Proceedings of the Thirty-Sixth International Conference on Automated Planning and Scheduling (ICAPS 2026)

Mitigating Deception and Interference in Online Goal Recognition Systems

Lorenzo Serina, Mattia Chiari, Matteo Olivato, Luca Putelli, Nicholas Rossetti, Ivan Serina, Alfonso Emilio Gerevini

Department of Information Engineering
Università degli Studi di Brescia
{name}.{surname}@unibs.it

Abstract

Online Goal Recognition (OGR) is the task of recognizing an agent's goal while that agent is executing a plan. Several OGR systems have been designed to observe and analyze an agent's actions in order to infer its goal. However, these systems generally assume that the agent is either unaware of being observed or is cooperating with the system. In this paper, we analyze two scenarios in which this assumption does not hold: Deception, where an agent deliberately hides its true goal by computing a deceptive plan; and Interference, where another agent tampers with the original plan. In particular, we evaluate the performance of the two state-of-the-art systems for OGR (ORL and CLERNET) in these two scenarios, gathering and extending different types of deceptive and interfering attacks. Moreover, we propose a framework (PAC-OGR) that mitigates the effect of the attacks by amending the manipulated plan and reasoning about the agent's behaviour. An experimental evaluation over several classical planning domains shows that PAC-OGR can be effectively integrated into existing OGR systems, making them more robust and reliable.

Introduction

Online Goal Recognition (OGR) systems are designed to infer an agent's objectives by analysing its actions. The state-of-the-art OGR systems, Online Recognition with Landmarks (ORL) (Vered et al. 2018) and CLERNET (Serina et al. 2025), assume keyhole recognition (Meneguzzi and Pereira 2021). In this paper, we challenge these assumptions considering two scenarios: *Deception* and *Interference*. In Deception, a scenario inspired by the work in (Price et al. 2023), the agent is supposed to be malevolent and executes a deceptive plan trying to hide its goal from the observer. We consider two kinds of deceptive plans: guided by landmarks (Hoffmann, Porteous, and Sebastia 2004) and by centroids (Pozanco et al. 2019). The landmarks for a goal are fluents that must hold at some point in the state trajectory of a valid plan reaching the goal. In a landmark-guided deceptive plan, the agent tries to achieve a *fake* goal that shares the highest number of landmarks with the real one and, from that goal, executes another plan towards the real goal. In a centroid-guided plan, the agent tries to reach an intermediate state (a centroid) that is equally distant from several possible goals, with the intention of not revealing its true goal (at least until the centroid state is crossed). In the Interference scenario, that we introduce in this work, an Intruder agent tampers with the observations made by the OGR system by adding fake executed actions, and by removing or replacing some of the actions that are truly executed by the observed agent.

Our first contribution is an evaluation of the robustness of the state-of-the-art OGR systems against different kinds of attacks. Moreover, we introduce and evaluate two integrable defence mechanisms which analyse both the observed actions (in pre-processing) and the predictions provided by the OGR systems (in post-processing). In the pre-processing phase, a custom algorithm analyses the preconditions and the effects of each observed action, trying to understand whether it has been tampered by the Intruder. The post-processing technique analyses the agent's behaviour in terms of the heuristic distances to each candidate goal after every observed action. A significant increase of the distance for a specific goal can indicate that the agent is no longer working towards that goal. This provides a hint that is exploited to revise the goal scores of the OGR system accordingly. An experimental analysis evaluates the robustness of state-of-the-art OGR systems and the performance of our new defence mechanisms over four well-known benchmark domains.

Background and Related Work

Goal recognition (GR) is the task of inferring an agent's objectives in an environment, based on a sequence of observed actions $\mathcal{O} = \langle o_1, \ldots, o_n \rangle$, where $o_i \in \mathcal{A}$ and $\mathcal{A}$ is the set of actions that the agent can perform (Ramírez and Geffner 2009). An *instance of the GR problem* in a given domain is then specified by: an initial state I of the agent and environment; the sequence of observations $\mathcal{O}$, and a set $\mathcal{G} = \{G_1, .., G_m\}$ of possible goals of the agent. The observations form a trace of the full sequence π of actions performed by the agent to achieve a goal $G^* \in \mathcal{G}$.

We consider the *online* setting of goal recognition (OGR), as formulated in (Vered et al. 2018), where the observation trace is complete, incrementally revealed action by action, and forms a prefix of the agent's plan. The candidate OGR solution is generated every time a new action is observed.

In this work, we focus on two state-of-the-art systems for OGR on discrete domains: ORL (Vered et al. 2018) and CLERNET (Serina et al. 2025). Following the approach in

(Pereira, Oren, and Meneguzzi 2017, 2020), ORL precomputes ordered sets of landmarks for each goal, and incrementally marks achieved landmarks as observations arrive. By estimating goal completion and pruning incompatible goals, it enables efficient and planner-free online goal recognition. Similarly to GRNet for offline goal recognition (Chiari et al. 2023), CLERNET is based on LSTM Networks and, whenever it observes an action performed by the agent, it processes the sequence of the actions observed so far to predict the single fluents in a goal. Then an algorithm aggregates these predictions to select the best candidate goal in $\mathcal{G}$.

Deceptive Planning generates a sequence of actions such that the observer is unable to determine the agent's goal. A first definition is introduced for path planning (Masters and Sardiña 2017), while its generalisation for classical planning is introduced in (Price et al. 2023). Most importantly, this paper employs two techniques exploiting planning knowledge: Landmark-based attacks, in which the Agent initially acts to achieve the deceptive goal that has the most similar landmarks (Hoffmann, Porteous, and Sebastia 2004) with the real one, and Centroid-based attacks, where it acts to reach a state close to several candidate goals. To the best of our knowledge, in this context, there is no study on defence mechanisms against Deceptive Planning.

Attacking OGR Systems

We study two types of attacks: our newly introduced *Interference*, and *Deception* (Price et al. 2023).

Interference

A standard OGR system considers two actors: an Agent performing actions to reach a goal, and an Observer that analyses the agent's behaviour to infer its goal. In our Interference scenario, we consider an additional malevolent agent, called the *Intruder*. We assume that the Intruder has full knowledge of the Agent's action models and of the environment (plan states), just like the Observer. The Intruder can tamper with the observations received by the Observer without directly interacting with the Agent. Therefore, while the Agent executes a plan π that achieves the correct goal G^*, the observation sequence $\mathcal{O}$ perceived by the Observer may differ from the true one. Since $\mathcal{O}$ is generated incrementally as the Agent performs actions, any corruption introduced at observation o_i persists for all subsequent observations. After each action is performed, the Intruder may attempt an attack with a certain probability p. We assume that p is fixed for any given OGR instance.

Insert attack. The Intruder injects a spurious observation (action) into $\mathcal{O}$. The inserted action is drawn from $\mathcal{A}$, but it may be inapplicable in the state where it is inserted.

Remove attack. The Intruder prevents some actions from being added to $\mathcal{O}$, resulting in an incomplete observation sequence. The Observer is unaware of these missing observations.

Replace attack. The Intruder substitutes a valid observation in $\mathcal{O}$ with a fake one. The fake action is drawn from $\mathcal{A}$, but it may be inapplicable in the state where it is inserted.

Deception

In the Deception scenario, the OGR framework involves two actors: the Agent and the Observer. Unlike the Interference scenario, here the Agent itself is malicious and actively attempts to hide its true goal from the Observer for as long as possible. To this end, the Agent applies goal-obfuscating strategies that influence its action choices. Note that, despite these strategies, the Agent must act within the constraints of the environment and cannot interact with or tamper with the Observer. Thus, the observation sequence $\mathcal{O}$ fully reflects the actions executed by the Agent. We consider two types of deceptive strategies: Landmarks and Centroid attacks.

Landmarks. The landmark-based attack was introduced in (Price et al. 2023). We adopt the *Most Similar Landmarks Goal* approach ($D\mathcal{L}_{Similar}$) as it has been shown to yield the strongest deception effects. The Agent initially executes a plan that achieves a wrong goal $G^- \in \mathcal{G} \setminus G^*$, chosen to maximize the landmark overlap of G^- and the true goal G^*. Specifically, for each candidate goal in $\mathcal{G} \setminus G^*$, we compute the set of its landmarks, and select G^- as the goal whose landmarks have the largest overlap with those for G^*. The malevolent Agent executes a plan reaching G^-, and then proceeds toward G^*, thereby misleading the Observer regarding the true target at least until it reaches G^-.

Centroid. A centroid C for an OGR task is defined as the state minimizing the average estimated cost (here the number of needed actions) of reaching each goal in $\mathcal{G}$ (Pozanco et al. 2019). In this attack, the Agent first navigates toward C, and then continues to the true goal G^*. This delays explicit commitment to G^* keeping the observation sequence consistent with multiple goal hypotheses. Following (Price et al. 2023), we adopt the $D\mathcal{C}$ strategy. To reduce computational cost and improve deception, we restrict the centroid computation to a subset $\mathcal{G}_\mathcal{C} \subseteq \mathcal{G}$ of fixed cardinality (in our experiments, $|\mathcal{G}_\mathcal{C}| = 3$). The subset is constructed by ranking goals according to their similarity with G^* and selecting the top $|\mathcal{G}_\mathcal{C}|$ goals. We use the landmark overlap measure introduced for Landmarks attack as similarity measure. Please note that G^* is always included in $\mathcal{G}_\mathcal{C}$.

Enhancing OGR Systems: PAC-OGR

To mitigate the adversarial effects of Interference and Deception attacks, we introduce the *PAC-OGR* (Pre-processing and Analysis for Clean OGR) framework. PAC-OGR acts as a tailored wrapper around the Goal Recognition system, preprocessing the observation stream and re-scaling the output scores using symbolic domain knowledge.

Pre-Processing In pre-processing, our defence mechanism filters out potentially tampered actions. This process is performed every time a new observation is added to $\mathcal{O}$. First, the observation sequence is validated using VAL (Howey, Long, and Fox 2004). If the sequence is found to be invalid, indicating tampering, we invoke a custom action-validation algorithm.

Given the initial state I and a sequence $\mathcal{O} = \langle o_1, \ldots, o_k \rangle$ of observed actions, we check the authenticity of each action o_i as follows. For every $o_i \in \mathcal{O}$, we compute two scores: P_i and E_i. P_i indicates the percentage of preconditions of

o_i that are satisfied in the state s_{i-1} generated by executing $\mathcal{O}_i^- = \langle o_1, \ldots, o_{i-1} \rangle$. Note that we construct s_{i-1} by applying all effects of the actions in $\mathcal{O}_i^-$ starting from I, without checking whether their preconditions are satisfied, and we calculate the percentage of preconditions of o_i that hold in that state. E_i indicates the percentage of effects of o_i that appear as preconditions for at least one subsequent action in $\mathcal{O}_i^+ = \langle o_{i+1}, \ldots, o_k \rangle$. The intuition behind these scores is that tampered actions — not belonging to the Agent's original plan — are unlikely to have their preconditions supported by prior actions, or produce effects that enable future actions. We compare these scores against two thresholds ϕ_p and ϕ_e to classify actions as follows. If $P_i > \phi_p$, we consider o_i authentic, as most of its preconditions are satisfied. Otherwise, we check E_i: if $E_i > \phi_e$, we consider o_i authentic, since its effects contribute to the remaining execution. If neither condition holds, o_i is *removed* from $\mathcal{O}$.

Post-Processing The post-processing component of PAC-OGR evaluates and re-scales the predictions made by the OGR system, which assigns a numeric score to each candidate goal entirely unaware of potential attacks. The idea is to examine the sequence of the states generated by the observed actions to access their coherence with each candidate goal, and weight the OGR predictions accordingly. To do this we exploit the function $f(s) = g(s) + h(s)$ used in standard heuristic search to evaluate a search state s. Specifically, for every candidate goal G, we compute $f^G(s_i) = i + h^G(s_i)$, where i is the length of the observation sequence $\langle o_1, \ldots o_i \rangle$ (i.e., the number of observed actions executed by the Agent so far), s_i is the state generated at the end of the sequence, and $h^G(s_i)$ is a heuristic estimate of the distance from s_i to G. Thus, for a goal G, we obtain the sequence $F_i^G = \langle f^G(s_1), \ldots, f^G(s_i) \rangle$.

Our method aims to penalize candidate goals for which the heuristic distances tend to increase over time, since this suggests that the Agent is moving away from them. However, the sequence of heuristic distances $\langle h^G(s_1), \ldots, h^G(s_i) \rangle$ alone would not always be informative enough, as it may exhibit plateaus or small fluctuations even when the Agent is actually pursuing G. By considering the f^G values instead, we also incorporate the length of the observation sequence, and a lack of heuristic-distance increase can be compensated by the (increasing) first term of f. The trend of the f_G values can be quantified by computing a linear regression over the sequence F^G in the form $y = ax + b$, and using the slope a as an indicator of goal deception: the larger a is, the less likely G is the true goal.

To ensure that a captures a reliable trend, the regression is performed only over the most recent portion of the sequence that exhibits sufficiently strong linear correlation, measured using Pearson's coefficient. Starting from the first observation, earlier values are iteratively discarded until the absolute value of the coefficient exceeds a threshold ϵ. Finally, we integrate this slope information with the score s_G^{OGR} provided by the OGR: for each candidate goal G, we compute its final score as $s_G = s_G^{OGR} \cdot \left(1 - \frac{\arctan(a)}{\pi/2}\right)$, where negative slopes a are set to 0 . Larger slopes thus penalize goals whose f^G trends make them increasingly implausible.

Experimental Evaluation

For our experimental evaluation, we use the two main SOTA OGR systems: ORL and CLERNET. We evaluate their robustness to Interference and Deception attacks, and PAC-OGR can enhance it. The experiments are conducted over four planning domains: BLOCKSWORLD, DEPOTS, DRIVERLOG, LOGISTICS. For each domain, we have a test set with 350 problems extracted from the test set provided in (Serina et al. 2025), onto which we applied the described attacks. For Interference attacks, we set $p = \{0.10, 0.15, 0.2\}$ and use an equal number of instances for each value. We evaluate the performance using the standard OGR metrics *Ranked First* (RF) and *Convergence* (CV) (Vered et al. 2018). For an OGR instance, RF is defined as the number of predicted goals (one per observation) correctly identified divided by the total number of observations. CV expresses how early the model can predict the correct goal, and it is calculated similarly to RF, but also considering whether the system does not predict another goal from that point onwards.

The thresholds ϕ_p and ϕ_e used to detect tampered actions are determined through an iterative optimization process testing different values. This process is performed on a dedicated validation set composed of 50 instances, where 30% of the observed actions have been tampered. The optimal parameter pair is selected by comparing the actions predicted as tampered against the ground truth, to maximize the $F_{0.5}$ score. The $F_{0.5}$ score is specifically chosen to penalize false positives, as the system considers false positives to be particularly harmful. In fact, misclassifying a valid action as tampered could exclude valid information, potentially causing performance degradation. For the post-processing component, h^G is computed as the average of admissible ($h_{\max}$ (Bonet and Geffner 2001), LM_{cut} (Helmert and Domshlak 2009)) and inadmissible heuristics h_{ff} (Hoffmann and Nebel 2001), h_{sa} (Keyder and Geffner 2007), in order to combine conservative lower bounds with more informative but optimistic estimates. ϵ is set to 0.95.

Results for Interference attacks The results of our robustness analysis are shown in the left and central parts of Table 1. These experiments demonstrate that ORL and CLERNET possess high inherent robustness, particularly against the Insert attack. For this type of attack, the average performance drop is modest: approximately 4 points for ORL (e.g., RF drops from 43.6 to 39.7) and roughly 2 points for CLERNET in terms of both metrics. The Delete attack causes comparable but slightly higher degradation. However, the Replace attack poses more issues, diminishing the RF and CV scores of several points for both ORL and CLERNET. Domains DEPOTS and DRIVERLOG are the most vulnerable. The worst result can be seen for the Replace attack in DRIVERLOG, where ORL dropped from 48.6 to 39.2 in RF. Conversely, LOGISTICS and BLOCKSWORLD are the most resilient domains.

The application of PAC-OGR delivers the most significant benefits against the Insert attack, yielding consistent gains across all domains. We observed average increases of approximately 1 point in both metrics for CLERNET and ORL (e.g., ORL CV rose from 37.4 to 38.6). The impact

Domain	System	No Attacks		Interference						Deception			
				Insert		Delete		Replace		Centroids		Landmarks	
		RF	CV	RF	CV	RF	CV	RF	CV	RF	CV	RF	CV
BLOCKS	ORL	36.8	34.7	34.5	32.2	**32.0**	**29.8**	**30.9**	**28.3**	25.7	23.9	11.5	7.0
	ORL + PAC-OGR	36.8	34.7	**34.7**	**32.5**	31.2	29.1	30.6	28.2	**31.1**	**27.3**	**19.5**	**11.5**
	CLERNET	43.2	42.0	42.4	39.6	**40.3**	**38.2**	**39.1**	**36.1**	33.4	**32.4**	17.3	11.5
	CLERNET + PAC-OGR	43.2	42.0	**42.8**	**40.3**	39.5	37.3	38.8	35.5	**36.0**	31.1	**20.3**	**13.2**
DEPOTS	ORL	45.3	44.0	40.6	38.1	**39.1**	**37.7**	37.4	34.4	31.2	29.0	11.3	5.9
	ORL + PAC-OGR	45.3	44.0	**42.0**	**40.1**	38.0	36.3	**37.5**	**35.4**	**40.0**	**33.3**	**15.4**	**8.2**
	CLERNET	48.3	43.5	46.1	39.4	**44.4**	**39.4**	**43.3**	36.2	39.1	**35.6**	**19.3**	**13.1**
	CLERNET + PAC-OGR	48.3	43.5	**47.2**	**42.0**	42.7	37.7	42.8	**37.0**	**41.4**	34.6	19.2	13.0
DRIVERLOG	ORL	48.6	46.7	42.8	39.8	**41.8**	**39.9**	39.2	35.9	32.6	28.9	13.3	7.2
	ORL + PAC-OGR	48.6	46.7	**44.0**	**41.6**	41.4	39.5	**39.7**	**37.3**	**38.3**	**29.3**	**21.5**	**12.1**
	CLERNET	58.3	52.3	56.2	48.8	**55.1**	**49.1**	51.4	43.6	39.5	32.3	24.8	16.8
	CLERNET + PAC-OGR	58.3	52.3	**57.1**	**49.9**	54.9	48.9	**52.4**	**45.3**	**42.2**	**32.3**	**25.9**	**18.2**
LOGISTICS	ORL	43.6	41.2	40.9	39.4	**39.4**	**38.4**	38.7	37.1	25.8	**21.3**	**8.4**	**3.9**
	ORL + PAC-OGR	43.6	41.2	**41.2**	**40.0**	38.9	38.0	**38.8**	**37.3**	**25.9**	20.0	7.9	2.6
	CLERNET	51.2	48.1	49.5	46.3	**46.9**	**44.1**	42.3	39.8	**36.2**	**26.1**	**13.1**	**5.8**
	CLERNET + PAC-OGR	51.2	48.1	**50.7**	**47.4**	46.5	43.6	**46.7**	**43.6**	35.1	23.4	12.8	5.5
All Domains	ORL	43.6	41.7	39.7	37.4	**38.1**	**36.5**	36.5	33.9	28.8	25.8	11.1	6.0
	ORL + PAC-OGR	43.6	41.7	**40.5**	**38.6**	37.4	35.7	**36.7**	**34.5**	**33.8**	**27.5**	**16.1**	**8.6**
	CLERNET	50.3	46.5	48.6	43.5	**46.7**	**42.7**	44.7	39.5	37.1	**31.6**	18.6	11.8
	CLERNET + PAC-OGR	50.3	46.5	**49.4**	**44.9**	45.9	41.9	**45.2**	**40.4**	**38.7**	30.4	**19.6**	**12.5**

Table 1: Experimental results comparing ORL and CLERNET without/with attacks, and considering the proposed defence methods. Best values between the standard and defence configurations are shown in **bold**.

of the Replace attack is more nuanced. While the global average indicates a slight overall improvement (e.g., ORL RF increased marginally from 36.5 to 36.7), domain-level analysis reveals that the defence is counter-productive in BLOCKSWORLD, especially for CLERNET. In this domain, our pre-processing component misclassifies some valid actions, causing information loss, and thus a slight performance decrease. For the Delete attack, PAC-OGR produces a slight performance drop in terms of both RF and CV. This is expected because the pre-processing component is designed only to remove tampered actions and not to infer missing ones, which results in possible information loss when the Delete attack is performed.

Results on Deception attacks The results of these experiments are shown in the left and right parts of Table 1. Both systems are vulnerable to deception attacks, as observed in (Price et al. 2023) for ORL. For the Centroid attack, both ORL and CLERNET exhibit an average loss of 15 points in terms of RF (from 43.6 to 28.8 and from 50.3 to 37.1) and CV (from 41.7 to 25.8 and from 46.5 to 31.6), as the agent moves towards states compatible with multiple goals. The Landmarks attack causes a much more pronounced performance drop for ORL and CLERNET in terms of RF (from 43.6 to 11.1 and from 50.3 to 18.6) and CV (from 41.7 to 6.0 and from 46.5 to 11.8), since the agent moves towards a fake goal, thereby misleading the system. However, PAC-OGR mitigates the impact of these attacks, enhancing the robustness of OGR systems. The defence technique applied to ORL leads to a remarkable improvement: for Centroid, from 28.8 to 33.8 RF, and from 25.8 to 27.5 CV; for Landmarks, from 11.1 to 16.1 RF, and from 6.0 to 8.6 CV. Both metrics increase consistently in BLOCKS, DEPOTS, and DRIVERLOG, with only a slight decrease observed in LOGISTICS. A smaller gain (1 point) can be seen for CLERNET in Landmarks, whereas for Centroid RF slightly increases (from 37.1 to 38.7) but CV decreases (from 31.6 to 30.4). Comparing the different domains, the best results are in DRIVERLOG, where both metrics improve, whereas the worst results are obtained in LOGISTICS.

Conclusions and Future Work

We have investigated the robustness of state-of-the-art OGR systems (ORL and CLERNET) under Interference and Deception attacks. Our experimental results show that both systems are robust to action tampering, whereas they lose performance against Deception, especially for the Landmarks attack. We have also introduced PAC-OGR, a defence mechanism that (i) filters out tampered actions; and (ii) leverages heuristic-based trend analysis to down-weight deceptive plans. Our experiments show that PAC-OGR enhances the robustness of both OGR systems, especially ORL. As future work, we aim to explore more adversarial settings and design neuro-symbolic strategies to further enhance the robustness of OGR systems, as in (Chiari et al. 2024).

Acknowledgments

This work was partially supported by project SERICS (PE00000014) and by project FAIR (B53C22003980006) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU, by the Italian Ministry of University and Research within the PRIMA 2024 programme project "Optimizing Water Resources in Coastal Areas using Artificial Intelligence" (AI4WATER – D53C25000510006), by the EU H2020 project AIPlan4EU (GA 101016442) and by the Climate Change AI project (No. IG-2023-174).

References

Bonet, B.; and Geffner, H. 2001. Planning as heuristic search. *Artif. Intell.*, 129(1-2): 5–33.

Chiari, M.; Gerevini, A. E.; Loreggia, A.; Putelli, L.; and Serina, I. 2024. Fast and Slow Goal Recognition. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, 354–362. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400704864.

Chiari, M.; Gerevini, A. E.; Percassi, F.; Putelli, L.; Serina, I.; and Olivato, M. 2023. Goal Recognition as a Deep Learning Task: The GRNet Approach. *Proceedings of the International Conference on Automated Planning and Scheduling*, 33(1): 560–568.

Helmert, M.; and Domshlak, C. 2009. Landmarks, Critical Paths and Abstractions: What's the Difference Anyway? In Gerevini, A.; Howe, A. E.; Cesta, A.; and Refanidis, I., eds., *Proceedings of the 19th International Conference on Automated Planning and Scheduling, ICAPS 2009, Thessaloniki, Greece, September 19-23, 2009*. AAAI.

Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *J. Artif. Intell. Res.*, 14: 253–302.

Hoffmann, J.; Porteous, J.; and Sebastia, L. 2004. Ordered Landmarks in Planning. *J. Artif. Intell. Res.*, 22: 215–278.

Howey, R.; Long, D.; and Fox, M. 2004. VAL: Automatic Plan Validation, Continuous Effects and Mixed Initiative Planning Using PDDL. In *16th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2004), 15-17 November 2004, Boca Raton, FL, USA*, 294–301. IEEE Computer Society.

Keyder, E.; and Geffner, H. 2007. Heuristics for Planning with Action Costs. In Borrajo, D.; Castillo, L. A.; and Corchado, J. M., eds., *Current Topics in Artificial Intelligence, 12th Conference of the Spanish Association for Artificial Intelligence, CAEPIA 2007, Salamanca, Spain, November 12-16, 2007. Selected Papers*, volume 4788 of *Lecture Notes in Computer Science*, 140–149. Springer.

Masters, P.; and Sardiña, S. 2017. Deceptive Path-Planning. In Sierra, C., ed., *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, 4368–4375. ijcai.org.

Meneguzzi, F.; and Pereira, R. F. 2021. A Survey on Goal Recognition as Planning. In Zhou, Z., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 4524–4532. ijcai.org.

Pereira, R. F.; Oren, N.; and Meneguzzi, F. 2017. Landmark-Based Heuristics for Goal Recognition. In Singh, S.; and Markovitch, S., eds., *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, 3622–3628. AAAI Press.

Pereira, R. F.; Oren, N.; and Meneguzzi, F. 2020. Landmark-based approaches for goal recognition as planning. *Artif. Intell.*, 279.

Pozanco, A.; E-Martín, Y.; Fernández, S.; and Borrajo, D. 2019. Finding Centroids and Minimum Covering States in Planning. In Benton, J.; Lipovetzky, N.; Onaindia, E.; Smith, D. E.; and Srivastava, S., eds., *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS 2019, Berkeley, CA, USA, July 11-15, 2019*, 348–352. AAAI Press.

Price, A.; Pereira, R. F.; Masters, P.; and Vered, M. 2023. Domain-Independent Deceptive Planning. In Agmon, N.; An, B.; Ricci, A.; and Yeoh, W., eds., *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, 95–103. ACM.

Ramírez, M.; and Geffner, H. 2009. Plan Recognition as Planning. In Boutilier, C., ed., *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, 1778–1783.

Serina, L.; Chiari, M.; Gerevini, A. E.; Putelli, L.; and Serina, I. 2025. Towards Efficient Online Goal Recognition through Deep Learning. In Das, S.; Nowé, A.; and Vorobeychik, Y., eds., *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2025, Detroit, MI, USA, May 19-23, 2025*, 1895–1903. International Foundation for Autonomous Agents and Multiagent Systems / ACM.

Vered, M.; Pereira, R. F.; Magnaguagno, M. C.; Kaminka, G. A.; and Meneguzzi, F. 2018. Towards Online Goal Recognition Combining Goal Mirroring and Landmarks. In André, E.; Koenig, S.; Dastani, M.; and Sukthankar, G., eds., *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, 2112–2114. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM.