

# Stochastic Goal Recognition Design

Christabel Wayllace

Washington University in St. Louis  
cwayllace@wustl.edu

## Abstract

Discovering the objective of an agent based on observations of its behavior is a problem that has interested both AI and psychology researchers for many years (Schmidt, Sridharan, and Goodson 1978; Kautz 1987). In AI, this problem is known as *goal recognition* or, more generally, *plan recognition*, and it has been used to model a number of applications ranging from software personal assistants (Oh et al. 2010; 2011a; 2011b) to security applications that recognize the plan or goal of terrorists (Jarvis, Lunt, and Myers 2005).

Traditionally, research in goal recognition focuses on finding efficient algorithms to infer the agent’s objective from online collected observations, however, a newly formulated problem takes another perspective. The problem, proposed by Keren, Gal, and Karpas (2014), is called *goal recognition design* (GRD) and is intended to reduce the complexity of the online goal recognition task by performing an offline modification of the underlying environment that the agent operates in. GRD aims to find the smallest set of modifications that *forces the agent to reveal its goal as early as possible*. This problem finds itself relevant in many of the same applications of goal recognition because, usually, the underlying environment can be easily modified.

Typically, a GRD problem has two components: (1) a measure of the efficacy of goal recognition and (2) a model of possible design changes one can make to the environment. In the seminal work by (Keren, Gal, and Karpas 2014), they proposed the *worst case distinctiveness* (*wcd*) metric as a measure of the goal recognition efficacy and *action removal* as design change. The *wcd* aims at capturing the maximum number of steps an agent can take without revealing its goal. Therefore, the objective is to find the set of removed actions that will cause a minimal *wcd*. The definition of a GRD problem is made under three key assumptions: (1) the agents in the system will act optimally (i.e., agents will move along a shortest path to its goal);(2) the actions of the agents are deterministic; and (3) the environment is fully observable (i.e., both states and actions of the agent are observable).

Since then, researchers have generalized the GRD problem to relax the first and third assumptions (Keren, Gal, and

	Generalizations			Metrics		Possible Designs	
	Subopt. Plans	Partially Obs. Env.	Stochastic Actions	<i>wcd</i>	<i>ecd</i>	Action Remov.	Sensor Refin.
Keren, Gal, and Karpas (2014)				✓		✓	
Son et al. (2016)				✓		✓	
Keren, Gal, and Karpas (2015)	✓			✓		✓	
Keren, Gal, and Karpas (2016a)	✓	✓		✓		✓	
Keren, Gal, and Karpas (2016b)	✓	✓		✓		✓	✓
Wayllace et al. (2016)			✓	✓		✓	
Wayllace, Hou, and Yeoh (2017)			✓	✓	✓	✓	
Wayllace et al. (2018)		✓	✓	✓	✓	✓	✓

Table 1: Properties of Current Goal Recognition Design Models

Karpas 2015; 2016a; 2016b) and have also proposed alternative algorithms to solve it (Son et al. 2016). Additionally, Keren, Gal, and Karpas (2016b) have proposed the refinement of sensors, which decreases the degree of observation uncertainty of agent actions, as a possible design change on the environment. Table 1 summarizes the generalizations, metrics, and possible designs of existing GRD models.

Our work aims to further extend the GRD problem with the objective to take into account the stochasticity and limitations present in the physical world, where agents do not control the outcomes of their actions, the observer is limited by the quantity and quality of the sensors, and / or where the agent is not optimal.

With this objective in mind we have proposed the *Stochastic GRD* (S-GRD) problem, where the outcomes of the agent’s actions are stochastic (Wayllace et al. 2016). Aside from this relaxation, we have proposed a new metric called *expected case distinctiveness* (*ecd*) (Wayllace, Hou, and Yeoh 2017), which weighs the possible goals based on their likelihood of being the true goal.

The framework has also been extended with the assumption that agent actions are no longer observable and agent states are only partially observable. This relaxation was motivated by practical applications such as agent navigation, where agent actions are unobservable yet its state (current location) can be (at least partially) observed. The partial observability of agent states is due to low sensor resolution (e.g., GPS) potentially causing that several nearby states to be indistinguishable from one another. Finally, we considered a new type of sensor refinement as a possible design to improve the observability of some states (Wayllace et al. 2018). Building on this last work, we submitted a paper that includes pruning of the search space for sensor refinement

and the generalization of the *ecd* metric for partially observable S-GRDs.

S-GRDs are modeled as a tuple  $T = \langle D, \mathbf{G} \rangle$  where the domain  $D$  uses MDP elements and  $\mathbf{G}$  represents the set of possible goals. Computing the different metrics requires using VI-like algorithms to solve an *augmented MDP* (Wayllace, Hou, and Yeoh 2017)—states and actions are coupled with the set of possible goals. The reason of using this type of MDP is that the set of possible goals for a particular state depends on the entire history of states visited to arrive to that state.

Partial observability requires the addition of a sensor model  $N$  represented as a partition of the set of states, such that each set in the partition contains states projecting the same observation. Sensor refinement is, therefore, modeled as a partition refinement.

Future work includes considering suboptimal agents, defining a new metric using entropy, and incrementing the number of possible starting states. We are also looking into modeling a cybersecurity application using our framework and currently we are using a realistic cybersecurity problem generator to evaluate our algorithms. The inclusion of suboptimality will allow us to model deceptive agents which is a more realistic assumption in this specific application.

## Research Plan

- **Progress to Date:** (1) Generalization of GRDs to S-GRD problems which handle stochastic action outcomes (Wayllace et al. 2016). (2) Generalization of S-GRDs to *Partially Observable* S-GRD (POS-GRD) problems (Wayllace et al. 2018). (3) Definition of a new metric (the *expected case distinctiveness* (Wayllace, Hou, and Yeoh 2017)) and generalization of the *wcd* (Wayllace et al. 2016; Wayllace, Hou, and Yeoh 2017) for S-GRDs. (4) Definition of a new type of design change for POS-GRDs: *sensor refinement* (Wayllace et al. 2018).
- **Spring 2019:** Generalize the S-GRD model to support non-optimal agents. By February 2019 I plan to have the algorithm design and a basic implementation.
- **Summer 2019:** (1) Generalize the model to multiple starting states. (2) Propose a new metric using entropy.
- **Fall 2019 - Spring 2019:** Write Dissertation.

## Personal Contribution

**Goal Recognition Design with Stochastic Agent Action Outcomes** (Wayllace et al. 2016): Brainstormed ideas for the model and algorithms, wrote the pseudocode, implemented algorithms, ran experiments, wrote mathematical proofs.

**New Metrics and Algorithms for Stochastic Goal Recognition Design Problems** (Wayllace, Hou, and Yeoh 2017): Brainstormed ideas for algorithms, wrote mathematical models and proofs, implemented algorithms, ran experiments, and wrote the first draft version.

**Accounting for Partial Observability in Stochastic Goal Recognition Design** (Wayllace et al. 2018): Brainstormed ideas for algorithms, wrote the mathematical definitions and

models, wrote pseudocode, implemented algorithms, ran experiments, wrote the drafts, and addressed corrections based on reviews.

## References

- Jarvis, P.; Lunt, T.; and Myers, K. 2005. Identifying terrorist activity with AI plan recognition technology. *AI Magazine* 26(3):73–81.
- Kautz, H. A. 1987. *A Formal Theory of Plan Recognition*. Ph.D. Dissertation, Bell Laboratories.
- Keren, S.; Gal, A.; and Karpas, E. 2014. Goal recognition design. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, 154–162.
- Keren, S.; Gal, A.; and Karpas, E. 2015. Goal recognition design for non-optimal agents. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 3298–3304.
- Keren, S.; Gal, A.; and Karpas, E. 2016a. Goal recognition design with non-observable actions. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 3152–3158.
- Keren, S.; Gal, A.; and Karpas, E. 2016b. Privacy preserving plans in partially observable environments. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 3170–3176.
- Oh, J.; Meneguzzi, F.; Sycara, K.; and Norman, T. 2010. AN-TIPA: An agent architecture for intelligent information assistance. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, 1055–1056.
- Oh, J.; Meneguzzi, F.; Sycara, K.; and Norman, T. 2011a. An agent architecture for prognostic reasoning assistance. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2513–2518.
- Oh, J.; Meneguzzi, F.; Sycara, K.; and Norman, T. 2011b. Probabilistic plan recognition for intelligent information agents: Towards proactive software assistant agents. In *Proceedings of the International Conference on Agents and Artificial Intelligence (ICAART)*, 281–287.
- Schmidt, C.; Sridharan, N.; and Goodson, J. 1978. The plan recognition problem: An intersection of psychology and artificial intelligence. *Artificial Intelligence* 11(1–2):45–83.
- Son, T. C.; Sabuncu, O.; Schulz-Hanke, C.; Schaub, T.; and Yeoh, W. 2016. Solving goal recognition design using ASP. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 3181–3187.
- Wayllace, C.; Hou, P.; Yeoh, W.; and Son, T. C. 2016. Goal recognition design with stochastic agent action outcomes. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 3279–3285.
- Wayllace, C.; Keren, S.; Yeoh, W.; Gal, A.; and Karpas, E. 2018. Accounting for partial observability in stochastic goal recognition design: Messing with the marauder’s map. *HSDIP 2018* 33.
- Wayllace, C.; Hou, P.; and Yeoh, W. 2017. New metrics and algorithms for stochastic goal recognition design problems. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 4455–4462.