Algorithms for Imitation Learning

- Language: English
- Examiner: Prof. Jim Mainprice
- Contact: jim.mainprice@ipvs.uni-stuttgart.de
- Type of course: Advanced seminar (Master students)
- Research page: https://ipvs.informatik.uni-stuttgart.de/mlr/HRM

Description

In this seminar, we will discuss the computational processes of imitation. This course is especially relevant for students interested in Artificial Intelligence, Machine Learning and Cognitive Science. Imitation learning is the problem of reproducing an observed behavior. Hence, these algorithms are especially useful for artifial agents Learning from Demonstrations. However, the study of the computational foundation of imitation learning inform us about how much knowledge we can extract from observation of a given phenomenon, and thus is an interesting philosophical and scientific endeavor in itself. Each session of this seminar will focus on a couple of algorithms and will highlight their strengths and weakness. On top of reinforcing their knowledge in imitation learning and AI, the participants should learn to study and criticize algorithms and strengthen their scientific communication skills.

Topics

We will discuss current theory and methods for imitation learning, for instance:

- Behavior Cloning
- Hidden Markov Models (HMMs)
- Gaussian Mixture Models (GMMs)
- Dynamic Movement Primitives (DMPs)
- Long Short-Term Memory (LSTMs)
- Markov Decision Processes (MDPs)
- Inverse Optimal Control (model based)
- Inverse Reinforcement Learning (model free)
- Hierarchical Approaches

Organization

The first lectures will provide a recap of the fundamentals and an overview of recent topics. In the subsequent weeks, students will present papers in the field of imitation learning. The seminar also includes a final project that will be based on a recent publication and demonstrate the approach/technique in simulation (e.g. using OpenAI Gym). A short report and final presentation are to be delivered at the end of the semester.

Prerequisites

Bachelor courses: *Artificial Intelligence* (Grundlagen der künstlichen Intelligenz). We assume familiarity with reinforcement learning and machine learning, in particular:

- Reinforcement Learning
- Definition of MDPs
- Policy and value iteration
- Machine Learning
- Classification and regression
- Fitting of linear and non-linear models
- Loss functions, Training/test error, overfitting