

Learning Physically Interactive Behaviours for Human-Robot Interaction

Project Manager
Vignesh Prasad

Principal Investigator
Prof. Jan Peters (PhD)

Project Term
2021 - 2021

Clusters
Lichtenberg II Cluster Darmstadt

Software
PyTorch

Institute
Intelligent Autonomous Systems,
Department of Computer Science,
Chair of Marketing and HR
Management

University
Technische Universität Darmstadt

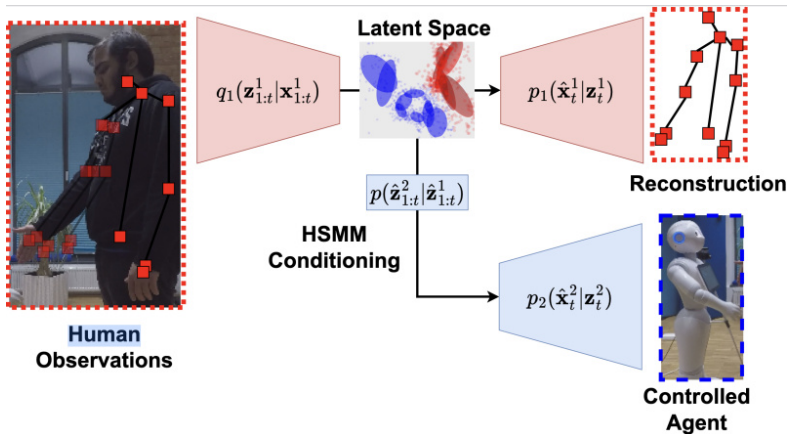


Figure 1: Overview of our approach, "MILD". We train VAEs to reconstruct the observations of the interactions agents ($x_{1:t}^1, x_{1:t}^2$) with an HSMM prior to learn a joint distribution over the latent space trajectories ($z_{1:t}^1, z_{1:t}^2$) of the interacting agents. During test time, the observed agent's latent trajectory $p(z_t^2 | z_{1:t}^1)$ conditions the HSMM to infer the controlled agent's latent trajectory \hat{x}_t^2 which is decoded to generate the agent's real-world trajectory.

Introduction

Modeling interaction dynamics to generate robot trajectories that enable a robot to adapt and react to a human's actions and intentions is critical for efficient and effective collaborative Human-Robot Interactions. This project explored Imitation Learning techniques to learn physically interactive behaviors for Human-Robot Interactions. The goal of this project was to enable robots to reactively respond to interactive actions being performed with a human in a timely manner that maintains spatial proximity to the human, for example waving, handshakes, high-fives, etc.

Methods

We propose a method that couples deep representation learning and probabilistic machine learning to address the problem of two-party physical Human-Robot Interactions. The developed approach uses Hidden Markov Models to model the joint latent space distributions of the interacting partners in the representation bottleneck of a Variational Autoencoder. Doing so enabled us to reactively generate suitable robot motions to be spatially and temporally coordinated with the human's motions. We train our model using demonstrations of humans performing interactions, which is then used to generate actions for a Humanoid Robot given the similarities in the Degrees of Freedom used to control the robot and the movement

capabilities of the human (Shoulder, elbow, and hand motions). While from a scientific standpoint, it is fairly straightforward, from an implementation perspective, it involves extensive tuning of various hyperparameters and input/output representations to ensure that the underlying distribution that is extracted from the data is meaningful and can be used to successfully solve the task at hand. Therefore, the cluster was used to perform extensive hyperparameter tuning by training various instances of the model to prune out the ones that do not learn meaningful distributions or do not solve the task successfully.

Results

Our experimental evaluations for learning Human-Robot Interactions from Human-Human Interaction demonstrations show that our proposed approach effectively learns meaningful underlying distributions that capture the nuanced interaction dynamics associated with such tasks. The final models developed in this project beat state-of-the-art approaches using recurrent neural networks for generating spatiotemporally coherent and accurate robot trajectories after observing the human interaction partner for a variety of interaction tasks on different robots. Although initial results proved ineffective in adequately learning underlying distributions either due to a lack of sufficient representation capacity of the deep learning model or due to collapsing into an uninformative distribution, sufficient hyperparameter tuning on the cluster enabled pruning out such models to yield the final set of model hyperparameters that consistently led to good performance.

Discussion

The models developed in this approach only evaluate the predictive performance on existing datasets. They would need further testing with the real robot and human interaction partners to gauge the true efficacy of the approach. Additionally, when running the predicted motions on the robot, it needs to be safe for both the human and the robot such that no unexpected forces are exerted either on the human or on the robot. Therefore, additional post-processing of the robot's motions, such as using Inverse Kinematics needs to be incorporated to equip a robot with the developed approach.

Publications

Prasad, V; Koert, D; Stock-Homburg, D; Peters, J. and Chalvatzaki, G. MILD: Multimodal Interactive Latent Dynamics for Learning Human-Robot Interaction. 472-479. In IEEE-RAS International Conference on Humanoid Robots (Humanoids) (2022)
<https://doi.org/10.1109/Humanoids53995.2022.10000239>

Last Update: 2024-05-17 18:39