

Inverse Optimal Control Adapted to the Noise Characteristics of the Human Sensorimotor System II

Project Manager
Matthias Schultheis

Researchers
Dominik Straub

Principal Investigator
Prof. Constantin Rothkopf

Project Term
2022 - 2023

Clusters
Lichtenberg II Cluster Darmstadt

Additional Software
JAX

Institute
Centre for Cognitive Science

University
Technische Universität Darmstadt



Introduction

Inverse optimal control (IOC) is the problem of inferring an agent's cost function and other properties of their internal model from behavior. It can be used to characterize behavior in sequential decision-making tasks. While IOC has been a fundamental task in artificial intelligence and machine learning, particularly reinforcement learning (RL) and robotics, it has widespread applicability in several scientific fields including behavioral economics, psychology, and neuroscience. Most existing work, however, is limited to fully-observable or linear systems, or requires the action signals to be known. In this project, we developed a probabilistic approach to inverse optimal control for partially-observable stochastic non-linear systems with unobserved action signals. By building on top of the previously proposed maximum causal entropy formulation, our approach unifies previous approaches to inverse optimal control. For evaluating the approach, we ran simulated experiments, applying different inverse optimal control methods to a large number of trajectories based on different parameter combinations. As the computation required many independent runs, which could be parallelized, the application of HPC was beneficial.

Methods

The inverse optimal control method is based on linearization of the dynamics to maintain a tractable likelihood. Internally, for solving the forward optimal control problem, iLQG and the

extended Kalman filter is used, yielding linear controller and filters. Via this formulation, an approximate likelihood in Gaussian form can be computed. For computing gradients in the optimization process, implicit differentiation is used.

Results

The method was successfully developed and evaluated on simulated data. It was shown that the true parameter values could be reliably estimated based on the noisy simulated data. Comparison to past IOC methods based on the maximum entropy formulation has shown that our method led to more accurate recovered parameter values. While our method developed in the past project period was limited to linear systems, the new method can handle non-linear systems via linearization, is more efficient through a gradient-based optimization, and generalizes past IOC methods.

Discussion

Through our evaluation, we show that our method can disentangle perceptual factors and behavioral costs despite the fact that epistemic and pragmatic actions are intertwined in sequential decision-making under uncertainty, such as in active sensing and active learning. The proposed method has broad applicability, ranging from imitation learning to sensorimotor neuroscience. This will enable answering novel scientific questions about how these quantities are affected by different experimental conditions, how they deviate from intended task goals and provided task instructions, or how they vary between individuals. This is particularly relevant to a computational understanding of naturalistic behavior, for which subjective utilities are mostly unknown.

Publications

Straub, D.; Schultheis, M.; Koepl, H.; Rothkopf, C. A.: "Probabilistic inverse optimal control with local linearization for non-linear partially observable systems." (2023) <https://doi.org/10.48550/arXiv.2303.16698>

Schultheis, M.; Rothkopf, C. A.; Koepl, H.: "Reinforcement learning with non-exponential discounting," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 35, 2022, pp. 3649–3662 <https://doi.org/10.48550/arXiv.2209.13413>

Last Update: 2023-08-10 12:53