Real-time goal recognition using approximations in Euclidean space

Douglas Tesch[b], Leonardo Amado[a], Felipe Meneguzzi[a]

[a]University of Aberdeen

[b]Pontifical Catholic University of Rio Grande do Sul

✉ douglas.tesch@edu.pucrs.br

Scan me for full paper

Motivation

- Very few online goal recognition approaches can work in discrete and continuous domains.
- Online goal recognition approaches often rely on repeated calls to a planner at each new observation, incurring high computational costs.
- Recognizing goals online in continuous space quickly and reliably is critical for any real-world trajectory planning problem (such as any application in robotics) since the real physical world is fast-moving.

Continuous and Discrete Domain Benchmarks

- The continuous domain experiments use simulated environments built on 29 benchmark scenarios from Moving-AI's [2] Starcraft maps[1].
- We use an openly available goal and plan recognition dataset [1] for the discrete domain experiments.

Online Goal Recognition

Given an initial state I and a set of goal hypotheses $g_n, n \in [1, \ldots, \mathcal{N}]$. Our approach searches for a trajectory $m_I^{g_n}$ that maximizes the probabilities of a sequence of observations belonging to the same goal.

$$m_I^{g_n R} = \underset{m_I^{g_n} \in M}{\text{argmax}}\, P(m_I^{g_n} \mid O)$$

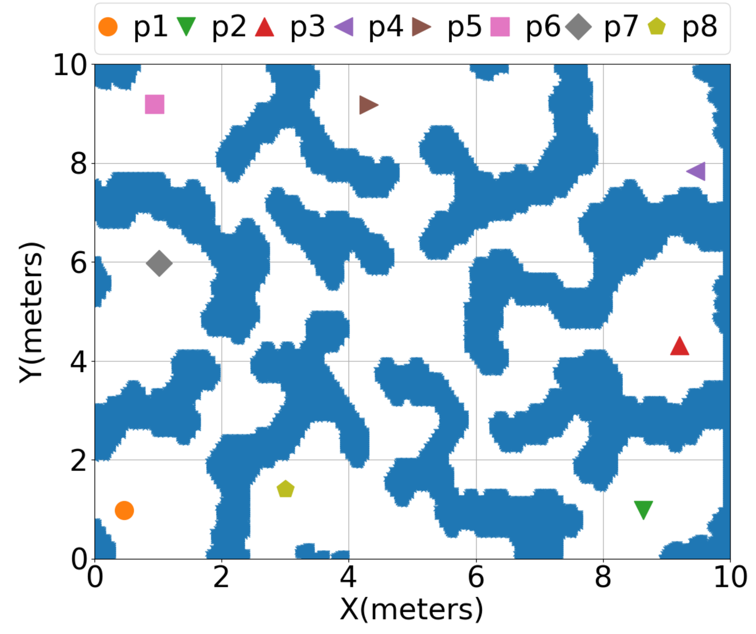

Figure 1: Starcraft's BigGameHunters map. Marks represent potential goal hypotheses positions.

Methodology

- We build a sequence of Cartesian points on X-Y axes (viapoints) representing a trajectory using RRT* (Rapidly Exploring Random Tree).

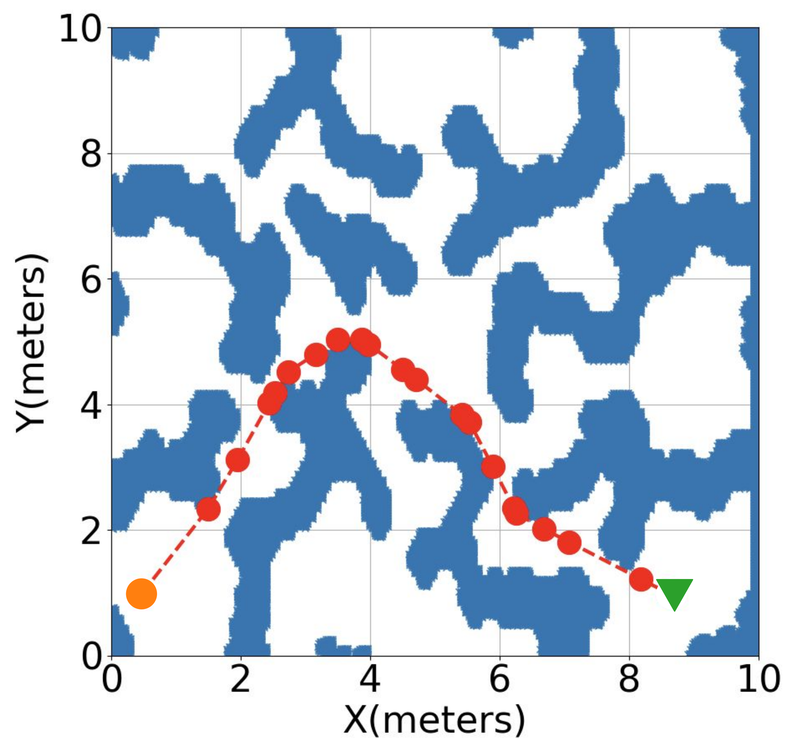

Figure 2: RRT^* output example with initial state at p_1 and goal states at g_2. Circles represent viapoints

- We use a 5th-degree polynomial model to link the viapoints and use Reinforcement Learning to obtain the model parameters.

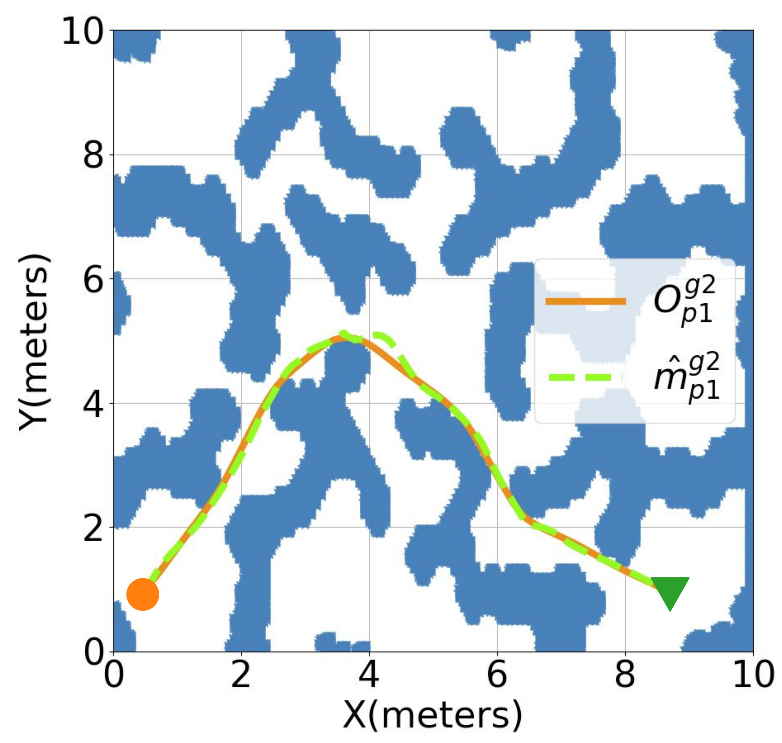

Figure 3: Contrasting approximate $\hat{m}^{g_2}_{p_1}$ and optimal $O^{g_2}_{p_1}$ trajectories.

- We use the Euclidean distance between the observations and the polynomial model to compute the probabilities of each goal hypothesis.

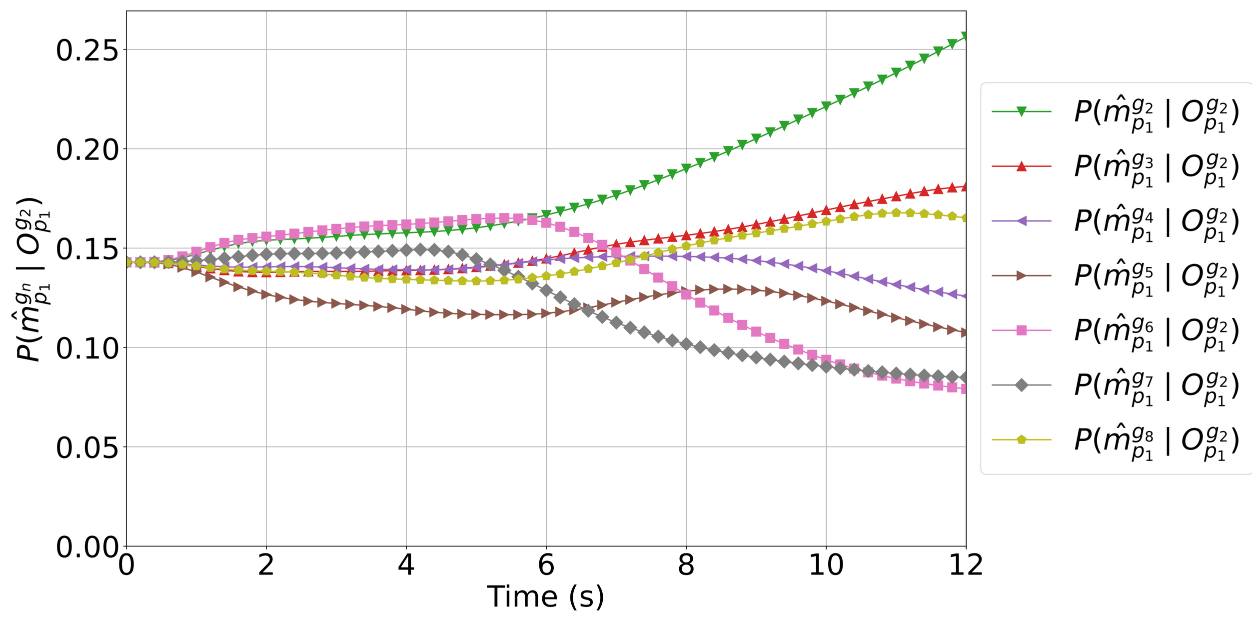

Figure 4: Conditional Probabilities $P(\hat{m}^{g_n}_{p_1} \mid O^{g_2}_{p_1})$ of all goals. Real goal hypothesis as g_2.

- We extend the approach to discrete domains by computing the euclidean distance directly from the STRIPS representation using

$$dist = \sqrt{\mid (o - m_I^{g_n}) \cup (m_I^{g_n} - o)) \mid}$$

Experiments Results and Conclusion

- Figure 5 shows our main results for continuous domains. The comparison against SoTA indicates that our approach has a superior PPV (Positive Predictive Value) when using TOP $k \geq 10$ solutions.

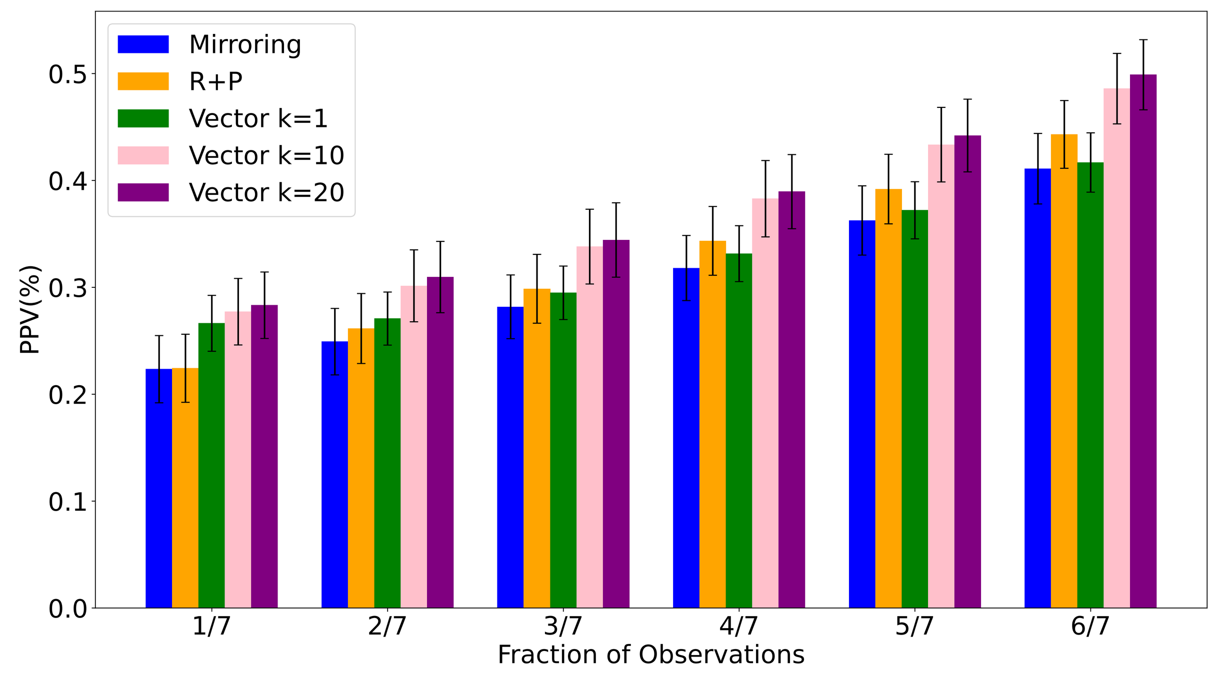

Figure 5: PPV percent and margin of error comparison from continuous domains over observations.

- Figure 6 shows our main results for discrete domains. The comparison against SoTA indicates that our approach performs similarly but with an inferior margin of error when using TOP $k \geq 20$ solutions.

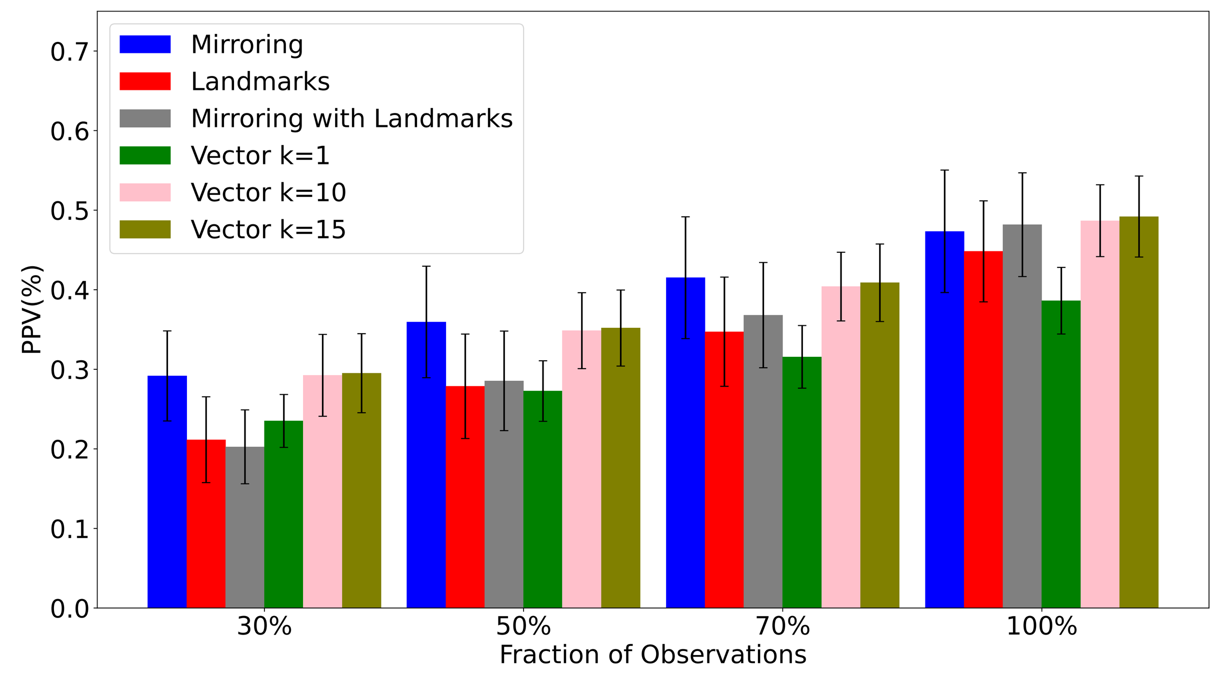

Figure 6: PPV percent and margin of error comparison from discrete domains over observations.

- Key contributions:
  - An efficient method of real-time goal recognition for continuous and discrete domains.
  - The method relies on a single call to the planner for each possible goal and uses a 5th-degree polynomial approximation at inference time for continuous domains to avoid costly optimal plan solutions.
  - Adding multiple solutions in the inference improves the overall recognition quality.

References

[1] Ramon Pereira, Nir Oren, and Felipe Meneguzzi. Landmark-based heuristics for goal recognition. In *Thirty-First AAAI Conference on Artificial Intelligence*, volume 31, 2017.

[2] Nathan R Sturtevant. Benchmarks for grid-based pathfinding. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(2):144–148, 2012.