Provably Safe Adaptive Reinforcement Learning for Autonomous Driving

Technical University of Munich

Department of Informatics

Chair of Robotics, Artificial Intelligence and Real-time Systems

Background

Safe reinforcement learning (RL) has shown strong potential for decision making in complex and uncertain environments [1]. However, for safety-critical systems, achieving high performance is not sufficient: the learned policy must also remain provably safe under changing environmental conditions. To address this challenge, this thesis focuses on extending a framework that combines learning-based action masking [3] with formally verified fail-safe maneuvers [2]. The method has already been tested in abstract benchmark tasks, to further demonstrate its applicability and effectiveness in more challenging scenarios, we will apply it to autonomous driving scenarios.

Description

This thesis aims to extend our current provably safe adaptive RL framework toward autonomous driving scenarios. The student will build on our existing framework and implement longitudinal and lateral action masking networks to prevent driving off-road and guarantee safe-distance with the preceding vehicles. The resulting method will be evaluated under various scenarios, including different road shapes and traffic conditions, and the adaptability will be assessed.

The project is closely connected to an ongoing research paper and is intended to provide experimental validation for further publication.

Tasks

- Get familiar with the existing code-base and our platform CommonRoad-RL.
- Adapt and integrate the action masking and fail-safe components.
- Design and conduct experiments under changing environments to evaluate safety, adaptability, and performance.
- Document the implementation, experiments, and findings in a written thesis.

References

[1] Hanna Krasowski, Jakob Thumm, Marlon Müller, Lukas Schäfer, Xiao Wang, and Matthias Althoff. Provably safe reinforcement learning: Conceptual analysis, survey, and benchmarking. *arXiv preprint arXiv:2205.06750*, 2022.

[2] Christian Pek and Matthias Althoff. Fail-safe motion planning for online verification of autonomous vehicles using convex optimization. *IEEE Transactions on Robotics*, 37(3):798–814, 2020.

[3] Mirco Theile, Lukas Dirnberger, Raphael Trumpp, Marco Caccamo, and Alberto L Sangiovanni-Vincentelli. Action mapping for reinforcement learning in continuous environments with constraints. *arXiv preprint arXiv:2412.04327*, 2024.

Supervisor:
Prof. Dr.-Ing. Matthias Althoff

Advisor:
Shuaiyi Li

Research project:

Type:
MA/GR/SA

Research area:
Safe Reinforcement Learning, Autonomous Driving

Programming language:
Python

Required skills:
Good programming skill of Python, familiar with git, experience with reinforcement learning is a plus, self-motivated working

Language:
English

Date of submission:
Flexible

For more information please contact us:

Phone:

E-Mail: shuaiyi.li@tum.de

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