

Imitation Learning for Lap Time Simulation

Project Manager
Siwei Ju

Principal Investigator
Dr. Oleg Arenz

Project Term
2022 - 2023

Software
TensorFlow

Institute
Intelligent Autonomous Systems

University
Technische Universität Darmstadt

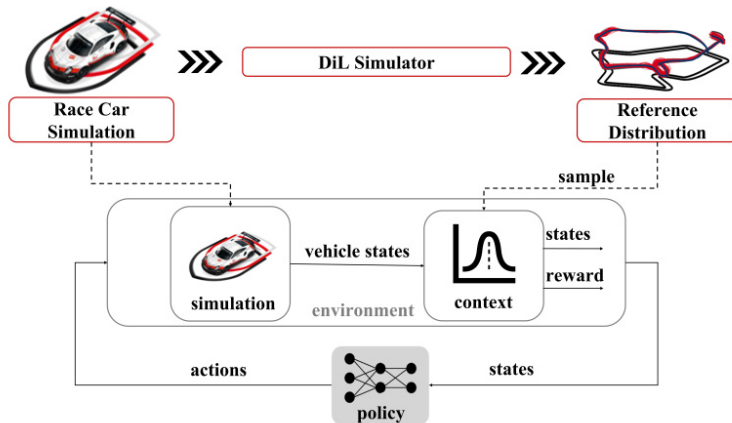


Figure 1: Our framework leverages the vehicle dynamics model to train and evaluate the agent while using the demonstrations generated by professional race drivers to provide the agent with context information during rollouts. These demonstrations are encoded into the reference distribution, allowing unlimited samples to be drawn for each rollout, enabling a probabilistic agent with human-like variance.

Introduction

In order to facilitate rapid prototyping and testing in the advanced motorsport industry, we consider the problem of imitating and outperforming professional race car drivers based on demonstrations collected on a high-fidelity Driver-in-the-Loop (DiL) hardware simulator for lap time calculation. We formulate a contextual reinforcement learning problem to learn a human-like and stochastic policy with domain-informed choices for states, actions, and reward functions. To leverage very limited training data and build human-like diverse behavior, we fit a probabilistic model to the expert demonstrations called the reference distribution, draw samples out of it, and use them as context for the reinforcement learning agent with context-specific states and rewards. In contrast to the non-human-like stochasticity introduced by Gaussian noise, our method contributes to a more effective exploration, better performance and a policy with human-like variance in evaluation metrics. Compared to previous work using a behavioral cloning agent, which is unable to complete competitive laps robustly, our agent outperforms the professional drivers.

Methods

Instead of conventional lap time calculation methods based on control theory with physical vehicle models, our work aims to develop a driver model based on imitation learning that can potentially facilitate the digital twin of DiL simulations. Our contribution is formulating race car driving as a contextual

reinforcement learning problem, and leveraging human demonstrations as context to enhance performance and efficiency, while learning human-like behavior and variance. We initialize the agent with behavioral cloning, and use reinforcement learning to improve the performance and robustness. Having only limited demonstrations, we fit them into a prior probability distribution called the reference distribution, and draw samples from it as the context during explorations. Using context-aware states and rewards informed by domain-knowledge, we train the policy by optimizing the expected return over all instances of the references, contributing to a more effective exploration and a stochastic policy with human-like variance, compared to the default step-based Gaussian noise for explorations. In practice, the demonstrations are generated by professional drivers in the DiL simulator.

Results

During the demonstration and training, the setting of the car is the same as in the professional races, without any driver assistance system such as traction control or anti-lock braking system (ABS), making it more challenging to control. Being trained and evaluated in the same simulation environment as in the DiL simulator with a simplified track model, our agent outperforms top-class professional race car drivers by 0.4 secs on average on a real racetrack in comparable scenarios. In addition, being optimized over the reference distribution, the agent learns a stochastic policy, exhibits realistic and human-like variance in terms of the driving line, the speed profile and the lap time. The agent is capable of completing laps even when faced with altered vehicle characteristics such as changes in power, grip, or balance, while predicting plausible lap time changes. The data generated by the agent can be used to predict metrics such as lap time, top speed, and driveability.

Discussion

For future work, developing a track- and setup-universal agent that can operate on various tracks and setups and reach the optimums without retrain is a compelling idea. Furthermore, more research on imitating individual driving styles of human drivers is of great interest to provide human-like driveability predictions.

Publications

Ju, S; van Vliet, P.; Arenz, O.; Peters, J. Digital Twin of a Driver-in-the-Loop Race Car Simulation With Contextual Reinforcement Learning. In IEEE Robotics and Automation Letters, vol. 8, no. 7, pp. 4107-4114 (2023) <https://doi.org/10.1109/LRA.2023.3279618>

Last Update: 2023-11-07 18:36