Deployed End-to-End Machine Learning on an Autonomous Racing Car

Proposed supervisor: Chris Holder - chris.holder@newcastle.ac.uk

School of Computing, Newcastle University

Fig 1: F1Tenth Autonomous Racing Vehicle

Autonomous motorsport is an emerging field combining the cutting-edge of engineering and computing. While enormous progress has been made recently towards autonomous road vehicles, translating this to the domain of racing, where vehicles operate at much greater speed at the edge of their dynamic limits, poses significant challenges.

F1Tenth is an open-source project [1] providing the software and hardware designs for small-scale autonomous racing vehicles (Fig 1). The School of Computing has constructed several F1Tenth vehicles that are used in research, student projects and outreach activities, however these currently rely on a relatively simplistic autonomy system.

A typical autonomous vehicle software stack comprises separate perception, planning, and control systems, however end-to-end approaches, where a single model translates sensor data directly to vehicle controls, have generated significant interest due to their potential to be more capable of adapting to the wide variety of unanticipated scenarios an autonomous system may encounter in the real world.

End-to-end autonomous driving models can be trained via either imitation – aiming to approximate the behaviour of a human expert driver – or reinforcement learning – aiming to maximise some reward metric, for example distance travelled without collision. However, such approaches are often restricted to simulation, due to the relative ease with which large quantities of data can be generated and the potential safety risks of working with real vehicles.

The overarching aim of this project is to train and deploy an end-to-end machine learning model on board an F1Tenth vehicle capable of successfully completing laps of previously unseen track layouts. The software developed, data collected, and models trained as part of this project will facilitate highly relevant research into areas such as AI interpretability, adaptability and efficiency, as well as engagement and outreach events aimed at educating the wider public about AI and its application in autonomous vehicles.

While the machine learning models trained for this project will be required to run onboard the vehicle's limited computing hardware, the training of multiple models using large scale datasets will require HPC resources.

[1] https://roboracer.ai/