A Framework to Infer Movement Planning from Observed Trajectories using Inverse Planning

Paulina Friemann (friemanp@cs.uni-freiburg.de)
Department of Computer Science
University of Freiburg, Germany

Joschka Boedecker
Neurorobotics Lab, Department of Computer Science
University of Freiburg, Germany

Andrew D. Straw
Institute of Biology I and Bernstein Center Freiburg
University of Freiburg, Germany

Navigation is an integral part of cognition, as it enables us to move around in an environment in a goal-directed manner. A plethora of algorithms have been developed for navigation and path-finding. From an outside perspective, it can be difficult to determine an agent's goals, or the algorithms and parameters that underlie its movement. This however is crucial to research fields in which no verbal access to the object of study is possible. As an example, understanding the underlying algorithms of insect navigation has increasingly gained attention from fields such as mobile robotics, due to the necessity of insects to use methods with high efficiency and low memory costs. However, interpreting insect navigation from a high-level standpoint is still a challenge.

Taking inspiration from a computational model on human goal inferences proposed by Baker, Saxe, and Tenenbaum (2009), we developed a framework based on Bayesian Inverse Planning to infer goals, algorithms and latent parameters from movement trajectories of a single agent. Whereas the original computational model focused only on the inference of goals given a trajectory, we extended the methodology to make inferences over any internal parameter or method, including the used planning algorithms. Our framework allows to be extended with different methods for planning, and uses these methods as grounds for comparison against movement trajectories.

Movement trajectories are either taken from experimental data of natural agents, such as insects, or sampled from the policies generated by the planning methods and parameters. The framework then uses Bayesian Inference to calculate the numerical, relative likelihoods of the different planning methods for a given movement trajectory.

As a proof of concept, we used the planning algorithm employed in Baker et al. (2009), which consists of an optimal planning approach based on value iteration, a method of Reinforcement Learning, extended by a determinism factor β, to introduce a stochastic noise into the trajectories. Furthermore, we embedded a greedy planning approach, likewise extended by a determinism factor β. To show the performance, we conducted two experiments in which we used randomly selected planning models, consisting of a planning type (greedy and optimal) and a value of β, to sample movement trajectories towards a goal in a simulated environment with fixed obstacles, depicted in Figure 1.

In each iteration, the sampled movement trajectories were compared against the policies produced by each planning model, and used as grounds for calculating the a posteriori likelihoods of the models. Analyses show that the determinism factor β, the planning method, and the combination are distinguishable using our framework. A graphic of this can be found in Figure 2.

The framework is extendable to include any variable or method that can influence a path, and to calculate their respective likelihoods. It allows numerical assessments of factors, models and goals, without having to make assumptions about others. Our

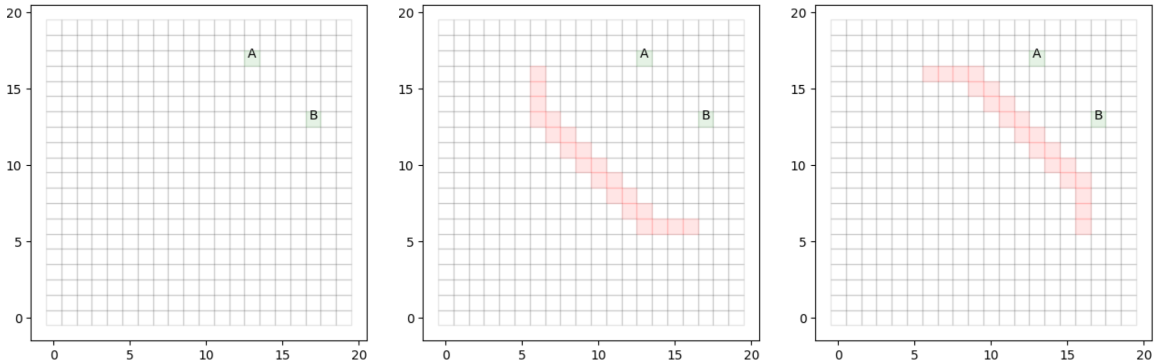

(a) Exp. 1. Starting positions and goal positions are sampled in each iteration.

(b) Exp. 2a. Convex obstacle between start and goals.

(c) Exp. 2b. Concave obstacle between start and goals.

Figure 1: Visualizations of the experimental area with two goals (A and B, whereas A is the goal which the agent pursues). Obstacles are depicted in red. In Experiments 2a and 2b, the starting position is the lower left corner.

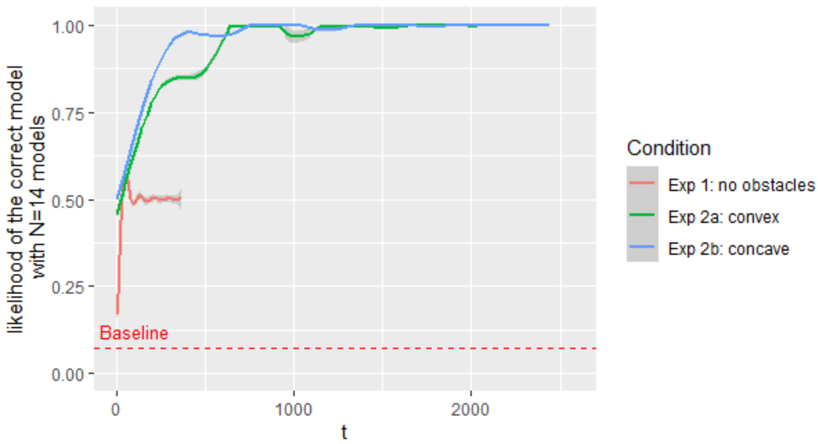

Figure 2: Average performance of the correct model over the time steps of trajectories. Depicted is the generalized additive model of the average likelihood. Note that in the case of no obstacles in Experiment 1, greedy and optimal planning methods produce the same policies. The baseline is a uniform distribution over all models, i.e., $1/|models|$, in this case 1/14.

intention is to use this framework for two purposes: to infer methods for and influences on planning in insects and to characterize performance of systems with reduced memory and high efficiency compared with state-of-the-art planning algorithms. Ultimately, we aim to gain insight into insect navigation and improve robotic systems that operate under strong resource constraints.

References

Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3), 329–349. doi: 10.1016/j.cognition.2009.07.005