Privacy Preserving Plans in Partially Observable Environments

Using Goal Recognition Design for Improved Privacy

Sarah Keren Avigdor Gal Erez Karpas

Faculty of Industrial Engineering and Management
Technion — Israel Institute of Technology

July 2016

Background : Goal Recognition Design

Offline design as a way to facilitate online goal recognition

Worst case distinctiveness (wcd) as a measure of model quality

Background : Goal Recognition Design

Offline design as a way to facilitate online goal recognition

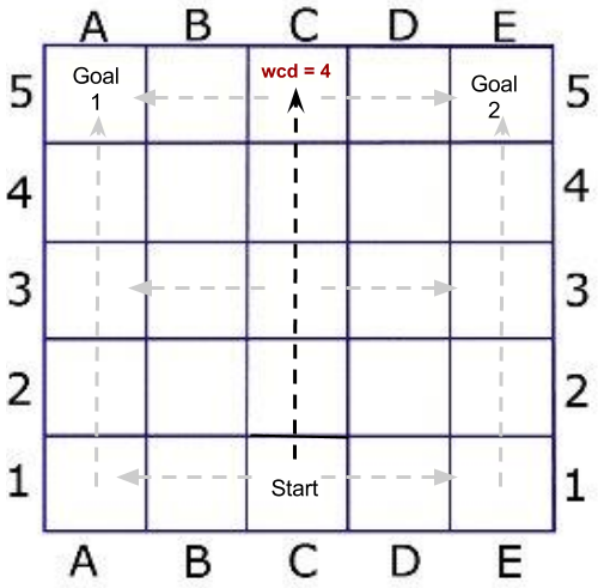

Worst case distinctiveness (wcd) as a measure of model quality

Background : Goal Recognition Design

Offline design as a way to facilitate **online** goal recognition

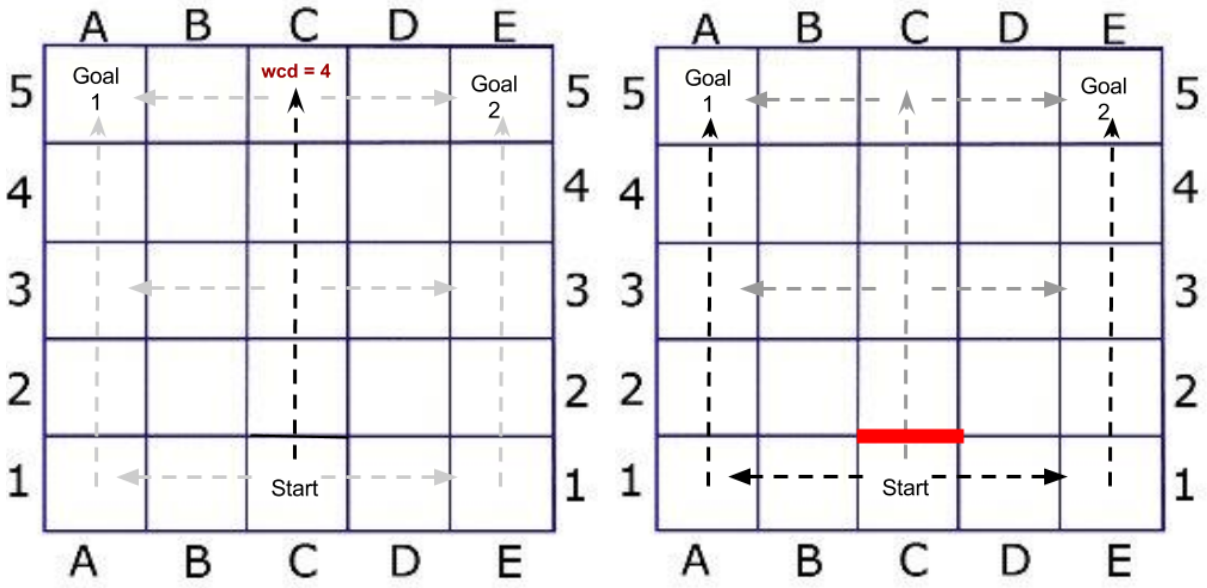

Worst case distinctiveness (wcd) as a measure of model quality

Offline design as a way to facilitate online goal recognition

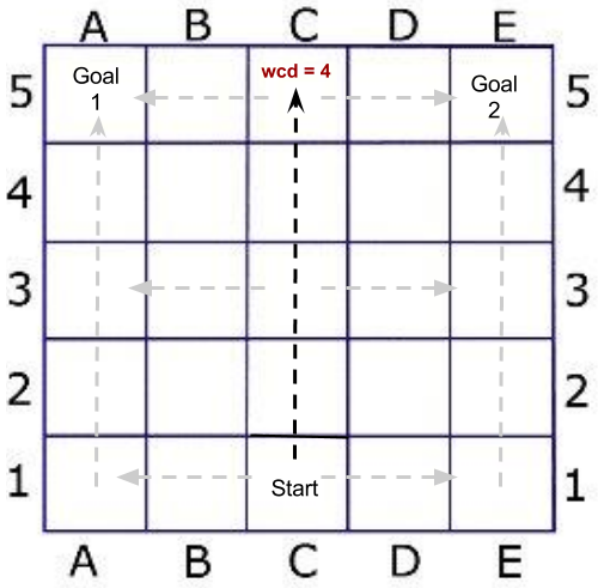

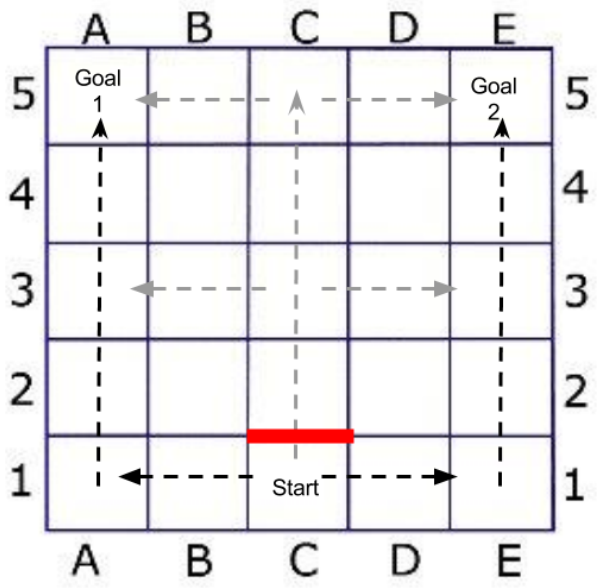

Worst case distinctiveness (wcd) as a measure of model quality

Background : Goal Recognition vs. Goal Recognition Design

Common to both

- STRIPS-like model:
  - Fluents F
  - Actions A with $a = \langle pre(a), add(a), del(a) \rangle$
  - Initial state $s_0 \subseteq F$
  - Set of possible goals $\mathcal{G}$
  - (Optional) sensor model which maps actions A to observation tokens

Background : Goal Recognition vs. Goal Recognition Design

Common to both

- STRIPS-like model:
  - Fluents F
  - Actions A with $a = \langle pre(a), add(a), del(a) \rangle$
  - Initial state $s_0 \subseteq F$
  - Set of possible goals $\mathcal{G}$
  - (Optional) sensor model which maps actions A to observation tokens

Goal Recognition — online

- Given a set of observations, what are the possible goals?
  - Generalizes plan libraries (Ramirez and Geffner, 2009 $\rightarrow$ 2016)

Background : Goal Recognition vs. Goal Recognition Design

Common to both

- STRIPS-like model:
  - Fluents F
  - Actions A with $a = \langle pre(a), add(a), del(a) \rangle$
  - Initial state $s_0 \subseteq F$
  - Set of possible goals $\mathcal{G}$
  - (Optional) sensor model which maps actions A to observation tokens

Goal Recognition — online

- Given a set of observations, what are the possible goals?
  - Generalizes plan libraries (Ramirez and Geffner, 2009 $\rightarrow$ 2016)

Goal Recognition Design — offline

- WCD: What is the maximum number of steps an agent can take before his goal is revealed?
- Reduce WCD: How can we modify the model to reduce WCD?

Extending the Goal Recognition Design Framework

Deterministic Environment

Compilation to classical planning (Keren et. al.):

- Optimal fully observable agents (ICAPS 014)
- Sub-Optimal fully observable agents (AAAI 2015)
- Some Actions are Non-observable (AAAI 2016)
- Arbitrary sensor model (IJCAI 2016) — right now

Compilation to ASP:

- Fully observable agents (Son et. al., AAAI 2016)

Stochastic Environment

Solution using MDP:

- Fully observable agents (Wayllace et. al., IJCAI 2016) — in 30 min.

Example

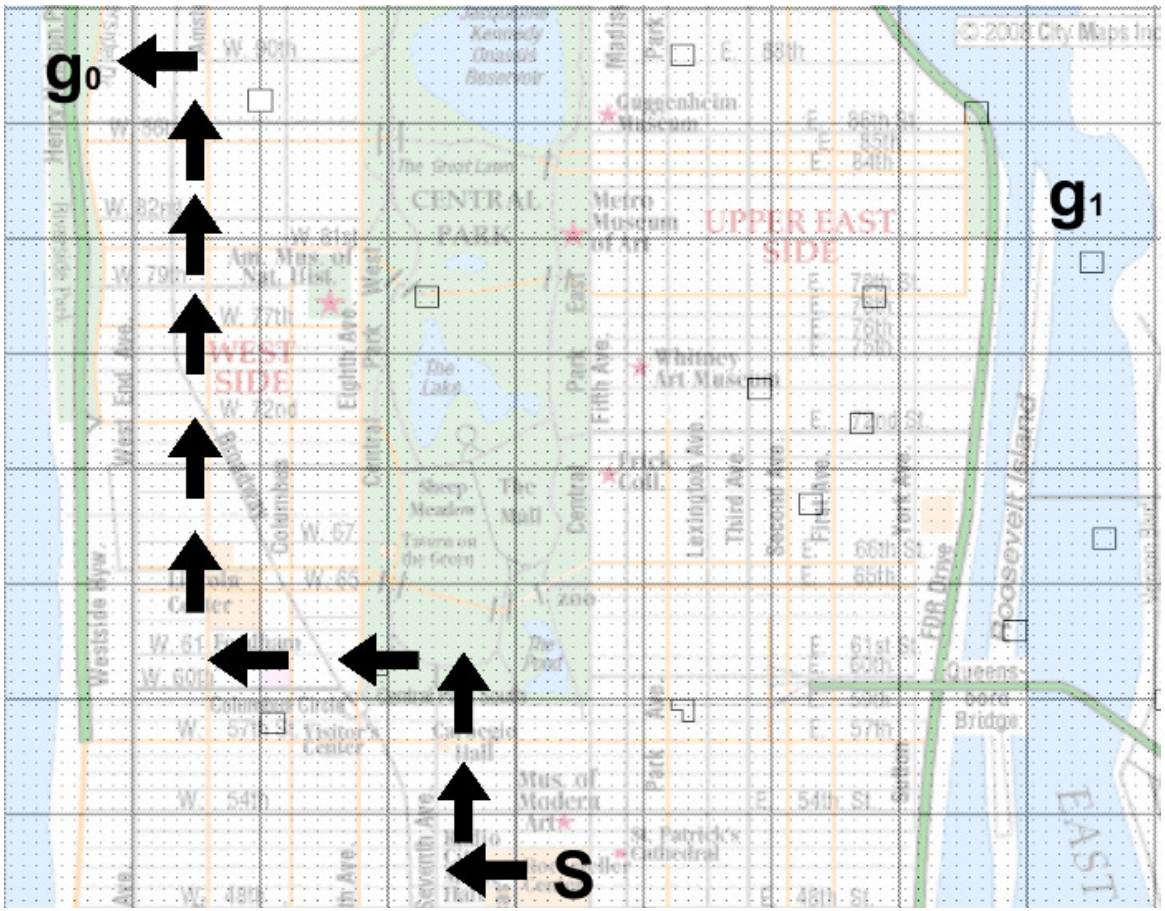

Full Observability

Example

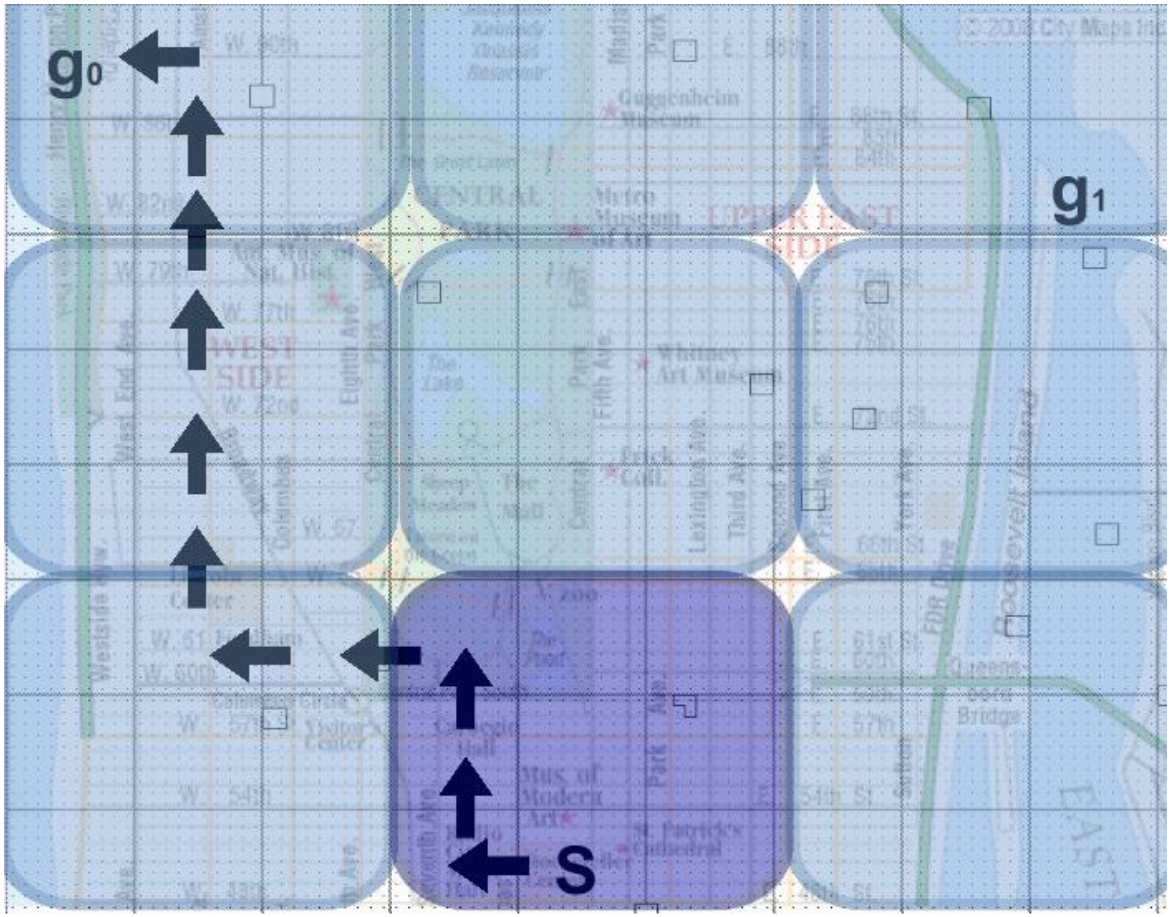

Coarse Sensors

Example

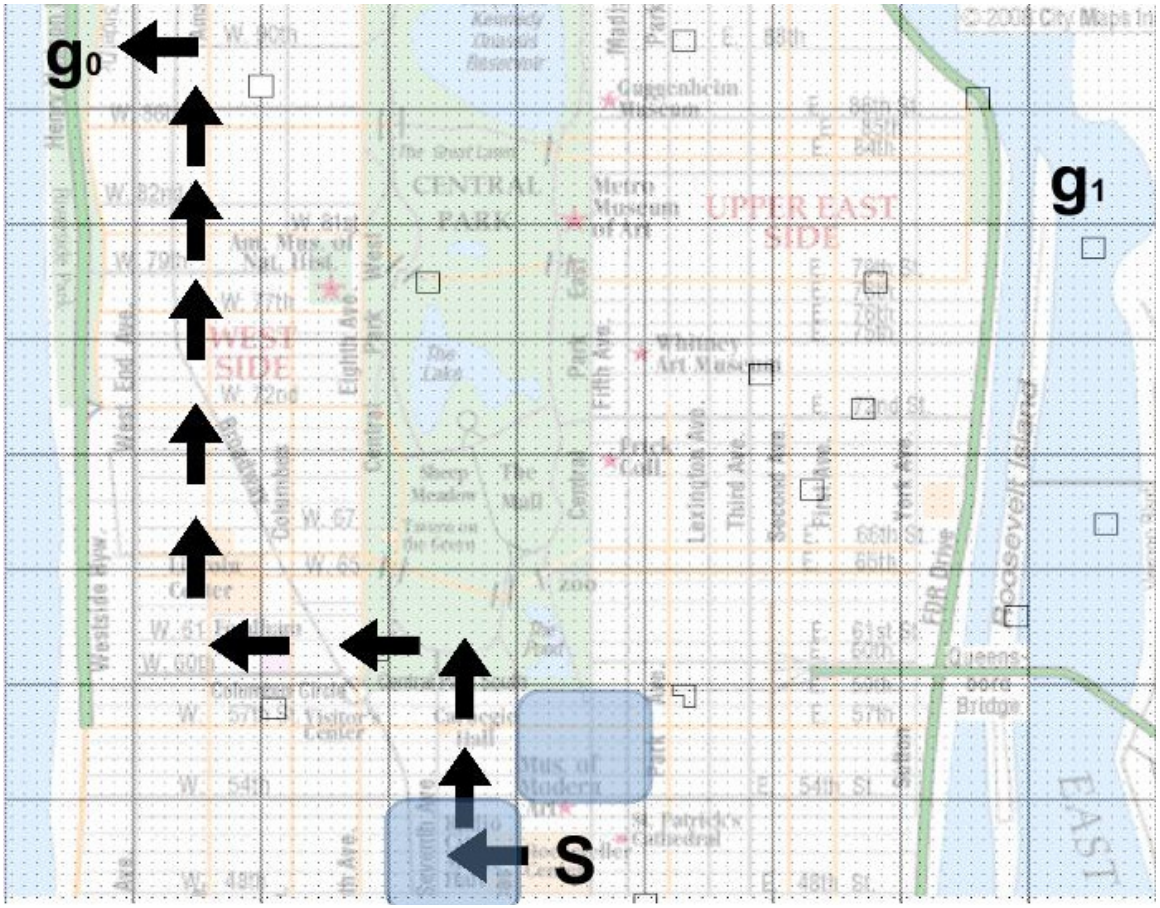

Noisy Sensors

Using Goal Recognition Design for Improved Privacy

Cloaking : How long can an agent keep his goal ambiguous ?

A user can choose a path that potentially maximizes its privacy

the wcd-path that allows him to stay ambiguous for at most wcd steps

Goal Recognition Design with Arbitrary Sensor Models

Sensor Model

Maps each action to a set of possible observation tokens.
The special token $o_\emptyset$ denotes non-observable action.

Goal Recognition Design with Arbitrary Sensor Models

Sensor Model

Maps each action to a set of possible observation tokens.
The special token $o_\emptyset$ denotes non-observable action.

Observable Projection

The observable projection of a path is a set of possible observation sequences, determined by the sensor model.

Goal Recognition Design with Arbitrary Sensor Models

Sensor Model

Maps each action to a set of possible observation tokens.
The special token $o_\emptyset$ denotes non-observable action.

Observable Projection

The observable projection of a path is a set of possible observation sequences, determined by the sensor model.

Non-distinctive Path

A path is non-distinctive if it has an observable projection, which is also the observable projection of a path leading to a different goal.

Goal Recognition Design with Arbitrary Sensor Models

Sensor Model

Maps each action to a set of possible observation tokens.
The special token $o_\emptyset$ denotes non-observable action.

Observable Projection

The observable projection of a path is a set of possible observation sequences, determined by the sensor model.

Non-distinctive Path

A path is non-distinctive if it has an observable projection, which is also the observable projection of a path leading to a different goal.

Worst Case Distinctiveness

The worst case distinctivenss (wcd) is the maximal non-distinctive path .

Calculating wcd: Compilation to Classical Planning

- Given GRD problem with two goals, we create classical planning problem with two agents each aiming at a separate goal
- Actions divided into
  - 'real' actions: change the state of the world
  - 'declare' actions: declare the observation token a 'real' action emits
- As long as both agents have declared the same observation sequence, they can get a discount when they declare the same observation token

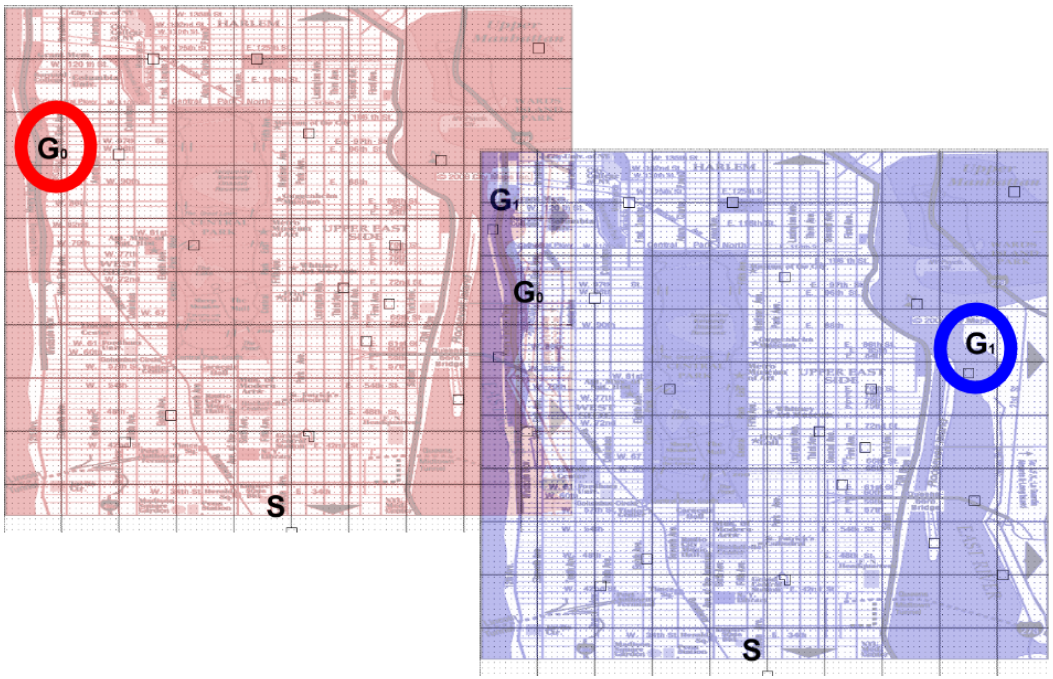

Empirical Evaluation : wcd

	LOGISTICS					BLOCKS WORLD					GRID-NAVIGATION			
	FULL	NO	POD-Obj	POD-Ac	POND	FULL	NO	POD-Obj	POD-Ac	POND	FULL	NO	POD	POND
wcd	1	1.2	1.2	13	13	5.3	6.1	6.1	8.5	8.5	2.8	3.02	3.09	3.18
time(LS)	2.85	—	—	—	—	4.9	—	—	—	—	0.3	—	—	—
time(LE)	35.1	83.75	—	—	—	72.4	74.1	—	—	—	0.3	0.24	—	—
time(CD)	263.8	107.1	94.7	117.3	397.3	82	103.3	96.1	113.2	373.5	0.63	0.64	0.48	1.33
% CD	0.8	0.9	0.9	0.85	0.7	1.0	1.0	1.0	1.0	0.75	1.0	1.0	1.0	1.0

Table 1: *wcd* Values, Running Time, and Coverage Ratio

- Measure effect non-deterministic partially observable sensor models have on the wcd value of a model and the efficiency of wcd calculation using the compilation.
- For each setting we manually created 5 sensor models : Fully observable (FULL), Non observable actions (NO), two versions of Partially observable deterministic (POD) and Partially observable non-deterministic (POND)
- For all domains, wcd increases with the decrease of observability and increase of uncertainty

Summary

- Extended Goal Recognition Design to handle arbitrary sensor models
- Allows us to find plans for privacy preserving agents
- Code and benchmarks available on our website:
  `http://ie.technion.ac.il/~sarahn/grd`

Thank You!