Deep Multi-Step Planning for Autonomous Driving

Background

Learning-based approaches, in particular reinforcement learning agents, have been applied to autonomous driving with promising results in recent years. However, despite optimizing for future cumulative reward, the output of such agents is typically limited to the current-step action, offering little transparency about long-term intentions. On the other hand, classical approaches (e.g. search, optimization) generally provide a multi-step solution, be it a sequence of control signals or an optimal trajectory for the vehicle to follow. This offers clear advantages in terms of explainability and verifiability that could advance the state of learning-based autonomous driving.

While clearly requiring more sophisticated network architectures than regular multilayer perceptrons, there are multiple conceivable ways of achieving deep learning-based autonomous driving with explicitly multi-step planning horizons - giving rise to a particularly interesting thesis topic.

Description

Two different candidate approaches stand out as worth exploring in particular:

1. Unlike vanilla neural networks, recurrent networks are intrinsically compatible with sequence-structured data. Through their internal memory states, they are capable of representing temporal interdependences in a principled manner. By recursively advancing some initial planning context through the network (without further external input signals), the resulting decoding procedure could be capable of intelligently planning a vehicle's motion in a multi-step fashion.
2. A constant-time planning algorithm can be attained by training a neural network to output a parameterized trajectory curve with a pre-defined mathematical structure. However, special care (i.e. clever architectural constraints) would be required to guarantee that the planned trajectories are feasibly given the dynamics model of the controlled vehicle as well as its current state.

In addition to designing a sensible input space and loss function for the learning problem, the thesis would revolve around choosing one of the two approaches suggested above (or another promising idea that catches your interest), and developing a trainable PyTorch trajectory planning module that implements your idea.

Tasks

1. Perform a literature review in order to identify the most promising approach(es) for multi-step neural trajectory planning.
2. Define an input space and a loss function for the learning problem.
3. Develop a PyTorch-based implementation of a planning module can be trained to output multi-step control sequences or trajectories.
4. Train, optimize and evaluate the model.
5. Wrap the model in an API for easy use in downstream applications.

Technische Universität München

Fakultät für Informatik

Lehrstuhl für Echtzeitsysteme und Robotik

Supervisor:
Prof. Dr.-Ing. Matthias Althoff

Advisor:
M.Sc. Eivind Meyer

Research project:
KoSi

Type:
MA

Research area:
Autonomous Driving
Deep Learning

Programming language:
Python

Required skills:
Machine Learning
Python (preferably with PyTorch experience)

Language:
English

Date of submission:
12. Dezember 2021

For more information please contact us:

E-Mail: eivind.meyer@tum.de