

Learn to Drive from Pixels only

Autonomous Driving with an RC Car using Reinforcement Learning and Evolution Strategies

Graduate

Matthias König

Introduction: Machines have been an essential part of the evolution of Humanity. The industrial revolution accelerated this process and marked the rise of autonomous machines. However, only simple, repetitive tasks were automated, and decision-making remained a task achieved solely by Humans. In the age of computers, autonomous decision-making has become more important. In the last decade, automated decision-making has progressed with the increasing computation power and new methods from machine learning. Automated decision-making is used in this work to learn autonomous driving. This work is a proof of concept to show that a radio-controlled car can learn to drive by itself by using only a camera. The car learned to drive on a track marked with blue and yellow cones. A big challenge was the transfer of the algorithm from the simulation to the real-world car. Since a simulation never exactly represents reality a domain gap has to be overcome.

Approach: A 3d simulation with photographic textures was built for training the algorithms. In fig. 1, an image of the simulation is shown. Two different methods were implemented to learn to drive the car. One approach used an object detector to localize the cones which marked the track. The reinforcement learning (RL) algorithms were trained with the detected positions of the cones. The domain gap was overcome by training the object detector with real-world data.

Another method used a hard-attention mechanism learned directly from the raw pixels. This approach attempted to keep the domain gap as small as possible by using photographic textures in the simulation. Further, the algorithm was trained in different simulations in parallel. Since the attention model was not differentiable, it had to be trained with evolution strategies. The attention model was very parameter efficient. With about 5000 trainable parameters, the model used more than 1000 times fewer parameters than comparable approaches.

Result: With both approaches, good driving behaviors could be learned in the simulation. The methods were evaluated on a real-world radio-controlled car. Fig. 2 shows an image from the camera of the real car with the RL algorithm which used the object detector. The algorithm with the object detector could drive multiple laps without leaving the track. Driving with the attention-based algorithm with the real-world car was a little shaky. However, it was also possible to drive multiple laps with the attention-based approach. In challenging lighting conditions, both approaches had some difficulty. They were caused partly by the car's camera lens, which was not optimal for this application. There is potential for both algorithms to improve performance.

Advisor

Hannes Badertscher

Co-Examiner

Gabriel Sidler, Teamup Solutions AG, Zürich, ZH

Subject Area

Data Science

Project Partner

Interdisciplinary Center for Artificial Intelligence (ICAI), OST Rapperswil

Fig 1. The simulation with the track and the car. Own presentation



Fig 2. The algorithm with the object detector. Own presentation



Fig 3. The algorithm with the attention mechanism. The white patches show where the attention is. Own presentation

