Applying Reinforcement Learning to Successfully Drive a Car Around a Track

Oskaras Klimašauskas, Gintautas Dzemyda

Institute of Data Science and Digital Technologies, Vilnius University, Vilnius, Lithuania

Introduction

This work addresses the problem of controlling a car to optimise route traversal using deep reinforcement learning techniques. Two methods are investigated to enable the car to successfully traverse the track. The first method uses ten beams to measure the distances from itself to the track boundaries. The second method relies on image analysis, where the car receives information from the environment through in-game visual frames. **The goal of this work is to evaluate how different input methods influence learning alongside effects of different exploration strategies.**

Observed inputs

4 Stacked frames of 50x50 images
One frame is captured once every 4 in-game frames.

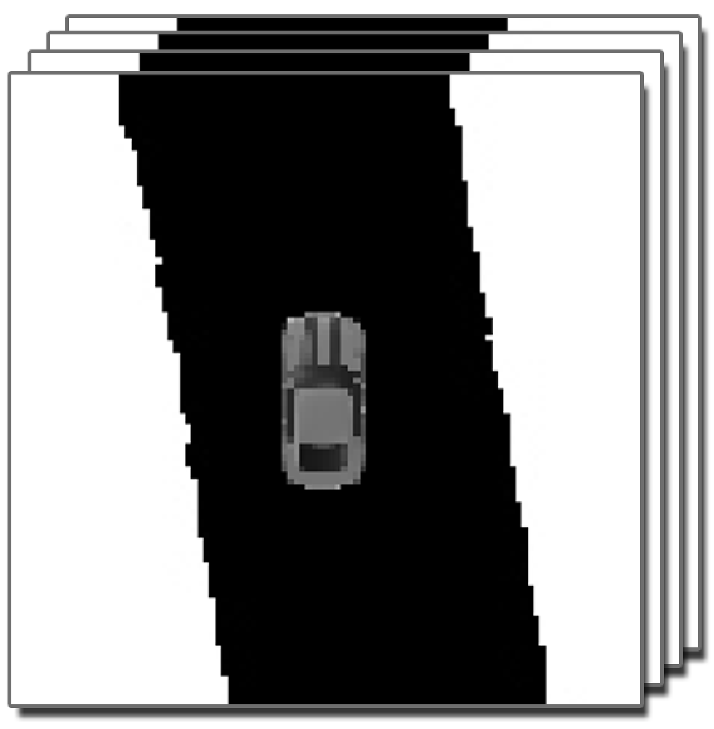

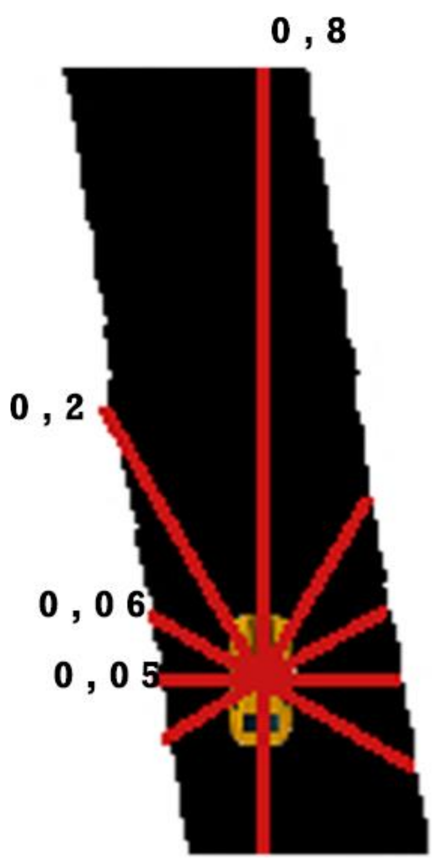

10 distance rays alongside velocity, reverse velocity, positive drift, negative drift.

Training progress

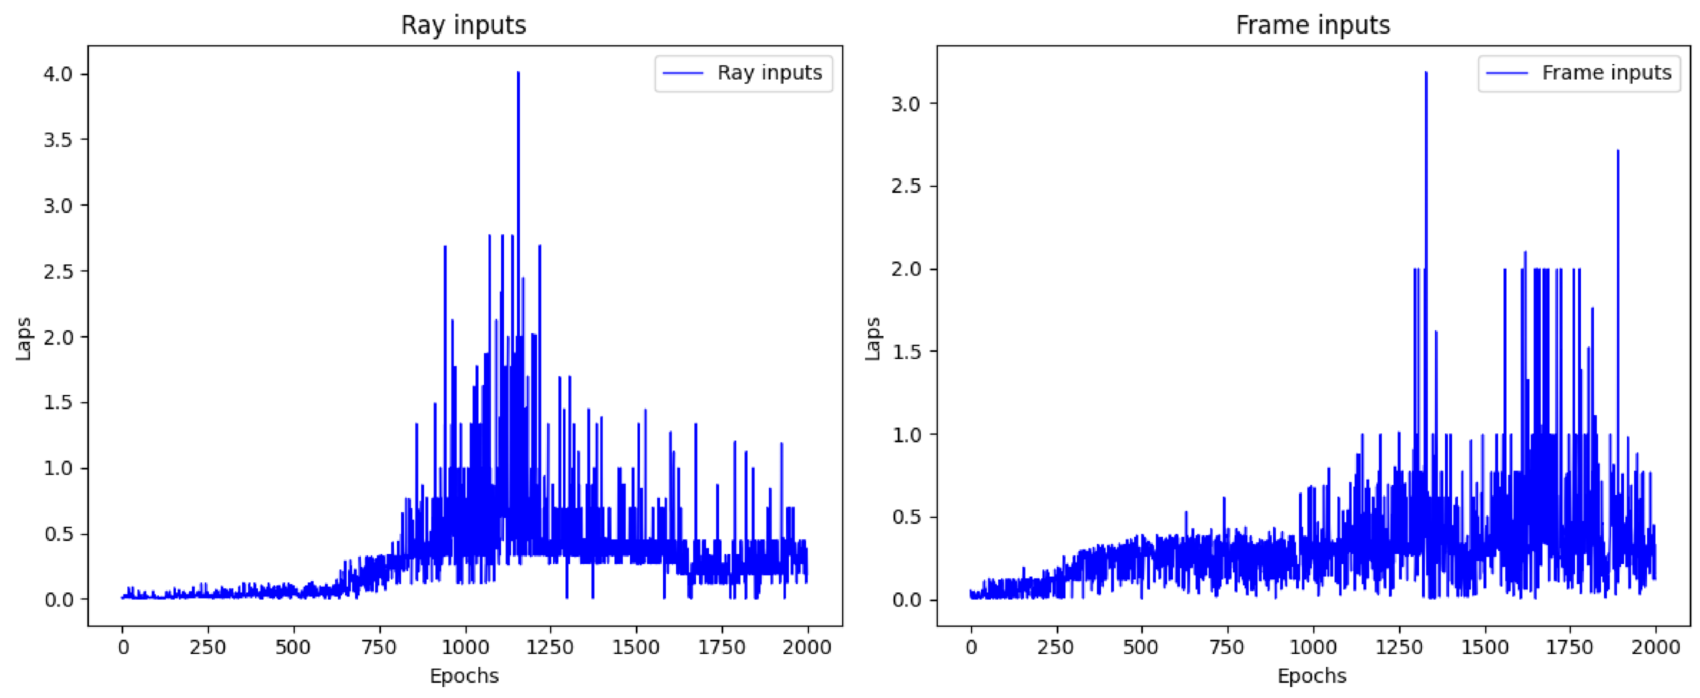

Ray inputs, with fewer features to learn, tend to overfit during training which leads to significantly worse performance over time.

Frame inputs often spend a lot of training in stagnation but generally perform better as training continues.

Trained environment

Both methods were tested on a black-and-white race track over 2000 epochs. The reward system was based on the distance the car traveled, with penalties applied for failing to drive forward. In addition to the base experiment, both methods were evaluated using interchangeable maps and random starting positions to assess their generalization capabilities.

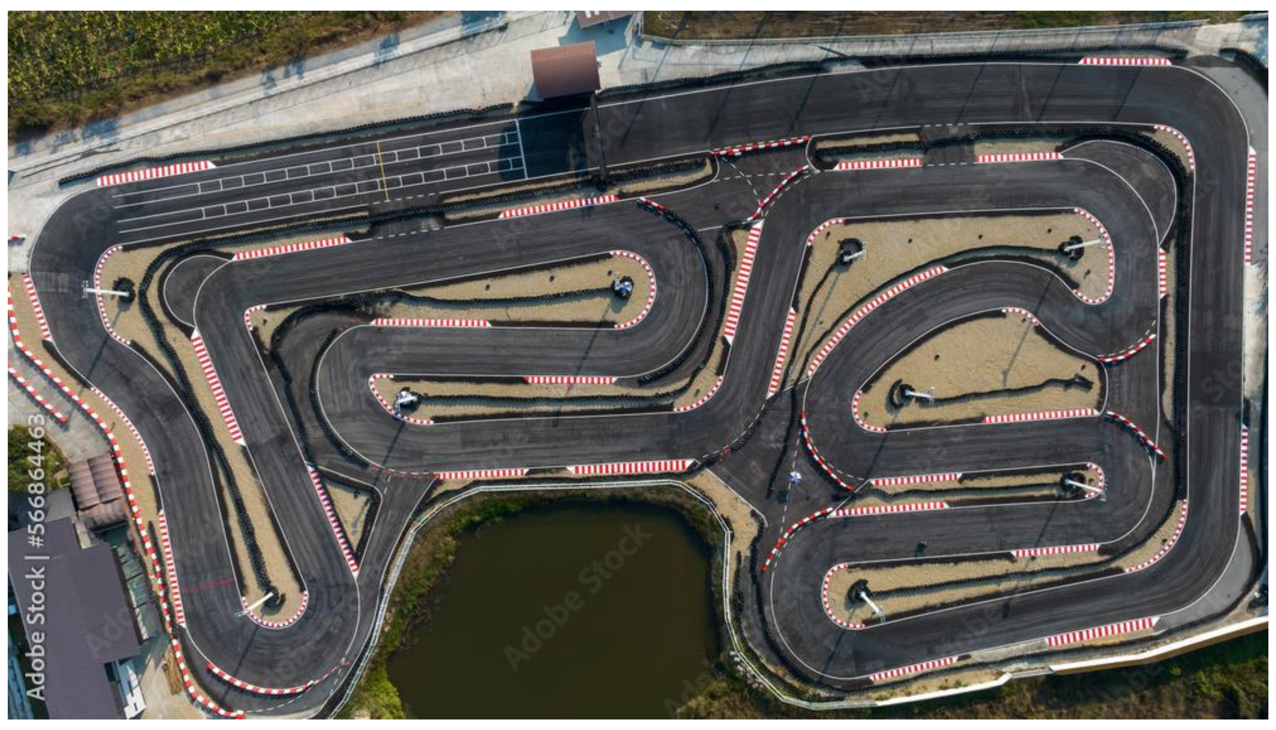

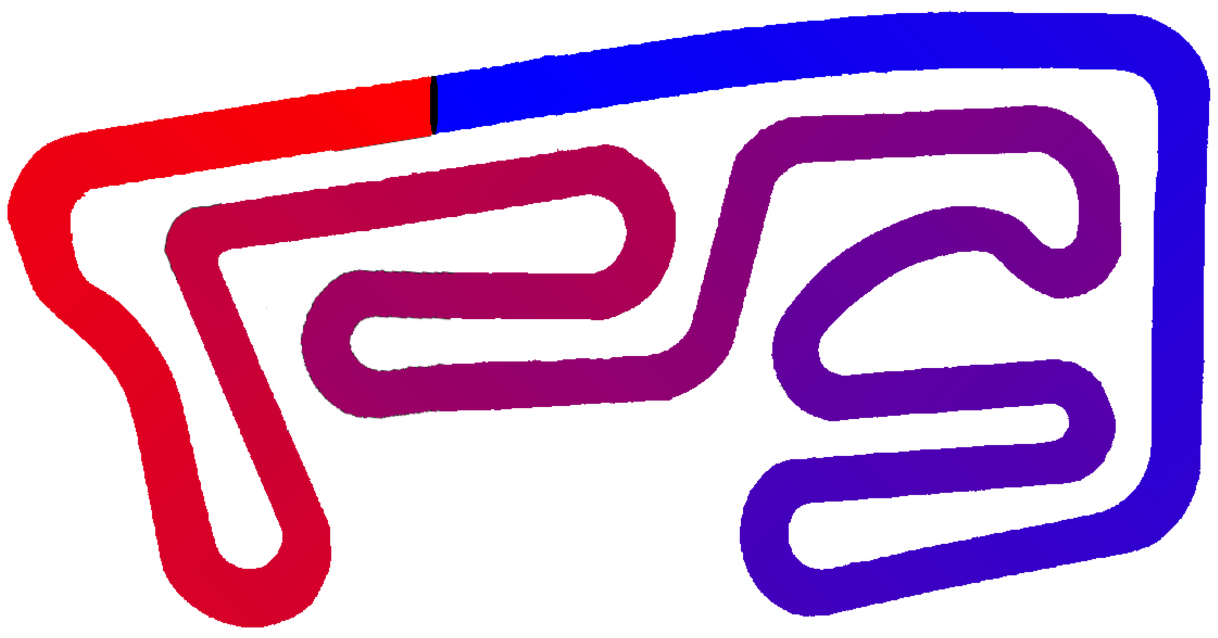

Results

1 – ray inputs
2d – frame inputs
mult – interchanging maps
rand – random starting pos.

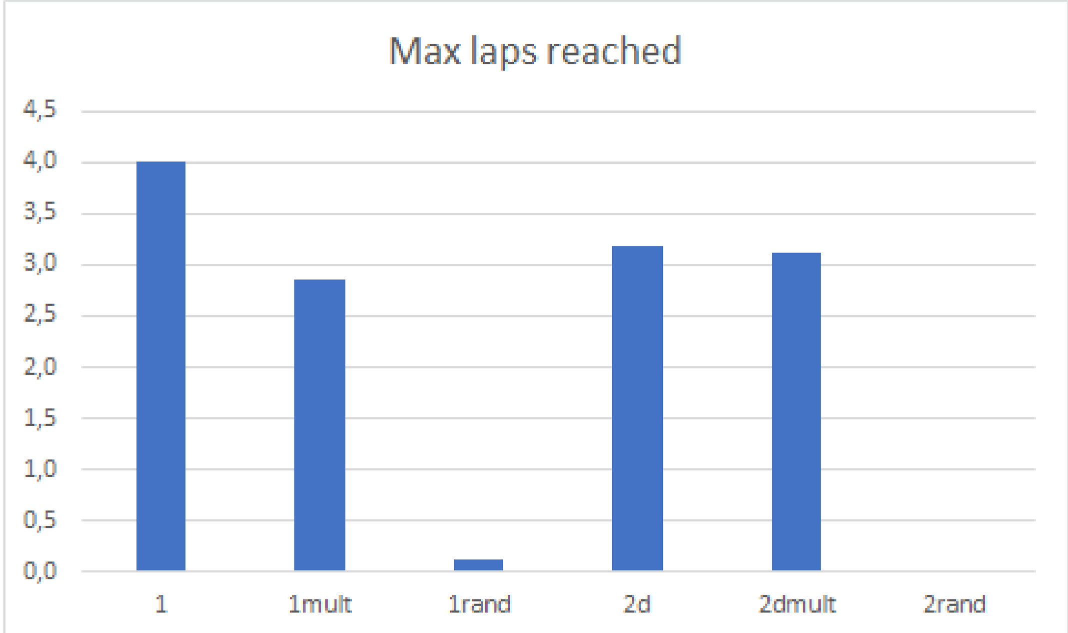

Randomized starting positions had failed the simulation completely as it resulted in too many states to explore.

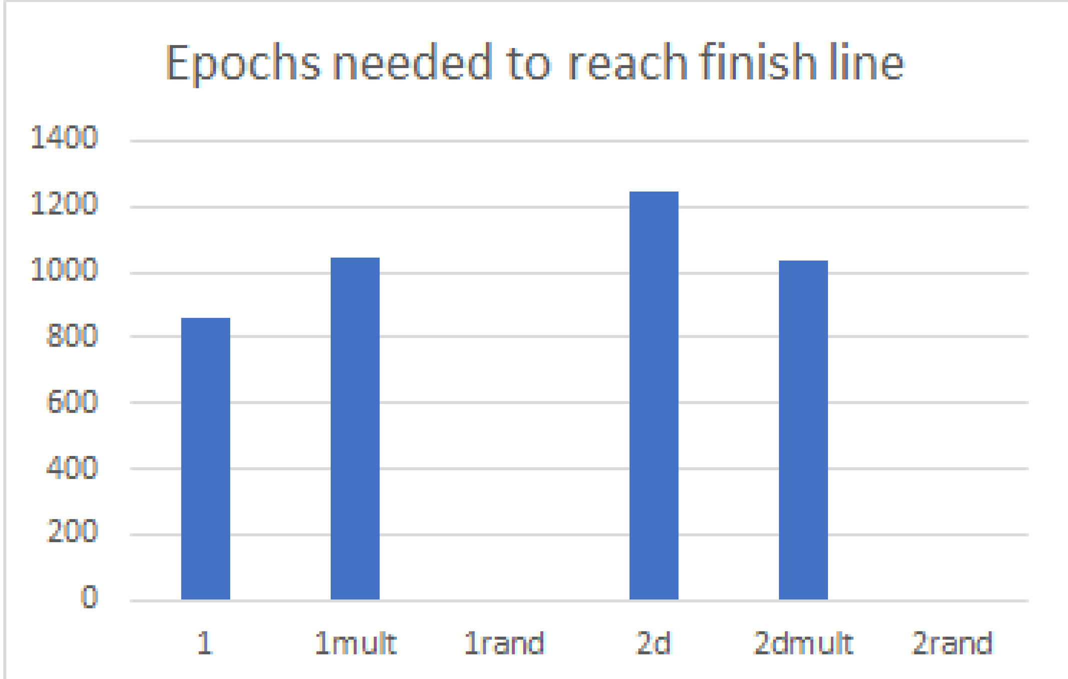

Stacking image frames performed slightly worse but had the advantage of not relying on in-game data such as velocity and distance.

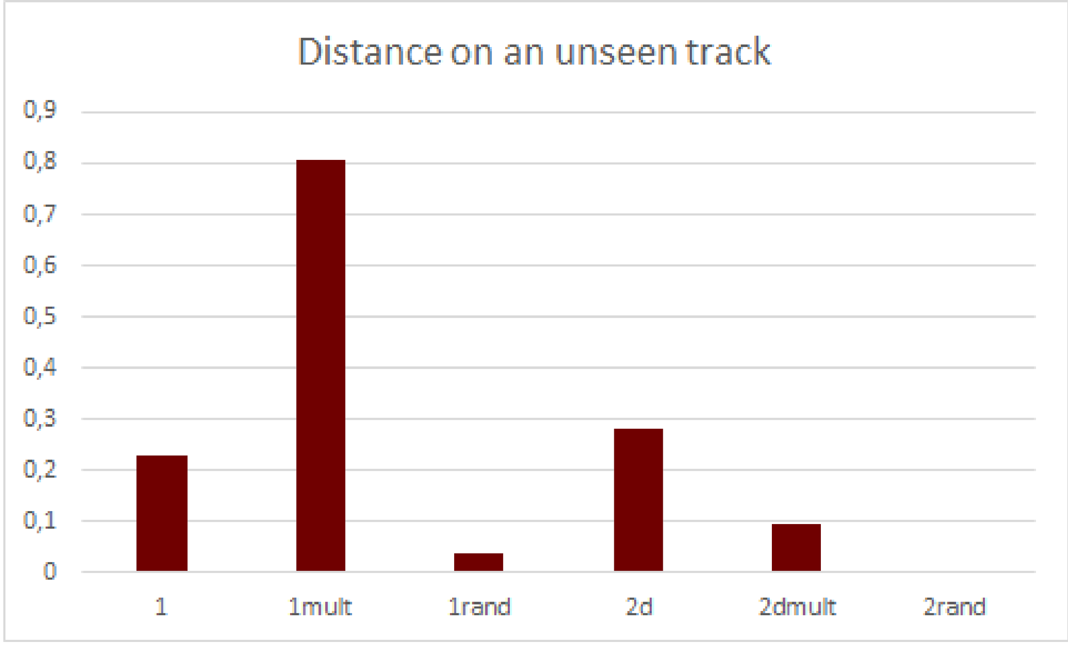

Interchanging maps with ray inputs proved to be the most effective approach for handling new maps that were not encountered during training.

Poster template by ResearchPosters.co.za