

# Multi-Agent Strategy Explanations for Human-Robot Collaboration

Ravi Pandya\*, Michelle Zhao\*, Changliu Liu, Reid Simmons, Henny Admoni

**Abstract**—As robots are deployed in human spaces, it is important that they are able to coordinate their actions with the people around them. Part of such coordination involves ensuring that people have a good understanding of how a robot will act in the environment. This can be achieved through explanations of the robot’s policy. Much prior work in explainable AI and RL focuses on generating explanations for single-agent policies, but little has been explored in generating explanations for *collaborative* policies. In this work, we investigate how to generate multi-agent strategy explanations for human-robot collaboration. We formulate the problem using a generic multi-agent planner, show how to generate visual explanations through strategy-conditioned landmark states and generate textual explanations by giving the landmarks to an LLM. Through a user study, we find that when presented with explanations from our proposed framework, users are able to better explore the full space of strategies and collaborate more efficiently with new robot partners.

## I. INTRODUCTION

Imagine you have just packed up your belongings to move into a new house. You and a friend have to move some large boxes out to the moving van, but it’s impossible to hold boxes and open doors along the way yourself. How would you decide the right strategy for the two of you to efficiently move all the boxes outside? You will likely briefly stop to discuss the different ways you might coordinate to hold the doors and move the boxes out, then settle on a plan before moving everything.

In this work, we take a step towards enabling robots with the same capability: *coordinating strategies for collaborative tasks* with people (Fig. 1). As robots start being deployed alongside people in social, home, and manufacturing settings, it is increasingly important for robots to be able to collaborate fluently with and around people. This is particularly true when there are multiple potential strategies for collaboration, since we need to make sure all agents take the same strategy to complete the task efficiently.

We examine how collaborative fluency can be facilitated by a robot through a combination of proactive and reactive processes that communicate strategies for coordination *a priori* but adapt to human strategies *in situ*. In this way, effective collaboration is not merely a reactive process where a robot purely adapts to the goals and behaviors of a human partner, but is a composite of reactive and proactive behaviors where the robot can leverage communication and suggestion to guide collaborations while remaining adaptive to the human in execution. Toward this goal, we examine how

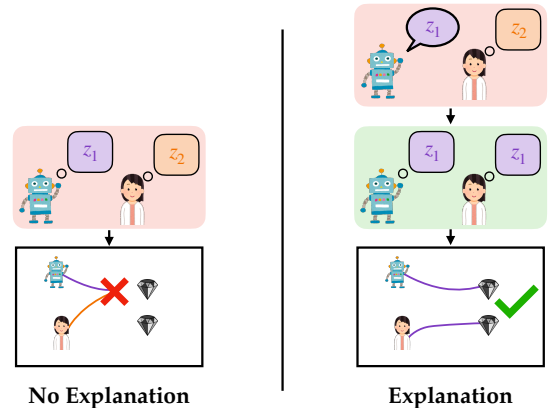


Fig. 1: Depiction of how strategy explanations can solve the strategy alignment problem. The task is for each agent to grab one diamond without collisions. **Left:** The human and robot start with mismatched strategies, and ultimately collide **Right:** The human and robot start with mismatched strategies, but the robot explains strategy  $z_1$ , aligning with the human and resulting in a successful rollout.

a robot can communicate particular strategies by distilling strategies into a few states.

The contributions of this work are:

- 1) A general method for generating visual strategy explanations for human-AI collaboration in games with multiple Nash equilibria;
- 2) A user study demonstrating that our method improves the ability of real users to collaborate efficiently with autonomous partners in a one-shot setting.

## II. RELATED WORK

**Machine Teaching:** One class of approaches for generating explanations of a robot’s behavior, called machine teaching [1], tries to generate a minimal set of examples that will allow the assumed learner model to understand how decisions were made. Originally applied to classification and regression tasks [2], [3], such techniques have recently been used to generate trajectory demonstrations [4], [5] for learners using inverse reinforcement learning [6]. Such approaches assume that a robot is either teaching a human how to perform a task or explaining how a robot would perform a task in isolation, but prior work lacks methods to generate multi-agent demonstrations for future collaborations.

**Policy Summarization:** A second class of explainability methods generates informative examples of a learned policy that may allow a human observer to understand its decisions, such as by finding states with high-entropy action distributions [7], by highlighting the “importance” of particular states [8] or by generating a reward map [9]. Other related work creates a graph of possible paths to a goal state and summarizes the set of necessary states to pass through [10].

\*contributed equally to this work.

Authors are with the Robotics Institute at Carnegie Mellon University, Pittsburgh, Pennsylvania, {rapandya, mzhao2, cliu6, rsimmons, hadmoni}@andrew.cmu.edu




Fig. 2: Pipeline for generating collaborative strategy explanations for a human. First, we identify strategies by clustering over the space of demonstrated examples. Then, we generate the set of landmark states to summarize each strategy cluster and generate textual descriptions with an LLM. Finally, the robot shows the human user the generated explanation before collaboration.

We note that the literature has focused on summarizing *single-agent* tasks, whereas our approach can be thought of as a multi-agent extension of policy summarization approaches for *collaborative* tasks.

### Human-Robot Collaboration in Multi-Strategy Tasks:

Much work has examined how robots can collaborate with human partners [11]. Recent work allowed robots to replan online to react to the strategies of other agents [12]–[14], sometimes doing explicit inference [15] and even considering safety [16]. Other work has considered online information gathering of the human’s strategy (or reward function) [17], [18], but these approaches rely on collision-avoidance behavior of humans to influence them. We note that these are largely *reactive* approaches to collaborating with humans, but we are interested in enabling robots to proactively choose collaborative strategies with human partners. Other work has been done in the space of proactively influencing multi-agent interactions [19], [20], but are either not tested with real users or are too brittle to interact fluently with them.

## III. MULTI-AGENT MULTI-STRATEGY GAMES

### A. Problem Formulation

We focus on a class of collaborative tasks where there may be multiple strategies (or equilibria) for completing the task. In particular, we consider  $N$ –player games represented by an HiP-MDP (Hidden Parameter Markov Decision Process)  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{Z}, \mathcal{T}, \mathcal{R} \rangle$  where  $\mathcal{S} \subseteq \mathbb{R}^n$  is the joint state space,  $\mathcal{A} \subseteq \mathbb{R}^m$  is the joint action space,  $\mathcal{Z} \subseteq \mathbb{R}^q$  is the hidden parameter space (or *latent strategy* space) where  $q \leq n$ ,  $\mathcal{T}$  is the transition dynamics, and  $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^N$  is the union of all agents’ reward functions. We assume that there are  $N - 1$  robot agents interacting with 1 human agent.

A joint policy  $\pi : \mathcal{S} \times \mathcal{Z} \rightarrow \mathcal{A}$  defines the actions all agents should take in a particular state conditional on the joint strategy  $z$ :  $a^t \sim \pi(s^t | z)$ , where  $a^t$  is the joint action of all agents and  $s^t$  is the joint state of the system. The joint policy is the union of all agents’ individual policies  $\pi_i$  that dictate the individual actions  $a_i^t$  of agent  $i$  under their strategy  $z_i$ :  $a^t = [a_1^t, \dots, a_N^t]^T = [\pi_1(s^t | z_1), \dots, \pi_N(s^t | z_N)]^T$ .

### B. Objective

Each agent has their own individual running reward function  $r_i(\cdot)$ . The agent’s overall objective is to maximize their sum of expected rewards:  $R_i = \sum_t r_i(s^t, a_i^t, s^{t+1})$ . The optimal solution to the HiP-MDP is a global Nash equilibrium,

i.e. one that finds the policies  $\pi_i^*$  and associated strategies  $z_i^*$  where the sum of the agents’ rewards are maximized so that no agent can change strategies to unilaterally improve their reward. Given an all agents’ strategy vectors  $\{z_1^*, \dots, z_N^*\}$ , we will know their rewards on the task  $R_i(z_1^*, \dots, z_N^*)$ , so we can mathematically say (with a slight abuse of notation) that a set of strategies is a Nash equilibrium if:

$$R_i(z_1^*, \dots, z_i^*, \dots, z_N^*) \geq R_i(z_1^*, \dots, z_i, \dots, z_N^*), \forall z_i \in \mathcal{Z}, \forall i \in [N]. \quad (1)$$

Solving for a global Nash equilibrium is generally intractable, so a joint policy  $\pi^*$  for all agents could be obtained through reinforcement learning [14], [19], imitation learning [21], or by solving a dynamic game [12], [13]. These kinds of games in the real world often have multiple classes of (local) Nash equilibrium strategies.

### C. Multi-Agent Strategy Alignment

We study the scenario where  $N - 1$  robots are collaborating with one human partner, which makes it difficult to roll out an optimal joint plan for all agents. The two main challenges here are that 1) humans may not be able to compute Nash equilibrium strategies, particularly in repeated games [22]–[25], and 2) many scenarios of interest have multiple different equilibria, so the human needs to know which one the agents may have decided on (this has been called the *strategy alignment* problem). All of the robot agents could coordinate with each other ahead of time, but they need some mechanism to coordinate with the human. Prior work uses online replanning to collaborate with humans [12], [15], but this turns the robot into a *reactive* agent instead of a *proactive* one. It may be useful to allow the human to lead the team’s strategy, but we want robots to *influence* the strategy that the team chooses. This can allow the human partner to choose strategies they may have otherwise been unfamiliar with.

## IV. METHOD: STRATEGY-CONDITIONED LANDMARKS

The full pipeline for our explanation generation method is illustrated in Fig. 2. We note that many collaborative tasks contain discrete modes, even if the full space of strategies  $\mathcal{Z}$  is continuous. We thus start by clustering  $\mathcal{Z}$ .

### A. Strategy Clustering

To cluster the space of potential strategies, we assume access to a dataset  $\mathcal{D}$  containing trajectories  $\xi_i$  of all  $N$  agents performing the task  $\mathcal{D} = \{\xi_1, \dots, \xi_l\}$  and their

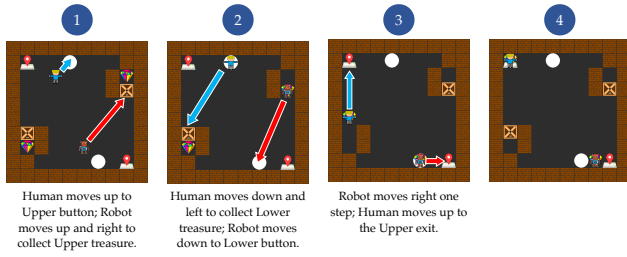


Fig. 3: An explanation of one strategy generated on the collaborative maze task. Each image is a strategy landmark state computed from our method Sec. IV and the text was generated from an LLM.

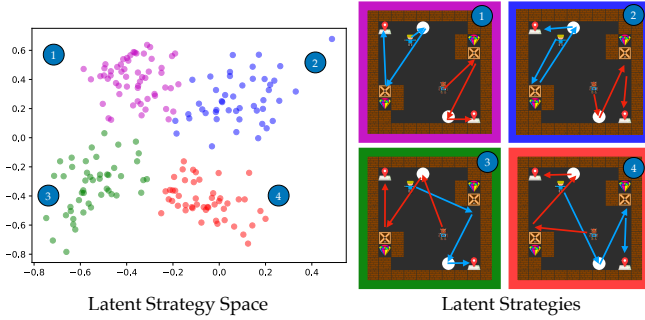


Fig. 4: **Left:** Clustered latent strategy space for maze navigation. **Right:** Resulting strategies from each cluster center are a different ordering of the human and robot picking each jewel vs. standing on each button.

associated strategy vectors  $\mathcal{L} = \{z_1, \dots, z_l\}$ . Such a dataset may come from humans collaborating on the task or from rolling out a policy  $\pi^*$  that solves the game multiple times (such as a differential game solver that outputs the Nash equilibrium policies for all agents).

We can then cluster the strategy vectors  $\mathcal{L}^1$  using k-means clustering. We use silhouette analysis [26] as a method for determining the number of strategy clusters. Resulting clusters of strategy vectors for the collaborative maze task (Sec. V) are visualized in Fig. 6. Ultimately, this gives us a discrete set of strategy cluster centers  $[z_1^*, \dots, z_k^*]$ .

### B. Strategy Landmark State Generation

Prior work [10] has shown that for MDPs that can be written as stochastic shortest-path problems, one can compute a representative subset of the state space called *landmark states* that every possible rollout of the policy must pass through in order to complete the task. We extend this idea to the HiP-MDPs we consider in this work by utilizing the previously discussed clustering of the latent space to generate *strategy-conditioned landmark states*.

**Definition 1 (Strategy-Conditioned Policy Landmark State):** Given an MDP  $\mathcal{M}$  and a strategy cluster center  $z_i^*$ , and a dataset of trajectories  $\mathcal{D} = \{\xi_1, \dots, \xi_l\}$ , a strategy-conditioned policy landmark state  $l$  (or simply *strategy landmark*) is a state (or cluster of states) that all trajectories  $\xi_j$  in the set of trajectories corresponding to the strategy cluster  $z_i^*$  (called  $\mathcal{D}_{z_i^*}$ ) pass through.

To find the set of strategy landmarks, we first discretize the state space then cluster the discretized trajectories from each strategy  $\mathcal{D}_{z_i^*}$ . This results in a set of landmark states

<sup>1</sup>In some scenarios, the latent strategy vectors may just be flattened trajectories, which works fine for our method.

for each strategy cluster:  $S_{land}^i = \{l_i^1, \dots, l_i^{k_i}\}$  where  $k_i$  a hyperparameter that determines the number of clusters and landmark states. One set of landmark states for the collaborative maze task (Sec. V) can be seen in Fig. 3.

### C. Textual Explanation Generation

During initial pilot studies, many users expressed that some amount of textual explanation would be useful for better understanding the strategies. We first wrote out the rules of each game and translated the strategy landmarks generated from Sec. IV-B to natural language. We then asked a large language model (ChatGPT) to generate descriptions for what happened between pairs of strategy landmarks  $(l_i^k, l_i^{k+1})^2$ . Text for one strategy is shown in Fig. 3.

### D. Baseline: Video Explanation

Since existing policy summarization methods do not explicitly deal with collaborative tasks, we construct an intuitive baseline of showing policy rollouts directly. For each strategy cluster  $\mathcal{D}_{z_i^*}$ , the video explanation baseline is a video of the trajectory corresponding to the strategy cluster center  $z_i^*$ . For example, videos for the collaborative maze task show rollouts of the trajectories depicted in Fig. 4 right.

### E. Modification: Landmark Video Explanation

To make the most fair comparison, we additionally augment our method with video explanations. Specifically, we freeze the video explanation from strategy cluster  $z_i^*$  at each landmark state  $l_i^j$  for 10 seconds and display explanation text.

## V. COLLABORATIVE MAZE TASK

For our first experimental domain, we use a 2D maze task built by previous work [14]. In this task, two agents need to collaborate to retrieve one treasure each and then take them to the exits. The agents need to coordinate their actions because the treasures are blocked by doors that can only be opened while one agent is standing on a nearby button. The task is shown in Fig. 4. The agents are rewarded when both agents reach different goals holding treasures.

### A. Co-GAIL

Prior work [14] collected a dataset  $\mathcal{D}$  of pairs of humans completing this task and used this data to train a policy (via their algorithm called Co-GAIL) that can output actions for the different strategies that humans exhibited. We utilize the latent strategy space learned by Co-GAIL to generate strategy clusters for strategy landmark generation (Sec. IV).

In an initial pilot study, we directly used the Co-GAIL model with real users. However, we found it was often too brittle to always finish the task (in strategy 2 in Fig. 4, the robot completed the game with only 25% of participants). In the original paper [14], the authors report 85-100% success rates with people, but only test with 4 human users repeatedly interacting with the system. Users likely became familiar with how to best play the game with this particular agent.

<sup>2</sup>In the future, it may be more scalable to use a vision-language model (VLM) to generate explanations so as to avoid manual task translation.




Fig. 5: A the three-part user study for testing our approach. First, users are assigned an explanation type: *landmarks*, *landmark videos*, *videos*, or *no explanation*. For each of two environments, users play  $N$  games with one explanation between each (or none in the *no explanation* case) and finally play one game with no prior communication with an agent that takes a random strategy.

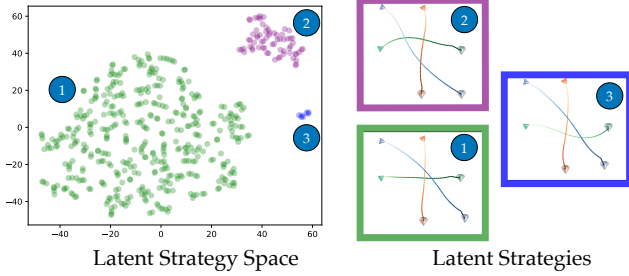


Fig. 6: **Left:** Clustered latent strategy space for social navigation task. **Right:** Resulting strategies from each cluster center. The agents start at the locations denoted by the triangles and the goal is to reach the diamond of the corresponding color without colliding.

### B. Robustly Acting with Human Collaborators

We created one robot planner for each strategy cluster that moves between the landmark states. Since the strategy clusters were still found using the latent codes output by the model trained with Co-GAIL, this approach allows us to take advantage of learning-based techniques while still robustly and reliably completing the task with new users. The planner  $\mathcal{P}$  takes in a pair of sequential landmark states for a particular strategy and outputs the robot’s action:  $a_R^t = \mathcal{P}(l_i^j, l_i^{j+1}, t)$ . Once the human and robot reach the strategy landmark state  $l_i^{j+1}$ , the planner moves to planning for  $l_i^{j+2}$ .

In this setting, we still want to measure how much the robot’s explanations are able to influence the person by measuring their adherence to the planned strategy. However, due to the nature of the task, the human and robot necessarily need to coordinate their actions in order to complete the task. We thus employ a basic Bayesian inference method over the first landmark state of each strategy to allow the robot to correctly respond to the actual chosen human strategy to complete the task:

$$b_R^{t+1}(l^1) = \frac{p(a_H^t | s_H^t; l^1) b_R^t(l^1)}{\sum_{\hat{l}^1} p(a_H^t | s_H^t; \hat{l}^1) b_R^t(\hat{l}^1)}. \quad (2)$$

where the assumed reward for the human in the likelihood function  $p(a_H^t | s_H^t; l^1)$  is the negative distance traveled.

## VI. SOCIAL NAVIGATION TASK

### A. Problem Setup

Our second experimental domain is a 2D social navigation task where  $N$  dynamic agents have different goals and need to reach them while avoiding collisions. Since we’re dealing with dynamical systems, it is natural to represent them as ODEs in continuous time. In practice, the systems are

integrated with a discrete time interval of  $t_s = 0.1s$  using a Runge-Kutta method (RK4). This allows the dynamical system to be in the same form as the HiP-MDP outlined in Sec. III-A. All  $N = 3$  agents have continuous unicycle dynamics  $\dot{s}_i = f(s_i) + g(s_i)a_i$ . The state  $s_i$  consists of the  $(x, y)$  position  $(p_x, p_y)$ , the velocity  $v$  and the steering angle  $\psi$ , and the control action is the acceleration and steering angle change  $a_i = [\dot{v}, \dot{\psi}]^T$ . The MDP state is  $s^t = [s_1^{t,T}, s_2^{t,T}, s_3^{t,T}]^T$ . The rewards are as detailed in [12].

### B. ILQGames

To solve the game posed in Sec. VI-A, we use the LQ-game solver proposed in [12] that solves for local Nash equilibria. Specifically, we use the implementation in Julia—iLQGames.jl [27]. Given an initial condition  $s^0$ , agent dynamics, and agent costs, the solver returns a trajectory  $\xi$  and control policy (strategy)  $z$  for all agents. We perturb the initial condition  $s^0$  500 times and run the solver on each initial condition to get a diverse dataset of multi-agent interactions  $\mathcal{D} = \{\xi_1, \dots, \xi_{500}\}$ ,  $\mathcal{L} = \{z_1, \dots, z_{500}\}$ . A silhouette analysis tells us there should be 3 clusters in the strategy space. We can see the resulting three strategy clusters in Fig. 6. The plotted latent strategy space shows a 2D t-SNE projection of the full latent strategies.

Equipped with a dataset  $\mathcal{D}$  of multi-agent trajectories and an associated set of cluster centers  $\{z_1^*, \dots, z_3^*\}$ , we generate explanations for each strategy in this task as detailed in Sec. IV. Here, it is easy to directly compute whether the human adhered to the planned strategy by checking which existing strategy cluster the rolled-out trajectory belongs to (since completion of the game does not depend on coordinated actions). We thus have the robot agents roll out their policies without inferring the human’s strategy online.

## VII. USER STUDIES

**Research Questions.** We conducted an in-person user study to answer the following research questions: **(RQ1)** How does the explanation type affect human performance and collaborative fluency in a zero-shot coordination task? **(RQ2)** Does the explanation type influence whether people adhere to the robot’s suggested strategy? **(RQ3)** Does robot-guided strategy exploration improve humans’ abilities to collaborate in subsequent games *without* prior strategy coordination?

**Experimental Setup:** We design a three-part study (Fig. 5) to answer all three research questions. All participants are given the same set of instructions for how to play each game before being guided through a series of collaborative tasks




Fig. 7: Task duration on study. Gray dotted: mean, black line: median.

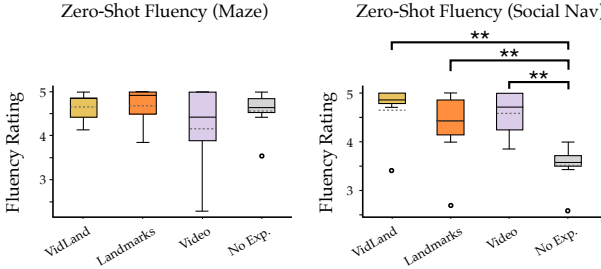


Fig. 8: Subjective fluency ratings

with the robot. Participants play the games using the joystick on a video game controller (Nintendo Switch Pro Controller).

In part 1 (P1), participants are given an explanation of one strategy by the robot (or no explanation in the baseline condition), then play one game with the robot. In part 2 (P2), the robot guides the participant through the rest of the strategies in a randomized order. Finally, in part 3 (P3), once the guided strategy exploration is complete, the participant plays the game with a new robot partner who executes a random strategy without explicitly communicating it. Participants experience this three-part study sequence for both the Maze (Sec. V) and Social Nav (Sec. VI) tasks.

**Independent Variables:** We vary the explanation type with 4 levels: *landmark images* (Sec. IV), *landmark videos* (Sec. IV-E), *videos* (Sec. IV-D), and *no explanation*.

**Objective Measures:** *Task performance* (relevant to P1 and P3) measures how long it takes for the human-robot team to complete the task. *Strategy adherence* (P2) is measured by the number of times the human chose to follow the robot’s suggested strategy. The no-explanation condition is removed from this analysis because the robot does not provide strategy guidance. *Strategy exploration* (P2) is measured by the number of unique strategies played by the team in each of the games prior to the final round with the unseen partner.

**Subjective Measures:** *Subjective fluency* (P1) is a measure of how well the human thinks the team coordinated their actions in the task, and is self-reported by the participant through seven Likert-scale questions drawn from [28] (Cronbach’s  $\alpha = 0.904$  for Maze, and  $\alpha = 0.913$  for Social Nav).

**Hypotheses.** We intuit that landmark-based explanations will improve participants’ understanding of the strategies, compared to video-only explanations which may contain noise. We expect improved task performance, adherence to the robot’s suggestion, and coordination with an unseen partner. **H1:** Participants guided by landmark images and landmark videos will complete the zero-shot coordination

task (in both Maze and Social Nav) faster than the other conditions (P1). **H2:** Participants guided by landmark images and landmark videos will adhere more to the robot’s guidance than the other conditions (P2). **H3:** Participants guided by explanations will execute more unique strategies than those not shown explanations (P2). **H4:** Participants guided by landmark images and landmark videos will be able to adapt to an unseen partner’s strategy faster than the other conditions in both environments (P3).

**Participants:** We recruited 49 in-person participants from the campus community and CMU’s Center for Behavioral and Decision Research participant pool, of which 49.0% self-identified as female, 44.9% as male, and 6.1% as non-binary. Participant ages ranged from 18 to 70 ( $M = 25.7$ ,  $SD = 9.4$ ). We removed participant data in each task if they experienced technical errors, giving us 43 participants, divided as: in the maze task, 14 in *landmark images* condition, 11 in *landmark videos*, 10 in *videos*, and 8 in *no explanation*; in the social navigation task, 13 in *landmark images*, 11 in *landmark videos*, 12 in *videos*, and 7 in *no explanation*.

**Collaborative Maze Results:** For each measure, we removed outliers using the interquartile range method. We used one-way ANOVAS for each measure, and used post-hoc tests with Bonferroni correction to examine main effects. In all figures, (\*) means  $p < 0.05$  and (\*\*) means  $p < 0.01$ .

**H1** was partially supported. We found a significant difference between explanation types on zero-shot coordination performance (Fig. 7), measured by the time taken to complete the game ( $F(3, 37) = 4.11, p = 0.013, \eta^2 = 0.25$ ). Post-hoc analysis showed that participants guided by landmark images ( $p_{bonf} < 0.01$ ) and landmark videos ( $p_{bonf} < 0.01$ ) completed the zero-shot coordination task significantly faster than those not shown any explanation, while those guided by videos showed no significant improvement over no explanation. There was no significant difference between conditions in subjective collaborative fluency (Fig. 8) in the zero-shot interaction ( $F(3, 38) = 2.06, p = 0.122, \eta^2 = 0.14$ ). However, we observe that the variance in subjective fluency and task duration was higher for video explanations.

**H2** was supported. Explanation type significantly affected adherence to the robot’s suggested strategies (Fig. 9,  $F(2, 32) = 6.795, p < 0.01, \eta^2 = 0.30$ ). Participants followed the robot’s guidance significantly more often when presented strategies in landmark image form than in video form ( $p_{bonf} < 0.01$ ).

**H3** was supported. Explanation type significantly affected how many unique strategies were actually performed by the team (Fig. 10,  $F(3, 39) = 7.367, p < 0.01, \eta^2 = 0.36$ ). Landmark videos ( $p_{bonf} < 0.01$ ) and landmark images ( $p_{bonf} < 0.01$ ) both had significantly more unique strategies compared to the no explanation condition. Participants without any robot guidance tended to not explore strategies; they repeatedly stuck with whichever strategy was performed first for games. Video demonstrations did not significantly influence people to explore the strategies in the games played. Although there was no significant difference between video and landmark video results, the means for both zero-




Fig. 9: Adherence to explained strategy by robot




Fig. 10: Number of strategies explored in repeated rounds.

shot task performance and number of strategies explored are higher for landmark videos and landmark images than for videos alone. This suggests that landmarks enhance participant understanding more effectively than videos.

**H4** was not supported. We did not find a significant difference in task performance with an unseen (random strategy) partner by explanation type (Fig. 11,  $F(3, 37) = 0.348, p = 0.79, \eta^2 = 0.03$ ). Since the starting state and environment do not change, practicing the task four times likely enabled participants to adapt to the random strategy partner.

#### Social Navigation Results:

**H1** was partially supported in the social navigation task. Explanation type significantly affected the duration of the first game (Fig. 7,  $F(3, 30) = 20.69, p < 0.01, \eta^2 = 0.67$ ). Participants who were provided video ( $p_{bonf} = 0.04$ ), landmark videos ( $p_{bonf} = 0.03$ ), and landmark images ( $p_{bonf} = 0.04$ ) all performed better than those not offered any strategy explanation. In other words, participants guided by explanations of any form (video, landmark videos, and landmark images) completed the zero-shot coordination task significantly faster than those not shown any explanation. Similarly, explanation type had an effect on subjective collaborative fluency ratings (Fig. 8,  $F(3, 35) = 7.07, p < 0.01, \eta^2 = 0.38$ ). The participants who were provided video ( $p_{bonf} < 0.01$ ), landmark videos ( $p_{bonf} < 0.01$ ), and landmark images ( $p_{bonf} < 0.01$ ) all found the collaboration more fluent than those not offered any strategy explanation.

**H2** was not supported. There was no significant difference in adherence to the robot’s suggested strategies by explanation type (Fig. 9,  $F(2, 33) = 0.056, p = 0.94, \eta^2 = 0.003$ ).

**H3** was supported. Explanation type had a significant effect on how many unique strategies were actually performed by the team (Fig. 10,  $F(3, 39) = 13.48, p < 0.01, \eta^2 = 0.51$ ). Landmark videos ( $p_{bonf} = 0.03$ ), landmark images ( $p_{bonf} = 0.03$ ), and videos ( $p_{bonf} = 0.03$ ) all significantly increased the number of unique strategies the team executed




Fig. 11: Task duration with new (random) partner.

compared to the no explanation condition.

**H4** was not supported. We did not find a significant difference in task performance with the unseen (random strategy) partner based on explanation type (Fig. 11,  $F(3, 33) = 1.30, p = 0.29, \eta^2 = 0.11$ ).

In the social navigation task, all explanations improved subjective fluency, task performance, and strategy exploration. Unlike in the maze task, each agent must only avoid—but does not rely on—the others. Landmarks and videos may provide similar benefits in such a simple task.

## VIII. CONCLUSION AND FUTURE WORK

We proposed a method for strategy alignment in multi-agent tasks when there are multiple discrete modes of collaboration. Specifically, we defined a framework for generating visual strategy explanations via strategy-conditioned landmark states. We evaluated the proposed method on two experimental domains where a robot (or multiple robots) used strategy explanations to induce efficient collaborations with a human partner. We ultimately found promising results; users tended to perform better in their first collaboration with a robot when presented with a strategy explanation, particularly one using landmarks. Users also explored the space of possible strategies more when guided by the robot’s explanations to do so.

Though we did not find that explanations prepared users for future collaborations, it suggests that in more complex tasks, robots may be able to communicate collaborative strategies efficiently. Researchers may consider explanation generation using home simulators [29], [30]. One challenge when scaling up may be determining appropriate strategies to explain when the set of discrete strategy modes is large. Additionally, future work may try extending to physical human-robot collaborations, where strategy alignment has effects on the *safety* of the interaction.

## ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE1745016 and DGE2140739 and additionally under Grant No. 2144489. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

## REFERENCES

- [1] X. Zhu, A. Singla, S. Zilles, and A. N. Rafferty, “An overview of machine teaching,” *arXiv preprint arXiv:1801.05927*, 2018.

- [2] X. Zhu, "Machine teaching: An inverse problem to machine learning and an approach toward optimal education," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 29, no. 1, 2015.
- [3] A. Shinohara and S. Miyano, "Teachability in computational learning," *New Generation Computing*, vol. 8, pp. 337–347, 1991.
- [4] S. H. Huang, D. Held, P. Abbeel, and A. D. Dragan, "Enabling robots to communicate their objectives," *Autonomous Robots*, vol. 43, pp. 309–326, 2019.
- [5] M. S. Lee, H. Admoni, and R. Simmons, "Machine teaching for human inverse reinforcement learning," *Frontiers in Robotics and AI*, vol. 8, p. 693050, 2021.
- [6] A. Y. Ng, S. Russell *et al.*, "Algorithms for inverse reinforcement learning," in *Icml*, vol. 1, 2000, p. 2.
- [7] S. H. Huang, K. Bhatia, P. Abbeel, and A. D. Dragan, "Establishing appropriate trust via critical states," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 3929–3936.
- [8] W. Guo, X. Wu, U. Khan, and X. Xing, "Edge: Explaining deep reinforcement learning policies," *Advances in Neural Information Processing Systems*, vol. 34, pp. 12 222–12 236, 2021.
- [9] S. Blakeman and D. Mareschal, "Generating explanations from deep reinforcement learning using episodic memory," *arXiv preprint arXiv:2205.08926*, 2022.
- [10] S. Sreedharan, S. Srivastava, and S. Kambhampati, "Tldr: Policy summarization for factored ssp problems using temporal abstractions," in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 30, 2020, pp. 272–280.
- [11] A. Ajoudani, A. M. Zanchettin, S. Ivaldi, A. Albu-Schäffer, K. Kose, and O. Khatib, "Progress and prospects of the human–robot collaboration," *Autonomous Robots*, vol. 42, pp. 957–975, 2018.
- [12] D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan, and C. J. Tomlin, "Efficient iterative linear-quadratic approximations for non-linear multi-player general-sum differential games," in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 1475–1481.
- [13] S. L. Cleac'h, M. Schwager, and Z. Manchester, "Algames: A fast solver for constrained dynamic games," *arXiv preprint arXiv:1910.09713*, 2019.
- [14] C. Wang, C. Pérez-D'Arpino, D. Xu, L. Fei-Fei, K. Liu, and S. Savarese, "Co-gail: Learning diverse strategies for human-robot collaboration," in *Conference on Robot Learning*. PMLR, 2022, pp. 1279–1290.
- [15] L. Peters, D. Fridovich-Keil, C. J. Tomlin, and Z. N. Sunberg, "Inference-based strategy alignment for general-sum differential games," *arXiv preprint arXiv:2002.04354*, 2020.
- [16] J. F. Fisac, A. Bajcsy, S. L. Herbert, D. Fridovich-Keil, S. Wang, C. J. Tomlin, and A. D. Dragan, "Probabilistically safe robot planning with confidence-based human predictions," *arXiv preprint arXiv:1806.00109*, 2018.
- [17] D. Sadigh, S. S. Sastry, S. A. Seshia, and A. Dragan, "Information gathering actions over human internal state," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 66–73.
- [18] 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, pp. 1–9.
- [19] A. Xie, D. Losey, R. Tolsma, C. Finn, and D. Sadigh, "Learning latent representations to influence multi-agent interaction," in *Conference on robot learning*. PMLR, 2021, pp. 575–588.
- [20] W. Z. Wang, A. Shih, A. Xie, and D. Sadigh, "Influencing towards stable multi-agent interactions," in *Conference on robot learning*. PMLR, 2022, pp. 1132–1143.
- [21] H. M. Le, Y. Yue, P. Carr, and P. Lucey, "Coordinated multi-agent imitation learning," in *International Conference on Machine Learning*. PMLR, 2017, pp. 1995–2003.
- [22] J. Wright and K. Leyton-Brown, "Beyond equilibrium: Predicting human behavior in normal-form games," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 24, no. 1, 2010, pp. 901–907.
- [23] F. Bornhorst, A. Ichino, O. Kirchkamp, K. H. Schlag, and E. Winter, "How do people play a repeated trust game?: Experimental evidence," *None*, 2004.
- [24] L. Bruttel and U. Kamecke, "Infinity in the lab. how do people play repeated games?" *Theory and Decision*, vol. 72, pp. 205–219, 2012.
- [25] S. Gächter, "Behavioral game theory," *Blackwell handbook of judgment and decision making*, pp. 485–503, 2004.
- [26] P. J. Rousseeuw, "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis," *Journal of computational and applied mathematics*, vol. 20, pp. 53–65, 1987.
- [27] L. Peters and Z. N. Sunberg, "ilqgames. jl: Rapidly designing and solving differential games in julia," *arXiv preprint arXiv:2002.10185*, 2020.
- [28] G. Hoffman, "Evaluating fluency in human–robot collaboration," *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 3, pp. 209–218, 2019.
- [29] X. Puig, K. Ra, M. Boben, J. Li, T. Wang, S. Fidler, and A. Torralba, "Virtualhome: Simulating household activities via programs," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 8494–8502.
- [30] X. Puig, E. Undersander, A. Szot, M. D. Cote, T.-Y. Yang, R. Partsey, R. Desai, A. W. Clegg, M. Hlavac, S. Y. Min *et al.*, "Habitat 3.0: A co-habitat for humans, avatars and robots," *arXiv preprint arXiv:2310.13724*, 2023.