Hessisches Kompetenzzentrum
für Hochleistungsrechnen

Learning to Plan

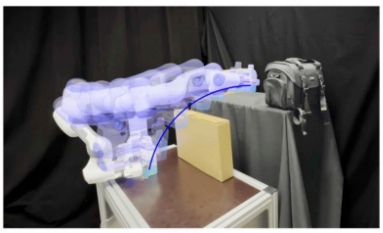
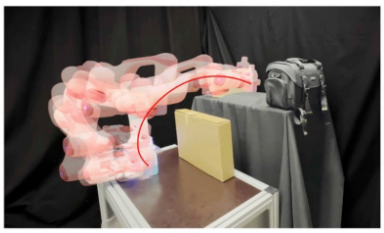



Figure 1: Execution of multimodal trajectories the Franka Panda robot. The motion in blue and red start and end at the same joint configurations, but it is possible to see two modes resulting from sampling trajectories with MPD.

Project Manager
Joao Carvalho

Principal Investigator
Prof. Jan Peters (PhD)

Project Term
2022 - 2023

Clusters
Lichtenberg II Cluster Darmstadt

Software
PyTorch

Institute
Intelligent Autonomous Systems

University
Technische Universität Darmstadt

Introduction

Motion planning is a crucial component of autonomous robot systems. It addresses the problem of finding a feasible, smooth, and collision-free path between a start and a goal point in a robot's configuration space, which can subsequently be executed by a lower-level controller. Following the imitation learning perspective, in this work, we learn the prior model from expert data and incorporate it into optimization-based motion planning. In particular, instead of explicitly sampling from a prior distribution as motion planning initialization, we propose to merge the prior sampling and motion optimization into one algorithm, by leveraging recent formulations in diffusion based generative models. This type of implicit model has shown impressive results in modeling multimodal and highdimensional data, such as image generation, having superior generative performance (with/out context guidance) compared to previous generative models. Indeed, these diffusion model properties are particularly well-suited for learning from demonstrations in robotics manipulation settings, where state space dimensions are usually large in manipulators (e.g. full state of position and velocity of Franka Emika Panda arm has 14 dimensions) and there exist thousands of trajectory samples in the expert dataset.

Methods

The method developed in this work is named: Motion Planning Diffusion (MPD). It can be divided in two parts – learning and planning. During the learning phase, we learn a trajectory generative model with a diffusion model using expert trajectories generated with an optimal motion planning algorithm. The diffusion model denoising function follows a Unet architecture that can handle large amounts of data and high-dimensional inputs. The output of this function is proportional to the log-probability of the probability density function of the data. During the planning phase, we formulate the motion planning problem as planning-asinference by sampling from a posterior distribution leveraging guidance in diffusion models. We use several costs from motion planning, e.g., collision-avoidance,

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

smoothness, endeffector orientation, and integrate them as guides in the classifier-guidance framework of diffusion models.

Results

We tested our algorithm against several baselines from simple to more complex environments. We found out that using a Diffusion Prior we can generate more collision-free trajectories than the baseline using a CVAE Prior. E.g., in the PointMass2D Dense environment, the success rate is 98% vs. 46%. Moreover, the diffusion model can produce more multimodal trajectories. Empirically, we observed that across different planning problems, the CVAE models tended to generate fewer modes than the ones from the diffusion. Observing the success rates of MPD and comparing them to CVAEPosterior, we see that using the diffusion model in combination with the cost gradients yields better results in terms of success rate and multimodality.

Discussion

In this work we proposed using diffusion models as priors for bootstrapping motion planning problems, via the planning-as inference perspective. We parameterize a trajectory with waypoints and construct a generative model over the whole trajectory. We train this model via supervised learning on motion plans generated with an optimal planner. At inference time, instead of sampling trajectories from the prior to only initialize an optimization-based motion planner, we propose to use the guidance properties of diffusion models to concurrently sample from the prior and bias these samples towards regions of low cost.

Publications

Carvalho J., et al, Motion Planning Diffusion: Learning and Planning of Robot Motions with Diffusion Models, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2023

Last Update: 2024-06-10 15:48