

Visuomotor Control of Pouring Liquids



Project Manager
Niteesh Midlagajni

Researchers
Carola Stork

Principal Investigator
Prof. Constantin Rothkopf

Project Term
2021 - 2022

Clusters
Lichtenberg II Cluster Darmstadt

Software
PyTorch

Additional Software
SPLisHSPLasH

Institute
Centre for Cognitive Science

University
Technische Universität Darmstadt

Introduction

Natural visuomotor control tasks such as pouring liquids into cups are trivial for humans but are challenging to model. Without the explicit knowledge of the physics governing the liquid dynamics, humans appear to solve the task effortlessly. It seems that we possess intricate cognitive mechanisms for reasoning about liquids, enabling us to predict the behaviour of liquids in a given environment. This understanding guides our choice of the optimal pouring action and duration based on the specific pouring task at hand. This project is about developing artificial agents that learn pouring behaviour using techniques from modern deep reinforcement learning and Imitation learning. Leveraging the comprehensive human pouring behaviour dataset we have collected, we aim to improve the learned policies and eventually conduct a comparison study.

Methods

We used a liquid simulator and adapted the environment to a setup similar to the physical dataset. For the Reinforcement learning environment, Gymnasium was used where the liquid simulator was integrated. For training of the agent, Twin Delayed DDPG (TD3) algorithm was used. The network architecture for both actor and critic networks was based on convolution networks and representations were designed such that neighbouring liquid particles could learn their spatial interactions. For imitation learning, we applied behavioural cloning (supervised learning) to train the actor network with pre-generated expert trajectories. We then used this pre-trained network as the actor network of the TD3 algorithm and

continued training it using reinforcement learning.

Results

For the baseline result where the agent learns from scratch, our agent did learn a valid pouring behaviour. However, the learned trajectories were very fast and the jug overturned at the end, resulting in a lot of spillage. Using behavioural cloning, the actor network on its own was able to learn a good pouring behaviour which followed the expert trajectories well. Integrating this pre-trained actor into the TD3 agent for further reinforcement learning training reduced the spillage of the agent by preventing the overturning behaviour observed previously.

Outlook

Despite being a complex problem, we showed that it is indeed possible to learn valid pouring behaviours using reinforcement learning and imitation learning. However, the performance achieved has not yet been adequate. Further work is required in order to improve the neural network architectures used. One possible way is to use special architectures that learn liquid particle interactions such as ConvSPF or DPI-net. Additionally, we are considering alternate state representation, where the agent learns directly from the image space data, and even augment the particle data with the image space data.

Last Update: 2023-11-06 16:40