TECHNISCHE UNIVERSITÄT MÜNCHEN
LEHRSTUHL FÜR INFORMATIONSTECHNISCHE REGELUNG
ORDINARIA: UNIV.-PROF. DR.-ING. SANDRA HIRCHE

M A S T E R ' S   T H E S I S

Learning Human Motion Models using Approximate Inverse Reinforcement Learning

Problem description:

Estimating the intent of an agent from demonstrations, is a critical task in many collaborative settings. Two popular approaches to tackle this problem is by means of inverse reinforcement learning (IRL) and dynamical movement primitives (DMP). Methods based on IRL assume the agent to behave optimally with regards to an intrinsic cost function, which is then inferred from demonstrations [1]. Even though the retrieved cost function depicts a rich representation of the agents behavior, IRL is computationally demanding and does generally not provide stability guarantees. Differently, DMP methods describe the convergence behavior of the agent at the disadvantage of losing descriptiveness regarding the agent's intent [2].
Based on the principal of inverse optimality, it is possible to approximately infer the cost function by an equivalent Lyapunov function [3]. Thereby, the expressiveness of the IOC approach and the stability guarantees of the DMP are retained. The underlying idea is to infer the value function of an agent directly and treat it as a control Lyapunov function. Thus, for some system, a closed-form control laws can be obtained directly and the Lyapunov function allows to make statements regarding the convergence properties of the close-loop system.
In this work a previously developed method shall be extended from low-dimensional, goal-directed tasks to be applicable in real-world settings. To this end methods that provide probabilistic guarantees, i.e., probably approximately correct (PAC) learning bounds, and stable Deep Neural Network representation will be researched. Finally, the developed method shall be evaluated on human-generated experimental data.

Tasks:

- Literature research on control Lyapunov function estimation for stochastic systems
- Development of a learning algorithm with stability guarantees
- Evaluation of proposed method on experimental data

Bibliography:

[1] W. Jin, D. Kulić, J. Feng-Shun Lin, S. Mou, and S. Hirche. Inverse optimal control for multiphase cost functions. *IEEE Transactions on Robotics*, 35(6):1387–1398, 2019.
[2] A. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal. Dynamical movement primitives: Learning attractor models for motor behaviors. *Neural Comput.*, 25(2):328–373, 2013.
[3] S. Tesfazgi, A. Lederer, and S. Hirche. Inverse Reinforcement Learning: A Control Lyapunov Approach. In *Accepted: IEEE 60th IEEE Conference on Decision and Control (CDC)*, 2021.

Supervisor: Samuel Tesfazgi
Start: As soon as possible

(S. Hirche)
Univ.-Professor