Autonomous F1Tenth Overtaking

Daniel van der Spuy
Supervised by Henry Williams and Emily Steiner

ROBOT LEARNING TEAM

1 Introduction

Overtaking is a crucial capability for autonomous robotic racing. This project explores using reinforcement learning (RL) to train F1TENTH race cars to execute safe and efficient overtaking maneuvers. RL models are developed in a custom simulation and tested on real cars.

Our approach compares state-of-the-art RL algorithms such as TD3 and SAC against conventional techniques like Follow-The-Gap to assess their effectiveness in dynamic overtaking.

Building on the CARES team's success in autonomous F1TENTH racing, this project enhances those capabilities by introducing competitive overtaking, further advancing high-speed autonomous competition.

2 Methods

Developed a simulation scenario where the RL car encounters other cars on the track during training with the following conditions:

- 2 non-RL cars spawning ahead of RL car
- The RL car is penalised for collisions
- Non-training cars use the Follow-The-Gap control method.

Spawning in training gave the following setup:

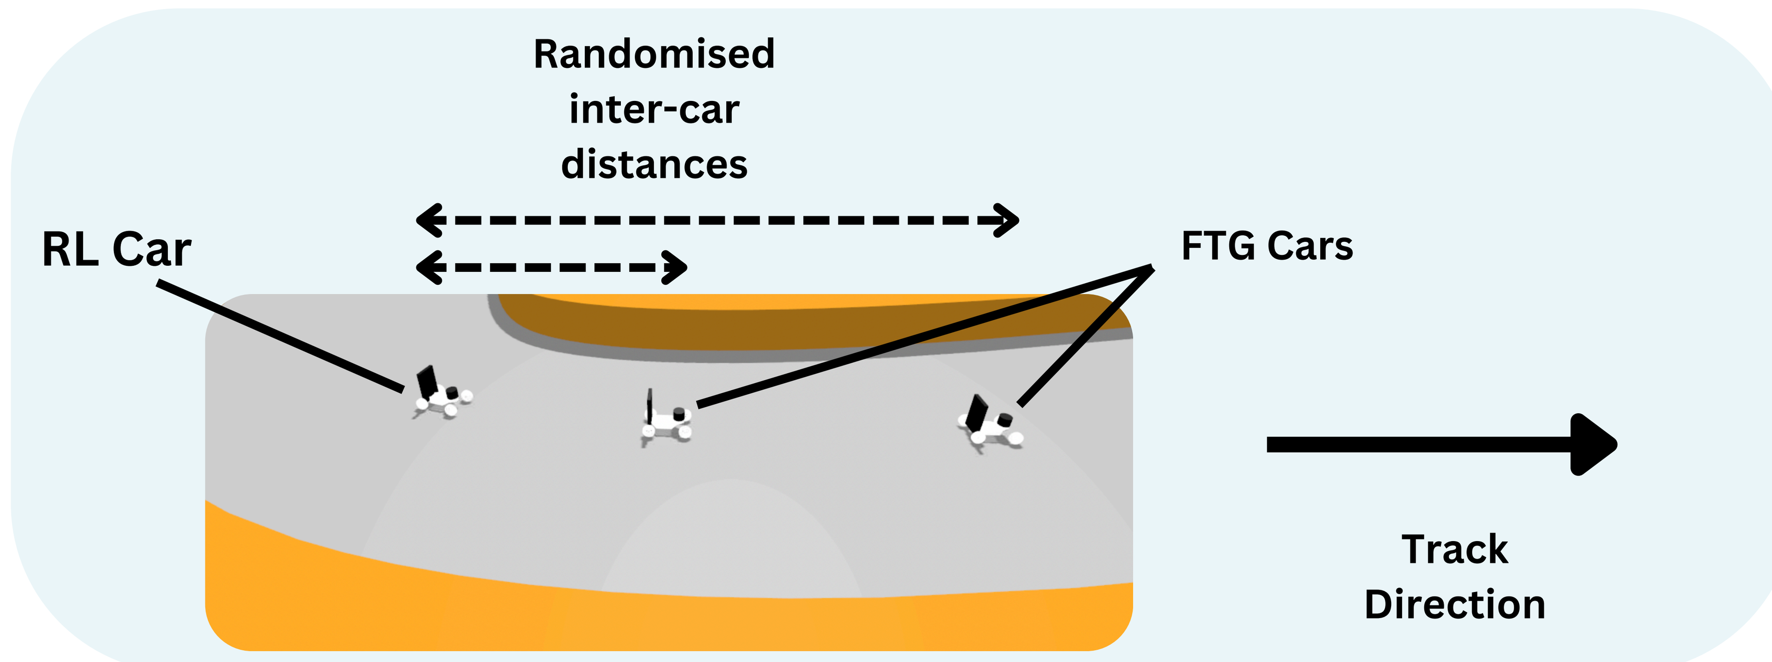

3 Simulation Results

Each model was evaulated 100 times in simulation to determine overtaking capabilities.

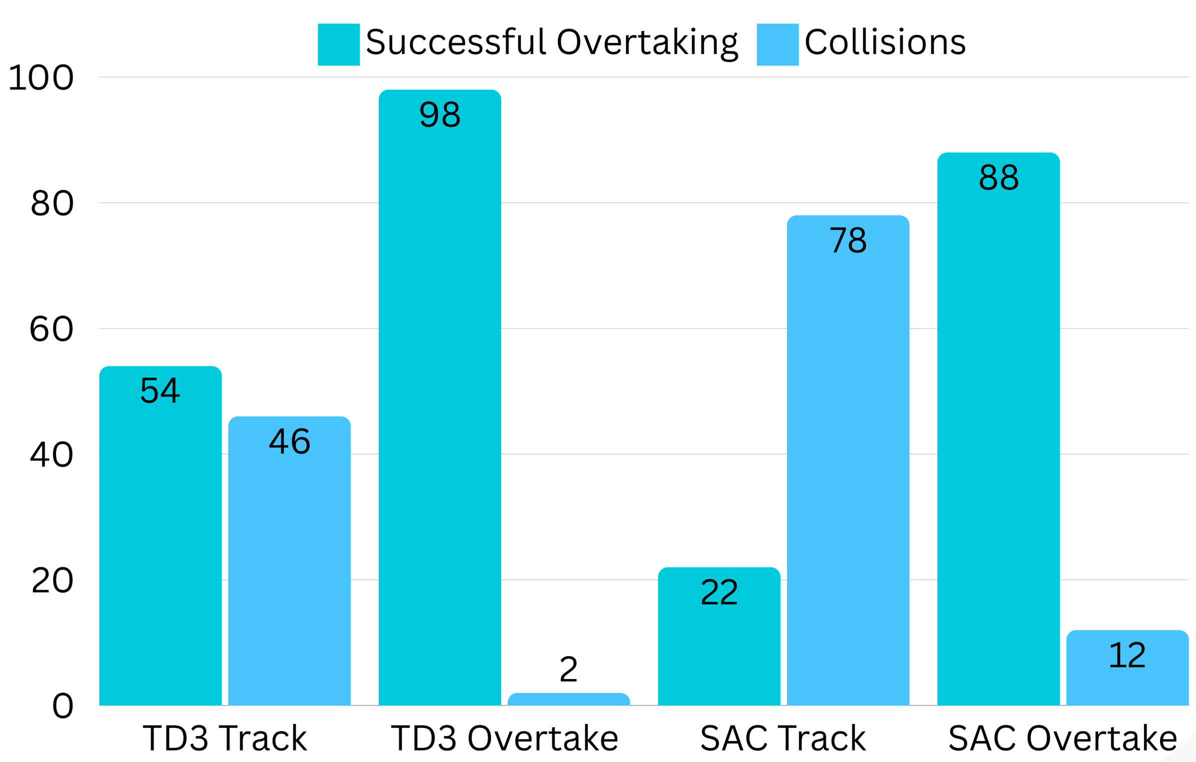

As seen in the graph, the models trained in the overtaking environment had a far higher rate of successful overtakes compared to the standard RL track following models.

Overall, TD3 performed better than SAC, although both saw a significant increase in the proportion of successful overtakes.

4 Real-World Results

The track is represented in the real-world by ducting for the cars to navigate through. Real-world testing has started as of poster publication but is not yet completed.

Preliminary results are shown below, where we can see that on the first testing day, the overtaking model performed significantly better than conventional and other RL control methods.

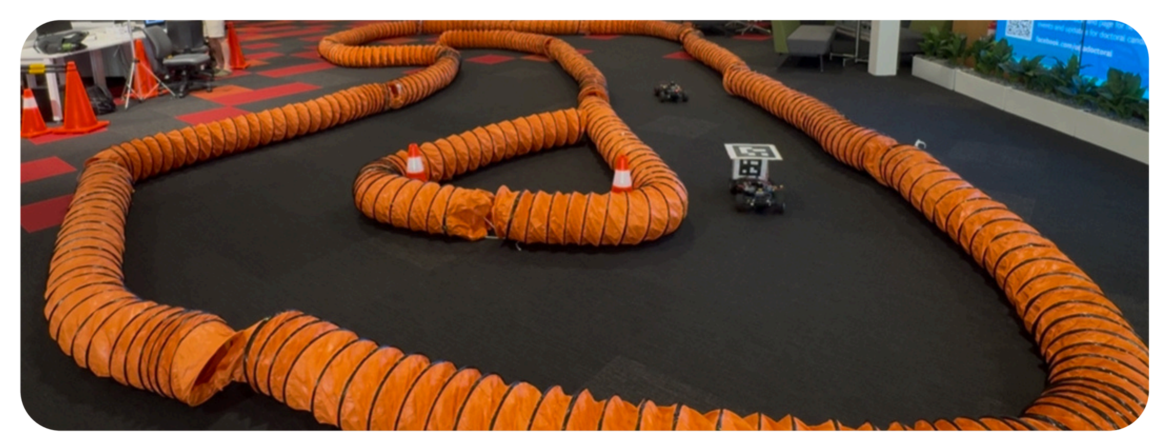

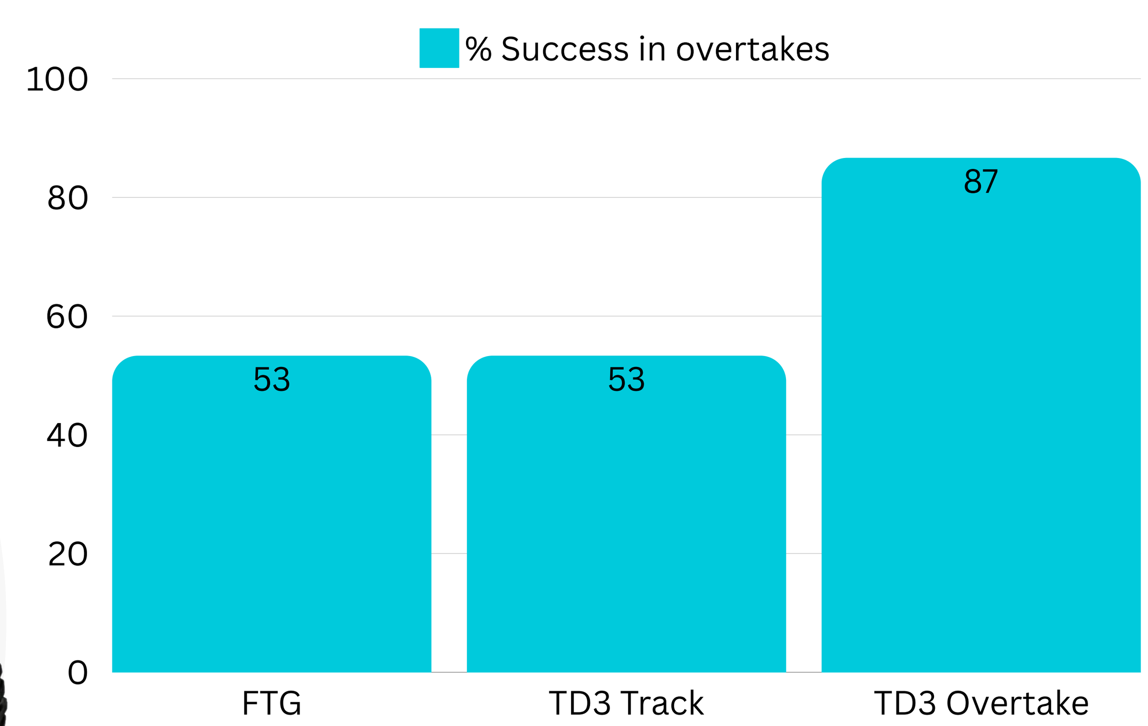

5 Discussion

The overtaking model has clearly learnt superior collision avoidance and overtaking skills compared to the Track RL algorithm. The behaviour observed resembles overtaking mechanisms similar to human driving, with careful driving around corners and opportunistic overtaking. The TD3 Overtaking model also was the only model to perform well racing another car running the same model - it is able to deviate from it's regular driving lines to execute overtaking.

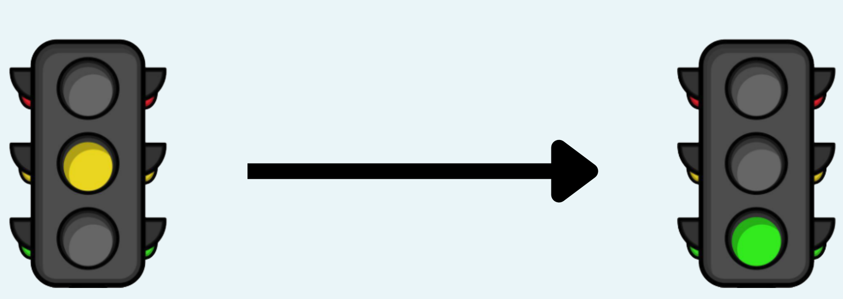

The model avoids collisions with other cars, slowing down to match the leading car's speed on tight corners where overtaking is not possible.

When the track opens up, or there is an opportunity for an overtake, the model reverts back to top speed and executes the overtaking move.

6 Future work

As shown by the results, there is potential for very effective overtaking using Reinforcement Learning methods.

In training, further improvements could be made to the overtaking scenario, such as randomising the speed of the non-learning cars to create a different challenge for the RL car every time the simulation is reset.

The reward function could also be tweaked to give the RL car a reward after successful overtaking - this could aid in encouraging the car to actively seek out overtaking of similar speed cars.

More generally, the simulator itself could be improved to add frictional forces and air resistance, decreasing the simulation-to-reality gap. This would align the performance in the real world more closely with what we see in the simulator, and would allow for more accurate RL training.

THE UNIVERSITY OF AUCKLAND
Te Whare Wānanga o Tāmaki Makaurau
NEW ZEALAND

CENTRE FOR AUTOMATION AND ROBOTIC ENGINEERING SCIENCE