

Super-Human Performance in Gran Turismo Sport Using Deep Reinforcement Learning

Florian Fuchs¹, Yunlong Song², Elia Kaufmann², Davide Scaramuzza², and Peter Dürri¹

Abstract—Autonomous car racing is a major challenge in robotics. It raises fundamental problems for classical approaches such as planning minimum-time trajectories under uncertain dynamics and controlling the car at its limits of handling. Besides, the requirement of minimizing the lap time, which is a sparse objective, and the difficulty of collecting training data from human experts have also hindered researchers from directly applying learning-based approaches to solve the problem. In the present work, we propose a learning-based system for autonomous car racing by leveraging high-fidelity physical car simulation, a course-progress-proxy reward, and deep reinforcement learning. We deploy our system in Gran Turismo Sport, a world-leading car simulator known for its realistic physics simulation of different race cars and tracks, which is even used to recruit human race car drivers. Our trained policy achieves autonomous racing performance that goes beyond what had been achieved so far by the built-in AI, and, at the same time, outperforms the fastest driver in a dataset of over 50,000 human players.

Index Terms—Autonomous Agents, Reinforcement Learning

SUPPLEMENTARY VIDEOS

This paper is accompanied by a narrated video of the performance: <https://youtu.be/Zeyv1bN9v4A>

I. INTRODUCTION

THE goal of autonomous car racing is to complete a given track as fast as possible, which requires the agent to generate fast and precise actions even when the vehicle reaches its physical limits. Classical approaches [1]–[3] generally decouple the racing problem into trajectory planning and vehicle control. From the trajectory-planning standpoint, the problem is about how to compute the minimum-time trajectory for a given race track. From the vehicle control standpoint, it captures the main difficulty and the most critical scenario in autonomous driving—how to safely control the vehicle during very extreme manoeuvres.

Manuscript received: October 15, 2020; Revised January 9, 2021; Accepted February 12, 2021.

This paper was recommended for publication by Markus Vincze upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by Sony R&D Center Europe Stuttgart Laboratory 1, the National Centre of Competence in Research (NCCR) Robotics through the Swiss National Science Foundation, and the SNSF-ERC Starting Grant.

¹F. Fuchs and P. Dürri are with Sony Europe B.V., Schlieren/Switzerland Branch (Florian.Fuchs@sony.com)

²Y. Song, E. Kaufmann, D. Scaramuzza are with the Robotics and Perception Group, University of Zurich, Switzerland (<http://rpg.ifi.uzh.ch>)

Digital Object Identifier (DOI): see top of this page.

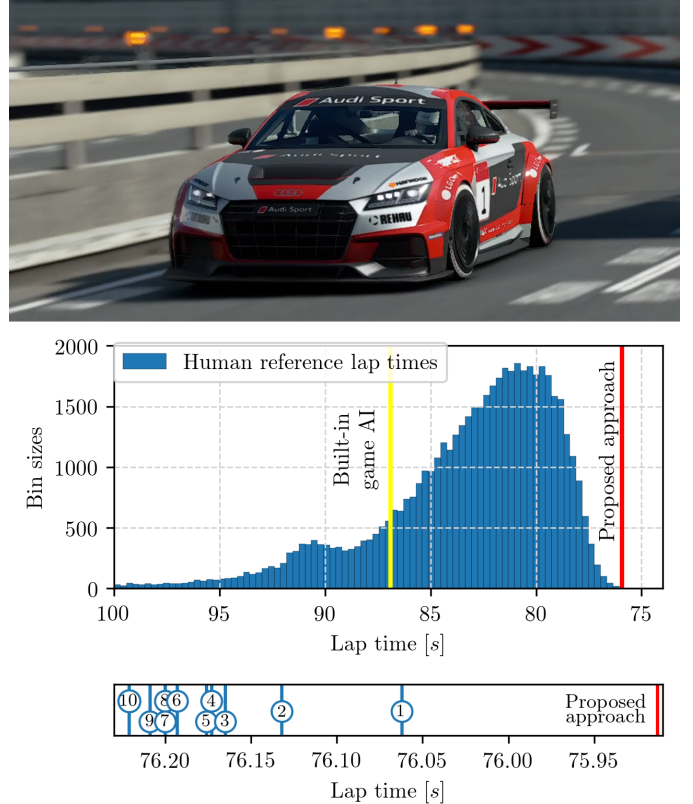


Fig. 1. **Top:** Our approach controlling the “Audi TT Cup” in the *Gran Turismo Sport* simulation. **Center:** Time trial lap time comparison between our approach (red line), the personal best lap times of over 50,000 human drivers from 70 countries (dark blue histogram), as well as the built-in non-player character (“game AI”, yellow line). Human lap times over 100 seconds are cut off to simplify the visualization. **Bottom:** Closer view at the lap times of the 10 fastest human drivers and our approach.

Classical approaches have been widely studied for addressing the autonomous car racing problem and have shown impressive results. In particular, much of the success relies on optimization-based techniques for trajectory generation and tracking, thanks to their capabilities of handling different constraints and nonlinear dynamics. However, several limitations do exist in this line of research, such as the requirement of expensive computations for the nonlinear optimization and the lack of flexibility in the objective (most optimizations require quadratic cost formulations). Overall, conventional systems for autonomous driving escalate in complexity as the model complexity increases and do not leverage the large volume of data generated by real-world driving.

The requirements of improving vehicle models using real-

world data and handling more complex cost formulations and general nonlinear dynamics motivated researchers to design learning-based systems. The first is related to dynamic model identification. In particular, multi-layer neural network models have shown to be helpful for high-speed driving when combining them with model predictive path integral control (MPPI) [1], [4], [5] or a simple feedforward-feedback control [6]. However, in order to run MPPI online in a receding horizon manner in a fast control loop, the optimization requires a highly parallel sampling scheme and heavily relies on graphics processing units (GPUs).

The second trend pertains to end-to-end vehicle control using neural networks, where a trained neural network policy can map high-dimensional raw observations directly to control commands. Hence, it has the advantage of forgoing the need for explicit state estimation and solving optimizations online, rendering fast and adaptive control performance. Combined with advanced dynamic model identification, neural network controllers have shown numerous achievements in many real-world robot applications [7]–[10].

In this work, we tackle the autonomous car racing problem in Gran Turismo Sport (GTS), which is a world-leading car simulator known for its realistic modeling of various race cars and tracks. GTS is a unique platform for studying autonomous racing systems, as it allows simulating a large number of realistic environments and benchmarking comparisons between autonomous systems and experienced human drivers. Such comparisons can offer invaluable insights into vehicle control during extreme manoeuvres.

Our approach makes use of model-free deep reinforcement learning and a course-progress-proxy reward to train a multilayer perceptron policy for the vehicle control. The key is formulating the minimum-time racing problem properly to maximizing the course progress and choosing meaningful low-dimensional state representations for efficient policy training. As a result, we demonstrate the first super-human control performance on the task of autonomous car racing. The strength of our controller is two-fold: 1) it does not rely on high-level trajectory planning and following (exactly what human drivers tend to do) 2) it generates trajectories qualitatively similar to those chosen by the best human drivers while outperforming the best known human lap times. Our findings suggest that reinforcement learning and neural network controllers merit further investigation for controlling autonomous vehicles at their limits of handling.

II. RELATED WORK

Prior work in the domain of autonomous racing can be grouped into three groups: (i) classical approaches relying on trajectory planning and following, (ii) supervised learning approaches, and (iii) reinforcement learning approaches.

Classical Approaches Classical approaches to autonomous car racing approach the problem by separating it in a chain of submodules consisting of perception, trajectory planning, and control. In particular, model predictive control (MPC) [2], [11]–[15] is a promising approach for controlling the vehicle at high speed. In [16], an MPC controller is combined with

learned system dynamics based on Gaussian Processes for the task of autonomous car racing. While being able to capture the complex dynamics of the car, the resulting controller requires to trade-off dynamical fidelity against computation time. Similarly, model predictive path integral control (MPPI) [1], [4], [5] is a more flexible approach that can be combined with complex cost formulation and neural network vehicle models. Both MPC and MPPI have shown impressive results in controlling physical vehicles at high speed in the real-world. Despite their successes, both approaches have limitations, such as the lack of flexibility in the cost function design or the requirement of highly parallel computations. Moreover, hierarchical approaches are also susceptible to failures of each submodule, where the operational states of each module are limited by a series of approximations and linearizations.

Imitation Learning Instead of planning trajectories and tracking them with a controller, imitation-based approaches directly learn a mapping from observation to control action in a supervised fashion. Learning such mapping requires labelled data, which is typically provided by human expert demonstrations or a classical planning and control pipeline. For example, ALVINN (Autonomous Land Vehicle in a Neural Network) [17] is one of the first autonomous driving systems that uses a neural network to follow a road. Similarly, a convolutional neural network (CNN) controller was trained by [18] for lane and road following. Pan et al. [19] use imitation learning for agile, off-road autonomous driving. Distilling human or algorithmic expert data in a learned policy is a promising approach to overcome the strict real-time constraint of classical approaches, but its performance is by design upper-bounded by the quality of the training data.

Reinforcement Learning Model-free reinforcement learning optimizes parametrized policies directly based on sampled trajectories, and hence, does not suffer from issues that have been hampering the aforementioned approaches, such as the necessity of solving nonlinear optimizations online or the dependence on labelled training data. For example, a number of studies [20]–[24] have demonstrated the success of using model-free deep RL for end-to-end driving and racing. Recently, a high-speed autonomous drifting system [24] was developed in simulation using the soft actor-critic (SAC) algorithm [25]. Here, off-policy training plays an important role in [21], [22], [24], in which the high sample complexity previously limiting the widespread adoption of deep RL methods in high-dimensional domains, was substantially reduced by combining offline training with an experience replay buffer. Despite the successful application of deep RL algorithms in real as well as simulated autonomous driving, there is, to the best of our knowledge, no work matching or exceeding the performance of human expert drivers regarding speed.

III. METHODOLOGY

Our main goal is to build a neural network controller that is capable of autonomously navigating a race car without prior knowledge about the car’s dynamics while minimizing the traveling time on a given track in the GTS environment. To achieve this goal, we first define a reward function that

formulates the racing problem and a neural network policy that maps input states to actions. We then optimize the policy parameters by maximizing the reward function using the SAC [25] algorithm. An overview of our system is shown in Figure 2.

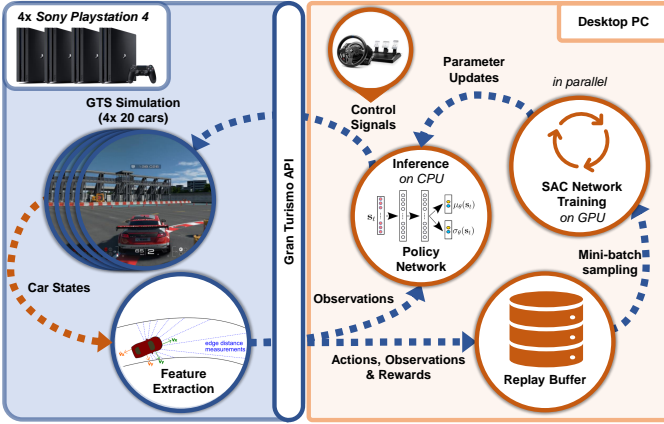


Fig. 2. System overview: we train a policy network to directly map observations, including a set of range finder measurements and the car’s velocity and acceleration, to control commands consisting of the car’s steering angle as well as the level of throttle and brake. We use a distributed sampling scheme to collect samples from 4 *PlayStation 4* simulating 20 cars each, and then, store the sampled trajectory in a fixed size first in, first out (FIFO) replay buffer. In parallel, we optimize the policy parameters using a soft actor-critic algorithm and data uniformly sampled from the replay buffer.

A. Minimum-Time Problem and Reward Function

We aim to find a policy that minimizes the total lap time for a given car and track. The lap time itself is a very sparse reward signal. Changes in the signal are therefore hard to attribute to specific actions of an agent. We design a proxy-reward based on the current course progress, which can be evaluated in arbitrary time intervals. To reduce the variance of the reward signal, we use an exponential discount for future rewards, with discount factor γ . Maximizing the course progress closely approximates minimizing the lap time when choosing a high enough discount factor. The new proxy-reward allows trading-off between an easily attributable but biased reward and a reward closer to the overall lap time goal, by adjusting γ and thus changing the weight of future rewards based on their temporal distance to an action and state. The construction of the progress reward can be seen in the top of Figure 3.

The use of an exponentially discounted future reward leads to a bias towards short-term rewards. In the GTS setting this bias reduces the incentive for an RL agent to brake, e.g. to prevent crashes. To counteract this short-term bias, we introduce a second reward term that penalizes wall contact relative to the car’s kinetic energy, leading to the final reward function of

$$r_t = r_t^{\text{prog}} - \begin{cases} c_w \|\mathbf{v}_t\|^2 & \text{if in contact with wall,} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $c_w \geq 0$ is a hyperparameter controlling the trade off between wall contact penalty and progress-based reward. Here $\mathbf{v}_t = [v_x, v_y, v_z]$ is a velocity vector that represents the current




Fig. 3. **Top:** The course progress cp_t at time t is constructed by projecting the car’s position on the track’s center line. We then define the progress-based reward r_t^{prog} as the increment in centerline-progress $r_t^{\text{prog}} = cp_t - cp_{t-1}$. **Bottom:** A subset of the observations fed to the policy and value networks.

linear velocity of the vehicle. A similar approach to incentivize collision avoidance was introduced in [26]. The introduction of the kinetic energy is justified through the energy-dependent loss in acceleration that takes place when hitting a wall. Without this additional wall contact penalty, we found the learned policies did not break and simply grinded along the track’s walls in steep curves. When using fixed valued wall contact penalties, we found the agent to either not react to the penalty or directly end up in a strategy of full braking and standing still to not risk any wall contact, depending on the strength of the penalty.

B. Policy Network

We represent the driving policy with a deep neural network. We make use of the SAC network architecture proposed in the original paper [25]. We use a policy network, two Q-function networks, and a state-value function network, each with 2 hidden layers with 256 ReLU nodes, resulting in a total of 599,566 trainable parameters. In the following, the features fed to the network and its prediction are explained in detail.

Input features We base the input feature selection on the results of a pre-study on GTS conducted with behavioral cloning, aiming to learn human-like driving based on a regression of expert’s actions on visited states. We iterated over combinations of subsets of features provided by GTS as well as additionally constructed features, which lead to the following feature selection:

- 1) The linear velocity $\mathbf{v}_t \in \mathbb{R}^3$ and the linear acceleration $\dot{\mathbf{v}}_t \in \mathbb{R}^3$.
- 2) The Euler angle $\theta_t \in (-\pi, \pi]$ between the 2D vector that defines the agent’s rotation in the horizontal plane and the unit tangent vector that is tangent to the centerline at the projection point. This angle is the only direct way for the

policy network to detect if a car is facing the wrong direction. 3) Distance measurements $\mathbf{d}_t \in \mathbb{R}^M$ of M rangefinders with a maximum range of 100m that measure the distance from the vehicle’s center point to the M edge points of its surrounding objects, such as the edge of the race track. The rangefinders are equally distributed in the front 180° of the car’s view. 4) The previous steering command δ_{t-1} . 5) A binary flag with $w_t = 1$ indicating wall contact and 6) N sampled curvature measurement of the course centerline in the near future $\mathbf{c}_t \in \mathbb{R}^N$. The curvature is represented through an interpolation of the inverse radii of circles fitted through centerline points provided by GTS.

We therefore represent the observation at a given time step t as a vector denoted as $\mathbf{s}_t = [\mathbf{v}_t, \dot{\mathbf{v}}_t, \theta_t, \mathbf{d}_t, \delta_{t-1}, w_t, \mathbf{c}_t]$. A subset of the features is illustrated in the bottom of Figure 3. To allow for a fair comparison with human drivers, we only use features that humans can either directly perceive or deduce from the GTS simulation.

Network prediction The output of the policy network $\mathbf{a}_t = [\delta_t, \omega_t]$ directly encodes the steering angle $\delta_t \in [-\pi/6, \pi/6]$ rad and a combined throttle-brake signal $\omega_t \in [-1, 1]$, where $\omega_t = 1$ denotes full throttle and $\omega_t = -1$ full braking. The combination of throttle and brake in a single signal is motivated by an analysis of human recordings, which revealed that fast strategies do not involve any simultaneous use of throttle and brake.

IV. EXPERIMENTS

We evaluate our approach in three race settings, featuring different cars and tracks of varying difficulty. We compare our approach to the built-in AI as well as a set of over 50,000 human drivers. To ensure a fair comparison to the human drivers, we constrain the policy to produce only actions that would also be feasible with a physical steering wheel. In the following, each race setting as well as the constraints imposed for the comparison with human drivers are explained.

Race settings We train separate agents for three experimental conditions for which human data from past online time trial competitions is available. The data, provided by GTS manufacturer *Polyphony Digital Inc.*, includes the personal best lap time and trajectory of every competitor, ranging from absolute beginners to world cup contenders. The competitions were restricted to fixed settings such as car, course, tires, and racing assistance settings. This allows for equal conditions when comparing human results with our approach. Figure 4 shows the used cars and tracks. Setting *A* and *B* use a track layout featuring a typical combination of curves and straight segments. The two settings differ in the cars used, with the “Audi TT Cup ’16” of setting *A* having a higher maximum speed, more tire grip, and a faster acceleration than the “Mazda Demio XD Turing ’15” of setting *B*. Setting *C* features the same car as setting *A* but a more challenging track layout which features a larger range of speeds due to the combination of long straight segments and very tight turns.

Leveling the playing field Whereas human players are limited by the physical inertia of their game controllers, our agent can, in principle, make arbitrarily large action changes.




Fig. 4. The tracks and cars used as reference settings to compare our approach to human drivers.

To allow for a fair comparison between our agent and human players, we estimated the maximal action change that a human can achieve within one frame with a *Thrustmaster T300* steering wheel and pedals. Based on that, we restricted the maximum change between frames during evaluation to 0.03 rad for the steering angle and 80% for each the throttle and brake range¹. We however found this restriction to not significantly influence the resulting lap times.

Robustness To detect potential shortcoming of our approach for a more general driving task, such as driving on multiple tracks or driving in the real world, we evaluate our agent’s robustness by testing its ability to generalize to a range of modified environments. We deploy the fastest agent trained in setting *A* on the following problem settings without any retraining: 1) transferring the policy to a different car (setting *B*), 2) transferring the policy to a different track (setting *C*), 3) changing the tire friction, 4) adding uniform noise to the agent’s observations, and 5) delaying the agents inference.

V. RESULTS

In this section we evaluate our approach on the three reference settings introduced in Section IV. We compare the lap times obtained by our approach with the ones achieved by the built-in AI of GTS and the fastest human drivers. Additionally, we analyze the resulting driving behaviour and compare it with driving strategies of professional human drivers. Due to the dynamic nature of our experiments, we encourage the reader to watch the supplementary videos.

A. Lap Time Comparison

Our approach outperforms the best human lap time in all three reference settings, overcoming the limitations of the currently built-in AI, which is outperformed by a majority of players. Table I shows the fastest achieved lap times for the three race settings.

On the first track our approach undercuts the fastest human lap times by 0.15 and 0.04 seconds for settings *A* and *B* respectively. We believe that the difference in margins to the best human drivers results from the speed differences between the two used cars, with the faster “Audi TT Cup” in

¹A similar fairness restriction was put on the agent proposed in [27], by limiting the number of actions an agent could take per minute.

TABLE I
TIME TRIAL COMPARISONS BETWEEN OUR APPROACH, HUMAN ONLINE COMPETITORS, AND THE BUILT-IN GTS AI FOR THE 3 RACE SETTINGS.

Driver	Metric	Setting A	Setting B	Setting C
Ours	Lap time [min]	01:15.913	01:39.408	02:06.701
Human players	Fastest lap [min]	01:16.062	01:39.445	02:07.319
	Median lap [min]	01:22.300	01:47.259	02:13.980
	# participants	52,303	28,083	52,335
Built-in GTS AI	Lap time [min]	01:26.899	01:52.075	02:14.252
	Slower than x% of humans drivers	82.6%	80.4%	53.9%

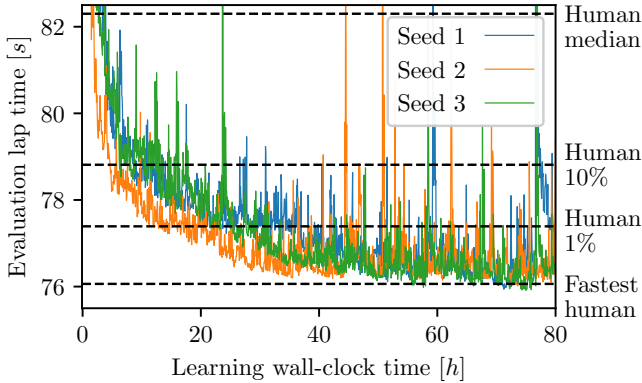


Fig. 5. Training progress for our approach learning on setting A. We evaluate the policy every 2 epochs of learning by letting the agent drive 2 consecutive laps. We take the lap time at the second lap as performance metric such that the fixed starting speed applied in the first lap does not distort the result.

setting A requiring a more agile strategy than the relatively slow “Mazda Demio” in setting B. While human players possibly struggle with the fast-paced setting, our approach does not suffer under the increased demands. Due to the similarities of the trajectories of setting A and B, we will in the following part only analyze the trajectories of setting A. Over 10 consecutive laps in setting A our agent’s lap time shows a standard deviation of 15 ms. A top 1% driver we invited for our study showed a standard deviation of 480 ms. In setting C our approach undercuts the best human time by 0.62 seconds. As in the other two settings, the increased margin can be explained by the more challenging track layout and faster car that render setting C the most difficult combination in our experiments.

Figure 5 shows the learning progress in setting A for three differently initialized neural network policies. With each random initialization our approach learns policies that achieve lap times faster than the fastest human reference lap after 56 to 73 hours of training, which in average corresponds to 2,151 training epochs and a total of 946,453 km driven. The 3 seeds result in similar learning progresses that mainly differ in the time and order in which they learn to master certain curves. Note that warm-starting the training process from a policy that is trained via imitation learning from human data does not improve final performance. Training the policy from scratch outperforms the behaviour-cloning policy after less than one our of training.

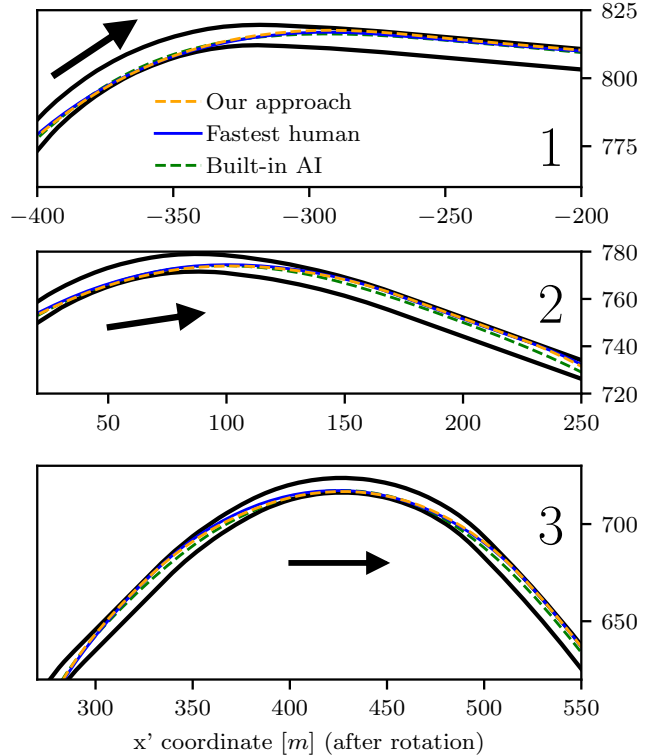


Fig. 6. Top-down view of driven paths for setting A. While the paths of our approach and the human driver are similar, the built-in GTS AI drives with a safety margin to the track’s walls, leading to tighter radii and a reduced speed.

B. Learned Driving Behavior

We analyze the models’ driving behavior at the epoch with the fastest evaluation lap time.

Out-in-out trajectory Our approach learned to make use of the whole width of the track to maximize its trajectories’ curve radii, driving a so called out-in-out trajectory. This allows the agent to drive higher speeds before losing traction. Figure 6 shows that learned behavior in a curve driving comparison between our approach, the fastest human, and the built-in AI for 3 curves in setting A. Our approach learned to drive curves similar to those of the fastest human, without having access to any human demonstration or using an explicit path planning mechanism. The top of Figure 7 shows the learned out-in-out trajectory for the steepest curve of setting A.

Anticipating curves Furthermore, our approach learned to detect curves early enough and assess their sharpness to decelerate to speeds that allow completing the curves without overshooting into the track’s walls, while at the same time not being overly cautious. This can be seen in the steepest curve of setting A in the top of Figure 7 and in a more extreme case in the hairpin of setting C in the bottom of the Figure. In the hairpin, the agent starts decelerating ~ 100 meters before the start of the curve.

Overall speed The top of Figure 8 shows a speed comparison between our approach, the fastest human lap, and the built-in AI in setting A. Our approach learned a control policy that closely matches and sometimes even improves on the




Fig. 7. Out-in-out driving behavior and early curve anticipation learned by our approach for setting A (top) and setting C (bottom). The brake signal is represented by the white bar marked in green. The segments are shown in the supplementary video from 0:24 to 0:30 and 2:50 to 3:04.

speed as well as the path of the fastest human reference lap. Even though path and speed of our approach and the fastest human are similar for setting A, our approach undercuts the fastest human lap time by 0.15 seconds. While this might not seem like a significant difference, this margin is similar to those between top contenders in real and simulated racing championships.

the human expert drives a wider trajectory on the incoming straight, our approach compensates on the outgoing straight, leading to a similar curve exit speed and curve completion time. The right side of Figure 9 shows the speed difference between our approach and the fastest human. Our approach achieves similar or higher speeds over most of the track, often driving tighter and faster curves than the human expert, which leads to an overall faster time.

In summary, these results indicate that our approach learns to drive autonomously at high-speed in different race settings, using different cars and driving on different tracks, including a challenging hairpin curve. In all settings our approach achieved lap times faster than those of all human reference drivers. Moreover, we find that our approach learns to drive trajectories that are qualitatively similar to those chosen by the best human players while maintaining slightly higher average speeds in curves by successfully executing later brake points.

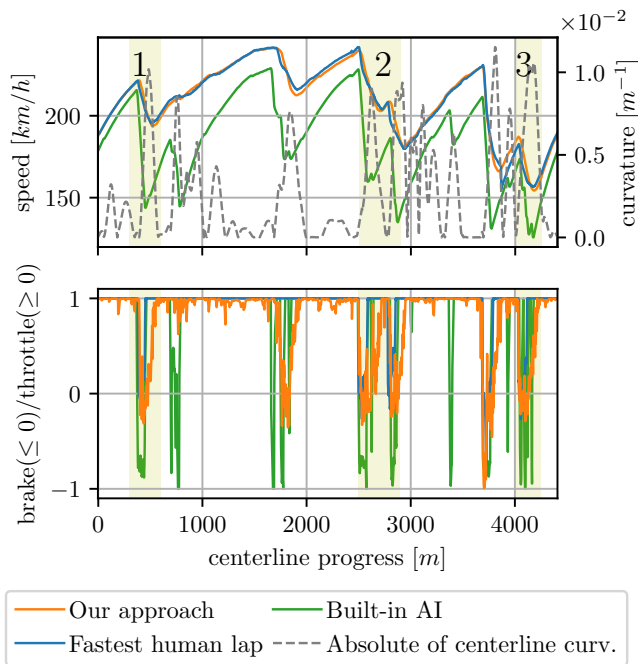


Fig. 8. **Top:** Comparison of speed between the fastest human, our approach, and the built-in GTS AI in setting A. **Bottom:** Comparison of combined throttle/brake signal.

The left side of Figure 9 shows the path and speed driven by our approach and the fastest human in the hairpin curve. While

C. Robustness

We analyze the robustness of our trained agent by modifying the test environment.

1) Transfer to new car: The agent still follows an out-in-out trajectory and is able to finish the track. In straight segments as well as slight curves, the agent is able to adjust to the change in dynamics and execute a feasible trajectory. In steep curves however, the agent is sometimes overshooting the track. This is most likely caused by the training car being easier to control than the test car, due to higher tire-surface friction caused by both the weight and the racing tires of the training car. For steeper curves, the agent is not able to extrapolate its behavior to the previously unseen dynamics of a car with less tire-surface friction.

2) Transfer to new track: The agent is able to drive wall-contact free in straight segments as well as slight curves. It is however not able to extrapolate its behavior to some of the

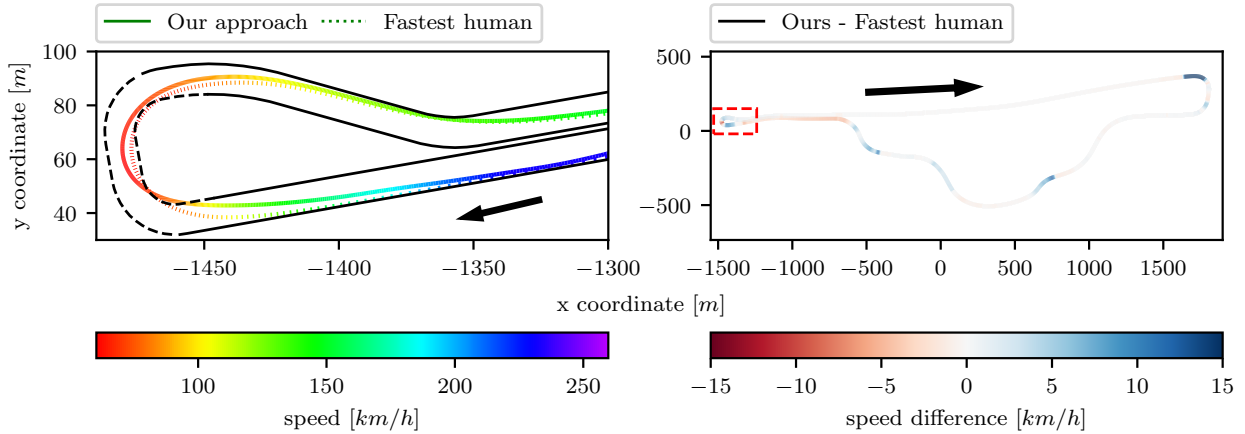


Fig. 9. **Left:** the trajectory driven by our approach and the fastest human trajectory in the hairpin segment in setting C. The solid track border is represented by walls on the track, while the dashed track border is simply drawn onto the track. To not get off the track, drivers are forced to brake early. **Right:** the speed difference between our approach and the fastest human over the whole track. Positive values indicate our approach driving faster in a segment.

unseen curve shapes. Since the agent gets stuck in the hairpin curve, it is not able to finish the track.

3) Increased tire friction: The agent is able to adjust to the change in friction by making corrections when deviating from the expected trajectory. In 2 curves, the agent over-estimates the outward force, leading to brief contact with the inner edge of the curve. This leads to a loss of 0.1 seconds on the baseline.

Lowered tire friction: The agent is able to make corrections to its trajectory to adjust to the change in dynamics and drive out-in-out paths for all but the 3 steepest curves. In the steepest curves, the agent takes a too aggressive path, based on the higher friction experienced in training, leading to contact with the outer edge at the end of the curve.

The two experiments indicate that our agent is able to make corrections to its path to a certain degree. However, for some track segments the corrections are not strong enough to handle the change in dynamics.

4) Noise in Observations: The top of Figure 10 shows the lap times in setting A for different degrees of noise added to the agent’s observation compared to the noise-free baseline performance. With up to 2% noise the agent drives trajectories close to our noise-free baseline, still placing in front of the fastest human. More noise in the agent’s observation leads to jitter in it’s action signal, resulting in a speed loss. With up to 9% noise our agent is still able to drive on the track without any contact with the edge of the track.

5) Delayed Inference: The bottom of Figure 10 shows the lap times in setting A when adding delays of various degrees to the agent’s inference compared to the non-delayed baseline performance. For delays up to 20 ms the agent performs close to the baseline, in some cases even slightly faster, outperforming the fastest human reference lap in each case. With a delay of 50 ms the agent loses 0.3 seconds on the baseline due to delay related deviations from the optimal line, but still achieves a lap time among the top 10 human drivers. With a delay of 100 ms our agent brakes too late before the steepest curve of the track, leading to wall contact and a loss of 1.4 seconds compared to our baseline. With a delay of 150 ms our agent is no longer able to drive a competitive

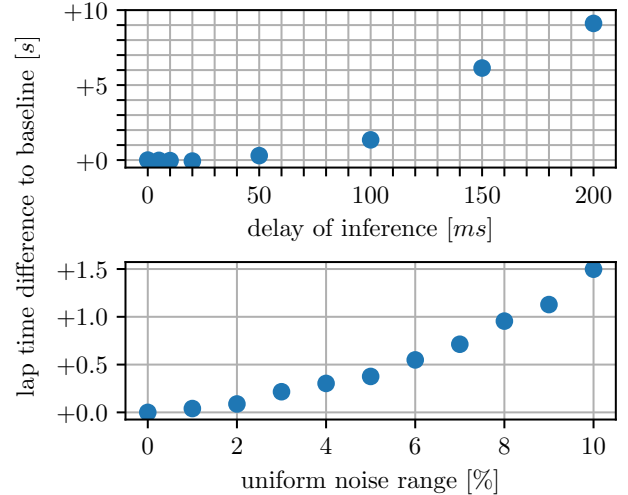


Fig. 10. **Top:** Loss in lap time when adding noise to the agent’s observation compared to undelayed baseline agent in setting A. **Bottom:** Loss in lap time when artificially delaying the agent’s inference.

racing line, often colliding with the edge of the track and losing over 6 seconds on our baseline. It’s possible that repeating the training process with the delay active is already sufficient to mitigate the decrease in performance. In addition, a predictive model such as the one proposed in [28] could be applied to counter the effect of the delay by estimating the real state based on the delayed state. In our comparison the human is disadvantaged by a relatively slow reaction time of approximately 0.2 seconds, compared to our agent, which is able to react in approximately 4 milliseconds, including both inference and communication with the PlayStation. We would argue, however, that this disadvantage is part of the human system and not part of the challenge of autonomous racing.

VI. CONCLUSION

In this paper, we have presented the first autonomous racing policy that achieves super-human performance in time trial settings in the racing simulator *Gran Turismo Sport*. The

benefits of our approach are that it does not rely on human intervention, human expert data, or explicit path planning, and it leads to trajectories that are qualitatively similar to those chosen by the best human players, while outperforming the best known human lap times in all three of our reference settings, including two different cars on two different tracks. This super-human performance has been achieved using limited computation power during both evaluation and training. Generating one control command with our policy network takes around 0.35 ms on a standard desktop PC CPU, and training, using 4 *PlayStation 4* game consoles and a single desktop PC, took less than 73 hours to achieve super-human performance.

Limitations of the present work include a) the restriction to single-player time trial races without other cars on the same track, and b) the constrained applicability of the learned control policies to a single track / car combination. We intend to address a) in future work by extending the observation space to allow the perception of other cars and modify the reward to disincentivize unfair driving behavior. To extend our approach to more track / car combinations we propose using more data-efficient RL algorithms, such as meta-RL, where the agent can adapt to new situations with only a small amount of new samples.

ACKNOWLEDGMENT

We are very grateful to Polyphony Digital Inc. for enabling this research. Furthermore, we would like to thank Kenta Kawamoto and Takuma Seno from Sony R&D Center Tokyo for their kind help and many fruitful discussions.

REFERENCES

- [1] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 1433–1440.
- [2] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1: 43 scale rc cars," *Optimal Control Applications and Methods*, vol. 36, no. 5, pp. 628–647, 2015.
- [3] E. Frazzoli, M. A. Dahleh, and E. Feron, "Real-time motion planning for agile autonomous vehicles," *Journal of guidance, control, and dynamics*, vol. 25, no. 1, pp. 116–129, 2002.
- [4] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic mpc for model-based reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, May 2017, pp. 1714–1721.
- [5] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [6] N. A. Spielberg, M. Brown, N. R. Kapania, J. C. Kegelman, and J. C. Gerdes, "Neural network vehicle models for high-performance automated driving," *Science Robotics*, vol. 4, no. 28, 2019.
- [7] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.
- [8] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, 2020.
- [9] J. Hwangbo, I. Sa, R. Siegwart, and M. Hutter, "Control of a quadrotor with reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2096–2103, 2017.
- [10] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in *Conference on Robot Learning*. PMLR, 2020, pp. 1101–1112.
- [11] T. Novi, A. Liniger, R. Capitani, and C. Annicchiarico, "Real-time control for at-limit handling driving on a predefined path," *Vehicle system dynamics*, 2019.
- [12] U. Rosolia and F. Borrelli, "Learning how to autonomously race a car: a predictive control approach," *IEEE Transactions on Control Systems Technology*, 2019.
- [13] J. Kabzan, M. d. I. Valls, V. Reijgwart, H. F. C. Hendriks, C. Ehmke, M. Prajapat, A. Bühler, N. Gosala, M. Gupta, R. Sivanesan *et al.*, "Amz driverless: The full autonomous racing system," *arXiv preprint arXiv:1905.05150*, 2019.
- [14] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, "Robust constrained learning-based nmpc enabling reliable mobile robot path tracking," *The International Journal of Robotics Research*, vol. 35, no. 13, pp. 1547–1563, 2016.
- [15] R. Verschueren, S. De Bruyne, M. Zanon, J. V. Frasch, and M. Diehl, "Towards time-optimal race car driving using nonlinear mpc in real-time," in *53rd IEEE conference on decision and control*. IEEE, 2014, pp. 2505–2510.
- [16] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, 2019.
- [17] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," in *Advances in neural information processing systems*, 1989, pp. 305–313.
- [18] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [19] Y. Pan, C.-A. Cheng, K. Saigol, K. Lee, X. Yan, E. Theodorou, and B. Boots, "Agile autonomous driving using end-to-end deep imitation learning," in *Robotics: science and systems*, 2018.
- [20] M. Jaritz, R. de Charette, M. Toromanoff, E. Perot, and F. Nashashibi, "End-to-end race driving with deep reinforcement learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 2070–2075.
- [21] M. Riedmiller, M. Montemerlo, and H. Dahlkamp, "Learning to drive a real car in 20 minutes," in *2007 Frontiers in the Convergence of Bioscience and Information Technologies*, Oct 2007, pp. 645–650.
- [22] A. Kendall, J. Hawke, D. Janz, P. Mazur, D. Reda, J. Allen, V. Lam, A. Bewley, and A. Shah, "Learning to drive in a day," in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 8248–8254.
- [23] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu, "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, 2020.
- [24] P. Cai, X. Mei, L. Tai, Y. Sun, and M. Liu, "High-speed autonomous drifting with deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1247–1254, 2020.
- [25] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International Conference on Machine Learning*, 2018, pp. 1856–1865.
- [26] G. Kahn, A. Villafior, V. Pong, P. Abbeel, and S. Levine, "Uncertainty-aware reinforcement learning for collision avoidance," *arXiv preprint arXiv:1702.01182*, 2017.
- [27] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev *et al.*, "Grandmaster level in starcraft ii using multi-agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
- [28] V. Firoiu, T. Ju, and J. Tenenbaum, "At human speed: Deep reinforcement learning with action delay," *arXiv preprint arXiv:1810.07286*, 2018.