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

**Florian Fuchs**[1,2], **Yunlong Song**[1], **Elia Kaufmann**[1], **Davide Scaramuzza**[1], **Peter Dürr**[2]

[1]Robotics and Perception Group
Depts. Informatics and Neuroinformatics, University of Zurich and ETH
[2]Sony Europe B.V.,
Schlieren/Switzerland Branch

**Abstract:** Autonomous car racing raises fundamental robotics challenges such as planning minimum-time trajectories under uncertain dynamics and controlling the car at its friction limits. In this project, we consider the task of autonomous car racing in the top-selling car racing game Gran Turismo Sport. Gran Turismo Sport is known for its detailed physics simulation of various cars and tracks. Our approach makes use of maximum-entropy deep reinforcement learning and a new reward design to train a sensorimotor policy to complete a given race track as fast as possible. We evaluate our approach in three different time trial settings with different cars and tracks. Our results show that the obtained controllers not only beat the built-in non-player character of Gran Turismo Sport, but also outperform the fastest known times in a dataset of personal best lap times of over 50,000 human drivers.

Video of the race performance: https://youtu.be/Zeyv1bN9v4A

## 1 Introduction

Autonomous driving at high speed is a challenging task that requires to generate fast and precise actions even when the vehicle is approaching its physical limits. Autonomous car racing, where the goal is to complete a given course in minimal time, features some of these difficulties of controlling a car close to its physical limitations. In recent years, competitions for real [1, 2] as well as simulated [3] autonomous cars showed promising results, but the demonstrated performances still lag behind what human experts achieve.

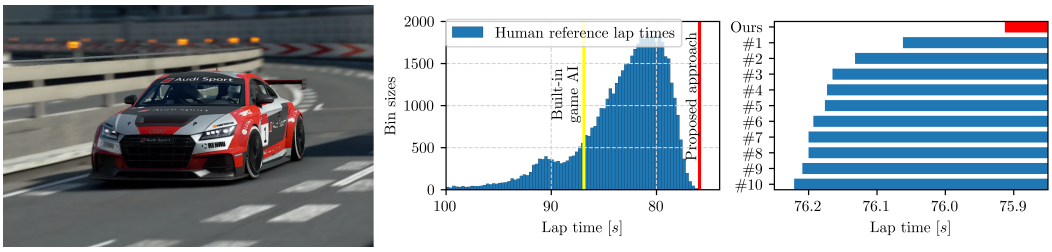


Figure 1: **Left:** Our approach controlling the "Audi TT Cup" in the *Gran Turismo Sport* simulation. **Center:** Time trial lap time comparison between our approach (red line), the personal best lap times of over 50,000 human drivers from 70 countries (dark blue histogram), as well as the built-in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

To our knowledge, the built-in non-player characters (NPC) included in modern car racing games are unable to compete with human expert players in fair comparisons. For example, the currently built-in NPC in *Gran Turismo Sport* (GTS) loses a total of 11 seconds compared with the fastest

human driver and is slower than 83% of all humans in one of our reference settings (see Figure 1). Other racing games apparently close the gap to human experts by granting an unfair advantage to the NPC, for example by increasing the engine power of the NPC's car; this, however, leads to frustration among human players who feel cheated [4]. A truly competitive autonomous control policy might thus increase player satisfaction and potentially even inspire better human performance.

We tackle the challenge of creating an autonomous agent for the *Gran Turismo Sport* (GTS) car racing simulator that is able of matching, or outperforming human experts in time trials. Gran Turismo Sport is the newest member of the Gran Turismo series, itself one of the highest-selling *PlayStation* series, with more than 80 million units sold [5]. Among racing games, GTS is known as a highly realistic driving simulation [6], modelling phenomena, such as the influence of tires' temperature and a car's current fuel level on traction. Therefore, similarly to real-world racing, the optimal trajectory (i.e., the trajectory leading to the fastest lap time) for a car in GTS depends not only on the geometry and properties of the track, but also on various (a priori unknown) physical characteristics and states of the car. Due to its similarity to real driving, and the relatively low price of training in GTS compared to training with actual race cars, GTS is also used to cast drivers for racing teams [7].

In contrast to previous approaches that relied on classical trajectory planning and control, our approach leverages reinforcement learning to train a deep sensorimotor policy that directly maps from observations to control commands. By maximizing a progress-based reward signal, our agents manage to drive autonomously on different tracks with different cars at high speed. In all tested settings, the agents achieve super-human performance by outperforming the fastest known times in a dataset of personal best lap times of over 50,000 human drivers.

## 2 Related Work

Prior work in the domain of autonomous racing can be grouped into three groups: (i) classical approaches relying on trajectory planning and control, (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 [8, 9, 10]. Considering the task of time-optimal trajectory planning, Timings and Cole [11] maximizes the distance traveled along a race track in a given fixed time, Velenis and Tsiotras [12] maximizes the vehicle speed at the final point of the maneuver, and Rucco et al. [13] directly minimizes the lap time. Model predictive control (MPC) [8] is a promising approach when considering solely the problem of controlling the vehicle at high speed given perfect state estimation and a sufficiently accurate dynamics model. Similarly, model predictive path integral control (MPPI) [10, 9] is a more flexible approach that can be coupled with complex cost criteria and deep neural networks to learn a dynamics model, however, it needs intensive parallel computation to solve the sampling-based online optimization. By separating the task of autonomous racing into submodules, above approaches are susceptible to failures of each submodule, rendering the entire system brittle.

**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) [14] 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 Bojarski et al. [15] for lane and road following. 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 an accurate dynamics model of the vehicle and its environment or the dependence on labelled training data. For example, a number of studies [16, 17, 18] have demonstrated the success of using model-free deep RL for end-to-end driving, where the task is to map images directly to control commands in order to drive a car safely.

Here, off-policy training plays an important role in both [17] and [18], 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.

## 3 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 [19] algorithm. An overview of our system is shown in Figure 2.

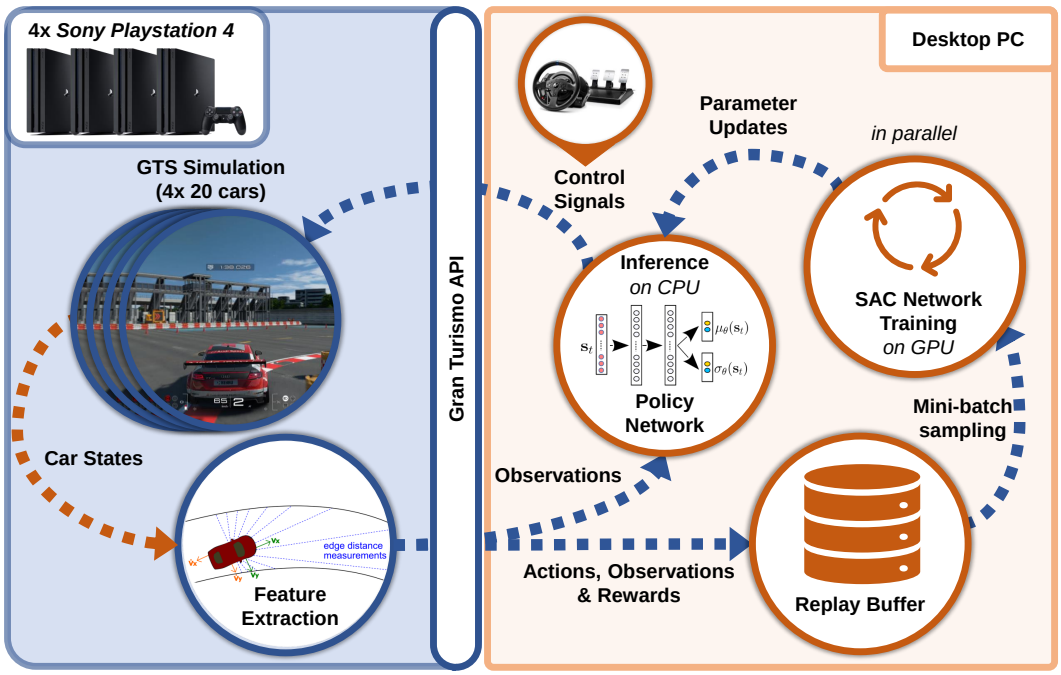


Figure 2: System overview: we train a policy network to directly map observations, including a set of range finder measurements and the car's velocity and acceleration, to control commands consisting of the car's steering angle as well as the level of throttle and brake. We use a distributed sampling 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.

### 3.1 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. Maximizing the course progress closely approximates minimizing the lap time when taking into account a long enough time horizon. 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 shortening or extending the agent's horizon respectively. We use an exponential discount for future rewards, with discount factor $\gamma$ controlling aforementioned trade-off. The construction of the progress reward can be seen in the left of Figure 3. The use of an exponentially discounted future

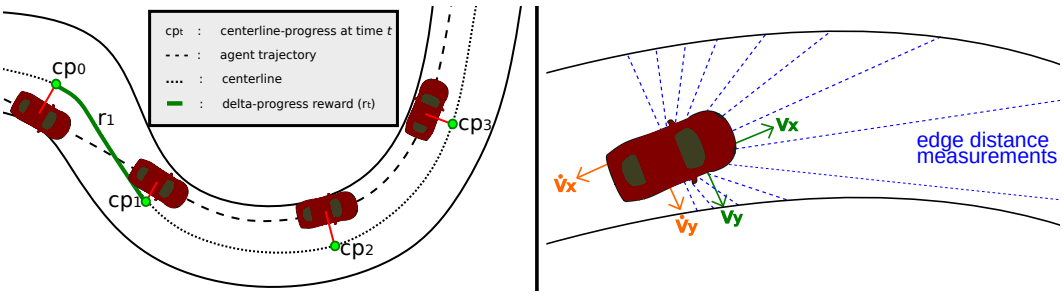


Figure 3: **Left:** The course progress $cp_t$ at time $t$ is constructed by projecting the car's position on the center line of the track. We then define the progress-based reward $r^{\text{prog}}$ as the increment in centerline-progress $r_t^{\text{prog}} = cp_t - cp_{t-1}$. **Right:** A subset of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

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

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

where $c_w \geq 0$ is a hyperparameter controlling the trade off between the wall contact penalty and the progress-based reward. Here $\mathbf{v}_t = [v_x, v_y, v_z]$ is a velocity vector that represents the current linear velocity of the vehicle. A similar approach to incentivize collision avoidance was introduced in [20]. 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.

### 3.2 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 [19]. 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** The neural network is fed a combination of input features consisting of the car's current linear velocity and acceleration, its orientation w.r.t. the track, a set of range finder measurements equally distributed towards the front of the car, the current steering angle, a binary flag indicating wall contact and the curvature of a receding horizon of the track, sampled at a set of equally-spaced points. A subset of the features is illustrated in Figure 3. The input features are concatenated to a single vector and fed to the network.

**Network prediction** The output of the policy network $\boldsymbol{a}_t = [\delta_t, \omega_t]$ directly encodes the steering angle $\delta_t \in [-\pi/6, \pi/6]\text{rad}$ 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.

## 4 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 NPC 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


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

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. 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 $0.03\,\mathrm{rad}$ for the steering angle and 80% for each the throttle and brake range[1]. We however found this restriction to not significantly influence the resulting lap times.

## 5 Results

In this section we evaluate our approach on the three reference settings introduced in Section 4. We compare the lap times obtained by our approach with the ones achieved by the built-in NPC 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 video.

### 5.1 Lap Time Comparison

Our approach outperforms the best human lap time in all three reference settings, overcoming the limitations of the currently built-in NPC, which is outperformed by a majority of players. Table 1 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 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*. 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

[1] A similar fairness restriction was put on the agent proposed in [21], by limiting the number of actions an agent could take per minute.

| Driver | Metric | Setting A | Setting B | Setting C |
|---|---|---|---|---|
| **Our approach** | 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 NPC | Lap time [min] | 01:26.899 | 01:52.075 | 02:14.252 |
| | Slower than x% of human players | 82.6% | 80.4% | 53.9% |

Table 1: Time trial comparisons between our approach, human online competitors, and the built-in GTS non-player character (NPC) for the 3 race settings.

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

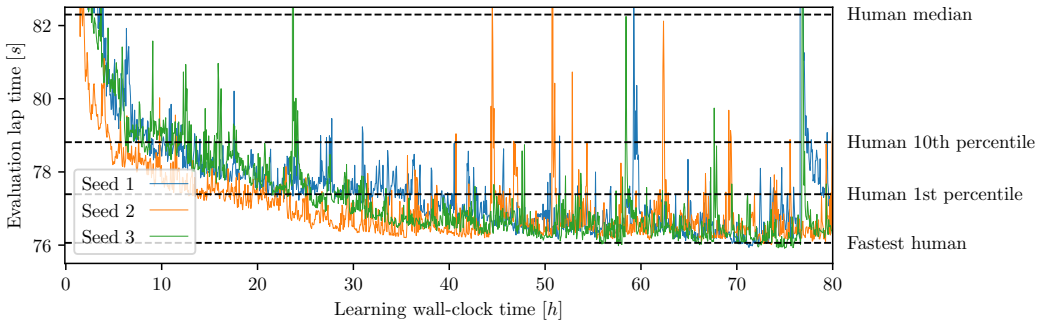


Figure 5: Training progress for our approach learning on setting *A*. We evaluate the policy every 2 epochs of learning by letting the agent drive 2 consecutive laps. We take the lap time at the second lap as performance metric such that the fixed starting speed applied in the first lap does not distort the result. During evaluation we remove the exploration noise by taking the mean of the policy distribution as action, which is an approximation for the maximum a posteriori action [19].

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.

### 5.2 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. The top of Figure 6 shows that behavior for the steepest curve of setting *A*. The right of Figure 7 shows a curve driving comparison between our approach, the fastest human, and the built-in NPC for 3 more 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.

**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 6 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 $\sim$100 meters before the start of the curve.

**Overall speed** The top left of Figure 7 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

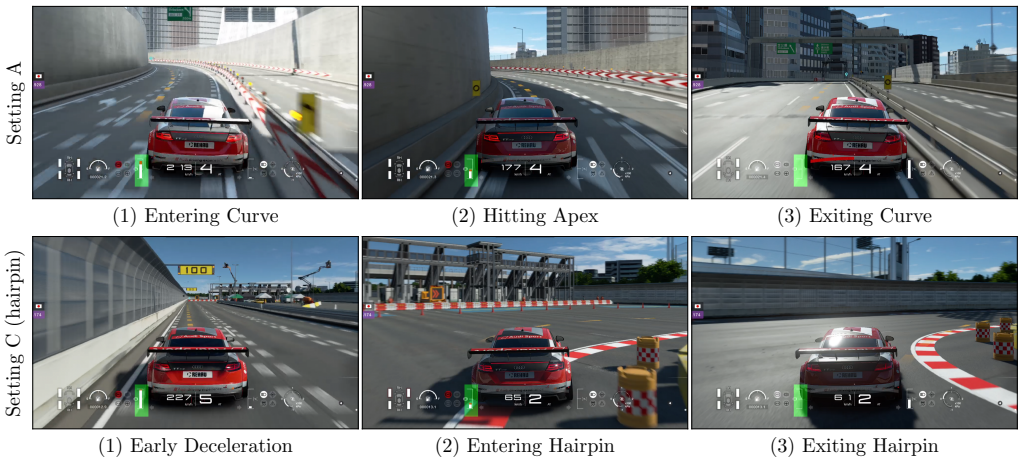


Figure 6: Out-in-out driving behavior and early curve anticipation learned by our approach for setting $A$ (top) and setting $C$ (bottom). The brake signal is represented by the white bar marked in green. The segments are shown in the supplementary video from 1:48 to 1:53 and 2:50 to 3:04.

matches and sometimes even improves on the 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.

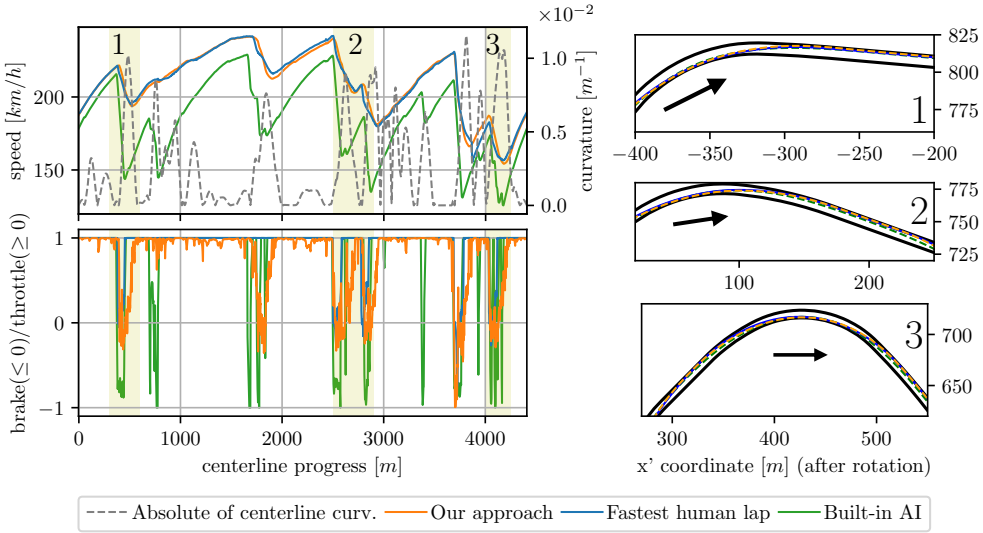


Figure 7: **Left:** Comparison of speed and the combined throttle/brake signal between the fastest human, our approach, and the built-in GTS NPC ("Built-in AI") in setting $A$. Note that the braking signal of our approach is stronger than that of the human reference lap, because the human uses a setting that passively brakes and this passive braking signal is not reflected in this plot. **Right:** Driven paths for setting $A$. While the paths of our approach and the human driver are similar, the built-in GTS NPC drives with a safety margin to the track's walls, leading to tighter radii and a reduced speed.

The left side of Figure 8 shows the path and speed driven by our approach and the fastest human in the hairpin curve. While 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 8 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.

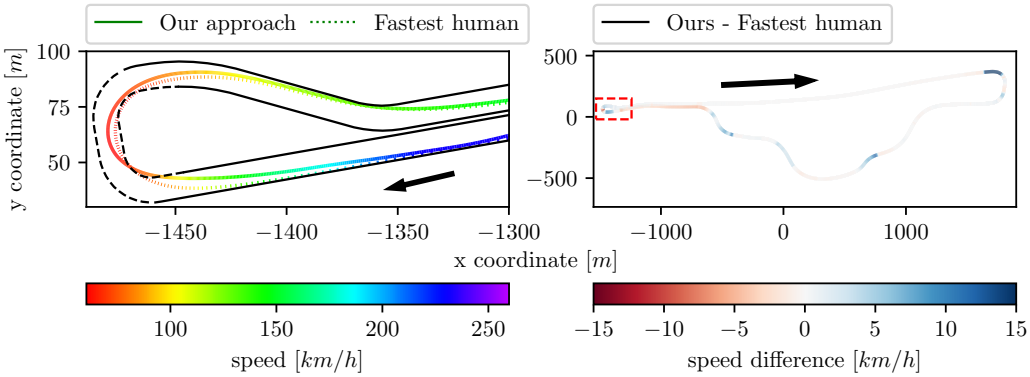


Figure 8: **Left:** the trajectory driven by our approach and the fastest human trajectory in the hairpin segment in setting *C*. The solid track border is represented by walls on the track, while the dashed track border is simply drawn onto the track. Cars are therefore not caught by a wall when overshooting the hairpin curve. 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. There is one outlier segment with values around +100 at the top right of the track. For ease of visualization the color scheme is however clipped at +15.

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 speed throughout curves.

## 6 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.

**Acknowledgments**

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. We are also very grateful to Polyphony Digital Inc. for enabling this research, and last but not least, 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] J. Kabzan, M. d. l. I. Valls, V. Reijgwart, H. F. C. Hendrikx, 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.

[2] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, et al. Stanley: The robot that won the DARPA Grand Challenge. *Journal of field Robotics*, 23(9):661–692, 2006.

[3] Amazon Web Services. AWS DeepRacer League. `https://aws.amazon.com/deepracer/league/`, 2019. [Online; accessed 01-June-2020].

[4] F. Forums. Unbeatable AI a bit cheaty pic in post. `https://forums.forzamotorsport.net/turn10_poststt118652_Unbeatable-Ai-a-bit-cheaty-pic-in-post.aspx`, 2018. [Online; accessed 01-June-2020].

[5] Sony Interactive Entertainment. Gran Turismo Franchise Sales Surpasses 80.4 Million. `sie.com/en/corporate/release/2018/180514.html`, 2018. [Online; accessed 01-June-2020].

[6] A. Ingram. Gran Turismo Sport - exploring its impact on real-world racing with Kazunori Yamauchi. `evo.co.uk`, 2019. [Online; accessed 01-June-2020].

[7] Sony Interactive Entertainment Inc. Gt academy. `gran-turismo.com/us/academy/`, 2017. [Online; accessed 01-June-2020].

[8] A. Liniger, A. Domahidi, and M. Morari. Optimization-based autonomous racing of 1: 43 scale rc cars. *Optimal Control Applications and Methods*, 36(5):628–647, 2015.

[9] 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)*, pages 1714–1721, May 2017. doi:10.1109/ICRA.2017.7989202.

[10] 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)*, pages 1433–1440, May 2016. doi:10.1109/ICRA.2016.7487277.

[11] J. P. Timings and D. J. Cole. Minimum maneuver time calculation using convex optimization. *Journal of Dynamic Systems, Measurement, and Control*, 135(3), 2013.

[12] E. Velenis and P. Tsiotras. Minimum Time vs Maximum Exit Velocity Path Optimization During Cornering. In *Proceedings of the IEEE International Symposium on Industrial Electronics, 2005. ISIE 2005.*, volume 1, pages 355–360, June 2005. doi:10.1109/ISIE.2005.1528936.

[13] A. Rucco, G. Notarstefano, and J. Hauser. An efficient minimum-time trajectory generation strategy for two-track car vehicles. *IEEE Transactions on Control Systems Technology*, 23(4): 1505–1519, 2015.

[14] D. A. Pomerleau. Alvinn: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pages 305–313, 1989.

[15] 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.

[16] 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)*, pages 2070–2075, May 2018. doi:10.1109/ICRA.2018.8460934.

[17] 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*, pages 645–650, Oct 2007. doi:10.1109/FBIT.2007.37.

[18] 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)*, pages 8248–8254, May 2019. doi:10.1109/ICRA.2019.8793742.

[19] 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*, pages 1856–1865, 2018.

[20] G. Kahn, A. Villaflor, V. Pong, P. Abbeel, and S. Levine. Uncertainty-aware reinforcement learning for collision avoidance. *arXiv preprint arXiv:1702.01182*, 2017.

[21] 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*, 575(7782):350–354, 2019.

[22] C. Berner, G. Brockman, B. Chan, V. Cheung, P. Debiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

[23] J. Hwangbo, I. Sa, R. Siegwart, and M. Hutter. Control of a quadrotor with reinforcement learning. *IEEE Robotics and Automation Letters*, 2(4):2096–2103, 2017.

[24] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529: 484–503, 2016. URL `http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html`.

[25] G. Barth-Maron, M. W. Hoffman, D. Budden, W. Dabney, D. Horgan, D. TB, A. Muldal, N. Heess, and T. Lillicrap. Distributional policy gradients. In *International Conference on Learning Representations*, 2018. URL `https://openreview.net/forum?id=SyZipzbCb`.

# Supplementary Material

## GTS Access

The GTS simulator runs on a *PlayStation 4*, while our agent runs on a separate desktop computer. We do not have direct access to the GTS simulator, neither do we have insights into the car dynamics modeled by GTS. Instead, we interact with GTS over a dedicated API via ethernet connection. The API provides the current state of up to 20 simulated cars and accepts car control commands, which are active until the next command is received. Whereas previous work in the area of RL often ran thousands of simulations in parallel to collect training data [22, 23, 24], we can only run one simulation per *PlayStation*, and for this work we make use of only 4 *PlayStations* during training. Moreover, the simulation runs in real-time and cannot be sped up for data collection during training nor can it be paused to create additional time for decision making. The simulator's state is updated at a frequency of 60 Hz, but, to reduce the load on the *PlayStations* when controlling 20 cars, we limit the command frequency to 10 Hz during training. During evaluation, we switch to one car, increasing the agent's action frequency to 60 Hz. We use a standard desktop computer with a *i7-8700* processor with a clock speed of 3.20GHz for inference and a *GeForce GTX 1080 Ti* graphics card for backpropagating through the models.

## Network Training

To train the networks we roll out 4x20 trajectories per epoch in parallel, with a length of 100 seconds each, by using all 20 cars available on each of the 4 *Playstations*. To minimize disturbance between cars during training, we initialize the position of agents equally distributed over the racing track with an initial speed of 100 km/h, which we found can accelerate training, since it allows the agents to faster approach the maximal feasible segment speeds, which are of interest for finding the fastest possible trajectory.

For the training process we make use of the TensorFlow-based SAC implementation by OpenAI[2]. We modified the code base to allow to asynchronously learn during roll outs and changed the default 1-step TD error to a 5-step TD error to stabilize training as also used in [25]. The combined training of all 4 networks takes approximately 35 seconds per epoch, which leaves the total epoch time at 100 seconds. Because of the real-time character of the GTS environment, an extensive hyperparameter search is not feasible. We therefore adapt most default hyperparameters from the OpenAI implementation, except for those listed in table 2.

| Hyperparameter | Value |
|---|---|
| Mini-batch size | 4,096 |
| Replay buffer size | $4 \times 10^6$ |
| Learning rate | $3 \times 10^{-4}$ |
| Update steps per epoch | 5,120 |
| Reward scale ($1/\alpha$) | 100 |
| Exponential discount ($\gamma$) for the "Mazda Demio" | 0.98 |
| Exponential discount ($\gamma$) for the "Audi TT Cup" | 0.982 |
| Wall contact penalty scale ($c_w$) | $5 \times 10^{-4}$ |
| Number of range finders ($M$) | 13 (every 15°) |
| Number of curvature mesurements ($N$) | 10 (every 0.2s) |

Table 2: Hyperparameters used for the SAC algorithm and the GTS simulation. Denser range finders lead to slower convergence but showed no additional improvements in the learned policy.

## Observation Space

The GTS simulator provides a rich amount of information about the agent's state and the racing environment, via a dedicated API. This includes the ground truth position, the linear velocity and acceleration, and others. 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

[2] github.com/openai/spinningup

by GTS as well as additionally constructed features, which lead to the following feature selection: 1) The linear velocity $\boldsymbol{v}_t \in \mathbb{R}^3$ and the linear acceleration $\dot{\boldsymbol{v}}_t \in \mathbb{R}^3$. 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 $\boldsymbol{d}_t \in \mathbb{R}^M$ of $M$ rangefinders 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 $\delta_{t-1}$. 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 $\boldsymbol{c}_t \in \mathbb{R}^N$. The curvature is represented through an interpolation of the inverse radii of circles fitted through centerline points provided by GTS, as can be seen in Figure 9. They are equally distributed from 1.0 to 2.8 seconds into the future from the car's current position, estimated with the car's current speed. We therefore represent the observation at a given time step $t$ as a vector denoted as $\boldsymbol{s}_t = [\boldsymbol{v}_t, \dot{\boldsymbol{v}}_t, \theta_t, \boldsymbol{d}_t, \delta_{t-1}, w_t, \boldsymbol{c}_t]$. 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. We apply z-score normalization to all features individually before feeding them to the policy and value networks, to account for large differences in feature scales. For example, the centerline difference angle $\theta_t$ lies within the range of $(-\pi, +\pi]$ $(rad)$, while the forward velocity $v_x$ can reach values up to 270 $(km/h)$. For feature x with mean $\bar{x}$ and standard deviation $\sigma$, we replace each feature instance with the normalized value $z = (x - \bar{x})/\sigma$.

### Action Space

We define the action vector $\boldsymbol{a}_t = [\delta_t, \omega_t]$ by two control signals: a steering angle $\delta_t \in [-\pi/6, \pi/6](rad)$ corresponding to the angle of the car's front wheels and a combined throttle-brake signal $\omega_t \in [-1, 1]$, where $\omega > 0$ denotes throttle and $\omega \leq 0$ represents braking. The magnitude of the throttle or the break is proportional to the absolute value $\|\omega\|$. For example, $\omega = 1$ is 100% throttle, $\omega = -1$ is 100% brake, and $\omega = 0$ is no throttle and no break. Combining the two signals reduces the complexity of the task. Since human expert recordings show that the best human strategies do not involve the simultaneous use of throttle and brake, we expect not to significantly restrict the agent by combining the two signals. To limit the outputs of the policy network to the valid range, we apply a *tanh* activation to the output layer as proposed in [19].

We use the automatic gearshift provided by GTS, since the used API does currently not allow to manually change gears. This option is also available to human players, but experienced players mostly use manual gearshift to have more control of when to shift.

### Analysis of the Results by a Domain Expert

To improve our understanding of the achieved results, we invited *Gran Turismo* domain expert TG (name omitted for reasons of anonymity), who has achieved top performance in several national and international competitions, to race in our reference settings and compare his performance against our approach. TG competed in two of our three reference settings and achieved lap times in the top 0.36 and 0.23 percentile of our human reference data.

When asked for his opinion on the policies' driving style, TG stated:

> *"The policy drives very aggressively, but I think this is only possible through its precise actions. I could technically also drive the same trajectory, but in 999 out of a 1000 cases, me trying that trajectory results in wall contact which destroys my whole lap time and I would have to start a new lap from scratch."*

### Improvements over the built-in NPC

In the 3 introduced reference settings the currently built-in NPC is outperformed by a majority of human players as can be seen in Table 1. The built-in NPC consists of a path following policy and a predefined reference trajectory. Due to the non-linearity of the GTS dynamics, a slight deviation from such a trajectory strongly changes the new optimal trajectory, which makes it practically impossible to pre-compute and then track an optimal trajectory. The built-in NPC solves this problem by using a cautious reference trajectory which includes a margin to the track's borders and allows

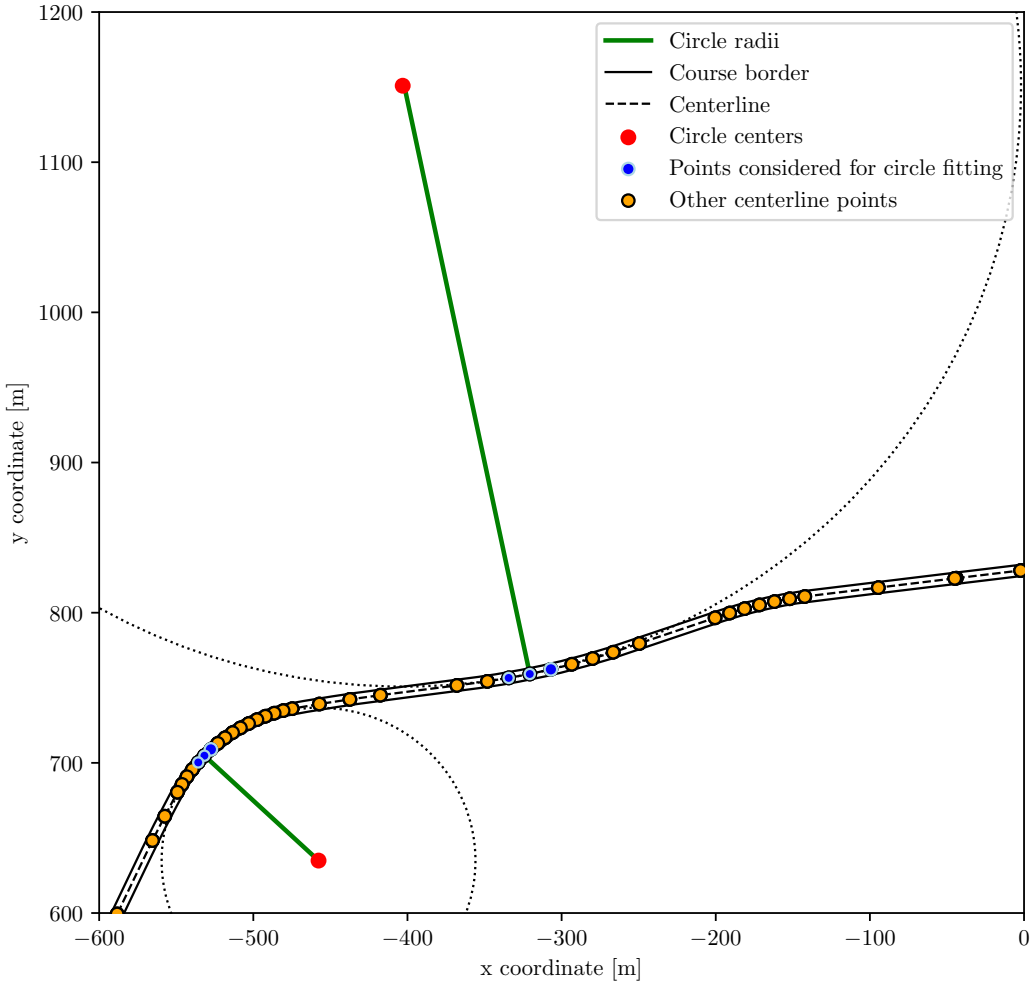


Figure 9: **Construction of the curvature measurement.** To represent the curvature of the centerline, we make use of centerline points provided by the GTS simulation. The points are distributed such that regions with higher curvature are represented by more points. This allows getting a precise curvature measurement from looking at only three neighboring points. We make use of that property by defining the curvature for each centerline point by the inverse radius of the circle given by that point and its two neighbors. Right curvature is represented by negative inverted radii, left curvature by positive inverted radii. We then interpolate between centerline points to end up with a continous representation of the curvature over the course progress.

recovery when deviating from the trajectory, as can be seen in Figure 7. This helps reducing the risk of contacting the track's side walls when deviating from the trajectory. However, it also leads to the trajectory's curves having smaller radii, forcing the NPC to decelerate to not lose traction, resulting in curve exit speeds up to 50 km/h slower than those of our approach and the fastest human. The bottom of Figure 7 shows how the built-in NPC brakes in segments where neither our approach nor the fastest human brake, to be able to follow its reference trajectory. By learning a flexible driving policy without any explicit restrictions, our approach was able to overcome the shortcomings of the previous built-in NPC and learn a driving policy more similar to that of a human expert driver regarding speed and path.