Reinforcement Learning for Automated Vehicle Control

Acknowledgements
In Austria the project was funded by the program Mobilit�at der Zukunft" and the Austrian Federal Ministry for Transport, Innovation and Technology (bmvit). The publication was written at VIRTUAL VEHICLE Research Center in Graz and partially funded by the COMET K2 { Competence Centers for Excellent Technologies Programme of the Federal Ministry for Transport, Innovation and Technology (bmvit), the Federal Ministry for Digital, Business and Enterprise (bmdw), the Austrian Research Promotion Agency (FFG), the Province of Styria and the Styrian Business Promotion Agency (SFG).

The aim of the research project is to use Reinforcement Learning (RL) to train an agent that is capable of navigating a robot in a complex real world environment. We use the mobile robotic test platform, SPIDER, developed at the Virtual Vehicle research labs as a test bed. The agent is trained to follow a predetermined path while avoiding stationary obstacles like trees or walls and moving obstacles like other vehicles or pedestrians along or near the path.

The learning process/the training of the agent is performed in a simulated virtual environment by a RL algorithm. The main concept of RL- algorithms is inspired by the human learning behavior: During the training the agent performs actions within the virtual environment, a driving simulation in our case, and earns feedback from the environment - either a reward or a punishment. The rewarding strategy we are using encourages the agent to steer towards way points on the path and to keep distance to any kind of obstacles. Basically the agent gets rewarded when he is moving in the direction of the next way point and gets punished if not. The distance to obstacles is measured using detection rings surrounding the agent. The more regions in those rings are blocked due to obstacles the more the agent gets punished. This allows it to evolve a strategy to earn as much rewards as possible and to find as a consequence a policy of how to follow a path without crashing into obstacles.

In real world driving the SPIDER observes the environment using four LIDAR sensors. The SPIDER's runtime environment, ROS, provides the environment data in the form of 2D cost maps, which encode the objects detected in the SPIDER's surrounding as 2D pixel images. These 2D pixel images and vectors connecting SPIDER's center of mass and way points on the given path are provided to the agent as input.

We have trained the agent due to the abovementioned methodology in a coarse-grained grid world environment renouncing on the in
uence of mass, acceleration
or deceleration forces in the vehicle control. The first experimental results show the feasibility of the approach: The agent is capable to follow a given path in a

variety of scenarios and to drive round stationary obstacles and to evade moving obstacles.
Currently we are extending the model to include the vehicle dynamics of the SPIDER robot and to use larger, more fine-grained pixel images as input of the agent. Both changes are refinements of the original simplified model and thus we are confident that the reinforcement learning approach will provide as an agent beeing capable to steer SPIDER in real world scenarios.