

# Peer-Aware Cost Estimation in Nonlinear General-Sum Dynamic Games for Mutual Learning and Intent Inference

Seyed Yousef Soltanian  
Arizona State University  
Tempe, AZ, United States  
ssoltan2@asu.edu

Wenlong Zhang  
Arizona State University  
Tempe, AZ, United States  
wenlong.Zhang@asu.edu

## ABSTRACT

Dynamic game theory is a powerful tool in modeling multi-agent interactions and human-robot systems. In practice, since the objective functions of both agents may not be explicitly known to each other, these interactions can be modeled as incomplete-information general-sum dynamic games. Solving for equilibrium policies for such games presents a major challenge, especially if the games involve nonlinear underlying dynamics. To simplify the problem, existing work often assumes that one agent is an expert with complete information about its peer, which can lead to biased estimates and failures in coordination. To address this challenge, we propose a nonlinear peer-aware cost estimation (N-PACE) algorithm for general-sum dynamic games. In N-PACE, using iterative linear quadratic (ILQ) approximation of dynamic games, each agent explicitly models the learning dynamics of its peer agent while inferring their objective functions and updating its own control policy accordingly in real time, which leads to unbiased and fast learning of the unknown objective function of the peer agent. Additionally, we demonstrate how N-PACE enables intent communication by explicitly modeling the peer’s learning dynamics. Finally, we show how N-PACE outperforms baseline methods that disregard the learning behavior of the other agent, both analytically and using our case studies. Codes available at: NPACE.

## KEYWORDS

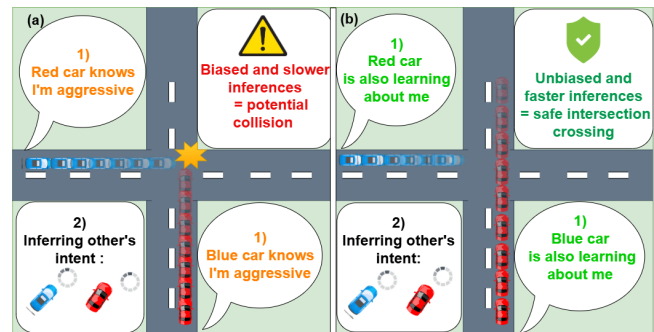
Dynamic Games; Intent Inference; Learning from a Learner; ILQgames

### ACM Reference Format:

Seyed Yousef Soltanian and Wenlong Zhang. 2026. Peer-Aware Cost Estimation in Nonlinear General-Sum Dynamic Games for Mutual Learning and Intent Inference. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/>

## 1 INTRODUCTION

General-sum dynamic games offer a powerful framework for modeling multi-agent interactions [23, 32], with applications in autonomous driving [9], assistive robots [21], and social navigation [11]. However, the process of finding the feedback Nash equilibrium (FBNE) solution of such games, which can be achieved by solving a set of Hamilton-Jacobi-Bellman (HJB) equations [5], is inherently complex, as traditional dynamic programming approaches suffer



**Figure 1: Demonstration of the core idea in N-PACE using a simulated general-sum game between two autonomous cars at an intersection. In (a), each agent assumes its peer knows its aggressiveness level, leading to biased inferences and possible collisions. In (b), N-PACE models the peer agent as a learning entity, resulting in a safe intersection crossing.**

from the *curse of dimensionality* [30]. This problem becomes even more challenging when agents do not know each other’s objective function and must perform inference online while making decisions, resulting in an incomplete-information differential game, a common scenario in human-robot interaction (HRI) [15, 21, 28, 36] and autonomous vehicle interaction [1, 39, 42].

Existing approaches to solve such games often rely on the assumption that one agent is an *expert* which knows the true intent of other agents [17, 19, 29, 32]. Despite success of these approaches in situations where the peer agent is indeed a complete information agent, in scenarios where all agents are learning simultaneously and inferring each other’s objective functions (which is defined as their *intent* in this paper), these assumptions can lead to biased estimates, poor coordination, and failure to achieve each agent’s goal [22, 34]. This is because while doing the inference, they fail to account for the fact that other agents can be learners, whose policies are conditioned on previously observed data [15]. It is even more important to consider such mutual learning in a competitive setting where an agent has no information about the others.

In this paper, we highlight the importance of accounting for the learning dynamics of peer agents while inferring their objectives, addressing critical limitations in existing approaches. We achieve this goal by developing a peer-aware cost estimation algorithm for general-sum nonlinear games (N-PACE), where each agent leverages prior knowledge about its peer agent’s learning dynamics (e.g., gradient-based learner, Bayesian learner) to condition its current inference of the other’s objective on their parameter learning,



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaaamas.org](http://www.ifaaamas.org)). <https://doi.org/10.65109/>

thereby removing the bias in their inferences. The N-PACE algorithm employs iterative linear quadratic (ILQ) approximation of the nonlinear general-sum game [11], both in policy update and intent inference phases. We also demonstrate how (ILQ) enables efficient approximation during the inference phase of our algorithm.

**Contributions.** This paper makes the following contributions. (1) We introduce the N-PACE algorithm for mutual intent inference in multi-agent interaction with continuous action-state and continuous intent space, where all agents are making online inferences. (2) We empirically and theoretically demonstrate the advantages of N-PACE in convergence, safety, and task completion, compared to baseline methods that assume the peer agent is an expert who knows the ego agent’s true intent. (3) We demonstrate how N-PACE, along with access to the learning dynamics of the peer agent, can be leveraged for intent communication and signaling to facilitate mutual learning and adaptation.

The closest works to ours are [11, 19, 22, 34], where [11] uses ILQGame for solving *complete information* multi-agent interactions. In [19], the ILQ approach is employed in a general-sum game where an “expert” demonstrates its intent to a learner agent knowing their learning dynamics, but it does not consider the situation when two learning agents (no experts) try to coordinate which is a more practical setting in multi-agent interactions. Finally, while [22, 34] study the problem of mutual learning and adaptation, the former’s problem setup is limited to a single-step optimization horizon cost optimization, and the latter considers only the linear quadratic setup. This work extends those to nonlinear general-sum games, which is a more practical multi-agent setup.

## 2 RELATED WORK

**Solving general-sum games.** Multi-agent interactions have been explored using various frameworks, including multi-agent reinforcement learning [6, 24], adaptive control [8], and game theory [4, 43]. A large portion of these interactions can be modeled as general-sum dynamic games [44, 45]. Traditional approaches for solving such games rely on adaptive dynamic programming to compute feedback Nash equilibria (FBNE) [37]. However, these methods are often hindered by the *curse of dimensionality* [30], making them impractical for high-dimensional problems and real-time implementation in robotics. To overcome these limitations, physics-informed neural networks (PINN) have been used to solve the HJB equation more efficiently [44]. Some recent work has focused on iterative linear-quadratic (ILQGames) approximations, which offer a more computationally efficient strategy for approximating the multi-agent interactions solution by finding the *local Nash equilibrium* of such games [11]. In this paper, we employ ILQGame for both the learning and control phases to find the approximated Nash equilibrium policies.

**Incomplete information general-sum games.** Solving incomplete information general-sum games is significantly more challenging than their complete-information counterparts [2, 7]. These games often require simplifications, such as modeling one agent (the ego agent) as a learner while assuming others are fully informed [16, 17, 29], treating the problem as an adaptive control problem [21], discretizing the intent space or relying on offline datasets for intent inference and belief updates [9, 13, 25]. Our

work is motivated by the lack of studies on incomplete information dynamic games where all agents are learners [14]. We will highlight the potential issues caused by these simplifications, such as modeling other agents as fully informed.

**Learning from a learner.** Online inference of an expert agent’s objective function has been extensively studied through inverse optimal control and inverse reinforcement learning [27, 31, 33]. While inverse non-cooperative games have been explored both offline and online [18, 26], existing works focus on learning from an expert system. Some works have considered *opponent shaping* and *learning from a learner* agent by incorporating their learning dynamics into offline multi-agent policy training approaches [10, 14], but they are inherently offline multi-agent policy training algorithms and not primarily designed for real-time intent inference and objective learning. Moreover, they are not built on game-theoretic setups where the future actions of one agent are conditioned on those of other agents. Motivated by these *learning from a learner* approaches and recent efforts in modeling human learning dynamics in human-robot interactions [15, 36], we propose a framework for mutual intent inference in systems where all agents are learners.

## 3 PROBLEM FORMULATION

**Notations and Assumptions.** We focus on a two-player setting with agents  $i$  and  $j$ . When referring to an arbitrary agent, we use the index  $k$  and denote their counterpart as  $-k$ . We consider the following system dynamics

$$s_{t+1} = f_t(s_t, a_t^i, a_t^j), \quad (1)$$

where  $s_t \in \mathbb{R}^n$  is the shared state vector at time step  $t$ , and  $a_t^i, a_t^j \in \mathbb{R}^m$  are the control inputs of agents  $i$  and  $j$ , respectively. The function  $f_t$  is assumed to be differentiable with respect to  $s_t, a_t^i$ , and  $a_t^j$ . We assume that all states and control signals  $s_t, a_t^i, a_t^j$  are observable to both players at each time step. For a fixed time horizon  $T$ , the running cost of each player  $k \in \{i, j\}$  is defined as  $g_t^k(s_t, a_t^i, a_t^j; \theta^k)$ , where  $g_t^k$  is a scalar function that is twice differentiable with respect to  $s_t, a_t^i, a_t^j$  and differentiable with respect to the *intent* parameter  $\theta^k$  for  $g_t^k$  (e.g., the targeted goal state or a weighting factor). Additionally, we use the notation  $\hat{\cdot}$  to denote the parameter estimate.

**General-sum game formulation.** We consider a receding horizon cost minimization problem for each agent, where at each time step  $t \in \{0, 1, \dots, T\}$ , each player  $k$  aims to minimize the cumulative cost:

$$J_t^k(s_t, a_t^i, a_t^j) = \sum_{\tau=t}^T g_\tau^k(s_\tau, a_\tau^i, a_\tau^j; \theta^k). \quad (2)$$

The Feedback Nash Equilibrium (FBNE) [3] is a solution concept for the above minimization problem, constrained by (1), where neither agent has an incentive to deviate from their policy, as doing so would increase their cumulative cost. In an FBNE, each agent’s policy is a function of the current state, ensuring that at every time step, the chosen control actions are optimal given the strategies of the other agents. We represent these policies using  $\pi_t^k(s_t; \theta^i, \theta^j)$  such that  $a_t^k = \pi_t^k(s_t; \theta^i, \theta^j)$ , emphasizing that they depend on both  $\theta^i$  and  $\theta^j$ , which is crucial for further formulation in our work.

**The incomplete information game.** The previously discussed FBNE solution was derived under the assumption that both agents

have full knowledge of  $\theta^i$  and  $\theta^j$ . However, in incomplete information games,  $\theta^j$  may be unknown to agent  $i$ , and  $\theta^i$  may be unknown to agent  $j$ . As a result, at each time step, they must rely on estimations  $\hat{\theta}_t^i$  and  $\hat{\theta}_t^j$ , respectively, and update these estimates over time. This leads to policies of the form  $\pi_t^i(s_t; \theta^i, \hat{\theta}_t^i)$  and  $\pi_t^j(s_t; \hat{\theta}_t^i, \theta^j)$ , which are not necessarily the same as FBNE strategies. To address this problem, the next section introduces how we use the ILQ method to obtain the policy  $\pi_t^k(s_t; \theta^k, \hat{\theta}_t^{-k})$  for each agent  $k$ , based on their current estimate of the other agent's parameter  $\hat{\theta}_t^{-k}$ . We then propose appropriate update rules for these estimates in Sec. 3.2.

### 3.1 ILQGame for Multi-Agent Interactions

The iterative linear quadratic regulator (iLQR) approach for locally optimal feedback control of nonlinear systems was first introduced in [20]. More recent work, such as [11], has demonstrated its efficiency in computing local Nash equilibrium solutions for general-sum dynamic games in real time. In our incomplete information game, where we need policies  $\pi_t^k$  to be computed as functions of the estimated parameters  $\hat{\theta}_t^i$  and  $\hat{\theta}_t^j$  in real time, we leverage the proposed ILQGame algorithm in [11] to solve the game for each agent  $k$  at every time step. The main reasons for using ILQGame in this work are its efficiency in real-time computations and, more importantly, its capability to parameterize the obtained policies as functions of  $\hat{\theta}_t^i$  and  $\hat{\theta}_t^j$ , which is crucial for having a differentiable layer in N-PACE as shown in Sections 3.2 and 3.3.

In the ILQGame approach, we begin with an initial admissible policy for each agent. In the forward pass, by applying these policies to the dynamics (1), we obtain a trajectory  $\zeta_t^0 = \{(s_\tau^0, a_\tau^0, r_\tau^0) \mid \tau \in \{t, t+1, \dots, T\}\}$ . Next, by linearizing the dynamics  $f_t$  and quadratizing the running costs  $g_t^k$  in the neighborhood of the trajectory  $\zeta_t^0$ , we construct a finite-horizon LQ game, which can be solved via coupled Riccati equations using the methods proposed in [3]. The linear feedback control solution of this LQ game is then used to update the agents' policies, as suggested in [11]. This iterative process continues until the system trajectory remains unchanged at each iteration and converges to a fixed trajectory, known to be the local Nash equilibrium trajectory  $\zeta_t^*$ .

In an incomplete information game, each agent learns the other's parameter  $\theta^{-k}$ , resulting in an estimate  $\hat{\theta}_t^{-k}$ . As a result, each agent independently solves a different ILQ problem to update their policy using  $\theta^k$  and  $\hat{\theta}_t^{-k}$ , leading to a distinct converged trajectory  $\zeta_t^{*k}(\theta^k, \hat{\theta}_t^{-k})$  at each step. At the converged trajectories, each agent's policy consists of a feedforward time-varying term  $\alpha_t$  and a feedback time varying term  $P_t$  in the form of

$$\begin{aligned} \pi_t^k(s_t; \theta^k, \hat{\theta}_t^{-k}) = & -\alpha_t(\theta^k, \hat{\theta}_t^{-k}) \\ & - P_t(\theta^k, \hat{\theta}_t^{-k})(s_t - s_t^{*k}(\theta^k, \hat{\theta}_t^{-k})), \end{aligned} \quad (3)$$

where  $s_t^{*k}(\theta^k, \hat{\theta}_t^{-k})$  represents states stored in the converged trajectory. In (3), we explicitly maintain and emphasize the dependence of each parameter on  $\theta^k$  and  $\hat{\theta}_t^{-k}$ . This not only highlights how agents update their policies based on new estimates but also serves as an introduction to the N-PACE algorithm.

**REMARK 1. (Extension to Other Methods)** The core idea in this paper is not limited to using ILQ for solving general-sum games. The

key requirement for the success of N-PACE is to have a policy explicitly parameterized by  $\theta^k$  and  $\theta^{-k}$ . In fact, alternative approaches, such as some reinforcement learning algorithms [24, 40], can also borrow the N-PACE core idea as long as they can generate policies parametrized by intent parameters.

### 3.2 Modeling the Learning Dynamics

In Sec. 3.1, we discussed how an agent  $k$  can update its policy at each step based on the current estimate  $\hat{\theta}_t^{-k}$  to solve a general-sum game. However, the question of how to learn and update  $\hat{\theta}_t^{-k}$  remains open. In this section, we propose a learning algorithm under the assumption that agent  $-k$  is an expert, for the sake of clarity and brevity. We then introduce approximation methods to enhance the efficiency of this learning process for real-time applications. Finally, in Sec. 3.3, we discuss how this learning approach can be extended to N-PACE for scenarios where all agents are learners.

**The learning dynamics  $h_t^k$ .** When agents can observe the states and control signals  $s_t, a_t^k, a_t^{-k}$  at each step, they can leverage this information to construct an error term for the online learning of the parameters  $\hat{\theta}_t^{-k}$ . An agent  $k$ , knowing that their peer is solving an ILQGame to generate policies, can attempt to model this policy using their current estimate  $\hat{\theta}_t^{-k}$  as  $\hat{a}_t^{-k} = \pi_t^{-k}(s_t; \hat{\theta}_t^{-k}, \theta^k)$  as given in (3). Consequently, they can define the following update rule as their learning dynamics  $h_t^k$  to refine the estimate of  $\hat{\theta}_t^{-k}$

$$\begin{aligned} \hat{\theta}_{t+1}^{-k} = & h_t^k(\hat{\theta}_t^{-k}, \theta^k, a_t^{-k}, s_t) = \hat{\theta}_t^{-k} \\ & - \alpha_k(\pi_t^{-k}(s_t; \hat{\theta}_t^{-k}, \theta^k) - a_t^{-k})^\top \nabla_{\hat{\theta}_t^{-k}} \pi_t^{-k}(s_t; \hat{\theta}_t^{-k}, \theta^k) \end{aligned} \quad (4)$$

**REMARK 2. (Expert Peer Assumptions)** Notice that in (4), when modeling the policy of the other agent  $-k$ , the parameter  $\theta^k$  is included in the policy model. This effectively assumes that the other agent has full knowledge of  $\theta^k$  and is treated as an expert, which can lead to incorrect and biased estimates.

**REMARK 3. (Flexibility of Learning Dynamics)** The proposed learning dynamics  $h_t^k$  follows a simple gradient descent algorithm with a learning rate  $\alpha_k$ , but it is not restricted to gradient-based methods in the N-PACE. It can also take the form of a Bayesian learner or even a neural network model, as explored in [36].

**Efficient approximation of the gradient.** The learning dynamics proposed in (4) require calculating the gradient of the policy model with respect to the unknown parameter  $\hat{\theta}_t^{-k}$ . The ILQ method, which generates the policy at each step, is an iterative process where the generated trajectories and policy improvement steps depend on the unknown parameter  $\hat{\theta}_t^{-k}$ . While this approach still relates the policy to  $\hat{\theta}_t^{-k}$ , the iterative nature of ILQ can make gradient calculation computationally expensive or even intractable for real-time implementation. Notably, in (3), as well as in the converged trajectory  $\zeta_t^{*k}(\theta^k, \hat{\theta}_t^{-k})$ , all parameters are explicitly expressed as functions of  $\hat{\theta}_t^{-k}$ .

To address this challenge, we adopt the approximation method proposed in [18]. Specifically, we ignore the dependency of the converged trajectory in ILQ on  $\hat{\theta}_t^{-k}$  and replace the gradient of the policy with the gradient of the linear policy obtained by solving the LQ game around the converged trajectory of iLQR. The intuition behind this approximation is that in ILQ, the policy derived from

solving the expanded LQ game around the converged trajectory serves as a good approximation of the final converged policy itself, as discussed and demonstrated in [18]. However, automatic differentiation (AD) methods can still be used to compute this gradient if there are no computational constraints. Alternatively, if a different game solver allows the policy to be directly expressed as a function of the unknown parameter  $\hat{\theta}_t^{-k}$  (e.g., an actor network receiving the parameter as a state input), one can bypass the need for this approximation method.

### 3.3 N-PACE for Multi-Agent Interactions

In Sec. 3.2, we proposed the learning dynamics  $h_i^k$  for each agent. However, the update rule in (4) assumes that agent  $-k$  is aware of  $\theta_k$  while generating its policy, effectively treating agent  $-k$  as an expert during estimation. This assumption is clearly invalid when all agents are learning. In N-PACE, we address this issue by ensuring that each agent  $k$  not only learns parameters of the other agent, but also accounts for the learning procedure of the other agent  $-k$ . This is feasible when agents have access to both learning dynamics,  $h_i^i$  and  $h_i^j$ . Consequently, they can replace  $\theta_k$  with  $\hat{\theta}_t^k$  in equation (4), eliminating the bias in the learning process. As a result, each agent, by having access to learning dynamics and tracking the learning of the other agent, arrives at the following system of learning equations at each step

$$\begin{aligned}\hat{\theta}_{t+1}^i &= h_i^i(\hat{\theta}_t^i, \hat{\theta}_t^j, a_t^i, s_t), \\ \hat{\theta}_{t+1}^j &= h_j^j(\hat{\theta}_t^j, \hat{\theta}_t^i, a_t^j, s_t).\end{aligned}\quad (5)$$

Details of N-PACE are shown in Algorithm 1. It should be noted that N-PACE does not require both agents to follow the same learning dynamics; what matters in N-PACE is that agents have knowledge of each other’s learning dynamics. An important property of N-PACE is that when both agents start with the same initial estimates for  $\hat{\theta}_0^i$  and  $\hat{\theta}_0^j$  and continuously track each other’s learning, they will always track the same estimations for both parameters. As a result, the learning dynamics (5) are effectively centralized, even though each agent performs the learning process independently. The situation where agents know others’ initial priors/estimates often happens in multi-agent settings when there is a strong prior on the parameters (e.g., a predefined level of rationality or aggressiveness in autonomous driving) or when agents are designed to coordinate despite having different unknown intentions (e.g., multi-UAV coordination without communication for cybersecurity). If the initial estimates or learning dynamics of each agent are uncertain (such as human-robot interactions), we refer the readers to works such as [36, 38, 41], which focus on data-driven modeling of these learning processes or latent representation of such learning behaviors. Even if agents misspecify each other’s initial estimates, N-PACE tends to attenuate the resulting bias compared to modeling the peer as a complete information agent that introduces a fixed bias into the learning procedure. Intuitively, for example, when the peer’s update map admits a stable fixed point, the coupled updates steer the estimates toward that fixed point, reducing the initial mismatch over time, trying to keep N-PACE robust to the initial mismatch. In contrast, methods that approximate the peer as a fully informed “expert” introduce a *structural* bias: the model mismatch is baked into the inference mechanism and therefore persists throughout

---

#### Algorithm 1 N-PACE for agent $i$

---

**Initialization:** Initialize  $\hat{\theta}_0^j$  and  $\hat{\theta}_0^i$ .

**for** each time step  $t < T$  **do**

**1. Policy Generation:**

a. Using (3), generate the policy as:

$$\pi_t^i(s_t; \theta^i, \hat{\theta}_t^j) = iLQSolver(s_t, \theta^i, \hat{\theta}_t^j)$$

b. Apply the control  $a_t^i = \pi_t^i(s_t; \theta^i, \hat{\theta}_t^j)$

**2. Prediction:**

a. Predict the  $j$ ’s policy as  $\pi_t^j(s_t; \hat{\theta}_t^j, \hat{\theta}_t^i)$  and model how  $j$ ’s predict your policy  $\pi_t^i(s_t; \hat{\theta}_t^i, \hat{\theta}_t^j)$  using the centralized solver  $iLQSolver(s_t, \hat{\theta}_t^i, \hat{\theta}_t^j)$

b. compute the gradient of each predicted policy with respect to their corresponding parameter  $\nabla_{\hat{\theta}_t^k} \pi_t^k$

**3. Learning:**

a. observe  $s_t, a_t^j$ , form the prediction errors.

b. Using the prediction and computed gradients apply the learning dynamics as:  $\hat{\theta}_{t+1}^i = h_i^i(\hat{\theta}_t^i, \hat{\theta}_t^j, a_t^i, s_t)$

$$\hat{\theta}_{t+1}^j = h_j^j(\hat{\theta}_t^j, \hat{\theta}_t^i, a_t^j, s_t)$$

**end for**

---


the interaction [22, 34]. A robustness analysis of N-PACE to misspecified learning dynamics and initial estimates has been provided in Sec 4.2. While N-PACE’s advantage in eliminating the inference bias of baseline methods is intuitive, the problem it addresses is generally non-convex; thus, any convergence guarantee necessarily depends on the stage-cost structure and the underlying dynamics. But to make the advantage theoretically more precise—relative to baselines that treat the peer as an “expert”—we analyze a class of dynamic games in which each agent’s policy is linear with respect to its own intent parameter  $\theta^k$ . This structure arises when the Hessian of the stage cost  $g_k(\cdot)$  with respect to states and control actions is independent of  $\theta^k$ , while the gradient terms depend linearly on  $\theta^k$ . Such a setting is common in goal-reaching tasks where the intent parameter directly represents the goal position.

**PROPOSITION 1 (N-PACE v.s TREATING THE PEER AS AN EXPERT).** *Suppose the stage costs  $g_t^k(s_t, a_t^i, a_t^j; \theta^k)$  in (2) have Hessians with respect to  $(s, a^i, a^j)$  that are independent of  $\theta^k$ , with  $\theta^k$  entering only linearly in the first-order terms, and on any compact operating set  $s_t \in \Omega$ , the policy of each agent admits the form  $\pi_t^k(s_t; \theta^k, \hat{\theta}_t^{-k}) = M_t^k(s_t, \hat{\theta}_t^{-k}) \theta^k + S_t^k(s_t, \hat{\theta}_t^{-k})$  with  $M_t^k, S_t^k$  bounded. Also, suppose each agent updates its peer parameter online by the gradient descent learning rule (4), then the convergence of intent estimation in N-PACE is guaranteed for sufficiently small learning rate  $\alpha$ , whereas methods that treat the peer as an expert/complete-information agent can fail to converge for any value of learning rate.*

**PROOF.** See Appendix. □

### 3.4 N-PACE for Intent Communication

The N-PACE algorithm also enables intent communication. Notice that the learning dynamics of agent  $-k$ , denoted as  $h_i^{-k}$ , are fundamentally conditioned on the actions  $a_t^k$  taken by agent  $k$ . As a result,



**Figure 2: Results of Case Study 1, the assistive lunar lander. (a) Trajectory comparison (b) and (c) Agents’ learning performance under different algorithms (d) final landing position error. (e) Final sum of intent learning error. (f) and (g) cumulative costs.**

each agent can modify its actions to influence the learning process of the other. To enable this, we propose to introduce an additional intent communication term in the form of  $\eta(\hat{\theta}_t^k - \theta^k)^2$  to the instantaneous cost function of each agent,  $g_\tau^k(s_\tau, a_\tau^i, a_\tau^j; \theta^k)$ , where  $\eta$  is the weighting scalar. This leads to the following optimization problem for each agent  $k$ :

$$\begin{aligned} \min_{a_t^k} J_t^{*k}(s_t, a_t^k, a_t^{-k}) &= \sum_{\tau=t}^T \left[ g_\tau^k(s_\tau, a_\tau^k, a_\tau^{-k}; \theta^k) + \eta \|\hat{\theta}_\tau^k - \theta^k\|^2 \right] \\ \text{s.t. } s_{\tau+1} &= f_t(s_\tau, a_\tau^k, a_\tau^{-k}), \quad \hat{\theta}_{\tau+1}^{-k} = h_1^k(\hat{\theta}_\tau^{-k}, \hat{\theta}_\tau^k, a_\tau^{-k}, s_\tau), \\ \hat{\theta}_{\tau+1}^k &= h_1^{-k}(\hat{\theta}_\tau^k, \hat{\theta}_\tau^{-k}, a_\tau^k, s_\tau). \end{aligned} \quad (6)$$

The influence problem above is computationally expensive and may be intractable for real-time implementation. Remember that the learning dynamics  $h_1^k, h_1^{-k}$  themselves require solving a general-sum game, using ILQgames in this case, for their prediction phase and computing the gradient of that prediction. To address this challenge, we propose an approximation method for solving the optimization problem in (6). Instead of considering the evolution of the other agent’s parameter learning throughout the entire horizon  $[t, T]$  in (6) to form the cumulative cost  $J^{*k}$ , we focus only on a single-step update of the other agent’s learning as  $\hat{\theta}_{t+1}^{-k} = h_1^k(\hat{\theta}_t^{-k}, \hat{\theta}_t^k, a_t^{-k}, s_t)$ , so we can modify the cumulative cost of agent  $k$  at each step as

$$J^{*k}(s_t, \hat{\theta}_t^{-k}, a_t^i, a_t^j) = J^k(s_t, a_t^i, a_t^j) + \eta \|\hat{\theta}_{t+1}^k - \theta^k\|^2, \quad (7)$$

where  $\hat{\theta}_{t+1}^k$  depends on  $a_t^k$  through the learning dynamics  $h_1^{-k}$ .

**LEMMA 1 (MONOTONE REDUCTION OF INTENT ESTIMATION ERROR UNDER INTENT COMMUNICATION).** *Assuming a fixed  $t$ , ego agent  $k$ , state  $s_t$ , peer action  $a_t^{-k}$ , current beliefs  $(\hat{\theta}_t^k, \hat{\theta}_t^{-k})$ , an admissible action domain  $\mathcal{A}_t^k \subseteq \mathbb{R}^{m^k}$ , a base objective  $C_t^k(a^k) = J^k(s_t, a^k, a_t^{-k})$ , and any learning dynamics  $h_1^{-k}$ . For any  $0 \leq \eta_1 < \eta_2$ , let action  $a_t^{k,\eta}$  be any solution of the optimization problem (7) at each step  $t$ , the parameter learning error  $E_t^k(a_t^{k,\eta}) = \|\hat{\theta}_{t+1}^k(a_t^{k,\eta}) - \theta^k\|^2$  satisfies  $E_t^k(a_t^{k,\eta_2}) \leq E_t^k(a_t^{k,\eta_1})$ . In particular,  $\eta \mapsto E_t^k(a_t^{k,\eta})$  is non-increasing.*

**PROOF.** See Appendix.  $\square$

**COROLLARY 1 (INTENT COMMUNICATION NEVER HURTS).** *Let  $a_t^{k,0} \in \arg \min_{\mathcal{A}_t^k} C_t^k$  be any no-teaching optimizer and, for  $\eta \geq 0$ , let  $a_t^{k,\eta} \in \arg \min_{\mathcal{A}_t^k} \Phi_t^{k,\eta}$  with  $\Phi_t^{k,\eta} := C_t^k + \eta E_t^k$ . Then for every  $\eta > 0$ ,  $E_t^k(a_t^{k,\eta}) \leq E_t^k(a_t^{k,0})$ . Moreover, if  $\arg \min_{\mathcal{A}_t^k} C_t^k \cap \arg \min_{\mathcal{A}_t^k} E_t^k = \emptyset$  (i.e. disjoint argmin sets), then there exists  $\eta^* > 0$  such that for all  $\eta > \eta^*$ ,  $E_t^k(a_t^{k,\eta}) < E_t^k(a_t^{k,0})$ .*

**PROOF.** The non-strict inequality is Lemma 1 with  $\eta_1 = 0$  and  $\eta_2 = \eta$ . For strictness, Define

$$C_\star := \min_{\mathcal{A}_t^k} C_t^k, \quad E_\star := \min_{a \in \arg \min C_t^k} E_t^k(a), \quad E_\star := \min_{\mathcal{A}_t^k} E_t^k(a).$$

Disjointness implies  $E_\star < E_0$ . Pick  $a_E \in \arg \min E_t^k$  and set  $\Delta C := C_t^k(a_E) - C_\star \geq 0$ ,  $\Delta E := E_0 - E_\star > 0$ . Then

$$\Phi_t^{k,\eta}(a_E) - \Phi_t^{k,\eta}(a_t^{k,0}) = \Delta C - \eta \Delta E.$$

Choosing  $\eta^* := \Delta C / \Delta E$  makes this negative for all  $\eta > \eta^*$ , hence any minimizer  $a_t^{k,\eta}$  satisfies  $E_t^k(a_t^{k,\eta}) \leq E_t^k(a_E) = E_\star < E_0 = E_t^k(a_t^{k,0})$ .  $\square$


We will show the effectiveness of this approximation for intent communication and empirical evidence validating Lemma 1 and Corollary 1 in Case Study 3 in Sec. 4.

## 4 CASE STUDIES

In this section, we show that explicitly modeling a peer’s learning dynamics with N-PACE improves task completion (Case Study 1) and safety (Case Studies 2–3) compare to baseline methods. We also evaluate the robustness of N-PACE to misspecified initial beliefs about the peer’s parameter learning (Case Study 2), and present both an ablation on initialization error and the effect of intent communication (Case Study 3). Across all case studies, the learning updates use gradient-descent-based Bayesian inference with a Gaussian belief. More details on baseline designs can be found in the extended version of this work [35]. (Simulation codes)

### 4.1 Case Study 1: Assistive Lunar Lander

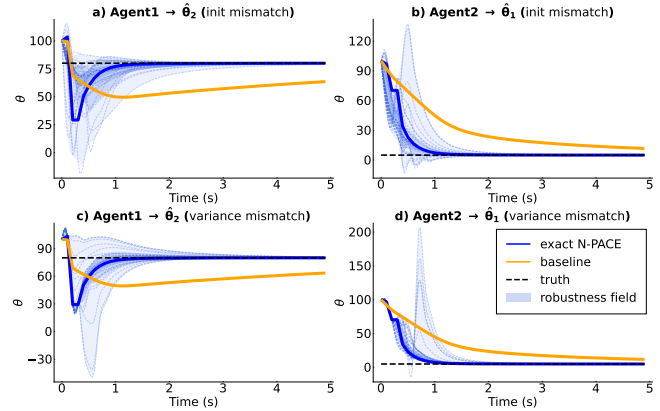
Our first example focuses on the importance of N-PACE in task completion. In this example, we study a shared control problem, an assistive lunar lander scenario in which a human agent controls only



**Figure 3: Trajectories of the lane-merging case study, the trajectories are plotted relative to the blue car coordinate (a) Safe lane merging under complete information (b) baseline methods, a risky behavior in merging from the aggressive driver (red). (c) Non-risky lane merging in N-PACE similar to the complete game, despite a 30% error in modeling the peer’s initial estimate.**

the angle of the lunar lander  $\phi$  by applying the torque  $T$ , while an autonomous agent controls the thrust force  $F$  along the direction of the lunar lander’s angle (assuming no gravity). Both agents cooperate to land the lunar lander at the final position  $(x_f, y_f)$ , but with asymmetric information: the human knows only the final desired  $x$ -position  $x_f$ , while the autonomous agent knows only the final desired  $y$ -position  $y_f$ . As a result, the human must learn the parameter  $\hat{y}_f$ , and the autonomous agent must learn  $\hat{x}_f$  based on observed interactions. The underlying nonlinear dynamics follow  $\dot{x} = v_x, \dot{y} = v_y, \dot{\phi} = \omega, \dot{v}_x = F \sin \phi, \dot{v}_y = T$  the running costs for the human  $g_h$  and the agent  $g_a$  are defined as  $g_a = 10(x - \hat{x}_f)^2 + 10(y - y_f)^2 + 10(\phi - \frac{\pi}{2})^2 + 10(v_x^2 + v_y^2 + \omega^2) + F^2$  and  $g_h = 10(x - x_f)^2 + 10(y - \hat{y}_f)^2 + 10(\phi - \frac{\pi}{2})^2 + 10(v_x^2 + v_y^2 + \omega^2) + T^2$ . The game has a time horizon of  $T = 5s$  with a sampling time of  $\Delta T = 0.1s$  that is used to discretize the above dynamics. In this example, we set  $x_f = 5$  and  $y_f = -5$ , while the lunar lander starts at the initial position  $(-5, 8)$  with a zero initial angle. We conducted three simulations for this case study: 1) a complete information game where both agents have full knowledge of  $x_f$  and  $y_f$  as a ground truth baseline, 2) a game theoretic scenario where each agent treats its peer as an expert, assuming the other agent has complete information, 3) a non-game-theoretic model predictive control (MPC) inspired by the baselines in [29] where each agent uses a history of observed data and fits a linear policy to its partner policy parameters (including the unknown goal position) and using this policy perform a receding horizon MPC, and finally 4) a scenario in which agents operate according to N-PACE.

Figure 2 shows a summary of the results, where Fig. 2 (a) presents the trajectory of the lunar lander under different algorithms, explicitly demonstrating the superiority of N-PACE compared to other approaches. It can be seen that the final position of the lunar lander in N-PACE is closer to the desired final position, and the entire trajectory aligns more closely with the complete information game trajectory (the local Nash equilibrium). In contrast, treating the other agent as an expert or the MPC-based baseline results in the lunar lander landing at a different position, emphasizing the importance of modeling the learning dynamics of the other agent for successful task completion. Figs. 2 (b) and (c) present the results of parameter learning performance for each agent under different algorithms, illustrating how N-PACE fast and successfully converges to the true final position goals in comparison to other baselines. Notably, treating the other agent as an expert can lead to a significant overshoot in the learning process. Intuitively, this may be due




**Figure 4: a) & b) Robustness of N-PACE in intent inference under mismatched assumptions about the peer’s initial belief. c) & d) robustness to wrong choice of initial variance in peer’s GD-based Bayesian inference. Envelopes representing 50 mismatched runs.**

to a high initial bias in assuming that the peer agent has complete information. 2 (f) and (g) interestingly show how N-PACE achieves a lower cumulative cost for each agent compared to other methods in this case study.

## 4.2 Case Study 2: Lane Merging

This example aims to study the N-PACE robustness to imperfect modeling of the peers’ learning dynamics. In our second case study, we consider a lane-merging task between two autonomous vehicles. The first car (blue) moves in a straight line, while the second car (red) attempts to merge into its lane. The blue car can only accelerate or decelerate  $a_1$ , whereas the red car can control both its acceleration along its direction of motion  $a_2$  and its steering angle velocity  $\delta$ . The game has a 5s horizon with a sampling time of  $\Delta T = 0.1s$  for dynamic discretization. During the game, Each agent is unaware of the aggressiveness level of others in their cost function and must learn it in real-time. The underlying dynamic is defined as  $\dot{x}_1 = v_1, \dot{v}_1 = a_1, \dot{x}_2 = v_2 \cos \phi_2, \dot{y}_2 = v_2 \sin \phi_2, \dot{\phi}_2 = \delta_2, \dot{v}_2 = a_2$ , and the running cost  $g$  for each car is  $g_{blue} = 0.1(x_1 - 25)^2 + 0.1v_1^2 + \theta_1 e^{-(d(t) - d_{safe})^2} + a_1^2$ , and  $g_{red} = 1y_2^2 + 10\phi_2^2 + 0.1v_2^2 + 0.1(x_2 - 25)^2 + \theta_2 e^{-(d(t) - d_{safe})^2} + a_2^2 + \delta_2^2$ , where  $d(t) = (x_1 - x_2)^2 + y_2^2$  is the distance between the two cars,  $d_{safe}$  is the safe distance they aim to maintain, and  $\theta_k$  is a parameter



**Figure 5: Comparison of  $\log(1 + \text{MSE})$  learning error across 300 Monte Carlo simulations of case study 3 with random values of aggressiveness.**

in the range  $[0, 100]$  that represents the aggressiveness of each agent (lower  $\theta_k$  corresponds to more aggressive behavior). Each agent initially assumes that the other agent is a non-aggressive, safety-conscious driver with  $\hat{\theta}_0^k = 100$ . However, both agents exhibit some level of aggressiveness: the first agent (the blue car) is highly aggressive with  $\theta_1 = 5$ , while the second agent (the red car) has a moderate aggressiveness level of  $\theta_2 = 80$ . To test robustness to misspecified initial beliefs, we introduced an initialization error  $e_{init} = \pm 70\%$  in each agent’s guess of the other’s initial estimation and randomly sampled 50 “initial parameter estimate of others” for each agent from this error bound. Equivalently, each agent  $k$  for each run uniformly samples from the range  $\hat{\theta}^k \in [30, 170]$  and sets it as the initial estimate of agent  $-k$  (although the true initial estimates are  $\hat{\theta}_0^k = 100$ ). Across all runs, as shown in figure 4 a) and b), N-PACE converged to the peer’s true intent and outperformed the baseline, indicating robustness to initialization mismatch. For higher error values, we observed N-PACE starting to show diverging behavior. Also, to better represent the effect of errors in the modeling of peers’ learning dynamics, we used an error  $e_{var} = \pm 55\%$  in how each agent models the initial variance in gradient-based Bayesian inference of the other agent as a way to model misspecified convergence speeds of the modeled learning dynamics. We randomly sampled 50 variances from this range and ran our simulations. Again, across all cases as shown in figure 4 c) and d) we observed N-PACE converging and outperforming the baseline. Indicating how N-PACE can lead to better intent inference despite the mismatch in modeling the initial estimate or learning dynamics of others. Qualitative trajectory comparison results and intent inference Results appear in Fig. 3 and Fig. 4

### 4.3 Case Study 3: Interaction Driving


In our third case study, to extensively study N-PACE and the intent communication approach proposed in (7), we conduct Monte Carlo simulations of an intersection driving game between two agents under incomplete information, similar to the scenario shown in Fig. 1. In this case study, two autonomous vehicles attempt to navigate through an intersection while being unaware of each other’s aggressiveness parameter,  $\theta_k$ . The agents are only able to control their acceleration along with the direction of their movement  $x, y$ . The dynamic of the game follows kinematic equations  $\dot{x}_1 = v_1, \dot{v}_1 = a_1, \dot{y}_2 = v_2, \dot{v}_2 = a_2$ , and each agent’s cost function is formed

as  $g_1 = (x_1 - 8)^2 + v_1^2 + 10 \theta_1 \exp(-(d(t) - d_{safe})) + a_1^2$  and  $g_2 = (y_2 - 8)^2 + v_2^2 + 10 \theta_2 \exp(-(d(t) - d_{safe})) + a_2^2$ . where  $d(t) = x_1^2 + y_2^2$  is the distance between two vehicles and  $d_{safe}$  is the safe distance they try to maintain. The aggressiveness of each agent is also defined same as Case Study 2. The game continues for 5 sec with a sampling time of  $\Delta T = 0.1$  sec, resulting in 50 samples.

In this example, each agent’s aggressiveness parameter,  $\theta_k$ , takes values in the range  $\theta_k \in [20, 50]$ , where a lower  $\theta_k$  indicates a more aggressive agent. In our Monte Carlo study, we sampled 300 aggressiveness values uniformly from  $[20, 50]$  for each agent. Using these 300 sampled values, we ran the game 300 times under five different scenarios: 1) Complete information, 2) The expert peer agent assumption, 3) An adversary min-max game baseline where each agent  $k$ , instead of trying to do intent inference, at each step decides conservatively based on  $\min_{a^k} \max_{\theta^{-k} \in [20, 50]} J^k(s_t, a_t^i, a_t^j)$  4) N-PACE, 5) N-PACE with intent communication, as proposed in Sec. 3.4, for different intent communication weighting values  $\eta = 0.1, 1, 10$ . The benchmarking results of these studies are shown in Table (1). As expected, under the complete information game scenario, no collisions were observed. The peer expert assumption resulted in 13 collisions. In contrast, N-PACE resulted in only 4 collisions. More interestingly, although the minmax approach was a conservative approach with a high average intersection crossing time, it still resulted in 2 collision cases, where on the other hand our intent communication approach (built on top of N-PACE) not only ensured safe crossing for all 300 scenarios, it also achieved the fastest crossing time among all incomplete-information scenarios, ensuring both safety and smooth interaction. This was also the case for  $\eta = 10$ , while only a single collision occurred for  $\eta = 0.1$ . The trajectories from these 300 simulations for different scenarios are shown in Fig. 6, where we plot the  $x$ -position of the first agent against the  $y$ -position of the second agent. These results highlight the improvements in both inference accuracy and safety assurance using the proposed intent communication framework built on top of N-PACE, although at the expense of increasing the average control effort (see table 1). Finally, Fig. 5 presents the log of the Mean Squared learning Error (MSE) of both agents across all 300 Monte Carlo simulations. We use  $\log(1 + \text{MSE})$  for better visualization and comparison. Interestingly, the results in Fig. 5 align well with the theoretical results in corollary 1 and lemma 1 that state increasing  $\eta$  always reduces the estimation error. These results highlight how intentional communication improves learning performance using higher  $\eta$ .

**Table 1: Case study 3 benchmark**

Method	Failure rate	Average control effort [ $m/s^2$ ]	Average time to cross [s]
Complete game	<b>0.0%</b>	5.00	1.497 ± 0.043
Expert peer assumption	4.33%	4.92	1.593 ± 0.044
N-PACE	1.33%	4.86	1.599 ± 0.054
N-PACE + $\eta = 0.1$	0.33%	5.42	1.521 ± 0.009
N-PACE + $\eta = 1.0$	<b>0.0%</b>	5.44	1.499 ± 0.029
N-PACE + $\eta = 10.0$	<b>0.0%</b>	5.45	<b>1.489 ± 0.035</b>
Minmax	0.66%	6.23	1.651 ± 0.165



**Figure 6: Trajectory  $(x, y)$  results of 300 simulations of the two-player incomplete information intersection game with random values of aggressiveness. Gray boxes indicate the collision area. (a) The complete information game results in no collisions. (b) The expert peer assumption method results in 13 collisions. (c) N-PACE results in only 4 collisions. (d) A conservative adversarial baseline still results in 2 collision cases. (e) N-PACE, combined with intent communication, achieves 0 collisions.**

## 5 CONCLUSIONS AND FUTURE WORK

In this paper, we presented N-PACE framework for solving general-sum dynamic games under intent uncertainty with continuous action-states and continuous intent space, which relies on treating the peer agent as a learning agent while inferring its cost parameters. We demonstrated the importance of accounting for the learning behavior of others and showed how this framework outperformed the approach of treating the peer agent as a complete-information. Additionally, we proposed an intent communication framework built on top of N-PACE for signaling in such settings. N-PACE assumes access to the learning dynamics of the peer agent. While such information is often available in many multi-agent systems, there are scenarios—such as human-robot interactions—where the learning dynamics may be unknown. The question of "learning the learning dynamics" presents an interesting research direction for future work, particularly in light of recent studies exploring the use of transformers for this purpose [15, 36]. Also, the mentioned works or cognitive research finding stating that humans can be modeled as Bayesian learners [12], pave the way for future applications of N-PACE in HRI. Finally, although N-PACE can conceptually be extended to more than two agents, as demonstrated in the extended version of this work in [35], scaling to large multi-agent systems remains computationally demanding and requires further research.

## APPENDIX

**PROOF OF PROPOSITION 1.** Let  $e_t^k := \hat{\theta}_t^k - \theta^k$  be the parameter estimation error of agent  $k$  using the N-PACE method, and  $\tilde{e}_t^k := \hat{\theta}_t^k - \theta^k$  be the parameter estimation error under the assumption of having an expert peer, for  $k \in \{i, j\}$ . The realized peer actions obey  $a_t^k = M_t^k(s_t, \hat{\theta}_t^{-k})\theta^k + S_t^k(s_t, \hat{\theta}_t^{-k}) + w_t^k$  (linear in  $\theta^k$ ) where  $w_t^k$  is a bounded variance measurement noise. Then the N-PACE and Expert peer assumption method error recursions are:

$$\begin{aligned}
 \text{(N-PACE)} \quad e_{t+1}^i &= (I - \alpha M_t^{i\top} M_t^i) e_t^i + \alpha M_t^{i\top} w_t^i, \\
 e_{t+1}^j &= (I - \alpha M_t^{j\top} M_t^j) e_t^j + \alpha M_t^{j\top} w_t^j, \\
 \text{(Expert assu)} \quad \tilde{e}_{t+1}^i &= (I - \alpha \tilde{M}_t^{i\top} \tilde{M}_t^i) \tilde{e}_t^i - \alpha \tilde{M}_t^{i\top} C_t^i \tilde{e}_t^j + \alpha \tilde{M}_t^{i\top} w_t^i, \\
 \tilde{e}_{t+1}^j &= (I - \alpha \tilde{M}_t^{j\top} \tilde{M}_t^j) \tilde{e}_t^j - \alpha \tilde{M}_t^{j\top} C_t^j \tilde{e}_t^i + \alpha \tilde{M}_t^{j\top} w_t^j,
 \end{aligned}$$

where  $\tilde{M}_t^k := M_t^k(s_t, \theta^{-k})$ , and the off-diagonal terms  $C_t^i, C_t^j$  include both  $M$ - and  $S$ -sensitivities via the mean-value form

$$\begin{aligned}
 C_t^i &= \int_0^1 \frac{\partial}{\partial \xi_t^i} \left( M_t^i(s_t, \xi_t^i(\tau)) \theta^i + S_t^i(s_t, \xi_t^i(\tau)) \right) d\tau, \\
 \xi_t^i(\tau) &:= \hat{\theta}_t^j + \tau(\theta^i - \hat{\theta}_t^j)
 \end{aligned}$$

and analogously for  $C_t^j$ ; thus  $(M_t^i(s_t, \theta^j) - M_t^i(s_t, \hat{\theta}_t^j))\theta^i + (S_t^i(s_t, \theta^j) - S_t^i(s_t, \hat{\theta}_t^j)) = C_t^i(\theta^j - \hat{\theta}_t^j) = -C_t^i \tilde{e}_t^j$  (and symmetrically for  $j$ ). If  $\|M_t^k\| \leq L$ ,  $\mathbb{E}\|w_t^k\|^2 \leq \sigma_k^2$ , and there is persistent excitation  $\lambda_{\min}(\mathbb{E}[M_t^{k\top} M_t^k]) \geq \mu_k > 0$ , then from N-PACE errors dynamics we have  $\mathbb{E}[\|e_{t+1}^i\|^2 | \mathcal{F}_t] \leq (1 - 2\alpha\mu_k + \alpha^2 L^4) \|e_t^i\|^2 + \alpha^2 L^2 \sigma_k^2$ , so any  $0 < \alpha < 2\mu/L^4$  gives contraction in mean square (a.s. convergence with diminishing stepsizes; bounded MSE floor with constant stepsizes). In contrast, expert assumption error dynamics contains the off-diagonal couplings  $-\alpha \tilde{M}_t^{i\top} C_t^i$  and  $-\alpha \tilde{M}_t^{j\top} C_t^j$  that arise from both the  $M$ - and  $S$ -term mismatches; unless these are uniformly dominated by the diagonal self-excitation (e.g., the symmetric part of  $\tilde{M}_t^{i\top} \tilde{M}_t^i$  exceeds  $\tilde{M}_t^{i\top} C_t^i$  for  $k \in \{i, j\}$ ), the one-step map need not be a contraction for any  $\alpha > 0$ , so the intent inference while assuming the peer agent is an expert is not guaranteed to converge and may even result in unstable parameter estimation/inference.  $\square$

**PROOF OF LEMMA 1.** Fix  $t$ , agent  $k$ , and data  $(s_t, a_t^k, \hat{\theta}_t^k, \hat{\theta}_t^{-k})$ . Let the admissible set  $\mathcal{A}_t^k$  be nonempty, and suppose minimizers  $a_t^{k,\eta} \in \arg \min_{a^k \in \mathcal{A}_t^k} \Phi_t^{k,\eta}(a^k)$  exist for the weights  $\eta$  under consideration (e.g., compact  $\mathcal{A}_t^k$  and continuity of  $C_t^k, E_t^k$  suffice). Recall the definitions:

$$\begin{aligned}
 C_t^k(a^k) &:= J^k(s_t, a^k, a_t^{-k}), \quad \hat{\theta}_{t+1}^k(a^k) := h_t^{-k}(\hat{\theta}_t^k, \hat{\theta}_t^{-k}, a^k, s_t), \\
 E_t^k(a^k) &:= \|\hat{\theta}_{t+1}^k(a^k) - \theta^k\|^2, \quad \Phi_t^{k,\eta}(a^k) := C_t^k(a^k) + \eta E_t^k(a^k).
 \end{aligned}$$

Take any  $0 \leq \eta_1 < \eta_2$ . By optimality of  $a_t^{k,\eta_2}$  for  $\Phi_t^{k,\eta_2}$  and  $a_t^{k,\eta_1}$  for  $\Phi_t^{k,\eta_1}$ ,

$$C_t^k(a_t^{k,\eta_2}) + \eta_2 E_t^k(a_t^{k,\eta_2}) \leq C_t^k(a_t^{k,\eta_1}) + \eta_2 E_t^k(a_t^{k,\eta_1}), \quad (8)$$

$$C_t^k(a_t^{k,\eta_1}) + \eta_1 E_t^k(a_t^{k,\eta_1}) \leq C_t^k(a_t^{k,\eta_2}) + \eta_1 E_t^k(a_t^{k,\eta_2}). \quad (9)$$

Add (8) and (9) and cancel the  $C_t^k$  terms on opposite sides to obtain

$$(\eta_2 - \eta_1) \left( E_t^k(a_t^{k,\eta_2}) - E_t^k(a_t^{k,\eta_1}) \right) \leq 0.$$

Since  $\eta_2 - \eta_1 > 0$ , it follows that  $E_t^k(a_t^{k,\eta_2}) \leq E_t^k(a_t^{k,\eta_1})$ . As  $\eta_1, \eta_2$  were arbitrary with  $\eta_1 < \eta_2$ , the map  $\eta \mapsto E_t^k(a_t^{k,\eta})$  is non-increasing.  $\square$

## ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1944833. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

## REFERENCES

- [1] Sunny Amatya, Mukesh Ghimire, Yi Ren, Zhe Xu, and Wenlong Zhang. 2022. When Shall I Estimate Your Intent? Costs and Benefits of Intent Inference in Multi-Agent Interactions. In *2022 American Control Conference (ACC)*. IEEE, 586–592.
- [2] Robert J Aumann, Michael Maschler, and Richard E Stearns. 1995. *Repeated games with incomplete information*. MIT press.
- [3] Tamer Başar and Geert Jan Olsder. 1998. *Dynamic noncooperative game theory*. SIAM.
- [4] Daan Bloembergen, Karl Tuyls, Daniel Hennes, and Michael Kaisers. 2015. Evolutionary dynamics of multi-agent learning: A survey. *Journal of Artificial Intelligence Research* 53 (2015), 659–697.
- [5] Alberto Bressan. 2010. Noncooperative differential games. a tutorial. *Department of Mathematics, Penn State University* 81 (2010).
- [6] Lorenzo Canese, Gian Carlo Cardarilli, Luca Di Nunzio, Rocco Fazzolari, Daniele Giardino, Marco Re, and Sergio Spanò. 2021. Multi-agent reinforcement learning: A review of challenges and applications. *Applied Sciences* 11, 11 (2021), 4948.
- [7] Pierre Cardaliaguet and Catherine Rainer. 2012. Games with incomplete information in continuous time and for continuous types. *Dynamic Games and Applications* 2 (2012), 206–227.
- [8] Fei Chen, Wei Ren, et al. 2019. On the control of multi-agent systems: A survey. *Foundations and Trends® in Systems and Control* 6, 4 (2019), 339–499.
- [9] Yi Chen, Lei Zhang, Tanner Merry, Sunny Amatya, Wenlong Zhang, and Yi Ren. 2021. When shall i be empathetic? the utility of empathetic parameter estimation in multi-agent interactions. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2761–2767.
- [10] Jakob N Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. 2017. Learning with opponent-learning awareness. *arXiv preprint arXiv:1709.04326* (2017).
- [11] David Fridovich-Keil, Ellis Ratner, Lasse Peters, Anca D Dragan, and Claire J Tomlin. 2020. Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games. In *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 1475–1481.
- [12] Thomas L Griffiths, Nick Chater, Charles Kemp, Amy Perfors, and Joshua B Tenenbaum. 2010. Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences* 14, 8 (2010), 357–364.
- [13] Haimin Hu, Kensuke Nakamura, Kai-Chieh Hsu, Naomi Ehrlich Leonard, and Jaime Fernández Fisac. 2023. Emergent coordination through game-induced nonlinear opinion dynamics. In *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 8122–8129.
- [14] Alexis Jacq, Matthieu Geist, Ana Paiva, and Olivier Pietquin. 2019. Learning from a learner. In *International Conference on Machine Learning*. PMLR, 2990–2999.
- [15] Kushal Kedia, Atiksh Bhardwaj, Prithwish Dan, and Sanjiban Choudhury. 2024. Interact: Transformer models for human intent prediction conditioned on robot actions. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 621–628.
- [16] Forrest Laine, David Fridovich-Keil, Chih-Yuan Chiu, and Claire Tomlin. 2021. Multi-hypothesis interactions in game-theoretic motion planning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 8016–8023.
- [17] Simon Le Cleac'h, Mac Schwager, and Zachary Manchester. 2021. Lucidgames: Online unscented inverse dynamic games for adaptive trajectory prediction and planning. *IEEE Robotics and Automation Letters* 6, 3 (2021), 5485–5492.
- [18] Jingqi Li, Chih-Yuan Chiu, Lasse Peters, Somayeh Sojoudi, Claire Tomlin, and David Fridovich-Keil. 2023. Cost inference for feedback dynamic games from noisy partial state observations and incomplete trajectories. *arXiv preprint arXiv:2301.01398* (2023).
- [19] Jingqi Li, Anand Siththaranjan, Somayeh Sojoudi, Claire Tomlin, and Andrea Bajcsy. 2024. Intent Demonstration in General-Sum Dynamic Games via Iterative Linear-Quadratic Approximations. *arXiv preprint arXiv:2402.10182* (2024).
- [20] Weiwei Li and Emanuel Todorov. 2004. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *First International Conference on Informatics in Control, Automation and Robotics*, Vol. 2. SciTePress, 222–229.
- [21] Yanan Li, Gerolamo Carboni, Franck Gonzalez, Domenico Campolo, and Etienne Burdet. 2019. Differential game theory for versatile physical human–robot interaction. *Nature Machine Intelligence* 1, 1 (2019), 36–43.
- [22] Changliu Liu, Wenlong Zhang, and Masayoshi Tomizuka. 2016. Who to blame? learning and control strategies with information asymmetry. In *2016 American Control Conference (ACC)*. IEEE, 4859–4864.
- [23] Dylan P Losey, Craig G McDonald, Edoardo Battaglia, and Marcia K O'Malley. 2018. A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction. *Applied Mechanics Reviews* 70, 1 (2018), 010804.
- [24] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems* 30 (2017).
- [25] Negar Mehr, Mingyu Wang, Maulik Bhatt, and Mac Schwager. 2023. Maximum-entropy multi-agent dynamic games: Forward and inverse solutions. *IEEE transactions on robotics* 39, 3 (2023), 1801–1815.
- [26] Timothy L Molloy, Jairo Inga Charaja, Sören Hohmann, and Tristan Perez. 2022. *Inverse optimal control and inverse noncooperative dynamic game theory*. Springer.
- [27] Timothy L Molloy, Jason J Ford, and Tristan Perez. 2020. Online inverse optimal control for control-constrained discrete-time systems on finite and infinite horizons. *Automatica* 120 (2020), 109109.
- [28] Stefanos Nikolaidis, Swaprava Nath, Ariel D Procaccia, and Siddhartha Srinivasa. 2017. Game-theoretic modeling of human adaptation in human-robot collaboration. In *Proceedings of the 2017 ACM/IEEE international conference on human-robot interaction*. 323–331.
- [29] Lasse Peters, Andrea Bajcsy, Chih-Yuan Chiu, David Fridovich-Keil, Forrest Laine, Laura Ferranti, and Javier Alonso-Mora. 2024. Contingency games for multi-agent interaction. *IEEE Robotics and Automation Letters* 9, 3 (2024), 2208–2215.
- [30] Warren B Powell. 2007. *Approximate Dynamic Programming: Solving the curses of dimensionality*. Vol. 703. John Wiley & Sons.
- [31] Nicholas Rhinehart and Kris M Kitani. 2017. First-person activity forecasting with online inverse reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision*. 3696–3705.
- [32] Wilko Schwarting, Alyssa Pierson, Javier Alonso-Mora, Sertac Karaman, and Daniela Rus. 2019. Social behavior for autonomous vehicles. *Proceedings of the National Academy of Sciences* 116, 50 (2019), 24972–24978.
- [33] Ryan Self, Moad Abudia, SM Nahid Mahmud, and Rushikesh Kamalapurkar. 2022. Model-based inverse reinforcement learning for deterministic systems. *Automatica* 140 (2022), 110242.
- [34] Seyed Yousef Soltanian and Wenlong Zhang. 2025. PACE: A Framework for Learning and Control in Linear Incomplete-Information Differential Games. *Proceedings of Machine Learning Research* vol 283 (2025), 1–15.
- [35] Seyed Yousef Soltanian and Wenlong Zhang. 2025. Peer-Aware Cost Estimation in Nonlinear General-Sum Dynamic Games for Mutual Learning and Intent Inference. *arXiv:2504.17129 [eess.SY]* <https://arxiv.org/abs/2504.17129>
- [36] Ran Tian, Masayoshi Tomizuka, Anca D Dragan, and Andrea Bajcsy. 2023. Towards modeling and influencing the dynamics of human learning. In *Proceedings of the 2023 ACM/IEEE international conference on human-robot interaction*. 350–358.
- [37] Kyriakos G Vamvoudakis and Frank L Lewis. 2011. Multi-player non-zero-sum games: Online adaptive learning solution of coupled Hamilton–Jacobi equations. *Automatica* 47, 8 (2011), 1556–1569.
- [38] Woodrow Zhuyuan Wang, Andy Shih, Annie Xie, and Dorsa Sadigh. 2022. Influencing towards stable multi-agent interactions. In *Conference on robot learning*. PMLR, 1132–1143.
- [39] Yiwei Wang, Yi Ren, Steven Elliott, and Wenlong Zhang. 2019. Enabling courteous vehicle interactions through game-based and dynamics-aware intent inference. *IEEE Transactions on Intelligent Vehicles* 5, 2 (2019), 217–228.
- [40] Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8 (1992), 229–256.
- [41] Annie Xie, Dylan Losey, Ryan Tolsma, Chelsea Finn, and Dorsa Sadigh. 2021. Learning latent representations to influence multi-agent interaction. In *Conference on robot learning*. PMLR, 575–588.
- [42] Yang Xing, Chen Lv, Huaji Wang, Hong Wang, Yunfeng Ai, Dongpu Cao, Efstathios Velenis, and Fei-Yue Wang. 2019. Driver lane change intention inference for intelligent vehicles: Framework, survey, and challenges. *IEEE Transactions on Vehicular Technology* 68, 5 (2019), 4377–4390.
- [43] Yaodong Yang and Jun Wang. 2020. An overview of multi-agent reinforcement learning from game theoretical perspective. *arXiv preprint arXiv:2011.00583* (2020).
- [44] Lei Zhang, Mukesh Ghimire, Zhe Xu, Wenlong Zhang, and Yi Ren. 2024. Pontryagin neural operator for solving general-sum differential games with parametric state constraints. In *6th Annual Learning for Dynamics & Control Conference*. PMLR, 1728–1740.
- [45] Lei Zhang, Mukesh Ghimire, Wenlong Zhang, Zhe Xu, and Yi Ren. 2023. Approximating discontinuous Nash equilibrium values of two-player general-sum differential games. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3022–3028.