Motion Planning Diffusion: Learning and Adapting Robot Motion Planning with Diffusion Models (Abstract Reprint)

Joo Carvalho[1,2], An Thai Le[1], Piotr Kicki[1], Dorothea Koert[1], Jan Peters[1,3,4]
[1]Intelligent Autonomous Systems Lab, Computer Science Department, Technical University of Darmstadt, Darmstadt, Germany
[2]Institute of Systems and Robotics, University of Lisbon, Lisbon, Portugal
[3]German Research Center for AI (DFKI), Research Department: SAIROL, Darmstadt, Germany
[4]Hessian.AI, Darmstadt, Germany

Abstract Reprint. This is an abstract reprint of the journal article by Carvalho, Le, Kicki, Koert, and Peters (2025).

Abstract

The performance of optimization-based robot motion planning algorithms is highly dependent on the initial solutions, commonly obtained by running a sampling-based planner to obtain a collision-free path. However, these methods can be slow in high-dimensional and complex scenes and produce nonsmooth solutions. Given previously solved path-planning problems, it is highly desirable to learn their distribution and use it as a prior for new similar problems. Several works propose utilizing this prior to bootstrap the motion planning problem, either by sampling initial solutions from it, or using its distribution in a maximum-a-posterior formulation for trajectory optimization. In this work, we introduce motion planning diffusion (MPD), an algorithm that learns trajectory distribution priors with diffusion models. These generative models have shown increasing success in encoding multimodal data and have desirable properties for gradient-based motion planning, such as cost guidance. Given a motion planning problem, we construct a cost function and sample from the posterior distribution using the learned prior combined with the cost function gradients during the denoising process. Instead of learning the prior on all trajectory waypoints, we propose learning a lower dimensional representation of a trajectory using linear motion primitives, particularly B-spline curves. This parametrization guarantees that the generated trajectory is smooth, can be interpolated at higher frequencies, and needs fewer parameters than a dense waypoint representation. We demonstrate the results of our method ranging from simple 2-D to more complex tasks using a 7-DOF robot arm manipulator. In addition to learning from simulated data, we also use human demonstrations on a real-world pick-and-place task. The experiment results show that diffusion models are strong priors for encoding multimodal trajectory distributions for optimization-based motion planning.

References

Carvalho, J.; Le, A. T.; Kicki, P.; Koert, D.; and Peters, J. 2025. Motion Planning Diffusion: Learning and Adapting Robot Motion Planning With Diffusion Models. *IEEE Transactions on Robotics*, 41: 4881–4901.