



# GRAIL: Goal Recognition Alignment through Imitation Learning

AAAI Track

Osher Elhadad  
Department of Computer Science, Bar  
Ilan University  
Ramat Gan, Israel  
oshereli2@gmail.com

Felipe Meneguzzi   
Department of Computer Science,  
Aberdeen University  
Aberdeen, Scotland, UK  
felipe.meneguzzi@abdn.ac.uk

Reuth Mirsky   
Department of Computer Science,  
Tufts University  
Medford, MA, USA  
reuth.mirsky@tufts.edu



## ABSTRACT

Aligning AI systems with human intentions requires understanding an agent’s goals from its behavior. Existing goal recognition methods typically rely on an *approximately optimal* goal-oriented policy representation, which may differ from the actor’s true behavior and hinder the accurate recognition of their goal. To address this gap, this paper introduces Goal Recognition Alignment through Imitation Learning (GRAIL), which leverages imitation learning and inverse reinforcement learning to learn one goal-directed policy for each candidate goal directly from (potentially suboptimal) demonstration trajectories. By scoring an observed partial trajectory with each learned goal-directed policy in a single forward pass, GRAIL retains the one-shot inference capability of classical goal recognition while leveraging learned policies that can capture suboptimal and systematically biased behavior. Empirical evaluations show that GRAIL outperforms standard reinforcement learning-based GR techniques in recognizing both suboptimal and biased behaviors in controlled, closed-set goal environments.

## KEYWORDS

Goal Recognition; Agent Modeling; Imitation Learning; Goal Alignment; Reinforcement Learning


### ACM Reference Format:

Osher Elhadad, Felipe Meneguzzi , and Reuth Mirsky . 2026. GRAIL: Goal Recognition Alignment through Imitation Learning: AAAI Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/ZCP56180>

## 1 INTRODUCTION

Goal Recognition (GR) aims to infer an agent’s underlying objectives from its observable behavior [39, 55]. Accurate and timely goal inference is critical for interactive and cooperative systems. Such systems include human-robot interaction where aligning with human intent ensures safe and supportive assistance [32, 49, 56] as well as multiagent scenarios, where predicting others’ goals facilitates adaptive coordination and strategic reasoning [10, 28, 46, 54].

By recognizing goals in real-time, AI systems can anticipate needs [52], proactively assist [42], and model discrepancies [53], or misalignments [11, 33, 35, 60]. However, existing GR approaches



**Figure 1: A navigation setup illustrating how ignoring systematic human biases like preference for park routes can lead to reduced inference accuracy in standard GR systems.**

often rest on a restrictive assumption that observed agents act (approximately) optimally with respect to a predefined objective [34]. Classical plan recognition methods [47], inverse planning [8, 9], and more recent reinforcement learning (RL)-based goal recognizers [3] build models that expect rational or reward-maximizing behavior. This leads to brittle performance when applied to agents whose actions are noisy, constrained, or systematically biased. Human behavior, for example, frequently deviates from optimality in meaningful ways, reflecting cognitive limitations, preferences, social norms, or learned heuristics. As Lindner and El-Assady [31] notes, and Masters and Sardina [33] illustrates, Boltzmann rationality models often cannot capture such behaviors. Thus, accommodating subrational or irrational actors requires new modeling assumptions, and existing GR systems often fail to align with real-world agents. Most GR methods also suffer from significant computational costs. Specifically, methods that rely on RL to train a policy per candidate goal, often requiring millions of interactions with the environment [3]. Similarly, methods that rely on planning algorithms may require planner invocations at inference time [48] for each observed trajectory, creating latency bottlenecks in interactive or time-sensitive applications.

Human navigation often deviates from shortest-path optimality, reflecting systematic preferences such as minimizing turns or following familiar landmarks [13]. Consider the scenario in Figure 1 where a person leaves home with two potential destinations: a grocery store to the northeast and their workplace to the southeast. Although the two main routes to the grocery, going east then north or north then east, take the same time, the person consistently prefers the path that passes through a scenic park located to the north. When observing this person walking directly east, an optimality-based goal recognizer would assign similar likelihoods to the grocery and workplace, since both are consistently equidistant throughout the observation. However, this overlooks



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

the person’s habit of going to the grocery via the park. By learning goal-directed behavior from previous walks, the GR system can capture that this person rarely begins by walking east when going to the grocery. Therefore, it correctly infers that the workplace is a more likely goal, even though both destinations are equally optimal under standard assumptions.

To correctly recognize goals in such cases, and to address the additional challenges mentioned above, we develop **GRAIL (Goal Recognition Alignment through Imitation Learning)**. GRAIL formulates GR as a collection of imitation learning (IL) problems, one for each candidate goal  $g$ . Unlike conventional GR methods that treat deviations from optimality as noise, our IL-based method handles the realities of human and agent behavior, that can be (1) suboptimal or (2) systematically biased-optimal. This allows GR systems to align with the actual behavior of agents that are not idealized optimal agents. This novel GR formulation not only enhances robustness to real-world agent behavior but also enables fast, planner-free GR at test time, making GRAIL suitable for interactive scenarios where low inference-time latency is critical. We show that a variety of IL algorithms can be repurposed for one-shot goal inference without any planner or reward modeling at test time.

Experiments across MiniGrid and PandaReach show that our imitation-based goal recognizers achieve substantial gains in recognition quality and practical efficiency. In systematically biased optimal MiniGrid tasks, all GRAIL variants achieve near-perfect GR, whereas the Q-learning baseline remains close to random. Under suboptimal behavior, GRAIL delivers substantial gains, particularly in the harder multi-goal settings. In PandaReach, our method substantially improves performance over the actor-critic baseline in noisy optimal conditions, while matching or only slightly underperforming compared to RL-based GR under clean optimal demonstrations. These results indicate that GRAIL scales gracefully with noise, preference-driven trajectories, and model mismatch in controlled, closed-set goal environments, while preserving lightweight, forward-pass-only inference suitable for real-time use. The full version appears in [20].

## 2 BACKGROUND

GRAIL employs two key techniques to solve *goal recognition* problems *Imitation Learning (IL)* and *Inverse Reinforcement Learning (IRL)*. To better place the contribution with respect to each of these domains, we refer the reader to Table 1. Goal Recognition is a *problem setting* whereby one can infer the goal of an observed agent given a partial observation sequence and a set of candidate goals. This is analogous, goal classification at inference time. In contrast, IL and IRL are *learning paradigms* [27, 41]. IL learns a policy from demonstration trajectories to mimic expert behavior, while IRL recovers a reward function (and often a derived policy) that explains demonstrated behavior under known dynamics. Like IL/IRL, GR may use demonstrations at training time, however, at inference time GR takes a partial trajectory plus candidate goals as input and outputs a goal label, whereas IL/IRL produce policies or rewards. Next, we provide the required background and definitions from each of these paradigms.

**DEFINITION 1 (GOAL RECOGNITION PROBLEM).** *A goal recognition problem is a tuple  $(S, A, G, O)$ , where:*

**Table 1: Role summary of GR, IL, and IRL. GR is a *problem*; IL/IRL are *learning paradigms* used here to build goal-directed policies for solving GR.**

	Object of study	Typical output	Typical use	Role in this paper
GR	Observed behavior & candidate goals	Goal label or posterior	Recognize intentions	Core task we solve
IL	Expert demos	Policy $\pi$	Learn behavior from examples	Learn goal-directed policy per goal
IRL	Expert demos & known dynamics	Reward $R$ (+ policy)	Recover preferences for planning/RL	Via AIRL, obtain goal-directed policy head per goal; no reward-based planning at test

- $S$  is the state space;
- $A$  is the action space;
- $G = \{g_1, \dots, g_K\}$  is a finite set of candidate goals; and
- $O = (o_1, \dots, o_T)$  is the observation sequence.

Observations can be of actions, states, or both. When observations are state-action pairs,  $O = (s_0, a_0, \dots, s_T)$  where  $T$  is the number of observations in the sequence. Solving a GR problem requires selecting  $g \in G$  that maximizes the likelihood of  $O$ . Definition 1 adapts earlier formalisations [3, 36] of GR. Most classical approaches treat this as an inverse-planning problem [3, 48]. They assume an optimal planner or a goal-directed RL policy and score each goal by how well its planner or policy could have produced the observations.

What follows positions this paper within the broader context of sequential decision-making. We start by briefly reviewing RL, which formulates decision-making as trial-and-error interactions with an environment to maximize cumulative reward. RL often makes an assumption that the environment dynamics is a *Markov decision process (MDP)*. An MDP is a tuple  $(S, A, T, \mathcal{R}, \gamma)$ , where:  $S$  is the state space;  $A$  is the action space;  $T(s' | s, a)$  defines the transition dynamics;  $\mathcal{R}(s, a)$  is the reward function; and  $\gamma \in [0, 1)$  is the discount factor. A solution to an MDP problem is a policy  $\pi(a | s)$  that seeks to maximize expected discounted return:  $\mathbb{E}_\pi \left[ \sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_t) \right]$ . While knowledge of the MDP allows planning algorithms to solve an MDP, RL methods bypass the need for knowledge of  $T$  and  $\mathcal{R}$ , and instead solves the MDP by sampling from the environment directly, or from a simulator. RL is well-suited for goal-directed behavior generation, but it assumes access to a sampling model of the reward function and often requires extensive interaction with the environment. Both assumptions are problematic for GR settings, where observed behavior is passive and the observed agent’s rewards are unknown.

*Inverse Reinforcement Learning (IRL)* aims to recover a reward function that explains observed agent behavior [6]. IRL algorithms use demonstrations to infer an agent’s preferences through a model of their reward, typically requiring full trajectories and known dynamics. In our work, we leverage IRL (specifically, AIRL) solely to obtain goal-directed policy heads.

*Imitation Learning (IL)* comprises a family of techniques where an agent learns to perform tasks by mimicking expert behavior, typically using demonstrations of state-action trajectories. Unlike IRL, which seeks to infer the underlying reward function, IL focuses on directly learning a policy that reproduces expert-like behavior, often without explicitly modeling or recovering the reward structure driving the expert’s decisions. Given  $N$  expert demonstrations  $D = \{\tau_i\}_{i=1}^N$ , where each trajectory  $\tau_i$  is a sequence of states and actions  $\tau_i = (s_0^i, a_0^i, \dots)$ , IL aims to learn a policy  $\pi(a | s)$  that minimizes divergence from the expert behavior distribution.

### 3 GOAL RECOGNITION ALIGNMENT THROUGH IMITATION LEARNING

We begin by formally extending Definition 1 to the problem of GR from demonstrations in Definition 2.

**DEFINITION 2 (GRD).** A *Goal Recognition from Demonstration (GRD) task* is a tuple  $(S, A, G, \mathcal{D}, \tau_{\text{obs}})$  where:

- $S$  is the set of all possible states;
- $A$  is the set of all actions;
- $G = \{g_1, \dots, g_K\}$  is a finite set of candidate goals;
- $\mathcal{D} = \{\mathcal{D}_g\}_{g \in G}$  is a set of demonstrations per goal.
- $\tau_{\text{obs}} = (s_0, a_0, \dots, s_T)$  is a new, partially observed trajectory.

Each  $\mathcal{D}_g = \{\tau_g^i\}_{i=1}^{N_g}$  is a set of trajectories  $\tau_g^i = (s_0^i, a_0^i, \dots, s_{T_i}^i)$  generated by an agent pursuing goal  $g$ . This formulation does not require all trajectories for a goal  $\{\tau_g^i\}_{i=1}^{N_g}$  to have a single starting point, nor a single end point, they should merely reflect the expected behavior of the goal-directed agent. The objective of a GRD task is to select the goal  $\hat{g} \in G$  that best explains  $\tau_{\text{obs}}$ .

The key insight behind GRAIL is that whether behavior is optimal, systematically biased, or noisy, it still leaves a distinctive *distributional fingerprint* on the trajectories. IL methods are ideally suited to capture these complex fingerprints. Once we have trained each  $\pi_g$ , we perform one-shot inference by evaluating a similarity score sim:

$$\text{sim}_g(\tau_{\text{obs}}) = \ell(\tau_{\text{obs}}; \pi_g),$$

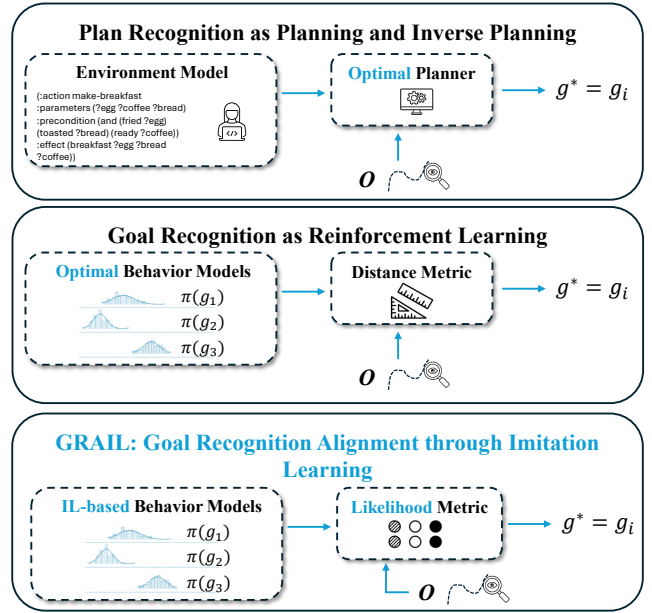
where  $\ell$  is a goal-agnostic scoring function. We can then infer  $\hat{g}$  by choosing the goal that maximizes this score.

$$\hat{g} = \arg \max_{g \in G} \text{sim}_g(\tau_{\text{obs}}).$$

GRAIL alternates between two stages. First, using *offline policy learning*, it trains a separate IL policy for each potential goal. Then, in an *online, one-shot inference* phase, it scores a new observation sequence against every trained policy to pick the most likely goal.

#### 3.1 Offline Policy Learning from Demonstrations

For each goal  $g \in G$ , we collect a dataset of demonstrations  $\mathcal{D}_g$  of an agent that aims to achieve  $g$ , and invoke our chosen IL routine to produce a goal-directed policy  $\pi_g(a | s)$ . By the end of this phase we possess a “bank” of  $K$  policies, each of which encodes the characteristic behavioral fingerprint of reaching its assigned goal, including any systematic biases or suboptimal quirks present in the demonstrations. In this paper, we explore three common IL algorithms at the backbone for GRAIL.



**Figure 2: A comparison between model-based GR approaches (top), the RL-based GR approaches (middle), and our proposed framework (bottom).**

*Behavioral Cloning (BC)* [7] treats imitation as supervised learning over state-action pairs. Training for this goal-specific policy maximizes the likelihood of expert actions.

*Generative Adversarial Imitation Learning (GAIL)* [26] Uses adversarial training to match the occupancy distributions of expert and policy rollouts. A discriminator distinguishes expert versus learner trajectories.

*Adversarial IRL (AIRL)* [22] Learns both a reward function  $\mathcal{R}_g(s, a)$  and a policy  $\pi_g$  such that the induced behavior explains expert data. The reward serves as a latent rationalization of actions.

#### 3.2 Online One-Shot Inference

When a new, partially observed trajectory  $\tau_{\text{obs}} = (s_0, a_0, \dots, s_T)$  arrives, we must quickly determine which  $\pi_g$  best explains it. Rather than invoking planning or rolling out policies in the environment, we compute for each goal a single similarity score  $\text{sim}_g(\tau_{\text{obs}}) = \ell(\tau_{\text{obs}}; \pi_g)$ , and then select  $\hat{g} = \arg \max_{g \in G} \text{sim}_g(\tau_{\text{obs}})$ .

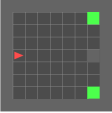
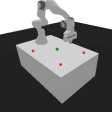
To implement a lightweight, one-shot scoring method uniformly applicable across all IL algorithms from our experiments, we use a similarity metric using *Negative Average MSE*.

$$\ell_{\text{MSE}}(\tau_{\text{obs}}; \pi_g) = -\frac{1}{T} \sum_{t=0}^{T-1} \|\pi_g(s_t) - a_t\|^2,$$

$$\text{sim}_g^{\text{MSE}}(\tau_{\text{obs}}) = \ell_{\text{MSE}}(\tau_{\text{obs}}; \pi_g).$$

This metric extends naturally to GAIL and AIRL by applying it to their policy heads  $\pi_g$ . Using this score, GRAIL performs fast inference over an entire trajectory in a single batch pass without requiring planning or RL rollouts.

Table 2: Comparison of domains, visualizations, and their specific characteristics.

Figure	Domain	Details
	MiniGrid	<p><i>Environments:</i> MiniGrid-9x9 environment with an obstacle at 7x4.</p> <p><i>Motivation:</i> Discrete navigation; comparison with GR baselines; interpretable and controllable (e.g., easy for designing suboptimal behavior).</p> <p><i>States:</i> Position <math>(x, y)</math> + direction (N,S,E,W) of agent (Discrete).</p> <p><i>Actions:</i> Turn, move forward, stay in place (Discrete).</p> <p><i>Reward:</i> <math>1 - 0.9 * (step_{count} / max_{steps})</math> for success, 0 otherwise (Sparse).</p>
	Panda-Gym	<p><i>Environments:</i> PandaReach environment.</p> <p><i>Motivation:</i> Realistic 3D robotic control scenario.</p> <p><i>States:</i> Robotic arm positions, velocities (Continuous).</p> <p><i>Actions:</i> Joint torques for robotic arm (Continuous).</p> <p><i>Reward:</i> Gradual reward based on distance to goal (Dense).</p>

### 3.3 Training versus Inference Regimes

BC-GRAIL trains one goal-conditioned policy per goal purely from offline demonstration trajectories, without interacting with the environment. GAIL-GRAIL and AIRL-GRAIL follow standard adversarial imitation learning. They alternate during training between updating a discriminator and a policy using environment rollouts, so environment interaction occurs at training time on the same order as other RL/IL methods. At inference time, all GRAIL variants perform GR without any planner calls or environment interaction. Thus, they only evaluate the learned policies on the observed partial trajectory and compute the scoring function.

## 4 EXPERIMENTAL DESIGN

*Environments.* We evaluate all methods on two domains with distinct characteristics: a discrete navigation environment - MiniGrid [16] and a continuous robotic manipulation environment - PandaReach [23]. Table 2 summarizes these differences appears. In the discrete MiniGrid task, the agent operates in a  $9 \times 9$  grid, starting at position (1, 4) facing right. The environment contains a fixed obstacle at position (7, 4), and up to six candidate goals are located at (7, 1), (7, 7), (5, 1), (5, 7), (7, 3), (7, 5). In the continuous PandaReach manipulation task, we use a simulation of a robotic arm to train policies to reach 3D targets, with up to four candidate goals.

In both environments, we manually defined the initial conditions (e.g., start position and orientation) and selected specific sets of candidate goals to construct controlled settings for GR (See Appendix A [20]). This manual configuration allows us to systematically generate expert behaviors under different assumptions: (i) *optimal* behavior, where the agent takes the shortest valid path to the goal; (ii) *systematically-biased* behavior, where the agent follows a structurally preferred route (e.g., favoring straight-line motion or a specific route like in the example in Figure 1); and (iii) *suboptimal* behavior, where the agent’s trajectory deviates due to noise or heuristic-driven choices. These settings allow us to evaluate how well GR methods generalize under realistic, non-idealized behavior. We elaborate further on the generation of these behavior regimes in the next Section.

To further support stable IL across all experiments and algorithms, we standardized the episode structure by using fixed-horizon rollouts instead of early termination upon goal reaching. Specifically, GAIL and AIRL require the input demonstrations to have a fixed-horizon. We seeded all experiments using the random number generator provided by Gymnasium [12].

*Baselines.* We compare GRAIL against two recent GR methods that frame inference as a matching problem against goal-directed policies. First, GRAQL [3] learns a separate Q-function  $Q_g(s, a)$  for each goal via Q-learning; at test time it treats the observed trajectory  $\tau_{obs}$  as inducing a “pseudo-policy” over the taken actions and scores each goal by the negative sum of Kullback-Leibler (KL)-divergences between this pseudo-policy and the softmax policy derived from  $Q_g$ . Second, DRACO [40] trains each goal-directed policy  $\pi_g$  via PPO (Proximal Policy Optimization) [50] on the true reward and then compares the empirical distribution of  $(s, a)$  pairs in  $\tau_{obs}$  to the occupancy measure of  $\pi_g$  using the 1-Wasserstein distance. Both baselines require per-goal RL training. All hyperparameters, training details, information on computational resources and efficiency can be found in Appendix A [20]).

*Behavioral Conditions.* To probe robustness, we collect two types of demonstration data. Under *optimal* conditions, MiniGrid trajectories are handcrafted to follow the shortest path to each goal, while PandaReach trajectories are produced by a converged SAC (Soft Actor-Critic) [25] agent. Under *suboptimal* conditions, we introduce systematic deviations: in MiniGrid we insert 180° turns (left and right) before each action without changing the overall path. In PandaReach, we use two types of action noise - Gaussian noise at levels of 0.1 and 0.3, and uniform noise at levels of 0.1 and 0.3 - and evaluated all experiments under low observability (2% of the trajectory). For MiniGrid, we generate 10 trajectories per goal (using 7 for IL training and 3 for testing its GR performance); for PandaReach we collect 200 per goal (150/50 train/test split). We evaluate GR performance, framed as a classification task, using standard metrics: Accuracy, Recall, Precision, and  $F_1$ -score.

Table 3:  $F_1$ -score (mean  $\pm$  std over 10 GR tasks) comparison for Suboptimal trajectories with 2, 4, 6 goals under 10%, 20%, and 30% observability. Bold indicates best performing method in each condition.

Goals	Obs	GRAQL	BC-GRAIL	GAIL-GRAIL	AIRL-GRAIL
2	10%	<b>0.8333 <math>\pm</math> 0.3416</b>	0.8 $\pm$ 0.3055	0.5674 $\pm$ 0.0735	0.5333 $\pm$ 0.0762
2	20%	<b>0.9 <math>\pm</math> 0.2108</b>	<b>0.9 <math>\pm</math> 0.2108</b>	0.7886 $\pm$ 0.0494	0.7683 $\pm$ 0.0682
2	30%	0.9333 $\pm$ 0.2	<b>1.0 <math>\pm</math> 0.0</b>	0.8652 $\pm$ 0.0306	0.8418 $\pm$ 0.0563
4	10%	0.4583 $\pm$ 0.2401	<b>0.5875 <math>\pm</math> 0.2286</b>	0.5107 $\pm$ 0.0835	0.4784 $\pm$ 0.0733
4	20%	0.6833 $\pm$ 0.2236	<b>0.725 <math>\pm</math> 0.2092</b>	0.7502 $\pm$ 0.0517	0.7223 $\pm$ 0.0662
4	30%	0.8125 $\pm$ 0.2734	<b>1.0 <math>\pm</math> 0.0</b>	0.8328 $\pm$ 0.0349	0.8012 $\pm$ 0.0589
6	10%	0.3067 $\pm$ 0.1374	0.4222 $\pm$ 0.1089	<b>0.4726 <math>\pm</math> 0.0768</b>	0.4450 $\pm$ 0.0652
6	20%	0.5111 $\pm$ 0.1764	0.6333 $\pm$ 0.1528	<b>0.7101 <math>\pm</math> 0.0531</b>	0.6833 $\pm$ 0.0617
6	30%	0.5722 $\pm$ 0.2636	<b>0.8444 <math>\pm</math> 0.1736</b>	0.7934 $\pm$ 0.0372	0.7651 $\pm$ 0.0571




Figure 3: Minigrad (9  $\times$  9 grid, 7  $\times$  7 free states) Q-learning visit counts per state (position + orientation) during training, shown as heatmaps with overlaid counts. Yellow indicates high visit frequency. (left) Training visits for goal at (7, 1). (right) Training visits for goal at (7, 7).

## 5 RESULTS

We tested GRAIL under various conditions to evaluate its performance, given systematically biased behavior (Section 5.1), suboptimal behavior (Section 5.2), and optimal behavior (Section 5.3). We further evaluate GRAIL’s efficiency in terms of training requirements (Section 5.4) and provide an ablation test for GRAIL’s inference metric (Section 5.5).

We averaged each reported result over 10 random seeds, with different initial environment states and policy initializations for each seed. We report mean  $\pm$  standard deviation for all metrics (accuracy, precision, recall,  $F_1$ -score).

### 5.1 Systematically Biased Optimal Behavior

In many scenarios, agents pursue goals along nominally optimal paths but exhibit consistent preferences for certain trajectories over equally short alternatives. We term this *systematically biased optimal behavior*, where standard GR methods that assume uniform optimality can fail to distinguish goals when the biased preference obscures the distinguishing segments of the trajectory. This is the situation in the use-case in Figure 1, which we translated to an experimental setup in Minigrad: the two goals, (7, 1) and (7, 7), are

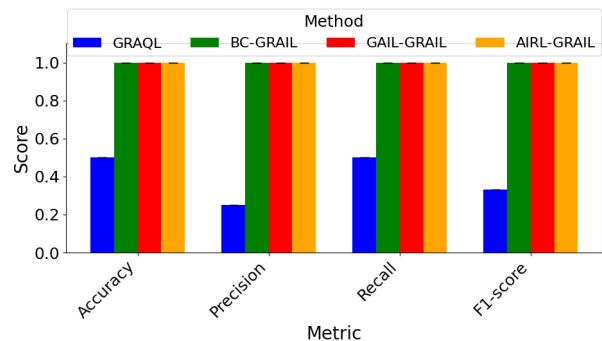


Figure 4:  $F_1$ -score comparison on Minigrad with systematically biased optimal demonstrations and 2 goals under 20-40% observability across 10 runs. GRAIL variants (green/red/yellow) achieve perfect inference, compared to GRAQL (blue).

placed in a 9  $\times$  9 obstacle world. Each goal admits multiple shortest-path solutions, yet Q-learning training visit-counts (Fig. 3) hint that the optimal routes (“east then north” for (7, 1), “east then south” for (7, 7)) dominate. We synthetically generated biased demonstrations in which the agent chooses its preferred optimal variant to be “north then east” for (7, 1), and “east then south” path for (7, 7).

We collect ten biased trajectories per goal and compared GRAQL against the GRAIL variants (BC-GRAIL, GAIL-GRAIL, AIRL-GRAIL) learned goal-directed policies. Each method was evaluated on tasks in which only the first 20-40% of each trajectory was observed, with performance quantified by Accuracy, Precision, Recall, and  $F_1$ -score.

As shown in Fig. 4, all three GRAIL methods recover the true goal with 100%  $F_1$  across the entire 20-40% observation range. In contrast, GRAQL’s  $F_1$  hovers around 0.3-0.5, effectively at chance, because the initial biased path segments (for (7  $\times$  7) path) coincide for both goals and the Q-functions assign them equal value. These results confirm that GRAIL, by directly modeling the biased distribution, can disambiguate goals even when RL-based inference cannot.

Table 4:  $F_1$ -score (mean  $\pm$  std over 10 GR tasks) comparison for Optimal trajectories with 2, 4, 6 goals under 10%, 20%, and 30% observability. Bold indicates best performing method in each condition.

Goals	Obs	GRAQL	BC-GRAIL	GAIL-GRAIL	AIRL-GRAIL
2	10%	<b>0.8 <math>\pm</math> 0.3055</b>	0.7333 $\pm$ 0.3266	0.7055 $\pm$ 0.0863	0.6604 $\pm$ 0.1092
2	20%	<b>0.9333 <math>\pm</math> 0.1491</b>	0.8667 $\pm$ 0.2309	0.9194 $\pm$ 0.0407	0.9047 $\pm$ 0.0793
2	30%	<b>1.0 <math>\pm</math> 0.0</b>	0.9333 $\pm$ 0.2	0.9649 $\pm$ 0.0230	0.9447 $\pm$ 0.0573
4	10%	<b>0.7167 <math>\pm</math> 0.2825</b>	0.6333 $\pm$ 0.2485	0.6412 $\pm$ 0.0973	0.6032 $\pm$ 0.0851
4	20%	<b>0.9 <math>\pm</math> 0.1833</b>	0.8167 $\pm$ 0.2191	0.8903 $\pm$ 0.0427	0.8453 $\pm$ 0.0701
4	30%	<b>0.9667 <math>\pm</math> 0.1</b>	0.9333 $\pm$ 0.1333	0.9482 $\pm$ 0.0261	0.9331 $\pm$ 0.0610
6	10%	0.5078 $\pm$ 0.1584	0.3733 $\pm$ 0.1839	<b>0.5772 <math>\pm</math> 0.0934</b>	0.5330 $\pm$ 0.0805
6	20%	0.8111 $\pm$ 0.2033	0.7111 $\pm$ 0.2129	<b>0.8462 <math>\pm</math> 0.0496</b>	0.8147 $\pm$ 0.0693
6	30%	0.8694 $\pm$ 0.1417	0.7889 $\pm$ 0.1262	<b>0.9188 <math>\pm</math> 0.0291</b>	0.8904 $\pm$ 0.0654




Figure 5:  $F_1$ -score comparison on PandaReach with optimal and noisy demonstrations (uniform and gaussian) and 4 goals under 2% observability across 10 runs.

## 5.2 Suboptimal Behavior

In real-world settings, agents often deviate systematically from shortest-path optima, whether through exploration, sensor noise, or habitual biases. We term this regime *suboptimal behavior*. We begin this evaluation with MiniGrid by varying task complexity. We adjust the number of candidate goals and the visibility of the agent’s trajectory to evaluate how each method generalizes as the problem becomes harder. In this domain, we evaluate GR with 2, 4, and 6 candidate goals across three partial observability levels (10%, 20%, 30% of the trajectory visible). Table 3 reports  $F_1$ -scores (mean  $\pm$  std) for the Q-learning-based GRAQL baseline and for BC-, GAIL-, and AIRL-GRAIL. GRAQL’s performance degrades as the number of goals increases: with 6 goals at 10% visibility, it achieves only  $0.31 \pm 0.14$ , while BC-GRAIL reaches  $0.42 \pm 0.11$ , GAIL-GRAIL  $0.47 \pm 0.08$ , and AIRL-GRAIL  $0.45 \pm 0.07$ . Remarkably, BC-GRAIL holds a lead in most of the scenarios, and achieves perfect recognition ( $1.00 \pm 0.00$ ) for 4 goals at 30% observability, highlighting the effectiveness of supervised cloning under sub-optimality.

In PandaReach, we compare GRAIL methods against the DRACO baseline in a four-goal continuous-control task under 2% observability and with Gaussian and uniform action noise at 10% and 30%.

Figure 5 illustrates that BC-GRAIL outperforms all other methods across noise levels, while GAIL-GRAIL shows particular strength as noise increases. AIRL-GRAIL and DRACO lag behind in these noisy settings. These results confirm that learning from suboptimal demonstrations enables BC-based policies to internalize systematic deviations, whereas approaches predicated on near-optimal behavior treat such deviations as noise.


Together, these results demonstrate that GRAQL and DRACO are advantageous only when the demonstrations are perfect, or when the goal set is very small. In all other scenarios: moderate observability, larger goal sets, or suboptimal/expert-imperfect data, GRAIL outperforms these methods. BC-GRAIL’s surprisingly strong performance under suboptimal trajectories further highlights the value of BC when behavior noise is structured.

## 5.3 Optimal Behavior

When demonstrations are noise-free shortest-path trajectories, we refer to this as the *optimal behavior* regime. Here, inference methods that assume idealized rationality should excel, given sufficient observation.

In the MiniGrid domain, Table 4 highlights several key patterns in  $F_1$ -score performance under optimal demonstrations. When only 10% of each trajectory is observable, GRAQL retains a modest advantage for small goal sets (2 and 4 goals), reflecting its reliance on precisely learned Q-values for very sparse data. As we increase observability to 20%, the gap narrows: GRAIL’s imitation-based methods rapidly catch up. Notably, for the six-goal scenario GRAIL ( $0.58 \pm 0.09$ ) and AIRL-GRAIL ( $0.53 \pm 0.08$ ) surpass GRAQL ( $0.51 \pm 0.16$ ), indicating that adversarial IL’s enforced coverage of diverse expert behaviors yields policies whose early-trajectory predictions generalize better across larger goal sets. By 30%, GAIL-GRAIL achieves  $0.92 \pm 0.03$  and AIRL-GRAIL  $0.89 \pm 0.07$  for six goals, whereas GRAQL lags at  $0.87 \pm 0.14$ . These results underscore that, while inference based on Q-learning excels in simple tasks, inference based on imitation-learning, especially those trained adversarially, scale more gracefully in data availability and goal-complexity.

In PandaReach, all GRAIL variants surpassed DRACO when using SAC-generated optimal trajectories. At just 2% observability (Figure 5), DRACO’s performance falls well below that of each



**Figure 6: GR inference time comparison on PandaReach with optimal and noisy demonstrations (uniform and gaussian) and 4 goals across 10 runs.**

GRAIL method. This shows that even under ideal conditions, imitation learning can match or outperform goal recognition based on actor-critic methods when only limited data is available.

Results in this optimal demonstration regime confirms that while RL-based GR methods can yield strong performance under extreme data scarcity and small goal sets, GRAIL’s policies scale more gracefully as observation coverage and problem complexity grow.

#### 5.4 Efficiency and Simulation Complexity

Beyond recognition accuracy, GRAIL delivers substantial efficiency gains both in training and inference. In the MiniGrid domain, BC-GRAIL requires only about 5s per goal, GAIL-GRAIL about 60 s, and AIRL-GRAIL about 240s, compared to roughly 300s for the GRAQL baseline. In PandaReach, BC, GAIL, and AIRL each train in under half the time needed by the PPO-based DRACO recognizer. BC-GRAIL achieves near-optimal recognition with just 7 demonstrations per goal in MiniGrid (and 150 in PandaReach), whereas GRAQL must visit on the order of  $2 \times 10^5$  state-action pairs to converge its Q-functions (see the heatmap in Fig. 3). At inference time, all GRAIL methods require only a handful of forward passes and are dramatically faster than DRACO (Fig. 6).

**Scalability to larger goal sets.** The detailed experiments reported above use relatively small closed sets of candidate goals (up to 6 in MiniGrid and up to 4 in PandaReach), consistent with standard GR formulations that assume a finite set of candidate goals or hypotheses [3, 36, 40]. We leave a more systematic large or infinite- $|G|$  study for future work.

#### 5.5 Inference Metric Ablation

We adopt negative average MSE between the policy’s predicted actions and the observed actions as our default goal score because it is simple, numerically stable, and applies well to both discrete and continuous action spaces. To situate this choice, we compare it against the Wasserstein-based distance proposed in DRACO [40], which was shown to be a strong metric for policy-based GR in

both MiniGrid and Panda-Gym, especially under partial observability and noisy observations, and to perform competitively with their Z-score distance. Earlier GR-as-RL work (GRAQL) [3] introduced three discrete, Q-based distances—MaxUtil, KL-divergence, and Divergence Point—and found that KL and MaxUtil generally dominate DP across a wide range of observability levels. Therefore we compare GRAIL to GRAQL with KL-divergence

In contrast to value-based distances such as MaxUtil, KL, Divergence Point, and Z-score, both our per-step MSE and DRACO’s Wasserstein distance operate directly on the action distributions of goal-directed policies and apply uniformly in discrete and continuous action spaces. We therefore focus our comparison on Wasserstein as a representative policy-based metric: DRACO empirically compares the GRAQL-style distances (MaxUtil, KL, and Divergence Point) against Wasserstein and Z-score, and reports that the Wasserstein- and Z-score-based variants achieve superior GR accuracy in most evaluated settings [40].

We measure both inference latency and recognition accuracy across all GRAIL variants using both our proposed negative averaged-MSE scoring function, and Wasserstein distance. Figure 6 plots the per-trajectory scoring time for each method under both MSE and Wasserstein metrics. Across the board, MSE-based scoring reduces latency, and this speed-up accumulates when evaluating large goal sets. Figure 5 compares  $F_1$ -scores under both similarity metrics in PandaReach, across four observation noise regimes (optimal, Gaussian noise at 0.1 and 0.3, and uniform noise at 0.1 and 0.3). Under perfect demonstrations, BC-GRAIL performs equally well with either metric, while AIRL-GRAIL sees a modest gain when using the MSE-based scoring. As noise increases, BC-GRAIL consistently benefits from MSE, showing improved robustness, and AIRL-GRAIL likewise enjoys moderate gains. By comparison, GAIL-GRAIL retains a clear advantage under the Wasserstein metric across noisy settings.

## 6 RELATED WORK

**Planning-based GR.** The foundational Ramirez and Geffner [48] framework casts GR as computing goal likelihoods via domain-independent planning, assuming known domain models. Subsequent work extended this to landmark-based inference [44, 45], with a comprehensive survey by Meneguzzi and Pereira [36]. Later work extends these approaches to leverage learning algorithms over the PDDL [4, 5, 17], and other representations [1, 2, 24, 38, 43]. All these works assume fully rational agents. Recent work has extended this to account for agents that plan under a bounded rationality model [33, 59]. Notice that across all of these environments, the recognizer’s inference process remains centered on constructing what it deems most likely, rather than modeling the observed agent’s behavior based on patterns learned from prior interactions.

**RL-based GR.** Goal Recognition as Q-Learning (GRAQL) [3] learns per-goal Q-functions via Q-learning and scores observations by comparing them to the softmax policies derived from these value functions, removing the need for explicit domain models. More recent work [21, 40] trains deep policy networks using actor-critic methods and performs recognition via Wasserstein distance, extending to continuous control domains. Both in planning-based approaches and RL-based ones, there is an assumption that the

actor’s behavior in the environment is consistent with the planner/learner’s policy.

GRAIL differs from these approaches by learning per-goal policies via IL rather than via planning or pure RL on environment rewards, thus it can better capture biases and suboptimal patterns that emerge in the actor’s behavior we aim to recognize. It performs planner-free, one-shot scoring of partial trajectories using these learned policies, designed to account for suboptimal and systematically biased agent behavior. GRAIL handles both discrete and continuous environments, demonstrating that imitation-based policy learning provides a flexible, data-driven alternative to planning- and RL-based GR methods.

## 7 CONCLUSION

This paper introduces GRAIL (Goal Recognition Alignment through Imitation Learning), a class of methods for inferring an agent’s goal from partial observations of its behavior given a finite set of candidate goals. The key idea is to learn one goal-directed policy per candidate goal using imitation learning methods such as Behavioral Cloning (BC), Generative Adversarial Imitation Learning (GAIL), and Adversarial Inverse Reinforcement Learning (AIRL), instead of relying on planners or environment reward functions. At test time, GRAIL performs planner-free, one-shot goal recognition (GR) by scoring how well each goal-directed policy explains the observed partial trajectory, without any environment rollouts. This yields a unified way to apply imitation learning (IL), including AIRL as a special case, to GR tasks in both discrete (MiniGrid) and continuous (PandaReach) control domains.

We evaluate GRAIL in settings with optimal, noisy, and systematically biased behavior. Across MiniGrid and PandaReach benchmarks, GRAIL variants often match or outperform RL-based GR baselines such as GRAQL and DRACO in recognition accuracy and  $F_1$ -score, particularly at low observability. An ablation on scoring metrics compares negative mean squared error (MSE) against Wasserstein score and shows that MSE achieves a favorable trade-off between accuracy and inference time while remaining simple and numerically stable.

Beyond GR accuracy and run-time, GRAIL offers several system-level attributes that are critical for practitioners. It provides *flexibility* in allowing different IL backends to be swapped in depending on data availability and domain characteristics. For example using BC in supervision-rich domains or GAIL and AIRL when requiring broader generalization. Second, GRAIL is inherently *modular* since after the offline training of per-goal policies, online inference requires only a single batch of forward pass, with no additional environment interaction or supervision. Moreover, GRAIL achieves *robustness* by aligning directly with observed behaviors, rather than penalizing every deviation from an idealized optimal planner/policy. Finally, because inference uses only learned policies and a simple scoring rule, the framework is easy to integrate into larger multiagent systems that need fast GR aligned with observed behavior.

### Limitations and Future Work

Despite these strengths, GRAIL has several limitations that point to directions for future work.

First, GRAIL assumes a closed, known set of candidate goals. This is the standard formulation in much of the GR literature, where the task is to classify observed behavior among a finite set of hypotheses. Open-set GR, where goals can lie outside a predefined inventory or must be discovered online, requires additional tools such as novelty detection, goal discovery, or dynamic hypothesis generation [51]. Extending GRAIL to handle unknown or changing goal sets is an important next step.

Second, our experiments assume fixed environment dynamics. We do not study shifts in dynamics, changes in action semantics, or transfer to entirely new environments. Understanding how to adapt or regularize goal-conditioned policies to ensure GRAIL remains effective under such shifts is an important future direction.

Third, the empirical study focuses on simulated domains with vector observations and synthetic models of suboptimality and bias. These settings make it possible to control noise and structure, but they do not capture complex, raw inputs from perception-heavy tasks, which are usually attributed to the complementary problem of *Activity Recognition* rather than *Goal Recognition* [39, 54].

Fourth, per-goal training scales linearly with the number of candidate goals. Although this is acceptable for the inventories considered here and is common in many GR approaches, very large goal sets would increase training cost and storage demands. Future work can investigate ways to learn general goal-conditioned policies that support zero-shot or few-shot transfer learning for new goals without retraining a policy from scratch [19, 51].

Finally, we explore only a limited set of policy architectures and scoring functions. More expressive policies, attention mechanisms, or sequence models might capture structured biases more effectively, and calibrated probabilistic scores could provide better uncertainty estimates over goals. Investigating these extensions, while retaining the simple and planner-free inference pipeline, is a natural progression of this line of work.

## Societal and Ethical Implications for AI Alignment

Interpreting an agent’s goal from behavior is a core challenge in AI alignment, but doing so without imposing unrealistic assumptions about optimality or rationality raises significant ethical and policy considerations. While GRAIL enables more behaviorally-grounded goal inference, particularly for humans whose actions reflect preferences and constraints beyond normative models, this capability comes with risks. Systems that infer goals from demonstrations must be designed with strong safeguards to prevent misuse, such as privacy violations, surveillance, or manipulation of inferred intent [30]. As GR becomes increasingly integrated into human-facing applications, we must uphold the following principles: (1) transparency in how algorithms generate inferences and communicate them to users [15, 57, 58]; (2) privacy protection by minimizing the collection and retention of sensitive behavioral data [14]; (3) fairness by preventing systematic biases in GR across different user groups; and (4) user agency by enabling individuals to understand, contest, or override inferred goals [18, 29, 37]. Although GRAIL presents a novel framework for aligned GR, its deployment must be governed by policies that ensure the protection of user values and prevent exploitative applications.

## REFERENCES

- [1] Leonardo Amado, Ramon Fraga Pereira, João Paulo Aires, Mauricio Cecilio Magnaguagno, Roger Granada, and Felipe Meneguzzi. Goal recognition in latent space. In *International Joint Conference on Neural Networks*, pages 1–8. IEEE, 2018.
- [2] Leonardo Amado, Gabriel Paludo Licks, Matheus Marcon, Ramon Fraga Pereira, and Felipe Meneguzzi. Using self-attention lstms to enhance observations in goal recognition. In *International Joint Conference on Neural Networks*, pages 1–8. IEEE, 2020.
- [3] Leonardo Amado, Reuth Mirsky, and Felipe Meneguzzi. Goal recognition as reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9644–9651, 2022.
- [4] Leonardo R. Amado, Ramon F. Pereira, and Felipe Meneguzzi. Robust Neuro-Symbolic Goal and Plan Recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Press, 2023.
- [5] Leonardo Rosa Amado, Ramon Fraga Pereira, and Felipe Meneguzzi. Combining lstms and symbolic approaches for robust plan recognition. In Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé, editors, *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems*, pages 1634–1636. ACM, 2021.
- [6] Saurabh Arora and Prashant Doshi. A survey of inverse reinforcement learning: Challenges, methods and progress. *Artificial Intelligence*, 297:103500, 2021.
- [7] Michael Bain and Claude Sammut. A framework for behavioural cloning. In *Machine intelligence 15*, pages 103–129, 1995.
- [8] Chris L Baker, Joshua B Tenenbaum, and Rebecca R Saxe. Goal inference as inverse planning. In *Proceedings of the annual meeting of the cognitive science society*, volume 29, 2007.
- [9] Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
- [10] Gagan Bansal, Besmira Nushi, Ece Kamar, Walter S Lasecki, Daniel S Weld, and Eric Horvitz. Beyond accuracy: The role of mental models in human-AI team performance. In *Proceedings of the AAAI conference on human computation and crowdsourcing*, volume 7, pages 2–11, 2019.
- [11] Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The complexity of decentralized control of markov decision processes. *Mathematics of operations research*, 27(4):819–840, 2002.
- [12] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [13] Geoffrey L Brown, Nidhi Seethapathi, and Manoj Srinivasan. A unified energy-optimality criterion predicts human navigation paths and speeds. *Proceedings of the National Academy of Sciences*, 118(29):e2020327118, 2021.
- [14] Jiannan Cai, Zhidong Gao, Yuanxiong Guo, Bastian Wibranek, and Shuai Li. Fedhip: Federated learning for privacy-preserving human intention prediction in human-robot collaborative assembly tasks. *Advanced Engineering Informatics*, 60:102411, 2024.
- [15] Tathagata Chakraborti, Sarath Sreedharan, Yu Zhang, and Subbarao Kambhampati. Plan explanations as model reconciliation: moving beyond explanation as soliloquy. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 156–163, 2017.
- [16] Maxime Chevalier-Boisvert, Bolun Dai, Mark Towers, Rodrigo Perez-Vicente, Lucas Willems, Salem Lahlou, Suman Pal, Pablo Samuel Castro, and J K Terry. Minigrid & miniworld: Modular & customizable reinforcement learning environments for goal-oriented tasks. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 73383–73394. Curran Associates, Inc., 2023.
- [17] Mattia Chiari, Alfonso Emilio Gerevini, Francesco Percassi, Luca Putelli, Ivan Serina, and Matteo Olivato. Goal recognition as a deep learning task: The GRNet approach. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 33, pages 560–568, 2023.
- [18] Alexandra Coman and David W Aha. AI rebel agents. *AI magazine*, 39(3):16–26, 2018.
- [19] Osher Elhadad and Reuth Mirsky. General dynamic goal recognition. *arXiv preprint arXiv:2505.09737*, 2025.
- [20] Osher Elhadad, Felipe Meneguzzi, and Reuth Mirsky. Grail: Goal recognition alignment through imitation learning. *arXiv preprint arXiv:2602.14252*, 2026.
- [21] Zihao Fang, Dejun Chen, Yunxiu Zeng, Tao Wang, and Kai Xu. Real-time online goal recognition in continuous domains via deep reinforcement learning. *Entropy*, 25(10):1415, 2023.
- [22] Justin Fu, Katie Luo, and Sergey Levine. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.
- [23] Quentin Gallouédec, Nicolas Cazin, Emmanuel Dellandrea, and Liming Chen. panda-gym: Open-Source Goal-Conditioned Environments for Robotic Learning. *4th Robot Learning Workshop: Self-Supervised and Lifelong Learning at NeurIPS*, 2021.
- [24] Christopher W Geib. Delaying commitment in plan recognition using combinatorial grammars. In *IJCAI*, pages 1702–1707, 2009.
- [25] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. Pmlr, 2018.
- [26] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.
- [27] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Christina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017.
- [28] Gal Kaminka, Jan Wendler, and Galit Ronen. New challenges in multi-agent intention recognition. In *Proceedings of the Fall Symposium on Intention Recognition for Collaborative Tasks*, 2001.
- [29] Daniel Kasenberg and Matthias Scheutz. Norm conflict resolution in stochastic domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [30] Henry Kautz. The third AI summer: AAAI robert s. engelmore memorial lecture. *AI magazine*, 43(1):105–125, 2022.
- [31] David Lindner and Mennatallah El-Assady. Humans are not boltzmann distributions: Challenges and opportunities for modelling human feedback and interaction in reinforcement learning. *arXiv preprint arXiv:2206.13316*, 2022.
- [32] Jean Massardi, Mathieu Gravel, and Éric Beaudry. Parc: A plan and activity recognition component for assistive robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3025–3031. IEEE, 2020.
- [33] Peta Masters and Sebastian Sardina. Goal recognition for rational and irrational agents. In *Proceedings of the 18th international conference on autonomous agents and multiagent systems*, pages 440–448, 2019.
- [34] Peta Masters and Mor Vered. What’s the context? implicit and explicit assumptions in model-based goal recognition. In *International Joint Conference on Artificial Intelligence 2021*, pages 4516–4523. Association for the Advancement of Artificial Intelligence (AAAI), 2021.
- [35] Malek Mechergui and Sarath Sreedharan. Goal alignment: re-analyzing value alignment problems using human-aware AI. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 10110–10118, 2024.
- [36] Felipe Rech Meneguzzi and Ramon Fraga Pereira. A survey on goal recognition as planning. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021, Canada.*, 2021.
- [37] Reuth Mirsky. Artificial intelligent disobedience: Rethinking the agency of our artificial teammates. *AI Magazine*, 46(2):e70011, 2025.
- [38] Reuth Mirsky and Ya’akov Gal. Slim: semi-lazy inference mechanism for plan recognition. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 394–400, 2016.
- [39] Reuth Mirsky, Sarah Keren, and Christopher Geib. *Introduction to symbolic plan and goal recognition*, volume 16. Springer, 2021.
- [40] Ben Nageris, Felipe Meneguzzi, and Reuth Mirsky. Goal recognition using actor-critic optimization. *arXiv preprint arXiv:2501.01463*, 2024.
- [41] Andrew Y Ng, Stuart Russell, et al. Algorithms for inverse reinforcement learning. In *Icml*, volume 1, page 2, 2000.
- [42] Jean Oh, Felipe Meneguzzi, and Katia Sycara. Probabilistic plan recognition for proactive assistant agents. *Plan, activity, and intent recognition*, 10:23, 2014.
- [43] Ramon Fraga Pereira, Mor Vered, Felipe Meneguzzi, and Miquel Ramirez. Online Probabilistic Goal Recognition over Nominal Models. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pages 5547–5553, 2019.
- [44] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. Landmark-based approaches for goal recognition as planning. *Artificial Intelligence*, 279:103217, 2020.
- [45] Luisa R de A Santos, Felipe Meneguzzi, Ramon F Pereira, and Andre Pereira. An LP-Based Approach for Goal Recognition as Planning. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*. AAAI Press, 2021.
- [46] Irina Rabkina and Kenneth D Forbus. Analogical reasoning for intent recognition and action prediction in multi-agent systems. In *Proceedings of the Seventh Annual Conference on Advances in Cognitive Systems*, pages 504–517. Cognitive Systems Foundation Cambridge, 2019.
- [47] Miguel Ramirez and Hector Geffner. Probabilistic plan recognition using off-the-shelf classical planners. In *Proceedings of the AAAI conference on artificial intelligence*, volume 24, pages 1121–1126, 2010.
- [48] Miquel Ramirez and Hector Geffner. Plan recognition as planning. In *Proceedings of the 21st international joint conference on Artificial intelligence*. Morgan Kaufmann Publishers Inc, pages 1778–1783, 2009.
- [49] Brian Scassellati. Theory of mind for a humanoid robot. *Autonomous Robots*, 12: 13–24, 2002.
- [50] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [51] Matan Shamir, Osher Elhadad, Matthew E Taylor, and Reuth Mirsky. Odgr: Online dynamic goal recognition. *arXiv preprint arXiv:2407.16220*, 2024.
- [52] Maayan Shvo and Sheila A McIlraith. Active goal recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9957–9966, 2020.

- [53] Sarath Sreedharan, Subbarao Kambhampati, et al. Handling model uncertainty and multiplicity in explanations via model reconciliation. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 28, pages 518–526, 2018.
- [54] Gita Sukthankar and Katia Sycara. Activity recognition for dynamic multi-agent teams. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(1):1–24, 2011.
- [55] Gita Sukthankar, Christopher Geib, Hung Hai Bui, David Pynadath, and Robert P Goldman. *Plan, activity, and intent recognition: Theory and practice*. Newnes, 2014.
- [56] Susanne Trick, Dorothea Koert, Jan Peters, and Constantin A Rothkopf. Multimodal uncertainty reduction for intention recognition in human-robot interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7009–7016. IEEE, 2019.
- [57] Alan R Wagner and Paul Robinette. An explanation is not an excuse: Trust calibration in an age of transparent robots. In *Trust in human-robot interaction*, pages 197–208. Elsevier, 2021.
- [58] Alan FT Winfield, Serena Booth, Louise A Dennis, Takashi Egawa, Helen Hastie, Naomi Jacobs, Roderick I Muttram, Joanna I Olszewska, Fahimeh Rajabiyazdi, Andreas Theodorou, et al. Ieee p7001: A proposed standard on transparency. *Frontiers in Robotics and AI*, 8:665729, 2021.
- [59] Tan Zhi-Xuan, Jordyn Mann, Tom Silver, Josh Tenenbaum, and Vikash Mansinghka. Online bayesian goal inference for boundedly rational planning agents. *Advances in neural information processing systems*, 33:19238–19250, 2020.
- [60] Simon Zhuang and Dylan Hadfield-Menell. Consequences of misaligned ai. *Advances in Neural Information Processing Systems*, 33:15763–15773, 2020.