



Name: Gergely Dániel Péter, BSc student, third year


Project type: thesis project

Topic: Training a Lane Keeping Agent using Reinforcement Learning

Supervisors: Dr. Balázs Bánhelyi, Dr. Gábor Gosztolya


The lane following problem:

- One of the core requirements for fully autonomous driving
- The goal is to teach a reinforcement learning agent to keep a vehicle moving at constant speed within its lane by appropriately changing the steering angle for as long as possible
- A custom learning environment was created built on the CARLA simulator, as it provides a quite realistic and customisable simulation environment
- The environment runs rollouts from random starting points on the „infinite highway” of the CARLA Town04 map



Experiments

- The PPO algorithm is used for training
- The Policy and Value Networks both work with the used steering value and some preprocessed images collected during the previous environment step to compute a steering action distribution and a state value estimation respectively
- The environment calculates the reward based on the following criteria: distance from lane center, difference between car orientation and lane direction, whether the chosen steer value decreases the previous metric
- Each simulation step lasts the same time (e.g.: 0.04 seconds) during which the steering value stays constant
- The image data that comes from the RGB sensor provided in CARLA is preprocessed using image cropping, edge detection techniques and binary thresholding to remove information not needed for lane following and to emphasize the positions of lane markings



Results & future work

Results:

- Best evaluation result so far was 73.86 seconds on average over 10 rollouts
- Noticably better performance when travelling next to solid lines compared to when both sides have dashed lines
- More data per training iteration gives better results at the start of and better average results over the training process
- Shadows don't seem to bother the agent despite sometimes being visible on the preprocessed images, but it sometimes „confuses” lamp posts with lane markings

Future work:

- The evaluation procedure needs revamping to be able to provide deeper insights and a more controlled way to measure the agents' performance for better fine tuning
- Restricting allowed steer value change to respect physical limitations
- Long term: Introducing additional complexity (car to follow, changing weather conditions etc.)

Education

