MT: Deep learning-based scene-aware behavior velocity planing for autonomous vehicles

Background

Autonomous driving has seen a spike of interest in recent years. Behavior planning, especially behavior velocity planning, is one essential part of the planning pipeline. Not hitting the front vehicle and keeping the speed limits are pretty much solved but how to dynamically adjust your velocity outside the traffic rules to, for example, an oncoming bus on a dense street, a large pedestrian crowd near a crosswalk, or implicitly taking occlusion into consideration ares still interesting questions. This master thesis is aimed at exploring the possibilities of deep learning in velocity-related behavior planning.

Figure 1: TUM research vehicle EDGAR.

Description

We are currently in the conception phase of this methodology and, thus, are in need of a literature research, a concept and a first prototype. As our vehicle EDGAR is currently driving LIDAR-only, we will use a point cloud scene representation to learn from; first from publicly available data sets such as nuscenes, and then, hopefully, from EDGAR.

With CommonRoad [1, 2], our chair is unique in the sense that we provide a complete, open-source scenario description and planning framework that is successfully running on a real car [3] (see this video). Our chair has deep practical knowledge in route planning[1], local planning[2] and is one of the leading chairs in the area of formally verified traffic rules and motion planning based on formal methods. However, in the area of velocity planning, especially non-traffic rule based velocity planning, we currently do not have a comparable knowledge base. The CommonRoad velocity planner[3] is the first step of changing that but it currently only considers the road geometry and the planning problem, thus not providing scene-aware behavior.

In many practical applications, such as Autoware, the behavior velocity planner is rule-based. Although this works reasonably well in many practical cases, it requires a substantial parametrization effort and may simultaneously produce overly conservative results in some cases, whilst not making use of implicit scene knowledge and driving overly confident in others. In research, fields like game theory and formal methods are explored. However, all of them rely on the current environment model, i.e. the perception and prediction, to be correct. And end-to-end learning is unpredictable and can result in dangerous behavior.

The approach of this thesis is aimed at enriching the currently existing framework [3] by a module that takes the original velocity plan and the velocity limits resulting from traffic rules as given and aims at finding a velocity within these limits to facilitate good behavior.

[1] pip install commonroad-route-planner
[2] pip install commonroad-rp
[3] pip install commonroad-velocity-planner, avaiable from August 2024

Technical University of Munich

Department of Informatics

Chair of Robotics, Artificial Intelligence and Real-time Systems

Supervisor:
Prof. Dr.-Ing. Matthias Althoff

Advisor:
Tobias Mascetta, M.Sc.

Research project:
i4Driving, CommonRoad

Type:
MT - Master Thesis

Research area:
Autonomous Driving - Planning

Programming language:
Python, PyTorch

Required skills:
Python, PyTorch, OOP, Git

Language:
English, German

Date of submission:
22. Juli 2024

For more information please contact us:

Phone: +49 (89) 289 18100

E-Mail: tobias.mascetta@tum.de

Website: www.ce.cit.tum.de/air/

Technical University of Munich

Department of Informatics

Chair of Robotics, Artificial Intelligence and Real-time Systems

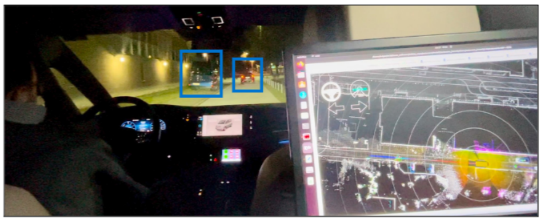

(a) Real vehicle view and RVIZ visualization. The other vehicles are highlighted for clarity.

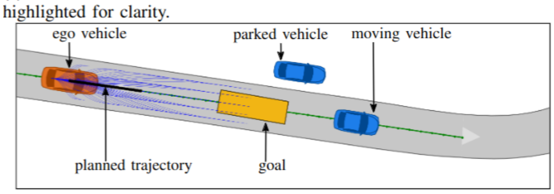

(b) Corresponding scenario in CommonRoad with planner output.

Figure 2: Exemplary planning on EDGAR with the CommonRoad Reactive Planner.

Prerequisites

Using a state-of-the-art point cloud backbone such as PointNet [4], you will develop your own architecture [5, 6, 7, 8] as well as implement, train and evaluate it. As this is a master thesis, advanced previous knowledge in applied deep learning is required, e.g., terminology (e.g. Feature, Layer, Head), layer types (e.g. Linear, Convolution, Batchnorm, Fully-Connected, ResNet Bottleneck, LSTM, Skip-Connection etc.), basic architecture types (e.g. GAN, AutoEncoder) and advanced PyTorch knowledge.

Tasks

Your tasks are:

1. Investigation current approaches in the literature on deep-learning based velocity planning
2. Familiarization with point cloud based deep learning architectures
3. Finding a promising data set
4. Definition of the term *good* velocity, perhaps based on concrete scenarios
5. Conceptualization of a promising network architecture
6. Iterative implementation, testing and evaluation of your approach
7. (Optional) If it works well and is fast enough, we will try to test it on EDGAR

Application

If you are interested in the topic, please send an email to the contact information provided on the right and attach a short **CV**, your current **grade report** and your **bachelor grade report**.

If you have extra-curricular projects or a student job, please attach information about it.

References

[1] Matthias Althoff, Markus Koschi, and Stefanie Manzinger. Commonroad: Composable benchmarks for motion planning on roads. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 719–726, 2017.

[2] TUM-IN06 and Matthias Althoff. Commonroad website, 2023.

[3] Gerald Würsching, Tobias Mascetta, Yuanfei Lin, and Matthias Althoff. Simplifying sim-to-real transfer in autonomous driving: Coupling autoware with the commonroad motion planning framework. In *2024 IEEE Intelligent Vehicles Symposium (IV)*, 2024.

[4] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

[5] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

[6] Aditya Khamparia and Karan Mehtab Singh. A systematic review on deep learning architectures and applications. *Expert Systems*, 36(3):e12400.

[7] Alex H. Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[8] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11784–11793, June 2021.

Technical University of Munich

Department of Informatics

Chair of Robotics, Artificial Intelligence and Real-time Systems