Elitenetzwerk Bayern

Einblicke in die Forschungsarbeit

MASTERARBEIT

„Deep Learning for Self-Driving Cars“

BENEDIKT JENIK

Elitestudiengang Software Engineering

University of Augsburg, December 2018

Deep Learning for Self-Driving Cars

Benedikt Jenik is an alumnus of the Elite Graduate Program „Software Engineering" at the University of Augsburg, Ludwig Maximilian University of Munich and the Technical University of Munich. There and at the Massachusetts Institute of Technology (MIT), he worked on deep learning research, with a focus on self-driving cars within the scope of his master's thesis.

How bleeding edge A.I. research is used in self-driving cars

Deep neural networks have revolutionized artificial intelligence research to a degree that has made it into mainstream news several times, like DeepMind's work in beating the world's best Go player, which was thought to be out of reach of algorithms for multiple decades. One particularly promising application of these deep learning algorithms is in the area of computer vision, meaning making computers detect and understand the contents of images. This research topic is especially applicable to developing and building self-driving cars, since they need to understand their surroundings using image and video data and is already widely applied commercially.

Benedikt Jenik built on that research, with a focus on not just showing that deep neural networks can be used for understanding scenes in the self-driving car context, but going further with also having the neural network learn to make decisions based on that understanding.

Picking the right abstraction level

An interesting perspective on the problem is that even just picking certain abstractions in the representation of a detection and scene understanding problem can blur the lines towards also already making decisions. One could for example structure the detection task as a pure segmentation problem, meaning just trying to find out which pixels belong to the drivable road and which to other objects like cars and pedestrians. This could be further abstracted in not just detecting 2D shapes, but 3D representations of those other cars, or even more exact lane information from the drivable area. This could happen by either a two-step process of the previously mentioned segmentation combined with further shape post processing algorithms, or in a single step where the neural net predicts the curvature of the road in front of it. This already significantly crosses that decision-making line as this information could be used to drive a car, by just steering accordingly, which has been shown to actually work surprisingly well. Jenik and his colleagues built on that work by showing that this way of making driving decisions, which is still too unreliable to be used in practice, is nevertheless useful as a secondary system: the difference in steering angle between a primary system like Tesla's Autopilot and this neural network based secondary system can be used to predict upcoming difficult or dangerous situations and warn the driver accordingly or even make him take over, in a semi-autonomous driving case. This was also deployed and successfully tested on-road in a real car - showing particular promise in busy highway traffic and construction zones.

Learning from crashing hundreds of times

Going up further along this abstraction hierarchy, one could even start to shift the task, from accurately describing the environment that could be used for driving decisions more towards actually having the neural network explicitly making these decisions while directly optimizing for some goal-metric. Jenik and colleagues did that for DeepTraffic, a simulation of dense highway traffic, where the algorithm was trained to efficiently navigate through random disorganized traffic, while optimizing for smoothness and speed. This was further extended as an online competition where deep learning researchers all over the world could submit their models, which got more than 17'000 submissions. In addition to providing insights into which deep learning models work particularly well, it was also shown that an agent doing well does not disturb other cars in the simulation, but actually affects traffic in a way that helps the other cars to go faster as well.

This type of reinforcement learning based, goal oriented neural network training can even be taken to more extremes, by making the algorithm figure out solutions in situations where humans would be unable to act quickly enough, like imminent crashes. An algorithm built by Jenik was able to learn to make its way through such situations using drifts and other complicated maneuvers, while learning from nothing but rewards or penalties after avoiding or getting into a collision in a simulation environment. Replicating these results in a 1:10 model scale is underway, with a promising start, and while it is clear that this should not be replicated with full-scale real cars, it nevertheless shows the potential of neural networks and may even help to discover collision avoidance strategies that could be implemented in the future.

Der 10 scale model car used in the collision research.
Rechte: Copyright 2017 Benedikt Jenik

More information:

https://arxiv.org/abs/1801.02805

https://arxiv.org/abs/1710.04459v1