Deep reinforcement learning for motion planning in man-made environments

Piotr Kicki

Examiner's Report

Prof. Ioannis Havoutis, University of Oxford

This report outlines the scope and content of the examined thesis, gives a brief summary of each chapter, and summarizes specific achievements and significant outcomes.

The thesis addresses the problem of learning-based motion planning for car-like vehicles and constrained kino-dynamic motion planning problems. The contributions include rapid path planning for car-like vehicles by building a solution from locally-defined polynomial segments generated by a neural network, and on-the-fly trajectory generation using a B-spline parametrization with a neural network generator. The latter has been effectively pre-trained on offline data and experiments and evaluations on a robotic air-hockey setup have been showcased. In all presented cases, neural networks are trained mainly under the reinforcement learning paradigm, where a range of loss functions have been evaluated to increase the efficiency of learning and quality of generated solutions.

Summary of the individual works

Chapters 2 and 3 focus on a machine learning-based approach to path planning for autonomous vehicles. Chapter 4 introduces a broader machine learning-based trajectory planning approach, focusing on robotic manipulation. Experimental evaluation of the proposed methods follows in Chapters 5 and 6, with both quantitative and qualitative analyses.

Chapter 2 outlines the implementation of learning-based planning, approximating the optimal planning function for local manoeuvrers in car-like vehicles. It formulates the planning problem as a Markov Decision Process (MDP) and suggests using a neural networks as planning policies.

Chapter 3 shifts from MDP to a bandit formulation to speed-up planning. It introduces B-spline path representations and a novel path construction procedure that incorporates boundary constraints directly into the solution.

Chapter 4 shifts focus to kino-dynamic motion planning for robotic manipulation. It presents a general approach for learning to plan near the constraint manifold and introduces a novel B-spline trajectory parametrization. The chapter proposes a machine learning-based planning algorithm capable of solving constrained kino-dynamic motion planning problems in milliseconds and replanning on-the-fly.

Chapter 5 provides the experimental evaluation of machine learning-based path planning for car-like vehicles from Chapters 2 and 3. The evaluation compares their capabilities with baseline motion planners on a dataset introduced in Chapter 2 and qualitative scenarios using the CARLA simulator.

Chapter 6 presents the experiments on constrained kino-dynamic learning-based motion planning. It emphasizes trajectory planning for robotic manipulators, demonstrating the applicability of the method from Chapter 4 in challenging tasks. The proposed approach is compared to state-of-the-art motion planners and evaluated on a real robotic setup of an air-hockey robotic platform.

The final Chapter 7 provides concluding remarks and discusses opportunities for future work.

Assessment of the contributions

Clearly these are substantial contributions to kino-dynamic motion planning using learning-based approaches, which have undergone rigorous scrutiny by reviewers before being published at prestigious venues. Furthermore, the thesis is generally well written, structured, and illustrated. Its descriptions are clear and correct.

In addition, strong evidence of theoretical and experimental evaluation was provided. Therefore, the reviewer believes that *the results and conclusions make a significant and substantial contribution to the field*. This is also well supported by the candidate's publication record in relevant top-tier robotics conferences and journals, for example, at the IEEE Transactions on Robotics (TRO), the IEEE International Conference on Robotics and Automation (ICRA) and the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), while much of the candidate's work has been presented to relevant national and international workshops.

Lastly, to address the evaluation criteria explicitly:

- The doctoral dissertation presents general theoretical knowledge of the Candidate in the discipline of kino-dynamic motion planing and machine learning as well as their ability to independently conduct scientific work.
- The subject of the doctoral dissertation is an original solution, or rather, a set of novel solutions, to a scientific problem.

Sincerely,

2024-03-12

Prof. Ioannis Havoutis
Associate Professor in Robotics
Dynamic Robot Systems Group
Oxford Robotics Institute
University of Oxford