

Domain-Independent Deceptive Planning

Adrian Price
Monash University, Australia
adrian.price@monash.edu

Ramon Fraga Pereira
University of Manchester, UK
ramon.fragapereira@manchester.ac.uk

Peta Masters
King’s College London, UK
peta.masters@kcl.ac.uk

Mor Vered
Monash University, Australia
mor.vered@monash.edu

ABSTRACT

We investigate *deceptive planning*, the problem of generating a plan such that an observer is unable to determine its ultimate goal. Most work in this area has focused on path and/or motion planning. However planning problems can be quite varied and challenging. We present domain-independent approaches for deceptive plan generation utilising the concepts of *landmarks*, *centroids*, and *minimum covering states*. We introduce new, domain-independent metrics to evaluate a plan’s deceptivity as a ratio between its deceptive quantity and cost; and we extensively evaluate the performance of our proposed approaches over widely different planning domains providing guidelines as to when to use each approach.

KEYWORDS

Planning; Deception; Plan Recognition; Domain-Independent

ACM Reference Format:

Adrian Price, Ramon Fraga Pereira, Peta Masters, and Mor Vered. 2023. Domain-Independent Deceptive Planning. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

The ability to deceive is a marker for human intelligence. It has been a central focus for *Artificial Intelligence* (AI) since Turing’s “Imitation Game” [51], that is, even before the term AI was coined [36]. *Deception* is often misunderstood. Although its first associations may be with crime and double-dealing, it is arguably only by knowing how and why particular actions have the potential to deceive that we can ensure they are performed in such a way as to be correctly understood [28, 31]. Moreover, deception in human affairs is also a social lubricant [56] whilst, in a security context, it is often regarded as synonymous with “privacy-preservation” [9, 29]. Unsurprisingly then, given deception’s multiplicity of uses in the human sphere, the more we interact with AI systems in our daily lives—as team members, personal assistants, driving aids, etc.—the more we expect them to be capable not only of detecting deception but also of deceiving/preserving privacy on our behalf.

In this paper, we investigate *deceptive planning*, which involves the generation of a plan such that an observer is unable to determine the plan’s purpose or assumes that its purpose is other than it is. Much previous work in this area has centred on *motion planning* and *path-planning* [e.g., 12, 32]. Even works whose definitions of

deception relate more broadly to *classical planning* (as defined in [15, Chapter 2]) ultimately provide examples and/or describe experiments that relate primarily to a path or motion-planning domain [e.g., 28, 35]. Planning problems, however, are more varied and challenging than the above examples suggest.

We offer new approaches to the generation of deceptive plans across a range of problems, utilising the concept of *relevant states* such as *landmarks* [22], *centroids* and *minimum covering states* [42]. Our core contributions is a range of *domain-independent* approaches to deceptive planning. We introduce *deceptive metrics*, which incorporate a novel definition of deceptive actions and quantify deception’s effectiveness as a ratio between its deceptive quantity (i.e., number of deceptive actions) and cost. We provide empirical evaluation of our approaches across multiple planning domains using two state-of-the-art online *goal recognition* approaches: *mirroring* [53] and *mirroring with landmarks* [55]. We analyse and compare approaches that provide exceptional deception, but are unrealistic for real world implementation owing to their excessive plan costs, with others that offer *relatively* good deception but are more economical. We also examine the impact of domain-specific factors, such as a limited number of landmarks, on the deceptiveness of plans.

In Section 2, we set out the technical background to our work. In Section 3, we present novel metrics for the evaluation and comparison of deceptive plans and twelve distinct approaches utilising the concepts of *landmarks*, *centroids* and *minimum covering states*. In Section 4, we provide an experimental evaluation and, in Section 5, we conclude the paper with discussion.

2 BACKGROUND AND RELATED WORK

We begin by setting out the technical framework on which we rely, then contextualise our work in relation to deceptive AI, landmarks, centroids and minimum covering states.

2.1 Technical Framework

We define our problems using the classical planning formalism [17, 46], assuming an environment which is *discrete*, *fully observable*, and *deterministic* [15, Chapter 2]. A *planning domain* \mathcal{D} is defined as $\langle F, A \rangle$ where: F is a set of *fluents* (i.e., environment properties); and A is a set of *actions* where every action $a \in A$ has a *positive cost*, denoted as $cost(a)$, and its own set of preconditions, add and delete lists: $Pre(a), Add(a), Del(a)$. We define a state S as a finite set of positive fluents $f \in F$ that follows the *closed world assumption* so that if $f \in S$, then f is true in S . We also assume a simple inference relation \models such that $S \models f$ iff $f \in S$, $S \not\models f$ iff $f \notin S$, and $S \models f_1 \wedge \dots \wedge f_n$ iff $\{f_1, \dots, f_n\} \subseteq S$. An action $a \in A$ is applicable

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

to a state S if and only if $S \models \text{Pre}(a)$, and generates a new state S' such that $S' \leftarrow (S \cup \text{Add}(a))/\text{Del}(a)$.

A *planning problem* \mathcal{P} is defined as $\langle \mathcal{D}, S_I, G \rangle$ where: \mathcal{D} is a planning domain as described above; $S_I \subseteq F$ is the *initial state*; and $G \subseteq F$ is a *goal state*. A *solution* to the planning problem \mathcal{P} is a *plan* $\pi = [a_1, \dots, a_n]$ that maps S_I into a state $S \models G$, that is, in which the goal state G holds. The main purpose of a planning problem is often to find a plan π with minimal cost (or length, where all actions have cost equal to 1). The cost of a plan $\pi = [a_1, a_2, \dots, a_n]$ is $\text{cost}(\pi) = \sum \text{cost}(a_i)$ and we say that a plan π^* is *optimal* if there exists no other plan π for \mathcal{P} such that $\text{cost}(\pi) < \text{cost}(\pi^*)$.

To evaluate the deceptivity of our plans, we use *goal recognition* (GR), the task of inferring an agent’s unobserved goal from its observed behaviour [37, 54]. GR is a well-established branch of AI with many real world applications such as assisting the elderly and disabled through recognition of their intent [16], daily living activity GR [18], recognition of fellow team members’ behaviours [23] and recognition of an adversarial agent’s intent [14]. While traditionally GR has been executed using plan libraries against which observed actions are matched [10], recent approaches accept a model of the domain as input and solve the GR problem using state-of-the-art planning [43, 44].

In this paper, we rely on model-based GR (*plan recognition as planning*) as defined in [44]. A *GR problem* \mathcal{P}_{GR} is formally defined as $\langle \mathcal{D}, S_I, \mathcal{G}, \text{Obs} \rangle$ where: $\mathcal{D} = \langle F, A \rangle$ and S_I are a planning domain and initial state, as above; \mathcal{G} is the set of possible goals, which includes the *real “hidden” goal* G_r (i.e., $G_r \in \mathcal{G}$),¹ and $\text{Obs} = [o_1, \dots, o_n]$ is a sequence of observations $o_i \in A$. Obs represents a valid (i.e., achievable, given the initial state S_I) but *partial* sequence of actions. That is, observable actions may be missing.

The solution to a GR problem \mathcal{P}_{GR} is identification of that goal $G \in \mathcal{G}$ which the GR system determines to be the *real goal* G_r . Contemporary GR techniques include planning and adapted heuristic functions [43, 44], standard planning graphs [13], top-K planning [50], landmarks [38, 39], learning and symbolic planning [2], linear programming [47] and mirroring [53, 55].

2.2 Deceptive Planning

Deceptive planning has been described as an inversion of the GR problem [6, 32]. While the aim of GR is to identify an agent’s goal by observing a sequence of its actions, the objective of a deceptive planner is to generate a sequence of actions such that an observer is *unable* to determine the agent’s goal.

Philosophers broadly agree that, from the deceiver’s point of view, *deception* involves the deliberate fostering or maintenance of false belief in the mind of the deceived [8]. Notions of belief and intentionality, however, are problematic in the context of machine behaviour, which has led some practitioners [e.g., 34, 48]) to prefer a definition based not on *intent* but on *behaviour*. An agent is deceptive “if its behaviour has *the potential to mislead*” [34, p.5].

Bell and Whaley’s general theory of deception [5] sets out two fundamental strategies. Every deception, they say, involves *dissimulation* (hiding the true) but it is by use of *simulation* (showing the false) that an observer can be made to believe something that is not true, a concept amplified in recent work on *extended goal*

recognition [33, 35]. In our work, we attempt to capture a similar intuition. While other planning paradigms typically treat ambiguity as deception [e.g., 12, 28, 34], we suggest that from the *observer’s* point of view, an ambiguous action—one that leaves the observer believing (correctly) that the real goal is one of a set of possible goals—is essentially truthful. Only when an action does *not* imply the real goal do we regard it as having the potential to mislead.

Much of the AI literature on deceptive planning focuses on path or motion-planning. In their study of deceptive robotic motion, Dragan et al. [12] propose three fundamental strategies: *exaggeration*, which over-emphasises actions that imply a false goal; *switching* between a false goal and the real goal, and *ambiguity*, the most economical of the methods but one which, under our definition, is not strictly deceptive at all. In the context of grid-based path-planning, Masters and Sardiña present multiple deceptive strategies, dependent for their implementation on the output of a probabilistic GR system [32]. They present measurements for deception at three levels of granularity: a *step, density*, a metric used to minimise the number of truthful steps in a path, and *extent* which—employing a concept of *path completion*—calculates how close to goal an agent can approach while remaining deceptive. In the context of path-planning, relative to a particular GR system, they identify a *radius of maximum probability*, within which deception cannot be achieved.

Approaching the problem of *Goal Obfuscation* in a *Classical Planning* setting but from a security perspective, Kulkarni et al. point out the limitations of dependence on any one particular GR system when the approach actually to be used by the observer is unknown. Their approach depends on ambiguity: an identical plan is generated no matter which goal is targeted, deviating only at the last possible moment to achieve their real intended objective [27].

Here, we demonstrate that, even without relying on ambiguity, it is possible to achieve *domain-independent* deceptive planning.

2.3 Landmarks

Landmarks are fluents (or actions) that must be achieved (or executed) at some point in all valid plans that achieve a goal state from an initial state [22]. *Landmarks* are often partially ordered based on the sequence in which they must be achieved. Given a planning problem $\mathcal{P} = \langle \mathcal{D}, S_I, G \rangle$, a formula ϕ_l is a landmark for \mathcal{P} iff ϕ_l is true at some point in all valid plans that achieve G from S_I .

Extracting landmarks and deciding their ordering has been shown to be PSPACE-complete [22], the same complexity as that for determining a plan’s existence [7]. Therefore, most landmark extraction algorithms [22, 26, 49] extract only a subset of landmarks for a given planning problem. Nevertheless, landmarks have been successfully employed to develop planners [45], planning heuristics [20], GR [38, 39], *counter-planning* [41], and action selection in *transparency planning* [30]. Here, we use landmarks to develop domain-independent approaches to deceptive planning. We denote a set of landmarks for a goal G as $\mathcal{L}(G)$. Figure 1a exemplifies the concept of landmarks in a GRID NAVIGATION problem: a domain where an agent can move between adjacent cells but may sometimes need to access a specific key to enter a restricted cell. In this example, from the initial cell (denoted by the robot) an agent must first pick-up the key, then open the lock, in order to achieve G .

By achieving landmarks associated with goals other than G_r , an observer can be made to believe that some *other* goal is real.

¹Note that the real goal G_r is **unknown** to the GR system in its role as the observer.




Figure 1: GRID NAVIGATION example with unit costs and 4-way navigation. (a) Key and padlock are landmarks required to achieve G ; (b) Relevant states (one example of each): C minimises average cost of reaching the goals, M minimises maximum cost, M_{\max} maximises minimum cost, and C_{\max} maximises the average cost; (c) Changes with G_r omitted.

2.4 Centroids and Minimum Covering States

Pozanco et al. [40, 42] introduced the concept of relevant states in planning: *centroids* and *covering states*. Figures 1b and 1c illustrate how they apply in the context of a GRID NAVIGATION problem.

Within a planning domain \mathcal{D} , a *centroid* C is a state that minimises the average cost of achieving each in the set of possible goals \mathcal{G} . Thus, referring to Figure 1b, the average cost to an agent of achieving the three goals G_0 , G_1 or G_2 is at its lowest $(1 + 3 + 5)/3$ when the agent is located at C . Similarly, a *max centroid* state C_{\max} maximises the average cost of achieving the set of goals \mathcal{G} . A *minimum covering state* M , meanwhile, is a state that minimises the maximum cost of achieving each goal in \mathcal{G} (in Figure 1b, M is 3 steps from each goal). Conversely, a *max minimum covering state* M_{\max} maximises the minimum cost of achieving each goal in \mathcal{G} . Note that many states may meet the criteria for a relevant state: Figure 1b provides one example of each.

While landmarks can be used to deceive (by directing the observer’s attention towards a false goal), centroids and minimum covering states can be used to confuse (by generating a condition where the observer does not know what to believe). The approach is related to that taken in [27]: by reference to any one of the relevant states, an identical initial plan would be generated no matter which goal the agent was targeting. Using relevant states, however, we can adapt the strategy to achieve a stronger deception by eliminating the real goal from the set on which the relevant states are based, as illustrated in Figure 1c. Under this approach, the observer is still confused but their attention is directed away from the real goal.

3 DOMAIN-INDEPENDENT DECEPTION

Let us first formally define the problem, which builds on [32] and the GR problem \mathcal{P}_{GR} set out in Section 2.1.

Definition 3.1. A domain-independent deceptive planning problem \mathcal{P}_D is a tuple $\langle \mathcal{D}, S_I, G_r, \mathcal{G} \rangle$ where:

- $\mathcal{D} = \langle F, A \rangle$ is a planning domain, F is a set of fluents, and A is a set of actions. Every action $a \in A$ has a positive cost: $\text{cost}(a)$;
- S_I is the initial state;
- G_r is the **real goal** that the agent aims to achieve and deceive from the observer;
- \mathcal{G} is the set of possible goals, including the real goal ($G_r \in \mathcal{G}$).

The *solution* to \mathcal{P}_D is a plan $\pi = [a_1, \dots, a_n]$ that maps S_I into a state $S \models G_r$, that is, a state in which the goal state G_r holds. The main purpose of deceptive planning is to generate a plan π (i.e., *deceptive behaviour*) that achieves the intended real goal G_r without revealing that intended goal to observers.

3.1 Quantifying Deceptiveness

The quantification of deception presents a dilemma: successful deception often leads away from the real goal but it must *achieve* its goal. To resolve the issue, we present a metric that considers both the *cost* and *quantity* of actions in the generated plan. Our formulation builds upon the concept of *truthful* and *deceptive steps* [32].

Masters and Sardiña [32] define a *truthful step* over path-planning as a step at which the probability of the real goal is greater than the probability of any other goal. We diverge from their original definition and rely instead on the *principle of rationality* to define a truthful action as one that *must* be included in an optimal plan that maps the *current* state into a state at which the real goal G_r holds.

To formalise this concept, given a plan $\pi = [a_1, a_2, a_3, \dots, a_n]$, we represent each action relative to the plan that incorporates it such that $\pi^1 = a_1, \pi^2 = a_2$, etc. On execution, each plan yields a trace, being an alternating sequence of states and actions $\sigma = [S_1, a_1, S_2, a_2, \dots, S_n]$. Now, let the set of *all* optimal plans that map the current state S into the real goal G_r be given by Π_{S, G_r}^* . With this in hand, a truthful action is given by Definition 3.2.

Definition 3.2. Given a known current state S , an action a is **truthful** iff $a = \pi^1$ for $\pi \in \Pi_{S, G_r}^*$. Otherwise, the action is **deceptive**.

Note that S in Definition 3.2 represents not the problem’s *initial* state but a plan’s *current* state. This requires that Π_{S, G_r}^* be *recomputed after every action*. If the first action in a plan is non-optimal but the recalculated plan from S_2 to the G_r is optimal, the plan has only a single deceptive action according to 3.2: if an action is part of an optimal plan to achieve the real goal *from the current state*, then that action is truthful. Thus, we use the principle of rationality to determine whether or not an action is deceptive.

One advantage of our definition is that it does not depend on observations, probabilities or the performance of any particular GR approach. Furthermore, whereas under the definition in [32], only when the real goal is unambiguously targeted is an action regarded as truthful, under our model, an action may be truthful simultaneously for *multiple* goals. That is, an action may increase ambiguity and confuse an observer but we argue it may be truthful nevertheless. If the observer can identify the real goal (whether singly or as one of a set) then they have not been deceived.

With these concepts properly in place, we can now specify a metric to capture how *much* of a given plan is deceptive. Under Definition 3.3, if **all** actions in the plan are deceptive, the deceptive *quantity* is 1; if none are deceptive, deceptive *quantity* is 0.

Definition 3.3. The **deceptive quantity** of a plan π is the ratio between the number of deceptive actions in a plan and the total number of actions in the plan. We denote the deceptive quantity of a plan π as $\mathbf{d}_{\text{Quantity}}(\pi)$, and formally define it as follows:

$$\mathbf{d}_{\text{Quantity}}(\pi) = \frac{|A_d|}{|\pi|}$$

where A_d is the set of all deceptive actions in π .

Observe that *deceptive quantity* performs a similar function to *deceptive density* in [32]. There, however, definitions are formulated relative to a deceiver wishing to minimise the number of truthful actions that may be observed. Our definition is relative to the observer and evaluates how often they may be deceived. Intuitively, the greater the *deceptive quantity* of a plan, the greater its “quality”; that is, the more likely it is to deceive.

We want to deceive but we also want to achieve our goal. As a counterbalance, therefore, we next present a metric to quantify the *deception cost* of a plan; that is, how much of a plan’s overall cost goes towards paying for its deceptive actions.

Definition 3.4. The **deception cost** of a plan π that maps S_I to G_r is a ratio between its additional cost (relative to an optimal plan $\pi^* \in \Pi^*(S_I, G_r)$) and the total cost of the plan. We denote the *deceptive cost* of a plan π as $d_{Cost}(\pi)$, and formally define it as follows:

$$d_{Cost}(\pi) = \frac{cost(\pi) - cost(\pi^*)}{cost(\pi)}$$

A plan with zero deceptive actions (i.e., where all actions occur in an optimal plan and are therefore truthful, c.f., Definition 3.2) has a deception cost of 0. If, however, a plan includes deceptive actions and costs twice as much as an optimal plan, then half its cost is paying for deception and its deception cost is, properly, 0.5.

Finally, by combining the metrics defined above, we can now provide a *score* that quantifies the *effective* deceptivity of a plan.

Definition 3.5. The **deception score** of a plan π is the ratio between its *deceptive quantity* and its *deception cost*. We denote the *deception score* of a plan π as $d_{Score}(\pi)$ and formalise it as follows:

$$d_{Score}(\pi) = \begin{cases} \frac{d_{Quantity}(\pi)}{d_{Cost}(\pi)} & cost(\pi) > cost(\pi^*) \\ 0 & cost(\pi) = cost(\pi^*) \end{cases}$$

The *deception score* increases with the number of deceptive actions that occur in π , but is always constrained by their cost. It tells us how much “bang”, if you will, we are getting for our buck and so provides a means of evaluating and comparing the efficacy of different deceptive approaches, which we set out next. Note that the special case in which if the cost of the plan with the deceptive actions and the cost of the optimal plan are the same, the deception score will be zero. Due to our definition of truthful steps, we would not expect this to happen, since even if a plan is ambiguous between goals, the actions will not be defined as deceptive as long as the plan does not deviate from the optimal plan.

3.2 Deceptive Approaches

We now present a baseline and three approaches to deceptive planning, each relying on a different concept for plan generation: *landmarks*, *centroids* and *minimum covering states*. We take as input a *deceptive planning problem* \mathcal{P}_D and, as output, return a plan π (i.e., a *deceptive behaviour*). To illustrate our approaches, we provide examples based on GRID NAVIGATION. In this domain, agents can move between (non-diagonal) adjacent cells but may need to collect keys to unlock padlocks to achieve certain restricted cells.

Simulation Deceptive Planning

Our *baseline* approach is SIMULATION, inspired by a widely used deceptive path-planning technique [25, 32] adapted here to work over non path-planning domains. Under this approach, a plan is




Figure 2: GRID NAVIGATION example to exemplify the landmark-based approaches. (a) Coloured keys can unlock padlocks of the same colour; (b) An optimal plan for G_r .

generated such that it appears to be targeting not G_r , but an alternative goal closest to the real goal in terms of cost.² This approach, as referenced in [32], is one of the two fundamental techniques described in Bell and Whaley’s general theory of deception which posits that deception can only be achieved by *simulation* (“showing the false”) or *dissimulation* (“hiding the true”) [5].

Landmarks-Based Deceptive Planning

We demonstrate our landmark-based approaches over the sample problem displayed in Figure 2a. In this example, some cells can only be achieved by using specific keys, e.g., to achieve the goal G_r from the initial state (indicated by the robot), the agent must first pick up the orange, blue, purple and red keys, and open the corresponding padlocks (in a specific order). Figure 2b shows an optimal plan π^* to achieve G_r from the initial state. The numbers on the cells correspond to the order in which the robot picks up keys and unlocks padlocks: first the red key, then the orange key, etc. Note that the robot does not need to pick up the green key or unlock the green padlock to achieve G_r . The cost of an optimal plan is $cost(\pi^*) = 24$, assuming a unified cost of 1 for every action.

In this example, the *Landmarks* are the locations of keys and, for each goal, this may constitute a partial or complete set of the existing keys: $\mathcal{L}(G_r) = [\text{red key, orange key, purple key, blue key}]$; $\mathcal{L}(G_1) = [\text{red key, orange key, purple key, green key}]$; $\mathcal{L}(G_2) = [\text{red key, orange key}]$.

Most Similar Landmarks Goal Approach. The first landmarks-based approach generates a deceptive plan to achieve an alternative goal that has the *most similar landmarks* to the real goal G_r . Once that goal has been achieved, an optimal plan to G_r is pursued. We denote this approach as $D\mathcal{L}_{Similar}$. We define the goal with the *most similar* landmarks as: $G_N = \max_{G_i \in \mathcal{G} \setminus \{G_r\}} \mathcal{L}(G_r) \cap \mathcal{L}(G_i)$. In landmark-based strategies, an inherent benefit of two goals being close together is that more landmarks are likely to be shared.

Figure 3a shows the plan adopted by an agent following the $D\mathcal{L}_{Similar}$ approach on the GRID NAVIGATION example. Since $\mathcal{L}(G_r)$ and $\mathcal{L}(G_1)$ share the most landmarks, $[\text{red key, orange key, purple key}]$, a plan is generated such that $\mathcal{L}(G_1)$ is achieved first, then $\mathcal{L}(G_r)$ is pursued (requiring the robot to backtrack in order to retrieve orange key).

Adopting our metrics, the cost of the plan $cost(\pi) = 48$ (again assuming a unified cost of 1 over all actions) and deceptive actions $|A_d| = 14$. Since the optimal plan had a cost of 24, the deception

²In navigational domains, it is convenient to think in terms of cost-distance. The closest goal in terms of cost-distance is the one that can be reached at the lowest cost.

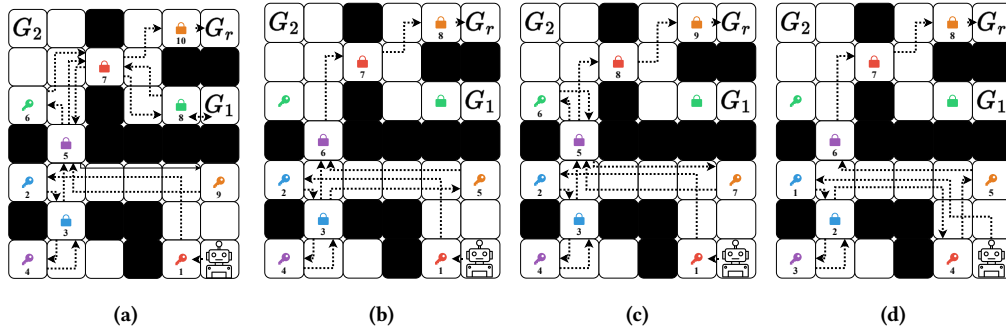


Figure 3: Examples of plans generated by our landmark-based deceptive approaches: (a) *Similar Landmarks* $D\mathcal{L}_{Similar}$; (b) *Shared landmarks* $D\mathcal{L}_{Shared}$; (c) *Combined Landmarks* $D\mathcal{L}_{Combined}$; (d) *Most Common Landmarks* $D\mathcal{L}_{MostCommon}$;

cost $d_{Cost}(\pi) = 0.5$, the deception quantity $d_{Quantity}(\pi) = 0.29$ and the plan’s overall deception score $d_{Score}(\pi) = 0.58$.

Shared Landmarks Approach. Our second landmarks-based approach again begins by finding the alternative goal that shares the most landmarks with the real goal but then deviates from the $D\mathcal{L}_{Similar}$ approach by creating a plan to achieve, not goal that has been identified, but rather the *shared landmarks* themselves. Only then does the agent continue to pursue the remaining landmarks of G_r . The shared landmarks are achieved according to their cost-distance from the initial state, i.e., closer shared landmarks are achieved first. This approach is denoted as $D\mathcal{L}_{Shared}$.

Figure 3b illustrates the plan adopted by an agent following the $D\mathcal{L}_{Shared}$ approach on the GRID NAVIGATION example. As shown, the agent achieves the landmarks shared by $\mathcal{L}(G_r)$ and $\mathcal{L}(G_1)$, $[\text{red}, \text{blue}, \text{purple}, \text{orange}]$, then the landmarks exclusive to $\mathcal{L}(G_r)$, $[\text{green}]$. Since shared landmarks are ordered by cost-distance from the initial state, the order of landmark achievement is $[\text{red}, \text{blue}, \text{purple}, \text{orange}, \text{green}]$. Once all landmarks have been achieved, the real goal is pursued optimally.

When analysing the deceptiveness of this plan, we find that the cost of the plan is lower, $cost(\pi) = 30$, and deceptive actions $|A_d| = 4$. The deception cost, calculated in relation to the cost of an optimal plan, is thus also lower, $d_{Cost}(\pi) = 0.2$. Deception quantity $d_{Quantity}(\pi) = 0.13$ (i.e., 13% of the plan is deceptive) and the overall deception score is higher: $d_{Score}(\pi) = 0.67$. We are getting more deception for less cost.

To amplify, in the $D\mathcal{L}_{Similar}$ approach, 50% of the cost goes towards achieving a plan that is 29% deceptive: the average deceptive step is 1.72% of the total plan cost. In the $D\mathcal{L}_{Shared}$ approach, 20% of the cost goes towards achieving a plan that is 13% deceptive: the average deceptive step costs only 1.54% of the total plan cost.

Combined Landmarks Approach. Here, we further exploit the power of landmarks. In this case, we divide them into three categories: *shared*, *closest exclusive* and *real exclusive*.

- (1) **Shared landmarks** \mathcal{L}_{\oplus} must be achieved for both the real goal G_r and an alternative goal G_N , being the goal with which G_r shares the most landmarks: $\mathcal{L}_{\oplus} = \mathcal{L}(G_r) \cap \mathcal{L}(G_N)$;
- (2) **Closest Exclusive Landmarks** $\mathcal{L}_{N\otimes}$ must be achieved for G_N but not G_r : $\mathcal{L}_{N\otimes} = \mathcal{L}(G_N) \setminus \mathcal{L}_{\oplus}$;
- (3) **Real Exclusive Landmarks**, $\mathcal{L}_{r\otimes}$ must be achieved for G_r but not G_N : $\mathcal{L}_{r\otimes} = \mathcal{L}(G_r) \setminus \mathcal{L}_{\oplus}$.

The order of achievement is $\mathcal{L}_{\oplus} \rightarrow \mathcal{L}_{N\otimes} \rightarrow \mathcal{L}_{r\otimes}$. By adopting this strategy, we aim to increase deceptiveness by achieving the landmarks common to most goals first and those unique to the real goal last. Within categories, landmarks are again ordered based on their cost-distance from the initial state. Figure 3c shows the plan adopted by an agent when following this approach. Over this domain, \mathcal{L}_{\oplus} is $[\text{red}, \text{blue}, \text{purple}, \text{orange}]$, $\mathcal{L}_{N\otimes}$ is $[\text{green}]$ and $\mathcal{L}_{r\otimes}$ is $[\text{yellow}]$. The order in which landmarks are to be achieved is $[\text{red}, \text{blue}, \text{purple}, \text{orange}, \text{green}, \text{yellow}]$, based first on the order of categories, then on the ordering of each landmark’s cost-distance from the initial state.

When analysing the deceptiveness of the plan under this approach, which we denote $D\mathcal{L}_{Combined}$, we find that $cost(\pi) = 36$ and deceptive actions $|A_d| = 7$. The deception cost $d_{Cost}(\pi) = 0.33$ and deception quantity is $d_{Quantity}(\pi) = 0.19$. This gives us an overall score of $d_{Score}(\pi) = 0.58$, the same as the $D\mathcal{L}_{Similar}$, slightly less effective than our shared landmarks approach $D\mathcal{L}_{Shared}$.

Most Common Landmarks Approach. Our last landmark-based approach relies on identifying the most landmarks common to *all* possible goals \mathcal{G} . Under this approach, denoted $D\mathcal{L}_{MostCommon}$, we achieve the landmarks of the real goal G_r ordered in such a way that the most common landmarks among all candidate goals are achieved first. Where landmarks have the same commonality, the landmark closer to the initial state is achieved first.

Figure 3d shows the plan taken by an agent when it follows the $D\mathcal{L}_{MostCommon}$ approach. Here, $\mathcal{L}(G_r)$ is $[\text{red}, \text{blue}, \text{purple}, \text{orange}, \text{yellow}]$. We first analyse the landmarks to identify their commonalities across all possible goals (including the real goal).

- Landmark red is common to G_r , G_1 and G_2 (3 goals);
- Landmark blue is common to G_r and G_1 (2 goals);
- Landmarks purple is common to G_r , G_1 and G_2 (3 goals); and
- Landmark orange is common to G_r (1 goal).

Under $D\mathcal{L}_{MostCommon}$, a plan is generated to achieve red and purple first, since they are required by the most goals. Their sub-ordering is determined by their cost-distance from the initial state. The plan then achieves blue and finally orange before the real goal is pursued optimally. Here, plan cost $cost(\pi) = 34$ and deceptive actions $|A_d| = 6$. Thus, the deception cost $d_{Cost}(\pi) = 0.29$ and deception quantity $d_{Quantity}(\pi) = 0.18$. This gives the $D\mathcal{L}_{MostCommon}$ approach our second highest overall deception score $d_{Score} = 0.6$.

Relevant States Based Deceptive Planning

We now present deceptive planning approaches which rely on centroids and minimum covering states [41]. We extract these *relevant states* using the searching algorithm proposed by Pozanco et al. [42], which searches through all reachable states by computing the cost of a plan to every possible goal by calling a planner. The relevant state (whether centroid or minimum covering state) is then found by looking for the best cost among all generated states by the searching algorithm.

Given a problem, \mathcal{P}_D , we extract centroids and minimum covering states with respect to three different set of goals:

- (1) The set of possible goals, \mathcal{G} , as originally defined in \mathcal{P}_D ;
- (2) A set of goals *excluding* the real goal, i.e., $\mathcal{G}_{NoG_r} = \mathcal{G} \setminus G_r$;
- (3) A set of goals comprising only the real goal G_r and the goal *closest* to the initial state S_I , that is, $\mathcal{G}_{ClosestG+G_r} = \min_{G \in \mathcal{G}} h(S_I, G) \cup G_r$, where h is a heuristic function that estimates the cost between two states. This combination enables us to get greater separation at minimal cost.

Centroid-Based Deceptive Planning. The three goal sets defined above allow us to extract four different *centroid* states using the same extraction algorithm, without any modification. These states are detailed below and are illustrated on the GRID NAVIGATION example depicted in Figure 4a:

- A *centroid* C for \mathcal{G} (pink cell \square);
- A *max centroid* C_{\max} for \mathcal{G} (dashed pink cell \square);
- A *centroid* C for \mathcal{G}_{NoG_r} (ignoring the real goal G_r), denoted as $C_{NoG_r \in \mathcal{G}}$ (light-blue cell \square);
- A *centroid* C for $\mathcal{G}_{ClosestG+G_r}$ (considering the closest goal from S_I), denoted as $C_{Closest}$ (dashed light-blue cell \square);

Our *centroid-based* approaches generate deceptive behaviours by computing an optimal plan π^* from the initial state S_I to one of the centroid states mentioned above, and another optimal plan π^* from such a state to the real goal G_r . As a result, we have four different approaches that rely on *centroids*, as follows: DC , DC_{\max} , $DC_{NoG_r \in \mathcal{G}}$, and $DC_{Closest}$. The deceptive behaviours generated by centroid-based approaches are plans that tend to approach most of the possible goals, i.e., plans that tend to have actions that aim to *minimise* or *maximise* the distance to the possible goals.

Figure 4a shows deceptive plans for the four centroid-based approaches. In this example, the cost of an optimal plan to achieve G_r is $cost(\pi^*) = 6$ (assuming a unified cost of 1). When analysing the deceptiveness of the plans for the four different centroid-based approaches, we have the following values for our deceptive metrics:

- DC : $|A_d| = 1$, $d_{Quantity} = 0.2$, $d_{Cost} = 0.25$, $d_{Score} = 0.8$;
- DC_{\max} : $|A_d| = 0$, $d_{Quantity} = 0$, $d_{Cost} = 0$, $d_{Score} = 0$;
- $DC_{NoG_r \in \mathcal{G}}$: $|A_d| = 1$, $d_{Quantity} = 0.2$, $d_{Cost} = 0.25$, $d_{Score} = 0.8$;
- $DC_{Closest}$: $|A_d| = 1$, $d_{Quantity} = 0.2$, $d_{Cost} = 0.25$, $d_{Score} = 0.8$;

Note that DC , $DC_{NoG_r \in \mathcal{G}}$, and $DC_{Closest}$ have different deceptive plans but the same scores under our metrics. As for DC_{\max} , the results are 0 for all metrics because the C_{\max} for this problem is an intermediate state along an *optimal plan* to achieve G_r . Recall that, under our model, a *confusing* plan is not strictly a *deceptive* plan: our DC_{\max} strategy results in a plan from S_I that passes through C_{\max} to the G_r , an optimal plan from S_I to G_r : it disguises the real goal without eliminating it from an observer’s consideration.




Figure 4: GRID NAVIGATION example. (a) Centroids deceptive plans; (b) Minimum Covering states deceptive plans.

Minimum Covering States Based Approaches. We extract different *minimum covering states* using the goal sets defined above:

- A *minimum covering state* M for \mathcal{G} ;
- A *max minimum covering state* M_{\max} for \mathcal{G} ;
- A *minimum covering state* M for \mathcal{G}_{NoG_r} , denoted $M_{NoG_r \in \mathcal{G}}$;
- A *minimum covering state* M for $\mathcal{G}_{ClosestG+G_r}$, denoted as $M_{Closest}$;

These *minimum covering states* are depicted in Figure 4b as: M using \mathcal{G} (brown cell \square); M_{\max} using \mathcal{G} (dashed brown cell \square); $M_{NoG_r \in \mathcal{G}}$ using $\mathcal{G}_{ClosestG+G_r}$ (yellow cell \square); $M_{Closest}$ using $\mathcal{G}_{ClosestG+G_r}$ (dashed yellow cell \square).

Similar to the centroids-based approaches, our minimum covering states-based approaches compute an optimal plan π^* from the initial state S_I to one of the minimum covering states set out above, and another optimal plan π^* from that state to the real goal G_r . Thus, we have four different minimum covering states-based strategies, as follows: DM , DM_{\max} , $DM_{NoG_r \in \mathcal{G}}$, and $DM_{Closest}$. These minimum covering states-based approaches generate plans that tend to make the observer’s recognition task more difficult due to the fact that the plans achieve intermediate states that are not fully optimal for *any* of the possible goals.

Figure 4b shows deceptive plans for the four minimum covering states-based approaches. The cost of an optimal plan to achieve G_r for this problem is $cost(\pi^*) = 6$ (assuming a unified cost of 1). When analysing the deceptiveness of the plans for these four approaches under our metrics:

- DM : $|A_d| = 0$, $d_{Quantity} = 0$, $d_{Cost} = 0$, $d_{Score} = 0$;
- DM_{\max} : $|A_d| = 0$, $d_{Quantity} = 0$, $d_{Cost} = 0$, $d_{Score} = 0$;
- $DM_{NoG_r \in \mathcal{G}}$: $|A_d| = 2$, $d_{Quantity} = 0.2$, $d_{Cost} = 0.4$, $d_{Score} = 0.5$;
- $DM_{Closest}$: $|A_d| = 0$, $d_{Quantity} = 0$, $d_{Cost} = 0$, $d_{Score} = 0$;

As before, the zero scores under our metrics occur when relevant states are part of an optimal plan for G_r . This implies that the plans may confuse but will not strictly deceive an observer. Note, however, that the example domain is grid-based which means there are very many alternative optimal plans from S_I to G_r . The domain is useful for illustration purposes (i.e., to demonstrate how each strategy works) but is not necessarily an ideal domain within which to demonstrate each strategy’s effectiveness.

4 EXPERIMENTS AND EVALUATION

We now set out our methodology and results from the experimental evaluation of the deceptive strategies enumerated above.³

³GitHub <https://github.com/AdrianPrice/Domain-Independent-Deception>

4.1 Domains, Benchmarks, and Planning Setup

We evaluated our approaches over three widely different domains to demonstrate the generality and breadth of their applicability; the `BLOCKS WORLD` domain: a domain that consists of a set of blocks with letters, a table, and a robot hand, where the goal is to find a plan that achieves a final configuration of blocks that corresponds to a certain word; the `LOGISTICS` domain, a domain that contains airports, trucks and air-planes, where the goal is to access and transport packages between locations; and finally, the `GRID NAVIGATION` domain, an adapted path-planning domain with additional properties (keys and padlocks) and actions, as used to illustrate our approaches throughout the paper.

Most existing planning algorithms rely on search-based problem solving using heuristic functions [4], e.g., `Fast-Downward` [19], one of the most well-known planners, a planner that employs different searching algorithms and a variety of planning heuristics.

For each domain, we constructed 10 deceptive planning problems of increasing difficulty. Each problem had 3 possible goals \mathcal{G} , including the real goal $G_r \in \mathcal{G}$.

We computed optimal plans for all approaches using the `A*` search algorithm from the `Pyperplan` planning library [1] with an admissible heuristic (i.e., `LMCUT` heuristic [20]) and extracted landmarks using the algorithm proposed by Hoffman et al. [22], also included in `Pyperplan`. For extracting centroids and *minimum covering states*, we used the sub-optimal search algorithm proposed by Pozanco et al. [42], which uses the `Fast-Downward` [19] planner with an inadmissible heuristic (`FAST-FORWARD` heuristic [21]).

4.2 Goal Recognition of Deceptive Plans

We evaluated the deceptiveness of our strategies not only using our own metrics but also by testing their ability to confound two state-of-the-art *online*⁴ GR systems using: *Mirroring* (M), a model-based recognition approach inspired by the method of GR used by humans [52]; and *Mirroring with Landmarks* (M+L), an online recognition approach which combines *Mirroring* with a generalised notion of landmarks [55]. In this way, we assess both the impact of our strategies and the usefulness of our model. In particular, we wished to evaluate the stability and performance of our landmarks-based deception approaches against a landmarks-aware GR system.

4.3 Experimental Evaluation

The results of our experimental evaluation are presented in Table 1. The first row, π^* , presents the results obtained over the *optimal plan* (hence a deception score of 0) for G_r . The second row presents the results obtained over our baseline approach denoted as `SIMULATION`. The remaining rows present our different deception generation approaches in three categories: *landmarks-based*, *centroid-based*, and *minimum-covering states*. The columns show average results over each of evaluation metric, again in three categories each corresponding to a different domain: `BLOCKS WORLD`, `LOGISTICS`, and `GRID NAVIGATION`.

Though not designed as optimal algorithms, we evaluated the relative efficiencies (generation time and cost) of our approaches, in addition to the degree of deceptivity they provide. Deception is

measured by d_{Score} and by the extent to which our plans deceived the two *online* GR approaches, *Mirroring* and *Mirroring with Landmarks*. To evaluate performance against the GR systems, we used an established metric which measures the ratio of the *real intended goal* that was *Not Ranked First*, $\neg R1st$ [52]⁵, i.e., in how much of the plan were the recognition algorithms deceived in that the real goal was *not* ranked first (as the recognised one by the GR approach).

We begin by looking at the results obtained over the `GRID NAVIGATION` domain, which mostly corresponds to traditional path-planning domains, in which the goal is to navigate from one location to the other. In terms of overall deception score (d_{Score}), the best result obtained is a ratio of 0.5 effective plan deceptivity. This is obtained by the *baseline*, commonly used path-planning deceptive approach, `SIMULATION` but also by two of our landmarks-based approaches, `DLSimilar` and `DLCCombined`. In terms of goal recognition, by the *Mirroring* approach M, again `SIMULATION` achieved the highest performance, effectively generating the most deceptive plan. However, when using the state-of-the-art recognition approach M+L, which combines *Mirroring* and *landmarks* for recognition, both landmark-based approaches achieve greater deception.

We now analyse the results over the `BLOCKS WORLD` domain. In terms of overall deception score, d_{Score} , the results are similar to the previous domain. The best result obtained was a ratio of 0.52 effective plan deceptivity, again obtained by the approaches `DLSimilar` and `SIMULATION`. However, when examining the effective recognition of the *Mirroring* approach M, `DLSimilar` achieved the best performance over the M+L recognition approach.

Finally, we analyse the results obtained over the `LOGISTICS` domain, arguably the most complex, real-world domain, containing different properties such as airports, trucks and air planes, and the aim is to get and transport packages between locations. In such complex domains, there are typically fewer optimal plans than the simpler path-planning domains. Over these types of domains, our deceptive approaches shine with the best d_{Score} achieved by `DMClosest` (0.85), `DCClosest` (0.83) and `DC` (0.81). In terms of goal recognition, again the `DLSimilar` and `SIMULATION` approach achieved the highest performance, although the `DLSimilar` proved to be, in this domain, much more efficient in terms of generation time.

Figure 5 allows us to see a clear trade off between plan cost and deception. Each graph shows how for each domain, and each deceptiveness evaluation metric, the cost of the plan (x -axis), influences the deceptiveness (y -axis) of that plan. Each approach category has its own shape representing it, with squares representing the baseline `SIMULATION` approach and π^* , circles representing the *landmark-based* approaches, triangles for the *centroid* approaches, and stars for the *minimum covering state* approaches.

The three deceptiveness evaluation metrics used in Figure 5 are $d_{Quantity}$, $M\neg R1st$ and $M+L\neg R1st$. Across all domains and all deceptiveness evaluation metrics there is a clear positive trend, that is, the more expensive a plan is, the more deceptive it is. It is important to note though, when choosing a deceptive planning approach for a real world case, there is more that goes into that thought than just deceptive quality. In the real world, there are often cost constraints enforced so that plans, including deceptive plans,

⁴In *online goal recognition*, observations are provided incrementally, and the aim is to recognise the agent's intended goal G_r , as soon as possible.

⁵Vered and Kaminka [52] defines *Ranked First* as the number of times a GR approach ranked the the intended goal G_r , as the most likely intended goal, indicating the general accuracy of the approach.

	GRID NAVIGATION							BLOCKS WORLD							LOGISTICS						
	$Time(s)$	$cost(\pi)$	$d_{Quantity}$	d_{Cost}	d_{Score}	M $\neg\mathcal{R}1st$	M+ \mathcal{L} $\neg\mathcal{R}1st$	$Time(s)$	$cost(\pi)$	$d_{Quantity}$	d_{Cost}	d_{Score}	M $\neg\mathcal{R}1st$	M+ \mathcal{L} $\neg\mathcal{R}1st$	$Time(s)$	$cost(\pi)$	$d_{Quantity}$	d_{Cost}	d_{Score}	M $\neg\mathcal{R}1st$	M+ \mathcal{L} $\neg\mathcal{R}1st$
π^*	0.876	6.8	0	0	0	0.72	0.31	7.52e-3	8.0	0	0	0	0.54	0.33	0.58	13.7	0	0	0	0.44	0.40
SIMULATION	4.18	14.0	0.26	0.51	0.50	0.92	0.83	0.0561	17.7	0.27	0.55	0.52	0.87	0.79	2.64	21.3	0.27	0.36	0.72	0.90	0.91
$D\mathcal{L}_{Similar}$	4.46	15.6	0.26	0.56	0.50	0.88	0.85	0.38	21.9	0.31	0.63	0.52	0.89	0.79	1.38	22.3	0.29	0.39	0.72	0.90	0.91
$D\mathcal{L}_{Shared}$	4.09	6.8	0	0	0	0.72	0.31	0.39	13.0	0.18	0.38	0.45	0.61	0.40	0.57	17.5	0.14	0.21	0.70	0.64	0.40
$D\mathcal{L}_{Combined}$	5.29	15.8	0.26	0.57	0.50	0.88	0.84	0.38	25.6	0.33	0.69	0.50	0.80	0.67	0.97	27.9	0.37	0.50	0.72	0.81	0.86
$D\mathcal{L}_{MostCommon}$	4.39	6.8	0	0	0	0.72	0.31	0.37	17.2	0.23	0.53	0.45	0.67	0.37	0.63	17.8	0.15	0.23	0.69	0.45	0.40
$DC_{Closest}$	75.0	6.8	0	0	0	0.72	0.31	2.53	10.4	0.09	0.23	0.30	0.61	0.32	150.0	18.4	0.23	0.26	0.83	0.65	0.56
$DC_{NoG, \in \mathcal{G}}$	57.3	11.6	0.19	0.41	0.40	0.85	0.73	2.99	12.6	0.16	0.37	0.35	0.60	0.57	151.0	22.9	0.29	0.40	0.74	0.83	0.78
DC	9.53	10.8	0.16	0.37	0.30	0.79	0.69	2.68	10.8	0.15	0.26	0.35	0.60	0.42	124.0	16.7	0.12	0.18	0.81	0.48	0.41
DC_{max}	4.85	11.2	0.17	0.39	0.35	0.80	0.67	7.69	12.8	0.18	0.38	0.35	0.70	0.50	13.8	15.8	0.08	0.13	0.64	0.45	0.35
$DM_{Closest}$	17.7	8.2	0.06	0.17	0.10	0.85	0.75	7.75	12.0	0.14	0.33	0.35	0.65	0.41	65.9	16	0.13	0.15	0.85	0.69	0.54
$DM_{NoG, \in \mathcal{G}}$	15.0	10.8	0.18	0.37	0.40	0.79	0.56	11.7	11.2	0.14	0.29	0.35	0.60	0.47	56.5	16.1	0.10	0.15	0.74	0.45	0.34
DM	8.77	11.0	0.15	0.38	0.30	0.76	0.52	7.22	14.6	0.18	0.45	0.35	0.60	0.42	36.8	15.4	0.09	0.11	0.70	0.51	0.44
DM_{max}	5.79	11.8	0.16	0.42	0.35	0.79	0.63	13.3	14.4	0.19	0.44	0.35	0.72	0.47	35.4	15.2	0.08	0.10	0.53	0.51	0.44

Table 1: Deception and recognition results over all deceptive planning approaches. The columns show the generation $Time$ (in seconds) for each approach (including landmark/relevant states extraction times), average plan cost $cost(\pi)$, deception quantity $d_{Quantity}$, deception cost d_{Cost} , deception score d_{Score} , and the deceptive ratio ($\neg\mathcal{R}1st$) of the plan obtained with *Mirroring* (M) and *Mirroring+Landmarks* (M+ \mathcal{L}) recognition approaches.

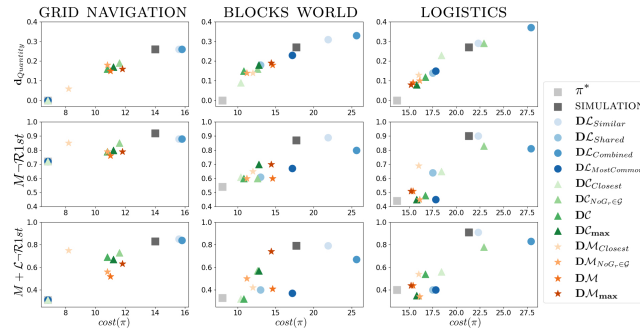


Figure 5: Plan cost plotted against $d_{Quantity}$, M $\neg\mathcal{R}1st$ and M+ \mathcal{L} $\neg\mathcal{R}1st$ values for all domains, showing a trend that the higher the plan cost, the higher the deception.

are completed (i.e. G_r is achieved) in a feasible amount of time. It can be clearly seen that while SIMULATION, $D\mathcal{L}_{Similar}$ and $D\mathcal{L}_{Combined}$ are often the most deceptive approaches, these approaches come at the cost of a very expensive plan. For deceptive plans involving hundreds or thousands of actions, it may make these approaches infeasible to be used for such large scale domains due to the massive overhead of the extra cost and relevant state approaches, *centroids* or *minimal covering states*, should be preferred.

Another important factor is the time taking to compute a valid plan. The $Time$ column in Table 1 shows the average time taken to generate the deceptive plans. This time includes, for *landmark-based* approaches, the extraction of landmarks, and for *centroid* and *minimal covering state* approaches, the extraction of the respective states. While all approaches are orders of magnitude slower than the time it takes to generate the optimal plan, *landmark-based* approaches are faster to generate when compared to the *centroid* and *minimal covering state* based approaches. This is due to the fact that for any planner (such as Pyperplan), computing a plan involves searching through an, often, extensive state-space until finding a node where all predicates from the goal state (\mathcal{G}) are true

[15]. When computing a plan for the *landmark-based* approaches, \mathcal{G} is often only a single predicate long, making the task of traversing the state-space simplistic in terms of time. While that simple search may need to be repeated multiple times (scaling with the amount of landmarks extracted), compared with the *centroid* and *minimal covering state* approaches, that not only take a sizeable amount of time to extract, but also require traversing deep into a state-space to satisfy a \mathcal{G} which may contain many predicates, *landmark-based* approaches are much more scalable for larger planning problems.

5 CONCLUSIONS

In this paper, we have presented two novel approaches to the domain-independent deceptive planning problem and have suggested multiple potential strategies for each, all capable of achieving high levels of deception under different conditions.

We introduced new metrics to evaluate a plan’s deceptivity, quantifying its effectiveness (i.e., ability to deceive) from the point of view of a human observer, as a ratio between the number of deceptive actions that occur during the course of a plan and their cost. We showed the applicability of our approaches over three planning domains and evaluated their efficiency and performance using plan cost and generation time as well as our proposed metrics and the ability of our plans to deceive two state-of-the-art GR systems.

Our results highlight the advantages of each deception approach against domain complexity. While for path-planning and grid navigation domains, the SIMULATION approach performs well, our *landmarks-based* approaches perform on par and have an advantage over more complex domains such as BLOCKS WORLD and LOGISTICS. While for the most complex, LOGISTICS, our *centroid* and *minimal covering states* based approaches performed best.

As future work, we intend to use a recent algorithm [24] for extracting centroids and minimum covering states, as well as using and adapting operator counting [11], and reconsidering the role of confusion and goal obfuscation techniques [3] in the generation of domain-independent deceptive plans.

ACKNOWLEDGMENTS

Peta Masters' research is supported by the UKRI Trustworthy Autonomous Systems (TAS) Hub (EP/V00784X/1).

REFERENCES

- [1] Yusra Alkhazraji, Matthias Frorath, Markus Grützner, Malte Helmert, Thomas Liebraut, Robert Mattmüller, Manuela Ortlieb, Jendrik Seipp, Tobias Springenberg, Philip Stahl, and Jan Wülfing. 2020. Pyperplan. <https://doi.org/10.5281/zenodo.3700819>. <https://doi.org/10.5281/zenodo.3700819>
- [2] Leonardo Rosa Amado, Ramon Fraga Pereira, and Felipe Meneguzzi. 2021. Combining LSTMs and Symbolic Approaches for Robust Plan Recognition. In *AAMAS*.
- [3] Sara Bernardini, Fabio Fagnani, and Santiago Franco. 2020. An Optimization Approach to Robust Goal Obfuscation. In *KR*.
- [4] Blai Bonet and Héctor Geffner. 2001. Planning as heuristic search. *Artificial Intelligence* 129, 1 (2001), 5–33.
- [5] J Barton Bowyer. 1982. Cheating: Deception in war & magic, games & sports, sex & religion, business & con games, politics & espionage, art & science. (1982).
- [6] Sviatoslav Braynov. 2006. Adversarial planning and plan recognition: Two sides of the same coin. In *Secure Knowledge Management Workshop*, Vol. 3. Citeseer, 67–70.
- [7] Tom Bylander. 1994. The Computational Complexity of Propositional STRIPS Planning. *Journal of Artificial Intelligence Research (JAIR)* 69 (1994), 165–204.
- [8] Thomas L Carson. 2010. *Lying and deception: Theory and practice*. Oxford University Press.
- [9] Tathagata Chakraborti, Anagha Kulkarni, Sarath Sreedharan, David E Smith, and Subbarao Kambhampati. 2019. Explicability? legibility? predictability? transparency? privacy? security? the emerging landscape of interpretable agent behavior. In *Proceedings of the international conference on automated planning and scheduling*, Vol. 29. 86–96.
- [10] Eugene Charniak and Robert P Goldman. 1993. A Bayesian Model of Plan Recognition. *Artificial Intelligence* 64, 1 (1993), 53–79.
- [11] Toby O. Davies, Adrian R. Pearce, Peter J. Stuckey, and Nir Lipovetzky. 2015. Sequencing Operator Counts. In *ICAPS*.
- [12] Anca D. Dragan, Rachel M. Holladay, and Siddhartha S. Srinivasa. 2014. An Analysis of Deceptive Robot Motion. In *Robotics: Science and Systems*.
- [13] Yolanda E.-Martín, María D. R.-Moreno, and David E. Smith. 2015. A fast goal recognition technique based on interaction estimates. In *IJCAI*.
- [14] Grady Fitzpatrick, Nir Lipovetzky, Michael Papisimeon, Miquel Ramirez, and Mor Vered. 2021. Behaviour Recognition with Kinodynamic Planning over Continuous Domains. *Frontiers in Artificial Intelligence* 4 (2021), 717003.
- [15] Hector Geffner and Blai Bonet. 2013. *A Concise Introduction to Models and Methods for Automated Planning*. Morgan & Claypool Publishers.
- [16] Christopher W Geib. 2002. Problems with Intent Recognition for Elder Care. In *Proceedings of the AAAI-02 Workshop "Automation as Caregiver"*. 13–17.
- [17] Malik Ghallab, Dana S. Nau, and Paolo Traverso. 2004. *Automated Planning - Theory and Practice*. (1st ed.). Elsevier. 635p.
- [18] Roger Leitzke Granada, Ramon Fraga Pereira, Juarez Monteiro, Rodrigo Coelho Barros, Duncan Ruiz, and Felipe Meneguzzi. 2017. Hybrid Activity and Plan Recognition for Video Streams. In *Workshops at the Thirty-First AAAI Conference on Artificial Intelligence*.
- [19] Malte Helmert. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research* 26 (2006), 191–246. <https://www.fast-downward.org>
- [20] Malte Helmert and Carmel Domshlak. 2009. Landmarks, Critical Paths and Abstractions: What's the Difference Anyway?. In *ICAPS*.
- [21] J. Hoffmann and B. Nebel. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research* 14 (may 2001), 253–302. <https://doi.org/10.1613/jair.855>
- [22] Jörg Hoffmann, Julie Porteous, and Laura Sebastia. 2004. Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research* 22, 1 (Nov 2004), 215–278.
- [23] Gal Kaminka, Mor Vered, and Noa Agmon. 2018. Plan Recognition in Continuous Domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [24] Erez Karpas. 2022. A Compilation Based Approach to Finding Centroids and Minimum Covering States in Planning. In *ICAPS*, Akshat Kumar, Sylvie Thiébaux, Pradeep Varakantham, and William Yeoh (Eds.).
- [25] Sarah Keren, Avigdor Gal, and Erez Karpas. 2015. Goal recognition design for non-optimal agents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29.
- [26] Emil Keyder, Silvia Richter, and Malte Helmert. 2010. Sound and Complete Landmarks for And/Or Graphs. In *ECAI*.
- [27] Anagha Kulkarni, Matthew Klenk, Shantanu Rane, and Hamed Soroush. 2018. Resource bounded secure goal obfuscation. In *AAAI Fall Symposium on Integrating Planning, Diagnosis and Causal Reasoning*.
- [28] Anagha Kulkarni, Siddharth Srivastava, and Subbarao Kambhampati. 2019. A unified framework for planning in adversarial and cooperative environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 2479–2487.
- [29] Zhengshang Liu, Yue Yang, Tim Miller, and Peta Masters. 2021. Deceptive Reinforcement Learning for Privacy-Preserving Planning. In *AAMAS*.
- [30] Aleck M. MacNally, Nir Lipovetzky, Miquel Ramirez, and Adrian R. Pearce. 2018. Action Selection for Transparent Planning. In *AAMAS*.
- [31] Peta Masters, Michael Kirley, and Wally Smith. 2021. Extended Goal Recognition: A Planning-Based Model for Strategic Deception. In *AAMAS*.
- [32] Peta Masters and Sebastian Sardiña. 2017. Deceptive Path-Planning. In *IJCAI*.
- [33] Peta Masters, Wally Smith, and Michael Kirley. 2021. Extended goal recognition: Lessons from magic. *Frontiers in Artificial Intelligence* 4 (2021).
- [34] Peta Masters, Wally Smith, Liz Sonenberg, and Michael Kirley. 2020. Characterising Deception in AI: A Survey. In *Deceptive AI*. Springer, 3–16.
- [35] Peta Masters and Mor Vered. 2021. What's the Context? Implicit and Explicit Assumptions in Model-Based Goal Recognition.. In *IJCAI*. 4516–4523.
- [36] John McCarthy, Marvin L Minsky, Nathaniel Rochester, and Claude E Shannon. 2018. A proposal for the Dartmouth summer research project on artificial intelligence (1955). *Reprinted online at http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html* (2018).
- [37] Felipe Rech Meneguzzi and Ramon Fraga Pereira. 2021. A Survey on Goal Recognition as Planning. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021, Canada*.
- [38] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. 2017. Landmark-Based Heuristics for Goal Recognition. In *AAAI*.
- [39] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. 2020. Landmark-Based Approaches for Goal Recognition as Planning. *Artificial Intelligence* 279 (Feb 2020).
- [40] Alberto Pozanco. 2020. GRS Readme. https://github.com/apozanco/GRS_0.1/blob/master/README.md
- [41] Alberto Pozanco, Yolanda E.-Martín, Susana Fernández, and Daniel Borrajo. 2018. Counterplanning using Goal Recognition and Landmarks. In *IJCAI*.
- [42] Alberto Pozanco, Yolanda E.-Martín, Susana Fernández, and Daniel Borrajo. 2019. Finding Centroids and Minimum Covering States in Planning. In *ICAPS*.
- [43] Miquel Ramirez and Hector Geffner. 2009. Plan Recognition as Planning. In *IJCAI*.
- [44] Miguel Ramirez and Hector Geffner. 2010. Probabilistic Plan Recognition using off-the-shelf Classical Planners. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.
- [45] Silvia Richter and Matthias Westphal. 2010. The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks. *Journal of Artificial Intelligence Research* 39 (2010), 127–177.
- [46] Stuart Russell and Peter Norvig. 2005. AI a Modern Approach. *Learning* 2, 3 (2005), 4.
- [47] Luisa R de A Santos, Felipe Meneguzzi, Ramon Fraga Pereira, and Andre Pereira. 2021. An LP-Based Approach for Goal Recognition as Planning. In *AAAI*.
- [48] Amanda Sharkey and Noel Sharkey. 2020. We need to talk about deception in social robotics! *Ethics and Information Technology* (2020), 1–8.
- [49] Malte Helmert e Matthias Westphal Silvia Richter. 2008. Landmarks Revisited. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (2008)*.
- [50] Shirin Sohrabi, Anton V. Riabov, and Octavian Udrea. 2016. Plan Recognition as Planning Revisited. In *IJCAI*.
- [51] Alan Mathison Turing. 1950. *Mind* 59, 236 (1950), 433–460.
- [52] Mor Vered and Gal A Kaminka. 2017. Heuristic Online Goal Recognition in Continuous Domains. *arXiv preprint arXiv:1709.09839* (2017).
- [53] Mor Vered, Gal A Kaminka, and Sivan Biham. 2016. Online Goal Recognition through Mirroring: Humans and Agents. In *The Fourth Annual Conference on Advances in Cognitive Systems*, Vol. 4.
- [54] Mor Vered, Reuth Mirsky, Ramon Fraga Pereira, and Felipe Meneguzzi. 2022. Advances in Goal, Plan and Activity Recognition. *Frontiers in Artificial Intelligence* 5 (2022).
- [55] Mor Vered, Ramon Fraga Pereira, Gal Kaminka, and Felipe Rech Meneguzzi. 2018. Towards Online Goal Recognition combining Goal Mirroring and Landmarks. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, 2018, Suécia*.
- [56] Aldert Vrij. 2008. *Detecting lies and deceit: Pitfalls and opportunities*. John Wiley & Sons.