Approximate Bayesian Decision Theory With Neural Networks

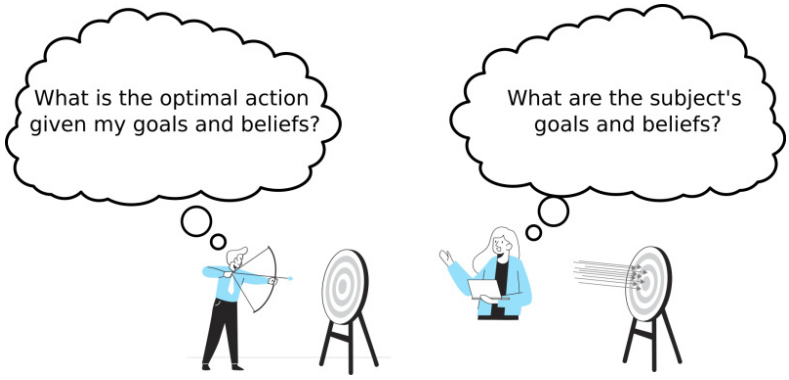


Figure 1: Left: Decision-making problem from the perspective of the subject. The subject needs to find the optimal action based on a cost function and their belief about the state of the world. Right: Inference problem about the subject's parameters from the perspective of the researcher. The researcher solves the inverse decision-making problem, i.e. they want to infer the posterior distribution over the parameters of the subject's perception-action system and cost function given a dataset of the subject's behavior.

Project Manager
Dominik Straub

Researchers
Tobias F. Niehues

Principal Investigator
Prof. Constantin Rothkopf

Project Term
2023 - 2024

Clusters
Lichtenberg II Cluster Darmstadt

Software
Python

Institute
Centre for Cognitive Science

University
Technische Universität Darmstadt

Introduction

Bayesian observer and actor models have provided normative explanations for many behavioral phenomena in perception, sensorimotor control, and other areas of cognitive science and neuroscience. They attribute behavioral variability and biases to different interpretable entities such as perceptual and motor uncertainty, prior beliefs, and behavioral costs. However, when extending these models to more complex tasks with continuous actions, solving the Bayesian decision-making problem is often analytically intractable. This is because the decision-making problem involves maximizing an expectation of a cost function over a probability distribution. Moreover, inverting such models to perform Bayesian inference about their parameters given behavioral data is even more difficult. Therefore, researchers typically constrain their models to easily tractable components, such as Gaussian distributions or quadratic cost functions, or resort to computationally expensive numerical methods.

Methods

To overcome these limitations, we amortize the Bayesian actor using a neural network, which is trained to approximate the optimal action. These parameters can include the actor's perceptual and motor uncertainty, their prior belief, and their cost function. The network is trained on a wide range of different parameter settings in an unsupervised fashion.
Because a forward pass through the neural network is fast and it is differentiable with respect to its inputs, we can use it as a stand-in for the optimal action in downstream applications. For

JUSTUS-LIEBIG-UNIVERSITÄT GIESSEN
TECHNISCHE UNIVERSITÄT DARMSTADT
UNIKASSEL VERSITÄT
Philipps Universität Marburg

example, the pre-trained neural network enables performing Bayesian inference of the Bayesian actor model's parameters given a dataset of behavioral data. We used the HPC to train neural networks to approximate optimal decisions for different cost functions. We then evaluated the inferences obtained using the neural network by running many simulated experiments in parallel.

Results

We first show on synthetic data that the inferred posterior distributions are in close alignment with those obtained using analytical solutions for the optimal action where they exist. Where no analytical solution is available, we recover posterior distributions close to the ground truth parameter values. Importantly, our method can be used to derive recommendations for experimental designs that allow to clearly identify parameters. Finally, we apply our method to empirical data from previously published studies. We show that it explains behavioral patterns such as undershots and regression to the mean.

Discussion

We have presented a new methodological framework for performing Bayesian inference about the parameters of Bayesian actor models. This approach is widely applicable in cognitive science, neuroscience, and other behavioral sciences to infer prior beliefs, perceptual and motor uncertainties and intrinsic costs of behavior from magnitude estimation tasks, production and reproduction tasks, and other sensorimotor tasks like throwing balls or shooting pucks. The approach is extensible to cost functions with other functional forms suited to particular tasks. Our code offers the possibility to implement new cost functions, train a neural network and then perform inference given behavioral data.

Publications

Niehues, T. F. Approximate Bayesian Inference of Parametric Cost Functions in Continuous Decision-Making, M.Sc. Thesis, Technical University Darmstadt, Germany

Last Update: 2024-10-21 19:53