i(.|s_t)$, since the dynamics are linear and the cost is quadratic, the expectation of the terms

inside the parenthesis in (41), depends only the first and second order moments of the policy π^i. Therefore, for any fixed first and second order moments of the policy π^i, the optimal policy π^{i^*} will be the distribution with maximum entropy subject to fixed first and second order moments. Leveraging the fact that with fixed first and second order moments, the maximum entropy distribution is a normal distribution (see for instance, Theorem 8.6.5 in [51]), we can conclude that the optimal policy for each agent i at any time step $\pi_t^{i^*}$ is normally distributed. Hence, it remains to find the mean and covariance of the optimal policy at any given time step. For every agent i, let the mean and covariance of agent i's policy be represented by μ_t^i and Σ_t^i at every time step t. The optimal solution to (41) can be found using dynamic programming. Let $V_t^{i^*}(s_t)$ denote the cost-to-go of agent i at Nash equilibrium at a given time step t

$$
V_t^{i^*}(s_t) := \min_{\pi^i} \mathbb{E}_{\pi^i, \boldsymbol{\pi}^{-i^*}} \sum_{k=t}^{T} \frac{1}{2} \Big(s_k^\mathsf{T} Q^i s_k + {l^i}^\mathsf{T} s_k + \sum_{j \in [N]} {a_k^j}^\mathsf{T} R^{ij} a_k^j \Big) - \sum_{k=t}^{T} \mathcal{H}(\pi_k^i(.|s_k)). \tag{42}
$$

We prove via induction that the cost-to-go for each agent at a given time step is a quadratic function of the state. We assume for time $t+1$, we have

$$
V_{t+1}^{i^*}(s_{t+1}) = \frac{1}{2} s^\mathsf{T} Z_{t+1}^i s_{t+1} + {\xi_{t+1}^i}^\mathsf{T} s_{t+1} + n_{t+1}^i, \tag{43}
$$

where Z_{t+1}^i, ξ_{t+1}^i, and n_{t+1}^i are the coefficients of the quadratic cost-to-go at time $t+1$. We will propagate the cost-to-go structure (43) backwards in time to prove that

$$
V_t^{i^*}(s_t) = \frac{1}{2} s^\mathsf{T} Z_t^i s_t + {\xi_t^i}^\mathsf{T} s_t + n_t^i. \tag{44}
$$

We prove this via dynamic programming. For time step t, we have

$$
V_t^{i^*}(s_t) = \min_{\pi_t^i} \mathbb{E}_{\pi_t^i, \boldsymbol{\pi}_t^{-i^*}} \frac{1}{2} \Big(s_t^\mathsf{T} Q^i s_t + \sum_{j \in [N]} {a_t^j}^\mathsf{T} R^{ij} a_t^j \Big) - \mathcal{H}(\pi_t^i(.|s_t)) + \mathbb{E}_{s_{t+1}} V_{t+1}^{i^*}(s_{t+1}). \tag{45}
$$

Using dynamics (21), the structure of cost-to-go function for time $t+1$ from (43), and the fact that the policies for every time step are normally distributed, one can verify that

$$
\mathbb{E} V_{t+1}^{i^*}(s_{t+1}) = \frac{1}{2} \Big[s_t^\mathsf{T} A^\mathsf{T} Z_{t+1}^i A s_t + \sum_{j \in [N]} {\mu_t^j}^\mathsf{T} {B^j}^\mathsf{T} Z_{t+1}^i A s_t +
$$
$$
s_t^\mathsf{T} A^\mathsf{T} Z_{t+1}^i \sum_{j \in [N]} B^j \mu_t^j + \sum_{j \in [N]} \left(B^j \mu_t^j \right)^\mathsf{T} Z_{t+1}^i + n_{t+1}^i + \mathrm{tr} \Big(
$$
$$
Z_{t+1}^i \sum_{j \in [N]} B^j \Sigma_t^j {B^j}^\mathsf{T} + Z_{t+1}^i \Big) + {\xi_{t+1}^i}^\mathsf{T} \Big(A s_t + \sum_{j \in [N]} B^j \mu_t^j \Big) \Big]. \tag{46}
$$

Moreover, we know that for a normal policy π_t^i with mean μ_t^i and covariance Σ_t^i, the entropy of the policy is

$$
\mathcal{H}(\pi_t^i(.|s_t)) = \frac{m^i}{2} \log(2\pi e) + \frac{1}{2} \log \det(\Sigma_t^i), \tag{47}
$$

where m^i is the dimension of the action space of agent i, and e is the Euler's number. Using (46) and (47), (45) which is the cost-to-go at time t, (45) can be rewritten as a minimization with respect to μ_t^i and Σ_t^i when we fix the mean and covariance of other agents' policies

$$
V_t^{i^*}(s_t) = \min_{\mu^i, \Sigma^i} \mathbb{E}_{\pi^i, \boldsymbol{\pi}^{-i^*}} \frac{1}{2} [s_t^\mathsf{T} Q^i s_t + \sum_{j \in [N]} {\mu_t^j}^\mathsf{T} R^{ij} \mu_t^j +
$$
$$
s_t^\mathsf{T} A^\mathsf{T} Z_{t+1}^i A s_t + \sum_{j \in [N]} {\mu_t^j}^\mathsf{T} {B^j}^\mathsf{T} Z_{t+1}^i A s_t +
$$
$$
s_t^\mathsf{T} A^\mathsf{T} Z_{t+1}^i \sum_j B^j \mu^j + \big(\sum_{j \in [N]} B^j \mu_t^j \big)^\mathsf{T} Z_{t+1}^i \big(\sum_{j \in [N]} B^j \mu_t^j \big)] +
$$
$$
{\xi_{t+1}^i}^\mathsf{T} \big(A s_t + \sum_{j \in [N]} B^j \mu^j \big) + n_{t+1}^i - \frac{1}{2} \log \det(\Sigma_t^i) - \frac{1}{2} m^i 2\pi e
$$
$$
+ \frac{1}{2} \sum_{j \in [N]} \mathrm{tr}(R^{ij} \Sigma_t^j + Z_{t+1}^i \big(\sum_{j \in [N]} B^j \Sigma_t^j {B^j}^\mathsf{T} \big) + Z_{t+1}^i). \tag{48}
$$

The objective function of (48) is a convex quadratic function of μ_t^i and a convex function of Σ_t^i. Taking the derivative of the objective function in (48) with respect to μ_t^i and Σ_t^i and setting the derivative equal to zero, one can verify that $\mu_t^{i^*}$ and $\Sigma_t^{i^*}$ have the structure of (24) and (25). By replacing $\mu_t^{i^*}$ and $\Sigma_t^{i^*}$ in (48) and rearranging the terms, one can see that $V_t^{i^*}(s_t)$ is also a quadratic function of the state s_t. This proves that at every time step t, $V_t^{i^*}(s_t) = \frac{1}{2} s^\mathsf{T} Z_t^i s_t + {\xi_t^i}^\mathsf{T} s_t + n_t^i$. By matching the coefficients of this quadratic function, recursions (28), (29), (30) and (31) are obtained. Note that we also need to verify the base case for our induction, i.e, the cost-to-go at the final time step $V_T^{i^*}(s_T)$ is quadratic. At the final time step T, the optimal policy $\pi_T^{i^*}$ is a zero mean normal distribution with covariance matrix $\Sigma_T^i = (R^i)^{-1}$. Plugging in this policy, we verify that in the final time step the cost-to-go is also quadratic in state. This proves the base case for our induction, which completes our proof that the optimal cost-to-go at any time step is quadratic in states. Now that we have proved that at every time step, $V_t^{i^*}(s_t)$ is of the form (44), we can obtain the recursions (26) and (27) on the matrix P_t^i and the vector α_t^i by replacing (24) and (25) in (45) and matching the coefficients.

REFERENCES

[1] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning." in *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.

[2] B. D. Ziebart, "Modeling purposeful adaptive behavior with the principle of maximum causal entropy," Ph.D. dissertation, Department of Machine Learning, Carnegie Mellon University, December 2010.

[3] S. Levine and V. Koltun, "Continuous inverse optimal control with locally optimal examples," in *Proceedings of the 29th International Coference on International Conference on Machine Learning*, 2012, pp. 475–482.

[4] D. Hadfield-Menell, A. Dragan, P. Abbeel, and S. Russell, "Cooperative inverse reinforcement learning," *arXiv preprint arXiv:1606.03137*, 2016.

[5] W. Zhan, L. Sun, D. Wang, H. Shi, A. Clausse, M. Naumann, J. Kümmerle, H. Königshof, C. Stiller, A. de La Fortelle, and M. Tomizuka, "INTERACTION Dataset: An INTERnational, Adversarial and Cooperative moTION Dataset in Interactive Driving Scenarios with Semantic Maps," *arXiv:1910.03088 [cs, eess]*, Sep. 2019.

[6] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical review E*, vol. 62, no. 2, p. 1805, 2000.

[7] R. E. Kalman, "When is a linear control system optimal?" *Journal of Basic Engineering*, vol. 86, no. 1, pp. 51–60, 03 1964.

[8] A. Y. Ng, S. J. Russell *et al.*, "Algorithms for inverse reinforcement learning." in *Icml*, vol. 1, 2000, p. 2.

[9] P. Abbeel and A. Y. Ng, "Apprenticeship learning via inverse reinforcement learning," in *Proceedings of the twenty-first international conference on Machine learning*, 2004, p. 1.

[10] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *International conference on machine learning*. PMLR, 2016, pp. 49–58.

[11] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 2016, pp. 4572–4580.

[12] S. Albrecht, K. Ramirez-Amaro, F. Ruiz-Ugalde, D. Weikersdorfer, M. Leibold, M. Ulbrich, and M. Beetz, "Imitating human reaching motions using physically inspired optimization principles," in *2011 11th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2011, pp. 602–607.

[13] K. Mombaur, A. Truong, and J.-P. Laumond, "From human to humanoid locomotion—an inverse optimal control approach," *Autonomous robots*, vol. 28, no. 3, pp. 369–383, 2010.

[14] P. Englert, N. A. Vien, and M. Toussaint, "Inverse kkt: Learning cost functions of manipulation tasks from demonstrations," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1474–1488, 2017.

[15] C. Awasthi, "Forward and inverse methods in optimal control and dynamic game theory," Master's thesis, the University of Minnesota, 2019.

[16] M. Menner and M. N. Zeilinger, "Maximum likelihood methods for inverse learning of optimal controllers," *arXiv preprint arXiv:2005.02767*, 2020.

[17] S. Natarajan, G. Kunapuli, K. Judah, P. Tadepalli, K. Kersting, and J. Shavlik, "Multi-agent inverse reinforcement learning," in *2010 Ninth International Conference on Machine Learning and Applications*. IEEE, 2010, pp. 395–400.

[18] K. Waugh, B. D. Ziebart, and J. A. Bagnell, "Computational rationalization: The inverse equilibrium problem," *arXiv preprint arXiv:1308.3506*, 2013.

[19] X. Lin, P. A. Beling, and R. Cogill, "Multi-agent inverse reinforcement learning for zero-sum games," *arXiv preprint arXiv:1403.6508*, 2014.

[20] X. Lin, S. C. Adams, and P. A. Beling, "Multi-agent inverse reinforcement learning for certain general-sum stochastic games," *Journal of Artificial Intelligence Research*, vol. 66, pp. 473–502, 2019.

[21] L. Yu, J. Song, and S. Ermon, "Multi-agent adversarial inverse reinforcement learning," *arXiv preprint arXiv:1907.13220*, 2019.

[22] J. Song, H. Ren, D. Sadigh, and S. Ermon, "Multi-agent generative adversarial imitation learning," *arXiv preprint arXiv:1807.09936*, 2018.

[23] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles," *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24 972–24 978, 2019.

[24] S. Lecleac'h, M. Schwager, and Z. Manchester, "Lucidgames: Online unscented inverse dynamic games for adaptive trajectory prediction and planning," *IEEE Robotics and Automation Letters*, 2021.

[25] F. Köpf, J. Inga, S. Rothfuß, M. Flad, and S. Hohmann, "Inverse reinforcement learning for identification in linear-quadratic dynamic games," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 14 902–14 908, 2017.

[26] S. Rothfuß, J. Inga, F. Köpf, M. Flad, and S. Hohmann, "Inverse optimal control for identification in non-cooperative differential games," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 14 909–14 915, 2017.

[27] J. Inga, E. Bischoff, F. Köpf, and S. Hohmann, "Inverse dynamic games based on maximum entropy inverse reinforcement learning," *arXiv preprint arXiv:1911.07503*, 2019.

[28] C. Awasthi and A. Lamperski, "Inverse differential games with mixed inequality constraints," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 2182–2187.

[29] L. Peters, D. Fridovich-Keil, V. Rubies-Royo, C. J. Tomlin, and C. Stachniss, "Inferring objectives in continuous dynamic games from noise-corrupted partial state observations," *arXiv preprint arXiv:2106.03611*, 2021.

[30] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, "The complexity of computing a nash equilibrium," *SIAM Journal on Computing*, vol. 39, no. 1, pp. 195–259, 2009.

[31] H. A. Simon, "A behavioral model of rational choice," *The quarterly journal of economics*, vol. 69, no. 1, pp. 99–118, 1955.

[32] ——, "Rational decision making in business organizations," *The American economic review*, vol. 69, no. 4, pp. 493–513, 1979.

[33] R. D. McKelvey and T. R. Palfrey, "Quantal response equilibria for normal form games," *Games and economic behavior*, vol. 10, no. 1, pp. 6–38, 1995.

[34] ——, "Quantal response equilibria for extensive form games," *Experimental economics*, vol. 1, no. 1, pp. 9–41, 1998.

[35] S. P. Anderson, J. K. Goeree, and C. A. Holt, "Noisy directional learning and the logit equilibrium," *Scandinavian Journal of Economics*, vol. 106, no. 3, pp. 581–602, 2004.

[36] S. Levine and V. Koltun, "Guided policy search," in *International Conference on Machine Learning*, 2013, pp. 1–9.

[37] ——, "Learning complex neural network policies with trajectory optimization," in *International Conference on Machine Learning*, 2014, pp. 829–837.

[38] T. Haarnoja, H. Tang, P. Abbeel, and S. Levine, "Reinforcement learning with deep energy-based policies," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 06–11 Aug 2017, pp. 1352–1361.

[39] J. Schulman, X. Chen, and P. Abbeel, "Equivalence between policy gradients and soft q-learning," *arXiv preprint arXiv:1704.06440*, 2017.

[40] S. Levine, "Reinforcement learning and control as probabilistic inference: Tutorial and review," *arXiv preprint arXiv:1805.00909*, 2018.

[41] T. Başar and G. J. Olsder, *Dynamic noncooperative game theory*. SIAM, 1998.

[42] D. H. Jacobson, "New second-order and first-order algorithms for determining optimal control: A differential dynamic programming approach," *Journal of Optimization Theory and Applications*, vol. 2, no. 6, pp. 411–440, 1968.

[43] Y. Tassa, N. Mansard, and E. Todorov, "Control-limited differential dynamic programming," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1168–1175.

[44] D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan, and C. J. Tomlin, "Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1475–1481.

[45] M. Wang, N. Mehr, A. Gaidon, and M. Schwager, "Game-Theoretic Planning for Risk-Aware Interactive Agents," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Las Vegas, NV, USA (Virtual): IEEE, Oct. 2020, pp. 6998–7005.

[46] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems." in *ICINCO (1)*. Citeseer, 2004, pp. 222–229.

[47] S. Yakowitz, "Algorithms and computational techniques in differential dynamic programming," *Control and Dynamical Systems: Advances in Theory and Applications*, vol. 31, pp. 75–91, 2012.

[48] J. Van Den Berg, S. Patil, and R. Alterovitz, "Motion planning under uncertainty using iterative local optimization in belief space," *The International Journal of Robotics Research*, vol. 31, no. 11, pp. 1263–1278, 2012.

[49] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for autonomous cars that leverage effects on human actions." in *Robotics: Science and Systems*, vol. 2. Ann Arbor, MI, USA, 2016.

[50] R. P. Bhattacharyya, R. Senanayake, K. Brown, and M. J. Kochenderfer, "Online parameter estimation for human driver behavior prediction," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 301–306.

[51] T. Cover and J. Thomas, *Elements of Information Theory*. Wiley, 2012.