An Adaptive Human Driver Model for Realistic Race Car Simulations

Stefan Löckel, Siwei Ju, Maximilian Schaller, Peter van Vliet, and Jan Peters

***Abstract*—Engineering a high-performance race car requires a direct consideration of the human driver using real-world tests or human-driver-in-the-loop simulations. Alternatively, offline simulations with human-like race driver models could make this vehicle development process more effective and efficient but are hard to obtain due to various challenges. With this work, we intend to provide a better understanding of race driver behavior from expert knowledge and introduce an adaptive human race driver model based on imitation learning. Using existing findings in the literature, complemented with an interview with a race engineer, we identify fundamental adaptation mechanisms and how drivers learn to optimize lap time on a new track. Subsequently, we select the most distinct adaptation mechanisms via a survey with 12 additional experts, to develop generalization and adaptation techniques for a recently presented probabilistic driver modeling approach and evaluate it using data from professional race drivers and a state-of-the-art race car simulator. We show that our framework can create realistic driving line distributions on unseen race tracks with almost human-like performance. Moreover, our driver model optimizes its driving lap by lap, correcting driving errors from previous laps while achieving faster lap times. This work contributes to a better understanding and modeling of the human driver, aiming to expedite simulation methods in the modern vehicle development process and potentially supporting automated driving and racing technologies.**

***Index Terms*—**

I. Introduction

THROUGHOUT more than 125 years of motorsports history, the fundamental goal of all participants did not change: reaching the best racing performance among competitors, which ultimately requires engineering a race car that fits its driver well. In fact, Milliken and Milliken already stated in 1995 that "it is the dynamic behavior of the combination of high-tech machines and infinitely complex human beings that makes the sport so intriguing for participants and spectators alike" [1]. Hence, for modern vehicle development in professional motorsports, a good understanding and modeling of the human (not necessarily lap time-optimal) driver are crucial to further improve the performance of the human-driver-vehicle-system. This objective is different from the motivation of robotic racing, where as-fast-as-possible synthetic drivers outperform human drivers [2]. However, the human decision-making process during racing is extremely complex and thus difficult to model, since:

1) many influencing factors exist;
2) vehicle dynamics are highly nonlinear and race cars are usually driven at the limits of handling, posing a difficult control task;
3) each driver exhibits an individual driving style;
4) human generalization and adaptation mechanisms are complex.

While challenges 1–3 have been successfully addressed in recent research with a framework that employs a deep neural network controller to capture these three aspects of human driving [3], [4], the problem of integrating human *adaptation* into a race driver model[1] remains unsolved. With this work, we intend to identify and better understand adaptation and learning techniques mastered by professional race drivers from related research and expert knowledge, contribute to the modeling of driver behavior by developing two methods to incorporate this behavior, and evaluate the proposed methodology within a realistic race car simulation environment as in the human-driver-in-the-loop (HDiL) simulator shown in Fig. 1.

A human-like race driver model could considerably extend and improve full vehicle simulations, ultimately enhance the resulting development efficiency and vehicle performance, while being much less expensive compared to HDiL simulations.

A. Problem Statement and Notation

In order to model human race driver behavior, we aim to learn a human-like control policy π^M which maps the current overall state $\mathbf{x}$, including vehicle state and situation on track, to the vehicle control inputs $\mathbf{a} = \begin{bmatrix}\delta & g & b\end{bmatrix}$ composed of steering wheel angle δ, throttle pedal position g and brake pedal

[1]A driver model represents a vehicle control policy aiming to mimic the behavior of the human race driver in order to support full vehicle simulations.

Manuscript received 21 July 2022; revised 3 October 2022 and 13 January 2023; accepted 30 May 2023. This work was supported by Dr. Ing. h.c. F. Porsche AG. This article was recommended by Associate Editor B.-F. Wu. *(Corresponding author: Stefan Löckel.)*

This work involved human subjects or animals in its research. The authors confirm that all human/animal subject research procedures and protocols are exempt from review board approval.

Stefan Löckel and Siwei Ju are with the Computer Science Department, Technische Universität Darmstadt, 64289 Darmstadt, Germany, and also with Dr. Ing. h.c. F. Porsche AG, 71287 Weissach, Germany (e-mail: loeckel@ias.tu-darmstadt.de).

Maximilian Schaller is with ETH Zürich, 8092 Zürich, Switzerland.

Peter van Vliet is with Dr. Ing. h.c. F. Porsche AG, 71287 Weissach, Germany.

Jan Peters is with the Computer Science Department, Technische Universität Darmstadt, 64289 Darmstadt, Germany.

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TSMC.2023.3285588.

Digital Object Identifier 10.1109/TSMC.2023.3285588

Fig. 1. Race car simulator at Porsche Motorsport [5]: Realistic visualization, a vehicle cockpit mounted on an actuated platform, and a high-fidelity vehicle dynamics model facilitate rapid testing of new vehicle configurations with the human driver in the loop. The vehicle model is developed in-house, has 14 degrees of freedom, and is validated using real-track data. It is accompanied by a high-resolution, laser-scanned track model. Details about the simulator can be found in [6]. This simulator is used to generate demonstration data from professional race drivers for our adaptive human driver model. Consequently, the simulator's vehicle model is taken to evaluate the human driver model, intending to support the future vehicle development process.

actuation b. This policy should be able to robustly maneuver a race car at the handling limits while being similar to the unknown internal driving policy π^E of human experts. At the same time, this expert policy is nondeterministic due to natural human imprecision and intentional adaptation, and able to generalize to new situations as, for example, new race tracks. In this work, we aim to approach the problem of modeling this behavior by:

1) identifying and understanding certain aspects of the most important adaptation and learning mechanisms through related work and expert interviews;
2) using these findings to considerably extend a data-based driver modeling approach;
3) evaluating the developed methods using data from professional race drivers and a state-of-the-art motorsports simulation environment.

Consequently, the resulting driver-specific control policy π^M should be able to generalize to unseen tracks and exhibit certain adaptation characteristics of the human driver. We thereby focus on the adaptation *result*, finishing laps with sufficient performance.

B. Related Work

This section discusses related work in all relevant fields, from methods to analyze or achieve optimal racing performance, to past work on the analysis, modeling, and imitation of human driver behavior, and research on the analysis of human adaptation behavior.

Optimal Racing Performance: To model the physics of a car in different driving situations, a variety of approaches with different complexity is available [1]. In classical control-based approaches, such vehicle models can be used to predict the driving behavior in standard maneuvers or to estimate the vehicle performance on a particular race track using lap time simulation approaches [7], [8], [9]. In the field of autonomous driving or racing, more recent research aims to achieve optimal performance with (data-driven) model predictive control (MPC) [10], [11], [12]. Furthermore, reinforcement learning can be used to train an agent that outperforms human drivers in simulated race environments [2], [13].

HDiL Simulation and Analysis: However, individual human driver behavior, being an important component of the vehicle-driver-entity, is often not sufficiently considered by these methods. This fact encourages motorsport teams to utilize HDiL simulation approaches, where the real driver operates the vehicle within a realistic simulation environment, facilitating faster prototyping and more realistic predictions of the true vehicle performance [6]. Furthermore, HDiL simulators enable the study of human driver behavior, for instance, perceptual and cognition skills of professional and nonprofessional race car drivers [14].

Modeling of Human Driver Behavior: Accordingly, a variety of related work describes car racing from the driver's perspective, analyzes racing techniques, driving lines, and the complex decision-making processes in greater detail, and contributes to a better understanding of the human driver in general [15], [16], [17]. Nevertheless, the task of modeling this behavior remains highly challenging. A number of approaches for building a driver model for different use cases mainly rely on conventional control architectures in partial driving scenarios [18], [19]. Using a cognitive architecture based on adaptive control, the driving behavior is modeled in a highway environment [20]. Some recently developed methods utilize imitation learning techniques to imitate human drivers: using supervised learning, random forests were trained to predict car control inputs from basic vehicle states [21] and it was shown that a feedforward neural network is able to track a driving line generated by a human [22]. Furthermore, methods based on (inverse) reinforcement learning were used to mimic drivers in highway driving scenarios [23], [24], [25], and were extended to imitate human behavior in a short-term race driving setting based on visual features [26]. By imitating a coach, reinforcement learning also enables end-to-end urban driving [27]. Besides that, research also targets specific human individuals [28], [29], [30] and hierarchical modeling [31]. These studies give insights into autonomous driving and driver modeling, but most of them are designed for urban driving and lack the ability to adapt when used for race car driving.

Probabilistic Modeling of Driver Behavior (ProMoD): Among the research on the modeling of human driver behavior, the ProMoD framework was demonstrated to be capable of completing full laps with a competitive performance by mimicking professional race drivers [3], [4]. The data-based and modular approach learns distributions of driving lines represented by probabilistic movement primitives (ProMPs) [32], [33] and trains a recurrent neural network on human race driver data in a supervised fashion. Furthermore, the driver identification and metrics ranking algorithm (DIMRA) was developed to classify individual driving styles using clustering algorithms and was later used as an evaluation method for the learned driver model [4].

Human Adaptation Behavior: Related to this topic, there seems to be a shift from linear and time-invariant models of human manual control to nonlinear and time-varying approaches that are apparent in current research trends [34]. In particular, adaptation over time is identified as a key aspect of human behavior that should and can be modeled by moving toward time-varying models. While the ProMoD framework is shown to work well in many situations, it is still lacking the functionality of a time-varying model, i.e., the ability to learn to drive on unknown tracks and to adapt and learn from gathered experience from driven laps. As such learning and adaptation aspects play fundamental roles in competitive motorsports, any robust and accurate driver modeling approach should be able to reflect them.

Human adaptation behavior w.r.t. adaption times for changing road types in a driving simulator is analyzed, yet not modeled in the work of [35]. Past research on modeling driver adaptation to sudden changes in the vehicle dynamics takes into account limb impedance modulation and updating of the driver's internal representation of the vehicle dynamics [36]. However, the latter work focuses exclusively on lateral dynamics with a first-principles approach without a superordinate objective such as lap time.

Among these approaches, ProMoD offers a solid foundation for this work, as the modeling approach is able to dynamically control a car in a race driving setting, mimicking individual driver behavior without achieving super-human performance. In this work, we considerably modify and extend ProMoD to model human driving adaptation—to the best of our knowledge, for the first time in the racing context. With the modular architecture, the driving policy adaptation remains interpretable. We considerably enhance the quality of a modern driver modeling approach, contribute to a better understanding of human race driver behavior, and aim to pave the way for more accurate vehicle simulations and, potentially, future autonomous racing.

II. METHODOLOGY

As a proper understanding of the human race driver is fundamental for modeling its learning techniques, we ground our methodology on key insights from literature, supplemented by findings of an expert interview with a professional race engineer[2] for LMP1[3] race cars. To derive modeling principles, we summarize these literature and expert insights into adaptation principles and select the most distinct of them with the help of a simple, questionnaire-based survey conducted with professional race drivers and expert motorsport engineers, as an extra layer of expertise. The adaptation principles identified in Section II-A are followed by a short summary of the recently presented ProMoD driver modeling framework in Section II-B. In Section II-C, we present a novel way to generalize the driver model to new tracks. Finally, Section II-D introduces a new method to optimize driving similar to a race driver based on previous laps.

[2]A race engineer works at the interface between the driver and the vehicle, trying to help the driver work with the vehicle and to find a vehicle setup tailored to the driver's needs.

[3]Le-Mans-Prototypes represent a top class of race cars used in different endurance racing series with races lasting up to 24 h.

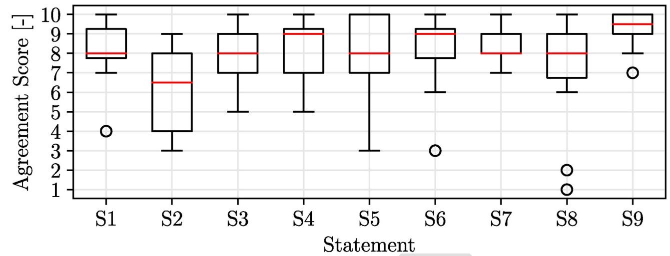

Fig. 2. Agreement levels of 12 experts with statements **S1**–**S9**. The experts were asked to choose an agreement level between 1 and 10 with step size 1. Red lines indicate the median. Boxes represent the interquartile range. The whiskers measure 1.5 times the interquartile range.

A. Adaptation Principles

Race drivers constantly pursue better racing performance in the presence of new tracks and modified vehicle setups. In this section, we aim to understand the most important principles for their adaptation behavior. We gather the following key insights from literature, extended with an expert interview[4] of a race engineer in the Appendix. We aggregate these two sources of insights into summarizing statements **S1**–**S9** detailed in the following, and finally conduct a simple, questionnaire-based survey to directly ask additional experts for their agreement with these statements. In Fig. 2, we measure the agreement of 12 additional experts (including drivers, race engineers, vehicle engineers, and tire engineers) with the 9 statements.

Objective (Delta) Lap Time: In order to (iteratively) optimize lap time [37], race drivers pay attention to the delta lap time, which is the difference between the current and the last (or best) lap time[9] (**S1**). Modifications to the vehicle setup and environmental changes are only considered a posteriori, which means that race drivers usually do not plan with them, but only react after experiencing them[9] (**S2**).

Risk Awareness: Race drivers are particularly risk-aware and constantly test for the vehicle limits [17], starting from a safe region and improving their driving incrementally[9] (**S3**).

Hierarchy: The choice of brake points heavily influences the speed profile of the entire corner [16], [38]. Subsequently, the speed profile heavily influences the driving line. Race drivers control brake points, speed profile, and driving line hierarchically, in this order[9] (which means that brake points are the main tuning knob) (**S4**).

Initialization—Driving on New Tracks: When starting on a new track, drivers tend to compare all new situations and corners to their experience from other tracks [16], [38], to get an initial guess of reasonable brake points and driving lines, which are subsequently refined[9] (**S5**). The initialization of brake points begins already before starting to drive, while the speed profile and driving line are initialized during the first few laps[9] (**S6**). After the first few laps, drivers are able to complete the lap with a close to competitive lap time[9] (**S7**).

Iteration—Adaptation Rules and Quantities: The general adaptation strategy seems to be similar for all drivers, where adaptation of the braking (brake points and peak brake pressure) is particularly important[9] (**S8**). By fine-tuning brake

[4]Findings from the expert interview are marked with this footnote. A summary of the interview is given in the Appendix.

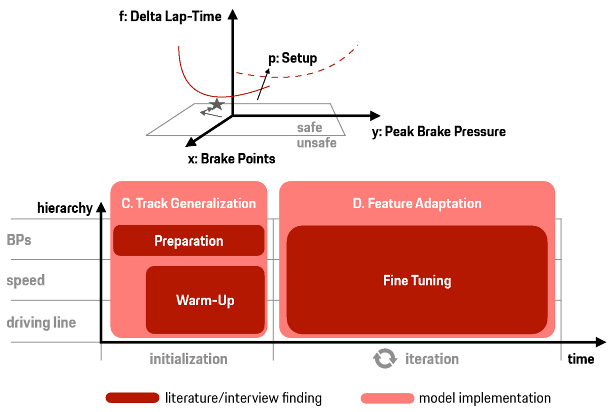

Fig. 3. Top: Iterative adaptation process visualized as an optimization problem. Bottom: Three phases of driver adaptation to solve the above optimization problem, arranged in hierarchy-time plane. Dark color denotes findings from the expert interview and related work, whereas light color signifies how the respective findings are implemented in the adaptive model.

points and peak brake pressure, drivers manage to achieve better performances[9] (**S9**).

Overall, the agreement level of the additional experts with the above statements is high. The lowest median agreement is 7 for **S2** on environmental changes viewed as a disturbance. The corresponding lower end of the interquartile range is 4, much less than 7 (or higher) for all other statements. Hence, we do *not* base our subsequent design choices on **S2**. Further, we observe outliers that might be connected to the diverse backgrounds of the 12 experts. Two tire engineers strongly disagree with **S8** on braking being particularly important for adaptation, applicative as a rule for all drivers. In contrast, both asked drivers strongly agree with this statement. Since drivers are the modeling target themselves, we decide to approve the main expert's statement that brake points are the key control variables. To summarize and simplify the problem, we set up the following qualitative model: Race drivers optimize delta lap time as a function of brake points, peak brake pressure, and other variables as visualized in Fig. 3. This function is parameterized through the vehicle setup. To solve this problem, the brake point variables are initialized in the *Preparation* phase in a safe region, i.e., such that the lap can be completed. Speed and driving line are initialized in hierarchical order during the *Warm-Up* phase. Afterward, drivers iteratively adapt and try out changes on all three hierarchical levels during *Fine-Tuning*.[5] Eventually, they arrive close to the optimizer shown as a star on the top of Fig. 3. This point usually lies close to the boundary of the safe set, as the driver will be operating the vehicle at the handling limits. As these generalization and adaptation capabilities are fundamental for professional race drivers, a driver model used for full vehicle simulations is required to have them as well. In the following, the basic ProMoD framework will be derived and subsequently extended with these skills.

B. ProMoD

The recently presented ProMoD framework combines knowledge and ideas from both race driver behavior and

[5]In the following, heuristically defined *control points* will be introduced for different vehicle states to directly adapt all three hierarchy levels.

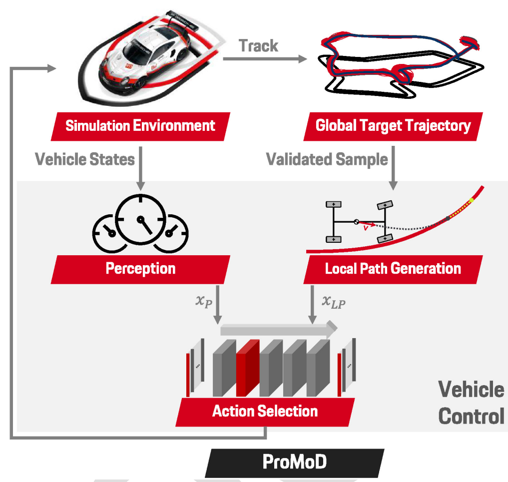

Fig. 4. Original ProMoD framework to imitate human race drivers in simulation [4]: *Global Target Trajectory* holds a distribution of potential target driving lines, relating to the driver's mental image of a driving corridor. *Local Path Generation* and *Perception* calculate a feature vector based on the current situation on track and a sampled target driving line. *Action Selection* maps the features to driver actions. Feeding back the predicted actions to the simulation environment closes the loop.

autonomous driving architecture. It consists of multiple modules as visualized in Fig. 4, where each of these modules represents fundamental steps in the decision-making process of a human race driver [3], [4].

Our novel generalization and adaptation methods are based on this architecture, which is summarized in the following.

Global Target Trajectory: Every driver keeps a mental image of the whole race track in their head, knowing approximately where to brake, to turn in, and to accelerate again in each corner. However, this imagined driving corridor is not precise, i.e., it incorporates variance, and additionally changes over time with gathered experience. Hence, we model the global target trajectory with a distribution over potential driving lines, which could be interpreted as a driving corridor, using ProMPs [32], [33]. For this purpose, both the spatial and the temporal information of every demonstrated driving line on a particular track is projected to a lower-dimensional weight space. We define a series of equally distributed radial basis functions (RBFs)

$$b_j(s) = \exp\left(-\frac{(s - c_j)^2}{2h}\right) \tag{1}$$

with function index $j \in \{1, 2, \ldots, N_{\mathrm{BF}}\}$, track distance s, constant width h, and c_j being the equally distributed centers of the functions. All basis functions are assembled into the basis function matrix $\boldsymbol{\Phi}_s \in \mathbb{R}^{N_s \times N_{\mathrm{BF}}}$, where the jth column contains $b_j(s)$ evaluated at N_s points, equidistant in terms of track distance. Subsequently, $\boldsymbol{\Phi}_s$ is aggregated into

$$\boldsymbol{\Psi}_s = \mathrm{diag}(\boldsymbol{\Phi}_s, \boldsymbol{\Phi}_s, \ldots, \boldsymbol{\Phi}_s) \in \mathbb{R}^{nN_s \times nN_{\mathrm{BF}}} \tag{2}$$

for n variables that the trajectory consists of. The weight vector

$$\boldsymbol{w}_i = \left(\boldsymbol{\Psi}_s^T \boldsymbol{\Psi}_s + \epsilon \boldsymbol{I}\right)^{-1} \boldsymbol{\Psi}_s^T \boldsymbol{\tau}_{s,i} \in \mathbb{R}^{nN_{\mathrm{BF}}} \tag{3}$$

is derived using ridge regression for each demonstration trajectory $\boldsymbol{\tau}_{s,i} \in \mathbb{R}^{nN_s}$ and regularization factor ϵ. By fitting a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{w}}, \boldsymbol{\Sigma}_{\boldsymbol{w}})$ over the N demonstration weights with mean $\boldsymbol{\mu}_{\boldsymbol{w}}$ and variance $\boldsymbol{\Sigma}_{\boldsymbol{w}}$

$$\boldsymbol{\mu}_{\boldsymbol{w}} = \frac{1}{N}\sum_{i=1}^{N} \boldsymbol{w}_i \in \mathbb{R}^{nN_{\mathrm{BF}}}, \tag{4}$$

$$\boldsymbol{\Sigma}_{\boldsymbol{w}} = \frac{1}{N}\sum_{i=1}^{N} \left(\boldsymbol{w}_i - \boldsymbol{\mu}_{\boldsymbol{w}}\right)\left(\boldsymbol{w}_i - \boldsymbol{\mu}_{\boldsymbol{w}}\right)^T \in \mathbb{R}^{nN_{\mathrm{BF}} \times nN_{\mathrm{BF}}} \tag{5}$$

we are able to describe the distribution of driving lines for a driver on a particular track efficiently. Subsequently, an arbitrary number of new driving lines which are similar to all demonstrations can be generated by sampling a weight vector from this distribution, $\boldsymbol{w}^* \sim \mathcal{N}\left(\boldsymbol{\mu}_{\boldsymbol{w}}, \boldsymbol{\Sigma}_{\boldsymbol{w}}\right)$, and using

$$\boldsymbol{\tau}^* = \boldsymbol{\Psi}_s \boldsymbol{w}^* \tag{6}$$

to retrieve a new driving line in the original formulation which could be subsequently used as *target trajectory*. While sampling this target trajectory all at once models the human driver's ahead-of-time plan based on experience and knowledge of the whole track, real-time planning based on the current state on the track is performed by ProMoD's *Local Path Generation* module.

Local Path Generation: For any situation on track, a human driver continuously plans the upcoming path a few seconds ahead. We use this module to mimic the path planning by calculating constrained polynomials and multiple preview features[6] based on the current vehicle state and the *target trajectory*. These local path features are denoted as $\mathbf{x}_{\mathrm{LP}}$.

Perception: In addition to the path planning features, each driver relies on additional information about their surroundings, such as visual information or experienced accelerations. These perception features, which mostly relate to basic vehicle states, are gathered inside this module and are denoted as $\mathbf{x}_P$.

Action Selection: The action selection process, i.e., the mapping from the current state (as described by the feature vector $\mathbf{x} = [\mathbf{x}_{\mathrm{LP}}\ \mathbf{x}_P]$) to human-like control actions $\mathbf{a}$, is learned using a recurrent neural network. It is trained on all available demonstration data for a particular driver, aiming to imitate its individual driving style and incorporating the dynamics of the action selection process.

This modular and hierarchical structure, compared with end-to-end learning such as in [13], increases interpretability when tuning the driving behavior. After the recurrent neural network, i.e., the action selector, is trained, it serves as a controller that drives the car by following the global reference trajectory. Subsequently, by modifying the global reference trajectory, the driver model can be adjusted for performance or generalization. Compared with the direct adaptation of the action selection policy (parameters of the recurrent neural network), the adaptation of the global reference trajectory has the following advantages: 1) fewer parameters to update; 2) an interpretable adaptation process; and 3) predictable and understandable results. In the following, we present methods to generalize and adapt this driver model in two different phases. Section II-C introduces *Track Generalization*, addressing the *Preparation* and *Warm-Up* steps identified from the literature and interview (see Fig. 3). Section II-D describes *Feature Adaptation*, modeling the iterative *Fine-Tuning*.

Algorithm 1 Estimating a Driving Line Distribution + Sampling

$\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa,dy}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa,dy} \leftarrow \textsc{BuildProMP}(\mathcal{D})$
$x'(s), y'(s), \kappa'(s) \leftarrow \textsc{BuildDrivingLine}(\mathcal{B}_{\mathrm{left}}, \mathcal{B}_{\mathrm{right}})$
$\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa'} \leftarrow \textsc{RidgeRegression}(\kappa'(s))$
$\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'} \leftarrow \mathbf{0}$
$\boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'} \leftarrow \textsc{EstimateVariance}(\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa'}, \boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa,dy}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa,dy})$

for $i \leftarrow 1, N_{\mathrm{samples}}$ **do**
  $\boldsymbol{w}_i^{*dy} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'}\right)$
  $x_i^*(s), y_i^*(s) \leftarrow \textsc{Reconstruct}(x'(s), y'(s), \boldsymbol{w}_i^{*dy})$
  $\Delta t_i^*(s) \leftarrow \textsc{EstimateSpeed}(x_i^*(s), y_i^*(s), \mathcal{P})$
end for

C. Track Generalization: Generate Driving Line Distributions

In order to generate first laps on a new, yet unknown track, it is required to learn a reasonable driving line distribution for the *Global Target Trajectory* module. All other modules of ProMoD are track-independent by definition and remain unmodified. Hence, we construct a driving line distribution for a new track based on its borders (assumed to be known) and prior knowledge from other tracks. Inspired by the results from Section II-A, we propose the methodology described in Algorithm 1. We utilize a novel ProMP description, conventional methods to fit driving lines based on geometric boundaries, and a method to estimate the variance of the driving line around the track based on experience from other tracks.

ProMPs on Demonstration Data: To encode prior knowledge from other tracks, we use all available driving line data from all known tracks $\mathcal{D}$ and calculate ProMPs with a modified representation as driving line distributions for each track separately. In particular, we take the vehicle positions in the Cartesian space for all laps on a given track and map them to a curvilinear description $x(s),\ y(s) \mapsto dy(s),\ \kappa(s)$ for each track. Thereby, dy represents the lateral deviation from a reference line and κ the line curvature, both based on the reference line distance s. While there is an overlap between the information in dy and κ, both representations are needed for subsequent calculations. Similar to the computation of RBF weights via ridge regression in the Cartesian space, driving lines are now represented by weight vectors $\boldsymbol{w}^{dy}$ and $\boldsymbol{w}^{\kappa}$ for dy, κ, and RBFs in the curvilinear space with equidistant discretization. Assuming a Gaussian distribution, we retrieve mean weight vectors $\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa}$, $\boldsymbol{\mu}_{\boldsymbol{w}}^{dy}$ and variances $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa}$, $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy}$ to describe the distribution of all available driving lines on a particular track. By iterating this process for all available tracks, we can aggregate all driving line information into $\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa,dy}$, $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa,dy}$. In the following, we estimate a driving line distribution for an unknown track by combining this stochastic information with a conventional path planning method.

Generate Mean Driving Trajectory: We start by estimating a mean driving trajectory which is only based on the given track

[6]Examples are a predicted lateral offset or a predicted speed difference from the target driving line. More details are given in [4].

boundaries $\mathcal{B}_{\text{left}}$ and $\mathcal{B}_{\text{right}}$. As the generation of a reasonable and collision-free path around the track is required, we decide to use Elastic Bands [39], [40]. While being computationally efficient and easy to interpret, this method exhibited reasonable driving line estimates with sufficient accuracy. The resulting trajectory is now taken as the reference and mean driving line for the new track. Similarly to the ProMP calculation on the available demonstration data, the curvature $\kappa'(s)$ of this Elastic Band driving line is projected to the lower-dimensional weight space and set as the mean curvature $\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa'}$ with $\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'} = \mathbf{0}$ by definition.

Variance Estimation: Using this mean trajectory and the existing corner information from other tracks, we estimate the variance with a sliding window approach. For this purpose, we are moving along the estimated mean driving line's curvature $\kappa'(s)$ and compare the current situation, described by a sequence of curvatures, to all situations on all known tracks as encoded in $\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa,dy}$, $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa,dy}$. By finding the most similar corner measured by the absolute difference between curvatures, we are now able to iteratively build $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'}$, which describes the variance of driving lines on the new track.[7]

Sampling and Reconstruction: Using the Elastic Band estimate $x'(s)$, $y'(s)$ for the mean driving line and the modified ProMP $\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'} = \mathbf{0}$, $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'}$ describing the lateral deviation from the mean, we are now able to sample new driving lines for the new track. In particular, we draw a sample weight vector $\boldsymbol{w}_i^{*dy} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'}\right)$ and retrieve the lateral deviation $dy_i^*(s)$ as $\boldsymbol{\Phi}_s \boldsymbol{w}_i^{*dy}$. Now, it is possible to construct a sample driving trajectory in the Cartesian space using

$$x_i^*(s) = x'(s) - \sin\big(\phi_i^*(s)\big)dy_i^*(s) \tag{7}$$

$$y_i^*(s) = y'(s) + \cos\big(\phi_i^*(s)\big)dy_i^*(s) \tag{8}$$

where ϕ_i^* is the mean heading angle of the vehicle and equals 0 when the vehicle drives purely into x-direction.

Speed Profile: In addition to the trajectory of the vehicle, ProMoD requires a speed profile for the *Local Path Generation* module. Since this velocity profile depends on the vehicle and its setup and is hard to estimate using the available demonstration data, we follow a more robust approach based on vehicle dynamics. For each sampled vehicle trajectory $x_i^*(s)$, $y_i^*(s)$, we utilize a conventional lap time estimation approach based on the vehicle performance envelope $\mathcal{P}$ to retrieve an approximate speed profile [7], [41].

Simulation: The sampled driving lines with corresponding speed profiles can now be used to reconstruct the original ProMP formulation within the previously presented ProMoD framework. Initializing with a reduced performance envelope $\mathcal{P}$ represents the *Preparation* phase on a new track and allows for safely simulating first laps. By iteratively expanding $\mathcal{P}$ and simulating the resulting driving lines and speed profiles, ProMoD is able to cautiously approach the vehicle limitations, aiming to mimic the *Warm-Up* phase. The complete process facilitates simulations on new tracks for which no demonstration data exists, enhancing our driver modeling framework with track familiarization abilities to generate first fast laps. After becoming familiar with a track, human drivers continuously optimize their performance, as shown in Section II-A. Hence, ProMoD should also be adaptable and learn from experience, which necessitates adaptation techniques.

[7]We use the curvature κ to find similar corners since it naturally describes the corner shape. The lateral deviation dy is used for sampling, as it allows for a more robust reconstruction.

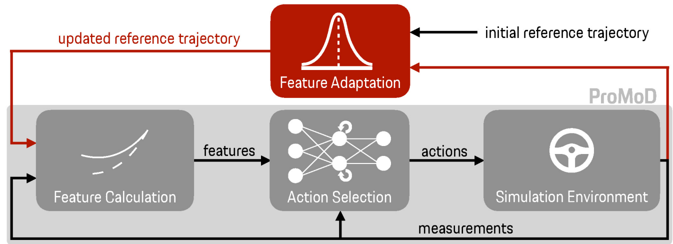

Fig. 5. Feature adaptation (red) extending the original ProMoD framework (gray), consisting of feature calculation (summarizes *Local Path Generation* and *Perception*), *Action Selection*, and the simulation environment. Every finished lap is analyzed and the reference trajectory is adapted correspondingly.

D. Feature Adaptation

Professional race drivers master the skill of continuously optimizing their performance by analyzing past laps and adapting accordingly. With an additional feedback loop as shown in Fig. 5, ProMoD is enabled to mimic this learning process to a certain extent. By only adapting the global target trajectory, which is used to compute local path planning features $\mathbf{x}_{\text{LP}}$, the behavior of ProMoD can be influenced. At the same time, ProMoD maintains its ability to imitate human drivers as the action selection module remains unchanged. In the following, we use *Conditioning* and *Scaling* to modify the global target trajectory while keeping it human-like:

Conditioning: Recall that the ProMPs for the global target trajectory are represented by a Gaussian weight distribution $p(\boldsymbol{w}) = \mathcal{N}\big(\boldsymbol{w} \mid \boldsymbol{\mu}_{\boldsymbol{w}}, \boldsymbol{\Sigma}_{\boldsymbol{w}}\big)$ with mean weight vector $\boldsymbol{\mu}_{\boldsymbol{w}}$ and covariance matrix $\boldsymbol{\Sigma}_{\boldsymbol{w}}$. We are now able to alter this distribution by conditioning the prior distribution to a new (algorithmically chosen) observation $\boldsymbol{x}_{s'}^* = \{\boldsymbol{y}_{s'}^*, \boldsymbol{\Sigma}_y^*\}$ at a specific location $s = s'$, as presented in [33]. Here, the control point $\boldsymbol{y}_{s'}^* \in \mathbb{R}^n$ is an algorithmically chosen target state (see Paragraph *Adaptation Process* for details) of the vehicle position and velocity to be reached at distance s', and variance $\boldsymbol{\Sigma}_y^* \in \mathbb{R}^{n\times n}$ is the confidence of this choice. The conditional distribution $p\big(\boldsymbol{w} \mid \boldsymbol{x}_{s'}^*\big)$ remains Gaussian with updated parameters

$$\boldsymbol{\mu}_{\boldsymbol{w}}^{[\text{new}]} = \boldsymbol{\mu}_{\boldsymbol{w}} + \boldsymbol{L}\big(\boldsymbol{y}_{s'}^* - \boldsymbol{\Psi}_s^T \boldsymbol{\mu}_{\boldsymbol{w}}\big), \tag{9}$$

$$\boldsymbol{\Sigma}_{\boldsymbol{w}}^{[\text{new}]} = \boldsymbol{\Sigma}_{\boldsymbol{w}} - \boldsymbol{L}\boldsymbol{\Psi}_{s'}^T \boldsymbol{\Sigma}_{\boldsymbol{w}} \tag{10}$$

where

$$\boldsymbol{L} = \boldsymbol{\Sigma}_{\boldsymbol{w}} \boldsymbol{\Psi}_{s'}\Big(\boldsymbol{\Sigma}_y^* + \boldsymbol{\Psi}_{s'}^T \boldsymbol{\Sigma}_{\boldsymbol{w}} \boldsymbol{\Psi}_{s'}\Big)^{-1} \tag{11}$$

relates the variances of the prior distribution and the new observation with $\boldsymbol{\Psi}_{s'} \in \mathbb{R}^{nN_{\text{BF}}\times n}$ representing the value of all basis functions at $s = s'$ [33].

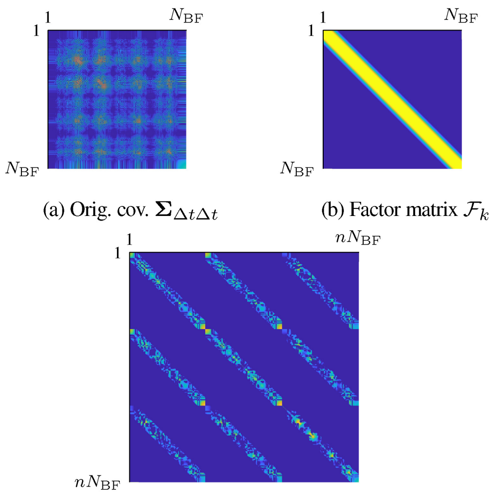

(a) Orig. cov. $\boldsymbol{\Sigma}_{\Delta t \Delta t}$ (b) Factor matrix $\mathcal{F}_k$

(c) Pre-processed cov. $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\text{masked}}$

Fig. 6. Masking the covariance matrix. (a) Part of the covariance matrix for a single variable ($\boldsymbol{\Sigma}_{\Delta t \Delta t} \in \mathbb{R}^{N_{\text{BF}} \times N_{\text{BF}}}$), where brighter colors indicate higher covariances. Far-off-diagonal correlations in the data potentially result from different vehicle setups in the demonstration data but are difficult to consider during conditioning. (b) Factor matrix for a single variable, where the elements on the diagonal are one, and off-diagonal entries are fading out to zero using bandwidth k. Here, k is selected such that distant and nonconsecutive turns cannot mutually influence each other. (c) Resulting matrix $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\text{masked}} \in \mathbb{R}^{nN_{\text{BF}} \times nN_{\text{BF}}}$ for three variables after masking, filtering out correlations over larger distances.

This procedure allows to move brake points or to shift apexes[8] by conditioning the prior distribution utilizing a set of rules derived from Section II-A. In the meantime, the correlations between different locations are taken into consideration by the covariance matrix which is learned from the data so that the whole trajectory is modified correspondingly. However, when using the prior variance without further consideration, conditioning at a specific turn potentially affects distant turns due to nonzero covariances in the data, as shown for $\boldsymbol{\Sigma}_{\Delta t \Delta t}$ in Fig. 6(a). As such a large effect across multiple turns is not considered to be human-like, we aim to reduce it by masking the original matrix using a factor matrix $\mathcal{F}_k \in \mathbb{R}^{N_{\text{BF}} \times N_{\text{BF}}}$ shown in Fig. 6(b). By multiplying $\mathcal{F}_k$ element-wise with each submatrix of $\boldsymbol{\Sigma}_{\boldsymbol{w}}$, we retrieve a masked matrix for conditioning

$$\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\text{masked}} = \begin{bmatrix} \mathcal{F}_k \circ \boldsymbol{\Sigma}_{xx} & \mathcal{F}_k \circ \boldsymbol{\Sigma}_{xy} & \mathcal{F}_k \circ \boldsymbol{\Sigma}_{x\Delta t} \\ \mathcal{F}_k \circ \boldsymbol{\Sigma}_{yx} & \mathcal{F}_k \circ \boldsymbol{\Sigma}_{yy} & \mathcal{F}_k \circ \boldsymbol{\Sigma}_{y\Delta t} \\ \mathcal{F}_k \circ \boldsymbol{\Sigma}_{\Delta tx} & \mathcal{F}_k \circ \boldsymbol{\Sigma}_{\Delta ty} & \mathcal{F}_k \circ \boldsymbol{\Sigma}_{\Delta t\Delta t} \end{bmatrix} \tag{12}$$

which effectively lowers the influence of conditioning on distant regions as shown in Fig. 6(c).[9] This matrix can then replace $\boldsymbol{\Sigma}_{\boldsymbol{w}}$ for effective local *Conditioning*.

Scaling: In order to fully utilize the vehicle's potential on straights, the speed profile can be adapted to influence the throttle actuation and braking behavior of ProMoD. Since the neural network performs trajectory tracking, aiming to minimize the control error between the reference speed and the actual speed, its output signals tend to fluctuate during intervals of full throttle. Therefore, if the actual velocity is larger than the reference velocity, ProMoD tends to accelerate less, even if the virtual driver is on a straight and expected to drive as fast as possible. This problem can be effectively solved by smoothly scaling the reference speed on long straights.

[8] An apex is defined as the closest point to the inner side of a corner, typically coinciding with the locally maximal curvature of the driving line.

[9] While the assumption of a fixed bandwidth k is not entirely human-like, it turned out to be sufficient to introduce the required adaptation characteristics. Future work may focus on finding a variable, distance-dependent masking to further enhance human likeness.

Algorithm 2 Adaptation Process

Input: $\boldsymbol{\mu}_{\boldsymbol{w}}^0$, $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{w}}^0$, envelope
$\boldsymbol{\Sigma}_{\boldsymbol{w}}^0 \leftarrow$ PROCESSVARIANCE($\hat{\boldsymbol{\Sigma}}_{\boldsymbol{w}}^0$)
$\boldsymbol{\tau}_0^{ref} \leftarrow$ CALCMEANTRAJECTORY($\boldsymbol{\mu}_{\boldsymbol{w}}^0$)
$\boldsymbol{\mathcal{I}}^{track} \leftarrow$ ANALYSETRACK($\boldsymbol{\tau}_0^{ref}$)
for $i = 0, 1, 2, \ldots$ **do**
  $\boldsymbol{\tau}_i \leftarrow$ SIMULATE($\boldsymbol{\tau}_i^{ref}$)
  $\boldsymbol{y}_s^* = \emptyset$
  if not ISCOMPLETED($\boldsymbol{\tau}_i$) **then**
    $\boldsymbol{y}_s^* \leftarrow \boldsymbol{y}_s^* \cup$ ANALYSEDL($\boldsymbol{\tau}_i, \boldsymbol{\mathcal{I}}^{track}$, envelope)
    if $\boldsymbol{y}_s^* == \emptyset$ or SLIPCHECK($\boldsymbol{\tau}_i, \boldsymbol{\mathcal{I}}^{track}$) **then**
      $\boldsymbol{y}_s^* \leftarrow \boldsymbol{y}_s^* \cup$ ADAPTSPEED($\boldsymbol{\tau}_i, \boldsymbol{\mathcal{I}}^{track}$)
    end if
  else
    $\boldsymbol{y}_s^* \leftarrow \boldsymbol{y}_s^* \cup$ CHECKINENVELOPE($\boldsymbol{\tau}_i$, envelope)
  end if
  $\boldsymbol{\mu}_{\boldsymbol{w}}^{i+1,0}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{i+1,0} = \boldsymbol{\mu}_{\boldsymbol{w}}^{i+1}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{i+1}$
  for $j = 1, 2, \ldots$, number of items in $\boldsymbol{y}_s^*$ **do**
    $\boldsymbol{\mu}_{\boldsymbol{w}}^{i+1,j}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{i+1,j} \leftarrow$ COND($\boldsymbol{\mu}_{\boldsymbol{w}}^{i+1,j-1}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{i+1,j-1}, \boldsymbol{y}_{s,j}^*$)
  end for

  $\hat{\boldsymbol{\tau}}_{i+1}^{ref} \leftarrow$ CALCMEANTRAJECTORY($\boldsymbol{\mu}_{\boldsymbol{w}}^{i+1}$)
  if ISCOMPLETED($\boldsymbol{\tau}_i$) **then**
    $\boldsymbol{\tau}_{i+1}^{ref} \leftarrow$ SPEEDSCALING($\hat{\boldsymbol{\tau}}_{i+1}^{ref}, \boldsymbol{\tau}_i, \boldsymbol{\mathcal{I}}^{track}$)
  else
    $\boldsymbol{\tau}_{i+1}^{ref} \leftarrow \hat{\boldsymbol{\tau}}_{i+1}^{ref}$
  end if
end for

Adaptation Process: The complete adaptation process, shown in Algorithm 2, is inspired by the insights from Section II-A and uses both introduced methods, *Conditioning* and *Scaling*, to continuously adapt ProMoD based on gathered experience. After simulating a lap, an initial check is done whether the lap was completed successfully. If this is not the case, the situation where the vehicle left the track is analyzed and the ProMP is conditioned using two subprocedures.

1) *Driving-Line Check and Adaptation:* As seen in Section II-A, the turn-in is the most important phase during cornering. Hence, the driving line is compared to the permissible driving corridor, represented by track borders or by the envelope of all demonstrations from the human drivers, and the largest deviation before the apex is found. Then, a new control point $\boldsymbol{y}_{s'}^*$ is added for *Conditioning* at this position, shifting the driving line distribution toward the permissible area.
2) *Velocity Adaptation:* If no valid adaptation is found or extreme tire slip occurs, a control point will be added to reduce the target speed shortly before the track was left.

In practice, ProMoD can eventually complete each critical corner when the target speed is low enough. Subsequently, the completed laps can be further adapted to improve the lap time and to keep the driving line in the envelope by:

1) Checking and reducing smaller deviations from the permissible driving corridor: Just like during real racing, ProMoD sometimes slightly exceeds the theoretically allowed driving corridor but still manages to complete

the lap. These situations are checked and additional control points are introduced for *Conditioning*.

2) Checking acceleration intervals and *Scaling* of the speed: As discussed before, sometimes ProMoD does not utilize the full vehicle potential during acceleration phases on straight lines. Hence, speed scaling is used to further increase the performance on already completed laps.

By introducing this process, we are able to encourage ProMoD to learn from the experience of previous laps, to correct mistakes, and to increase performance, matching the requirements illustrated in Fig. 3.

III. Evaluation

In this work, we use data of professional race drivers gathered from the HDiL simulator shown in Fig. 1 to train and evaluate our driver model. All rollouts of our driver model are simulated using the same in-house developed vehicle model of a high-performance race car, guaranteeing realistic vehicle dynamics and facilitating comparability to the human demonstrations. The task of driving the simulated race car is highly challenging as its only driver assistance system is *Traction Control*. In order to safeguard intellectual property, all plots in this article are shown normalized.

A. Track Generalization

We evaluate the presented track generalization method of our ProMoD framework on two race tracks, Motorland Aragón (AGN) and the Yas Marina Circuit in Abu Dhabi (ABD), and exclude demonstration data from these tracks during training. For each track, we initially estimate driving line distributions according to the methodology presented in Section II-C and draw N_{samples} driving lines from these distributions. When using these driving line samples for simulation on the corresponding unknown tracks, ProMoD is capable of completing full laps on the respective race track, as visualized in Fig. 7 for ABD.

For AGN, the track generalization method achieves comparable results considering the similarities of the resulting driving line and driver action distributions with the human driver. Furthermore, we compare the performance of ProMoD and the human driver on both tracks with equal vehicle setups. Fig. 8 visualizes the resulting lap time distributions, normalized to the median lap time of the human driver on each track, respectively.

Here, ProMoD is able to achieve lap times close to those of the human driver, with a slightly increased median due to small deviations in the expected speed profiles as visible in Fig. 7(a) between reference distances 0.1 and 0.2. These deviations result from the herein utilized conventional lap simulation approach [7], [41] that marginally underestimates the available acceleration potential and hence permissible speed of the vehicle in dynamic situations. This is a reasonable limitation, as the track generalization method is mainly intended to safely finish first laps on a new track with a close to competitive performance. In contrast to baseline machine learning models and also to conventional lap simulation approaches [7], [41] that rely on simplified vehicle models and do not consider human characteristics, extensive evaluations of the basic ProMoD framework in earlier research [3], [4] already demonstrated that the framework can robustly mimic human driving styles in a variety of settings. These findings are underlined by an extended evaluation of the adapted ProMoD model in the following section.

B. Feature Adaptation

The feature adaptation process is tested on two different tracks, the Silverstone Circuit (SVT) and Motorland Aragón (AGN), as these tracks turned out to be particularly difficult to finish for the driver model and, hence, are a suitable environment to demonstrate the applicability of our method. We start with an evaluation of the local effects of *Conditioning* and *Scaling* by showing the executed adaptations, the resulting changes in terms of driving line, and the selected actions of the driver model. Subsequently, we test the complete adaptation process on both tracks, showing that the method is able to pass previously unfinished turns and to improve lap time.

Local Effect—Adaptation: The local effects of adaptation are presented in Figs. 9 and 10, visualizing adaptations of the driving line and the speed profile, as well as the resulting action signals and driven lines. Here, ProMoD fails initially at Turn (T) 6/7 of SVT due to considerably exceeding the vehicle potential as shown in Fig. 9(b). In order to adapt the speed profile effectively, three control points are used to set the lower peak speed value, resulting in earlier braking and consequently helping to avoid the mistake and pass the turn. At the same time, with the purpose of reducing the curvature and avoiding corner-cutting, the driving line is pulled outwards around fifty meters before the first apex as shown in Fig. 9(a). After two iterations of simultaneously adapting both the speed profile and the driving line, ProMoD succeeds in this turn. Note that such intermediate iterations are part of our modeling *algorithm* and not part of the adaptation model itself, that is resembled by the final iterate of speed profile and driving line.

Local Effect—Scaling: Scaling is particularly useful on straights if ProMoD initially does not fully utilize the vehicle potential due to a modified vehicle setup and a too conservative prior target speed definition. Its effect becomes apparent when observing the throttle actuation signal. With a higher reference speed, the model tends to utilize full throttle more often on long straights, as shown in Fig. 11. Consequently, the fluctuations of the throttle signal in those intervals are eliminated, and the lap time is improved by about 0.2 s.

Adaptation Process: The developed adaptation process for ProMoD has been successfully tested on SVT and AGN as visualized in Fig. 12. While it requires four iterations to complete SVT, ProMoD needs more iterations for AGN since it fails at more locations. On both tracks, the learning speed is slower compared to a race driver, but ProMoD ultimately succeeds in completing a lap after less than 20 iterations, with at most five iterations for a problematic turn. To indicate the adaptation progress, the lap progress and the portion of the

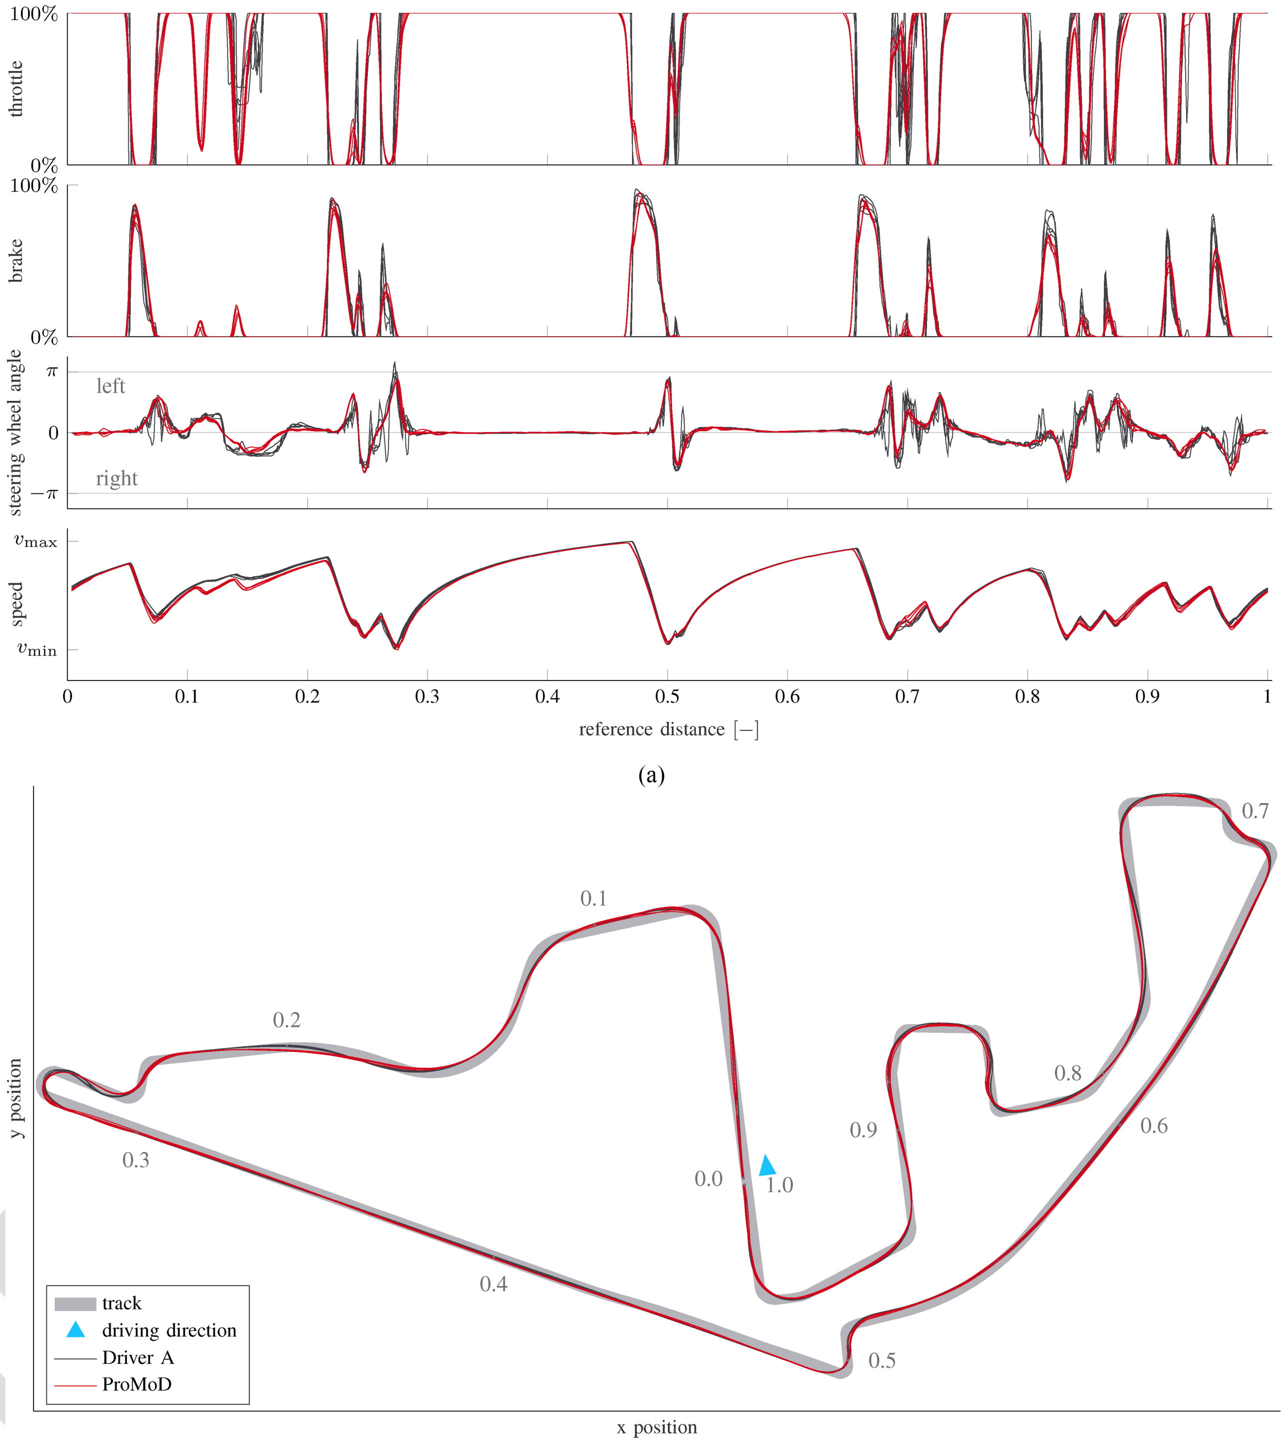

Fig. 7. Track Generalization results on ABD: We compare five laps of the human driver (dark gray) to five laps of the track-generalized ProMoD framework (red) with an identical vehicle setup. (a) Comparison of the driver actions and the resulting speed profiles over the normalized track reference distance. Here, ProMoD is able to approximately reproduce the throttle, braking, and steering activity of the real driver considering the braking points, actuation speeds, and amplitudes. The velocity profile shows small deviations after the first corner where ProMoD does not fully utilize the vehicle potential due to a slightly over-conservative speed profile estimation in this region. (b) Resulting simulated driving lines around the track (light gray) where numbers indicate the reference distance. The position of the start/finish line and the driving direction is indicated by the bright blue triangle. Here, ProMoD is able to generalize and approximately follows the demonstrations of the human driver even though they were not used during training for this race track. Some deviations are present at particularly challenging locations (e.g., the hairpin corner on the left), which, however, do not prevent ProMoD from finishing the lap with reasonable performance. These deviations may be reduced by using adaptation methods to learn from the gathered experience on the track.

lap with full throttle are plotted over the number of iterations, corresponding to the objective of finishing laps and optimizing the lap time, respectively, while imitating the human drivers.

DIMRA: Finally, we use DIMRA to evaluate the adapted model regarding the similarity of its driving style to that of the target human driver [4]. In Fig. 13, each marker

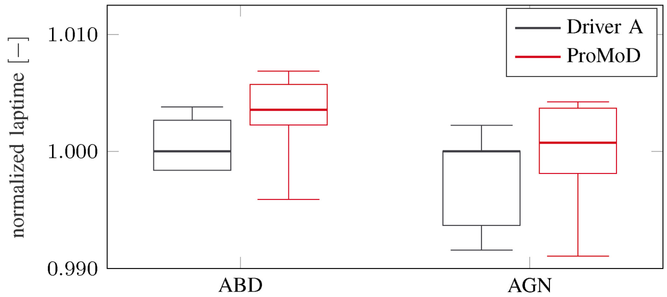

Fig. 8. Lap time comparison for track generalization on race tracks ABD and AGN: Times are normalized to the median demonstration lap time of the corresponding track. The whiskers correspond to the minimum/maximum values, the boxes indicate the upper/lower quartiles, and the thick central line shows the median value. Here, ProMoD is able to finish laps on unknown race tracks, less than 0.5% slower than the human driver in the median and at a competitive pace for its fastest laps. The slightly slower median lap time might be a result of a yet nonoptimal speed profile or driving line distribution.

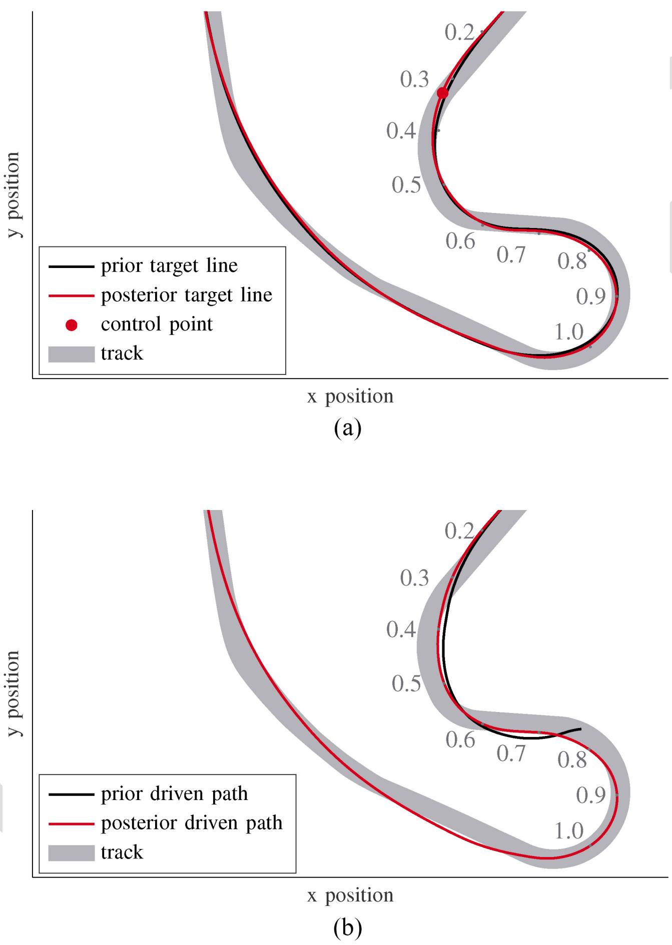

Fig. 9. Adaptation of the target line for T6/7 on SVT and the resulting driven paths. (a) Prior (black) and posterior (red) target lines. The posterior target line is pulled outwards before the first apex using a control point at corner entry, as ProMoD initially exceeded the vehicle potential and left the track. (b) Resulting lines driven by ProMoD. After simultaneous adaptation of the target line and the velocity profile, ProMoD is able to successfully finish this turn.

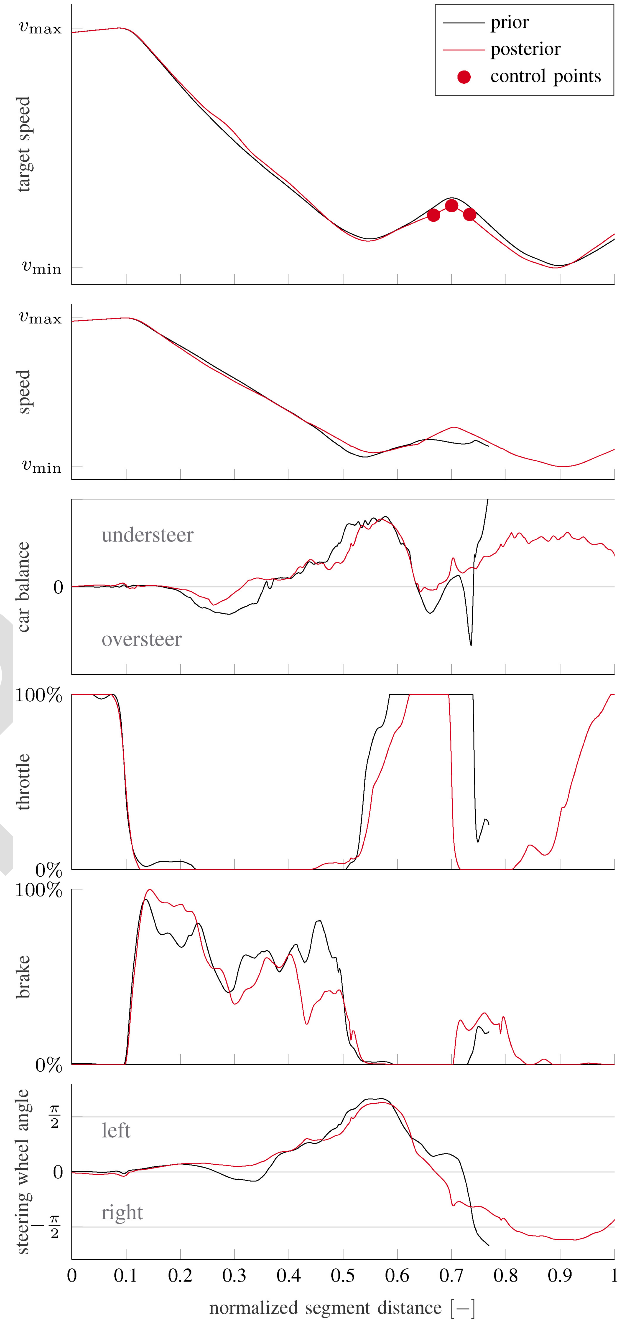

Fig. 10. Target speed and resulting vehicle states and driver actions over the normalized segment distance before and after adaptation (two iterations) of the target speed profile for T6/7 on SVT: Via three control points, the target speed profile is adapted while its general shape is preserved. The car balance refers to the dynamic driving state. When operated close to the friction limit (e.g., while cornering), the car balance typically assumes an oversteer (over-rotating, negative values) or understeer (under-rotating, positive values) state [1]. Before adaptation, at normalized segment distance 0.25, the vehicle oversteers and ProMoD is able to recover the vehicle by countersteering, at the cost of losing speed. However, at distance 0.65, ProMoD largely exceeds the grip potential, sliding over both axles which forces the vehicle off the track [see Fig. 9(b)]. After adapting the speed profile and driving line, ProMoD is able to keep the vehicle safely on track. Via *Action Selection*, ProMoD automatically increases the braking force during the first turn-in, accelerates later, and lifts the throttle and brakes earlier for the following turn.

represents a single lap with three metrics characterizing the individual driving style: throttle speed, brake speed, and the time of simultaneously pressed brake and throttle pedals. This plot indicates that after adaptation, the driver model remains capable of mimicking the individual characteristics of a specific driver while considerably differing from the others.

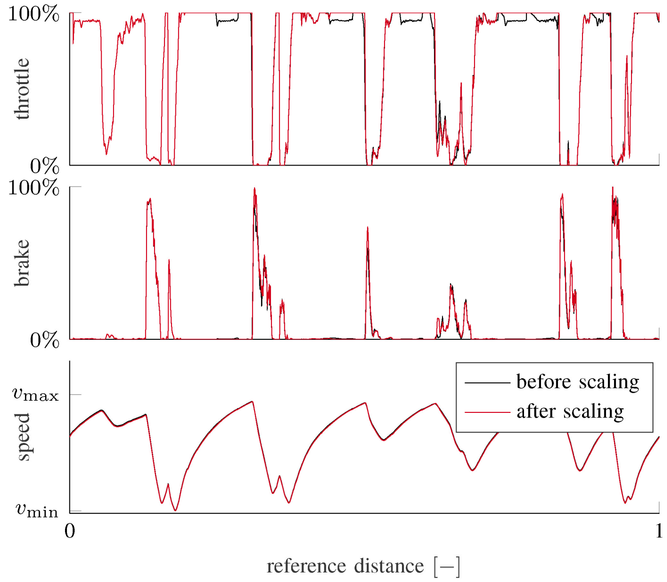

Fig. 11. Effect of *Speed Scaling* on straights: After scaling, ProMoD effectively utilizes the longitudinal potential of the vehicle and uses full throttle on most straights. For intervals where ProMoD would fail in subsequent turns due to the increased speed, scaling is prevented.

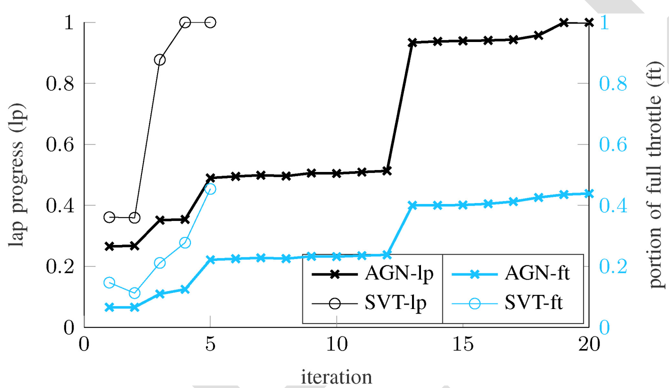

Fig. 12. Adaptation progress of ProMoD on AGN and SVT: For both tracks, ProMoD succeeds in completing a previously unfinished lap within 20 iterations, shown by lap progress (*lp*). The portion of full throttle is denoted by *ft*, where average expert values are 0.6152 and 0.5289 on SVT and AGN, respectively. Additional iterations can be used to further increase performance.

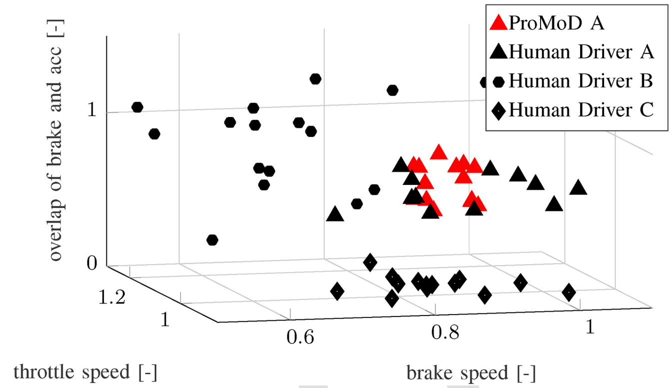

Fig. 13. Top three DIMRA driving style metrics of ProMoD and human drivers on SVT. ProMoD accurately mimics the individual driving style of driver A while still being distinguishable from two other professional race drivers.

IV. Conclusion

In this article, we collect insights into the general adaptation behavior and the learning processes of professional race drivers and derive new methods to extend ProMoD, an advanced modeling method for race driver behavior. With the purpose of understanding driver behavior in general and identifying the most important adaptation processes, this work starts with key insights from related work and experts inside and outside of the cockpit. Based on this acquired knowledge, we develop a novel method that estimates human-like driving line distributions for unknown tracks. These distributions can be used to simulate complete laps with almost competitive performances and human-like driver control inputs in a professional motorsport driving simulator. Subsequently, we present a feature adaptation method that allows ProMoD to learn from the gathered experience of previous laps. We demonstrate the model's ability to continuously learn from mistakes and to improve driving performance in terms of lap completion and time. This work contributes to the modeling and a better understanding of driver behavior, paving the way for advanced full-vehicle simulations with consideration of the human driver and potentially future autonomous racing.

Due to its modular architecture, ProMoD might be extended in various ways in future research. For feature adaptation and optimization, new methods may be introduced such as generating a more human-like masking matrix. Besides that, the neural network of the *Action Selection* module could be adapted to learn from experience using reinforcement learning techniques, or real track data may be used to provide more demonstration data. In order to better understand and model the efficient and complex adaptation process of human race drivers, approaching our modeling problem from the perspective of behavioral science is worth to be explored. On top of the development of the new adaptation methods, additional performance criteria related to the human adaptation process over subsequent laps could be defined for a more holistic assessment of the adaptation methods and improvement of the model. Furthermore, human-like qualitative feedback, which is based on encountered problems during driving, could help to further support the vehicle development process. In addition, our driver model may be extended to a multiagent environment with opponents on the race track, facilitating a more accurate prediction of true racing performance and potentially optimizing full racing strategies. Finally, ProMoD might be applied to similar use cases with the target of modeling human behavior in dynamic environments with small stability margins.

Appendix

Expert Interview

Is there a universal adaption rule that applies to all drivers and tracks?

Indeed, it turns out that adaption strategies are very similar across different drivers, tracks, and vehicles, in spite of the individual driving behavior, the various layouts of the tracks and the continuously modified vehicle setups. The driver's main goal is to "brake as late as possible, and accelerate as early as possible." The resulting driving line, the turn-in, and the on-throttle behavior are seen as a consequence of pursuing that goal.

How do drivers drive their first laps on a new track?

When faced with a new track, what a driver would do can be divided into three phases: 1) preparation; 2) warm-up; and 3) subsequent fine tuning.

1) *Preparation:* Drivers come to a new track with a memorized "database of corner information," collected from their prior experience, simulator sessions, statistical data, etc. First, drivers characterize each new corner by comparing it with those in their memory and assemble a first guess of the driving line. Since every corner is unique, this first guess is usually a rough approximation. At this point, it is helpful to consult other drivers to improve the initial guess. Finally, they set brake points, utilizing signs in the environment such as brake markers. Having concretized all prior information and exchanged opinions with fellow drivers of specific positions for hitting the brake pedal, the drivers start their first laps on a new track.
2) *Warm-Up:* Race drivers are particularly talented in assessing risk. They usually start off with a slow and safe speed profile, which they adapt from lap to lap to higher velocities. This process can take very few iterations. For example, one driver managed to reach a competitive lap time on the Le Mans circuit surprisingly after only five laps.
3) *Fine Tuning:* After warming up, drivers are able to complete the lap with a close to competitive lap time, which they then try to improve incrementally. Usually, drivers do not reach a global optimum but are aware of how to improve. High- and changing-speed corners are the most difficult ones, where spinning should be prevented, as it is extremely difficult to control.

Which quantities do race drivers adapt and how? Do they pay attention to specific metrics?

Although the goal of improving lap time is sound and clear, the real optimization process is indeed very complicated, and many factors have to be taken into consideration. The following three aspects are most critical during optimization.

1) *Delta Lap Time:* The adaption behavior of race drivers is result-oriented. They are not paying much attention to the exact speed values at local points around the track, but rather to the lap time difference to the previous or best lap. The association with the optimization problem is visualized on the top of Fig. 3.
2) *Brake Point:* Hitting the brake is where the corner starts. It is the most crucial tuning knob, not only because it influences the speed profile, but also since it is the source of any issues arising throughout the following corner. I.e., all issues should be traced back to the brake point, and cannot be locally analyzed.
3) *Peak Brake Pressure:* The driver attempts to predict the future state of the car when making decisions. In the presence of slip, however, uncertainty about the vehicle state is introduced, eventually leading to wrong predictions by the driver. Therefore, slip management is crucial during cornering, with the maximum brake pressure helping to anticipate imminent slip.

How do race drivers behave when the vehicle setup is modified? Will they preadapt their strategy according to the setup?

It is extremely complicated to analyze the car and the behavior of the driver simultaneously. Therefore, when new vehicle setups are tested, the drivers do not and are not expected to have much idea of what has been adapted on the car. Sometimes, race engineers would do blind tests in order to isolate the influences of the modified setups from those of the drivers.

References

[1] W. F. Milliken and D. L. Milliken, *Race Car Vehicle Dynamics*. Warrendale, PA, USA: Soc. Autom. Eng., 1995.

[2] P. R. Wurman et al., "Outracing champion Gran Turismo drivers with deep reinforcement learning," *Nature*, vol. 602, no. 7896, pp. 223–228, 2022.

[3] S. Löckel, J. Peters, and P. van Vliet, "A probabilistic framework for imitating human race driver behavior," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 2086–2093, Apr. 2020.

[4] S. Löckel, A. Kretschi, P. van Vliet, and J. Peters, "Identification and modelling of race driving styles," *Veh. Syst. Dyn.*, vol. 60, no. 8, pp. 2890–2918, 2022.

[5] Porsche AG. "Simulator." Nov. 2019. [Online]. Available: https://presse.porsche.de/prod/presse_pag/PressResources.nsf/Content?ReadForm&languageversionid=1043793&hl=christophorus-393-misc

[6] K. Fritzsche, F. Strecker, and M. Huneke, "Simulation-based development process for Motorsport and series production," *ATZ Worldwide*, vol. 119, no. 9, pp. 60–63, 2017.

[7] B. Siegler, "Lap time simulation for racing car design," Ph.D. dissertation, Dept. Comput. Sci., Univ. Leeds, Leeds, U.K., 2002.

[8] D. P. Kelly, "Lap time simulation with transient vehicle and tyre dynamics," Ph.D. dissertation, School Eng., Cranfield Univ., Cranfield, U.K., 2008.

[9] G. Perantoni and D. J. N. Limebeer, "Optimal control for a formula one car with variable parameters," *Veh. Syst. Dyn.*, vol. 52, no. 5, pp. 653–678, 2014.

[10] U. Rosolia, X. Zhang, and F. Borrelli, "Data-driven predictive control for autonomous systems," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 1, no. 1, pp. 259–286, 2018.

[11] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 3363–3370, Nov. 2020.

[12] T. Novi, "A control system framework for autonomous vehicles at the limits of handling," Ph.D. dissertation, School Eng., Univ. Florence, Florence, Ialy, 2019.

[13] F. Fuchs, Y. Song, E. Kaufmann, D. Scaramuzza, and P. Dürr, "Superhuman performance in gran turismo sport using deep reinforcement learning," *IEEE Robot. Autom. Lett.*, vol. 6, no. 3, pp. 4257–4264, Jul. 2021.

[14] P. M. van Leeuwen, S. de Groot, R. Happee, and J. C. F. de Winter, "Differences between racing and non-racing drivers: A simulator study using eye-tracking," *PLoS ONE*, vol. 12, no. 11, Nov. 2017, Art. no. e0186871.

[15] R. Bentley, *Ultimate Speed Secrets: The Complete Guide to High-Performance and Race Driving*, Motorbooks Int., Minneapolis, MN, USA, 2011.

[16] M. Krumm, *Driving on the Edge: The Art and Science of Race Driving*. London, U.K.: Icon, 2015.

[17] J. C. Kegelman, "Learning from professional race car drivers to make automated vehicles safer," Ph.D. dissertation, Dyn. Design Lab., Stanford Univ., Stanford, CA, USA, 2018.

[18] R. A. Hess and A. Modjtahedzadeh, "A control theoretic model of driver steering behavior," *IEEE Control Syst. Mag.*, vol. 10, no. 5, pp. 3–8, Aug. 1990.

[19] L. König, *Ein virtueller Testfahrer für den querdynamischen Grenzbereich*, vol. 43. Heidelberg, Germany: expert Verlag, 2009.

[20] D. D. Salvucci, "Modeling driver behavior in a cognitive architecture," *Human Factors J. Human Factors Ergonom. Soc.*, vol. 48, no. 2, pp. 362–380, 2006.

[21] P. Cichosz and Ł. Pawełczak, "Imitation learning of car driving skills with decision trees and random forests," *Int. J. Appl. Math. Comput. Sci.*, vol. 24, no. 3, pp. 579–597, 2014.

[22] H. Wei, W. Ross, S. Varisco, P. Krief, and S. Ferrari, "Modeling of human driver Behavior via receding horizon and artificial neural network controllers," in *Proc. IEEE Conf. Decis. Control*, 2013, pp. 6778–6785.

[23] A. Kuefler, J. Morton, T. Wheeler, and M. J. Kochenderfer, "Imitating driver Behavior with generative adversarial networks," in *Proc. IEEE Intell. Veh. Symp.*, 2017, pp. 204–211.

[24] R. Bhattacharyya et al., "Modeling human driving behavior through generative adversarial imitation learning," 2020, *arXiv:2006.06412*.

[25] Z. Huang, J. Wu, and C. Lv, "Driving behavior modeling using naturalistic human driving data with inverse reinforcement learning," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 8, pp. 10239–10251, Aug. 2022.

[26] Y. Li, J. Song, and S. Ermon, "InfoGAIL: Interpretable imitation learning from visual demonstrations," in *Proc. NIPS*, 2017, pp. 3815–3825.

[27] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "End-to-end urban driving by imitating a reinforcement learning coach," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 15222–15232.

[28] Z. Deng et al., "A probabilistic model for driving-style-recognition-enabled driver steering behaviors," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 52, no. 3, pp. 1838–1851, Mar. 2022.

[29] S. Schnelle, J. Wang, H. Su, and R. Jagacinski, "A driver steering model with personalized desired path generation," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 47, no. 1, pp. 111–120, Jan. 2017.

[30] B. Shi et al., "Evaluating driving styles by normalizing driving behavior based on personalized driver modeling," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 45, no. 12, pp. 1502–1508, Dec. 2015.

[31] Y. Wang et al., "Imitation learning of hierarchical driving model: From continuous intention to continuous trajectory," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 2477–2484, Apr. 2021.

[32] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, "Probabilistic movement primitives," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 2616–2624.

[33] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, "Using probabilistic movement primitives in robotics," *Auton. Robots*, vol. 42, no. 3, pp. 529–551, 2018.

[34] M. Mulder et al., "Manual control cybernetics: State-of-the-art and current trends," *IEEE Trans. Human–Mach. Syst.*, vol. 48, no. 5, pp. 468–485, Oct. 2018.

[35] A. Ronen and N. Yair, "The adaptation period to a driving simulator," *Transp. Res. F Traffic Psychol. Behav.*, vol. 18, pp. 94–106, May 2013.

[36] R. Deborne, R. Gilles, and A. Kemeny, "Modeling driver adaptation capabilities in critical driving situations," in *Proc. SAE World Congr. Exhibit.*, 2012, pp. 1–10.

[37] F. Braghin, F. Cheli, S. Melzi, and E. Sabbioni, "Race driver model," *Comput. Struct.*, vol. 86, nos. 13–14, pp. 1503–1516, 2008.

[38] R. Bentley, *Speed Secrets: Professional Race Driving Techniques*. Minneapolis, MN, USA: MotorBooks Int., 1998.

[39] S. Quinlan and O. Khatib, "Elastic bands: Connecting path planning and control," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1993, pp. 802–807.

[40] S. K. Gehrig and F. J. Stein, "Collision avoidance for vehicle-following systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 8, no. 2, pp. 233–244, Jun. 2007.

[41] B. Siegler, A. Deakin, and D. Crolla, "Lap time simulation: Comparison of steady state, quasi-static and transient racing car cornering strategies," *SAE Trans.*, vol. 109, no. 6, pp. 2575–2581, 2000.

AQ4

Stefan Löckel received the B.Eng. degree in mechatronics from DHBW Mannheim, Mannheim, Germany, in 2015, the M.Sc. degree in electrical engineering and control from TU Kaiserslautern, Kaiserslautern, Germany, in 2017, and the Ph.D. degree in computer science from TU Darmstadt, Darmstadt, Germany, in 2022.

He is currently employed with Dr. Ing. h.c. F. Porsche AG, Weissach, Germany, where he works on machine-learning models for the development of future sports cars.

Siwei Ju received the B.Sc. degree in mechanical engineering from the Beijing Institute of Technology, Beijing, China, in 2018, and the M.Sc. degree in automation engineering from RWTH Aachen, Aachen, Germany, in 2021. She is currently pursuing the Ph.D. degree with Dr. Ing. h.c. F. Porsche AG, Weissach, Germany, and Intelligent Autonomous System, Computer Science Department, TU Darmstadt, Darmstadt, Germany. AQ5

Her research focuses on race driver modeling using reinforcement learning and imitation learning.

Maximilian Schaller received the B.Sc. degree in mechanical and process engineering from TU Darmstadt, Darmstadt, Germany, in 2019, and the M.Sc. degree in robotics, systems and control from ETH Zürich, Zürich, Switzerland, in 2022.

He is currently affiliated with the Electrical Engineering Department, Stanford University, Stanford, CA, USA. His research focuses on optimization for control and learning for time-series modeling.

Peter van Vliet received the M.Sc. degree in mechanical engineering from the Delft University of Technology, Delft, The Netherlands, in 2005, specializing in systems and control.

Since 2012, he has been with Dr. Ing. h.c. F. Porsche AG, Weissach, Germany, working on a variety of topics around high-performance road and race cars. His current interests are research activities in the field of machine learning, learning control systems and data driven, and advanced method development.

Jan Peters has been a Full Professor (W3) for Intelligent Autonomous Systems with the Computer Science Department, Technische Universität Darmstadt, Darmstadt, Germany, since 2011, and, at the same time, he has been the Head of the Research Department on Systems AI for Robot Learning, German Research Center for Artificial Intelligence (Deutsches Forschungszentrum für Künstliche Intelligenz, DFKI), Kaiserslautern, Germany, since 2022. He is also a Founding Research Faculty Member of the Hessian Center for Artificial Intelligence, Darmstadt. AQ6

Dr. Peters has received the Dick Volz Best 2007 U.S. Ph.D. Thesis Runner-Up Award, the Robotics: Science and Systems—Early Career Spotlight, the INNS Young Investigator Award, the IEEE Robotics and Automation Society's Early Career Award, as well as numerous best paper awards. In 2015, he received an ERC Starting Grant. He was an ELLIS Fellow and a Program Director in 2020 and an AAIA Fellow in 2021.

AUTHOR QUERIES

AUTHOR PLEASE ANSWER ALL QUERIES

PLEASE NOTE: We cannot accept new source files as corrections for your paper. If possible, please annotate the PDF proof we have sent you with your corrections and upload it via the Author Gateway. Alternatively, you may send us your corrections in list format. You may also upload revised graphics via the Author Gateway.

If you have not completed your electronic copyright form (ECF) and payment option please return to the Scholar One "Transfer Center." In the Transfer Center you will click on "Manuscripts with Decisions" link. You will see your article details and under the "Actions" column click "Transfer Copyright." From the ECF it will direct you to the payment portal to select your payment options and then return to ECF for copyright submission.

Carefully check the page proofs (and coordinate with all authors); additional changes or updates WILL NOT be accepted after the article is published online/print in its final form. Please check author names and affiliations, funding, as well as the overall article for any errors prior to sending in your author proof corrections.

AQ1: Please supply index terms/keywords for your paper. To download the IEEE Taxonomy, go to http://www.ieee.org/documents/taxonomy_v101.pdf.
AQ2: Please confirm or add details for any funding or financial support for the research of this article.
AQ3: Please provide the department name for Dr. Ing. h.c. F. Porsche AG and ETH Zürich.
AQ4: Current affiliations in the biographies of Stefan Löckel and Maximilian Schaller do not match the First Footnote. Please check and correct where needed.
AQ5: Please specify the field of study for the Ph.D. degree of the author S. Ju.
AQ6: Please provide the educational background for the author J. Peters.

An Adaptive Human Driver Model for Realistic Race Car Simulations

Stefan Löckel, Siwei Ju, Maximilian Schaller, Peter van Vliet, and Jan Peters

***Abstract*—Engineering a high-performance race car requires a direct consideration of the human driver using real-world tests or human-driver-in-the-loop simulations. Alternatively, offline simulations with human-like race driver models could make this vehicle development process more effective and efficient but are hard to obtain due to various challenges. With this work, we intend to provide a better understanding of race driver behavior from expert knowledge and introduce an adaptive human race driver model based on imitation learning. Using existing findings in the literature, complemented with an interview with a race engineer, we identify fundamental adaptation mechanisms and how drivers learn to optimize lap time on a new track. Subsequently, we select the most distinct adaptation mechanisms via a survey with 12 additional experts, to develop generalization and adaptation techniques for a recently presented probabilistic driver modeling approach and evaluate it using data from professional race drivers and a state-of-the-art race car simulator. We show that our framework can create realistic driving line distributions on unseen race tracks with almost human-like performance. Moreover, our driver model optimizes its driving lap by lap, correcting driving errors from previous laps while achieving faster lap times. This work contributes to a better understanding and modeling of the human driver, aiming to expedite simulation methods in the modern vehicle development process and potentially supporting automated driving and racing technologies.**

AQ1 ***Index Terms*—**

I. Introduction

THROUGHOUT more than 125 years of motorsports history, the fundamental goal of all participants did not change: reaching the best racing performance among competitors, which ultimately requires engineering a race car that fits its driver well. In fact, Milliken and Milliken already stated in 1995 that "it is the dynamic behavior of the combination of high-tech machines and infinitely complex human beings that makes the sport so intriguing for participants and spectators alike" [1]. Hence, for modern vehicle development in professional motorsports, a good understanding and modeling of the human (not necessarily lap time-optimal) driver are crucial to further improve the performance of the human-driver-vehicle-system. This objective is different from the motivation of robotic racing, where as-fast-as-possible synthetic drivers outperform human drivers [2]. However, the human decision-making process during racing is extremely complex and thus difficult to model, since:

1) many influencing factors exist;
2) vehicle dynamics are highly nonlinear and race cars are usually driven at the limits of handling, posing a difficult control task;
3) each driver exhibits an individual driving style;
4) human generalization and adaptation mechanisms are complex.

While challenges 1–3 have been successfully addressed in recent research with a framework that employs a deep neural network controller to capture these three aspects of human driving [3], [4], the problem of integrating human *adaptation* into a race driver model[1] remains unsolved. With this work, we intend to identify and better understand adaptation and learning techniques mastered by professional race drivers from related research and expert knowledge, contribute to the modeling of driver behavior by developing two methods to incorporate this behavior, and evaluate the proposed methodology within a realistic race car simulation environment as in the human-driver-in-the-loop (HDiL) simulator shown in Fig. 1.

A human-like race driver model could considerably extend and improve full vehicle simulations, ultimately enhance the resulting development efficiency and vehicle performance, while being much less expensive compared to HDiL simulations.

A. Problem Statement and Notation

In order to model human race driver behavior, we aim to learn a human-like control policy π^M which maps the current overall state $\mathbf{x}$, including vehicle state and situation on track, to the vehicle control inputs $\mathbf{a} = \begin{bmatrix} \delta & g & b \end{bmatrix}$ composed of steering wheel angle δ, throttle pedal position g and brake pedal

Manuscript received 21 July 2022; revised 3 October 2022 and 13 January 2023; accepted 30 May 2023. This work was supported by Dr. Ing. h.c. F. Porsche AG. This article was recommended by Associate Editor B.-F. Wu. *(Corresponding author: Stefan Löckel.)* AQ2

This work involved human subjects or animals in its research. The authors confirm that all human/animal subject research procedures and protocols are exempt from review board approval.

Stefan Löckel and Siwei Ju are with the Computer Science Department, Technische Universität Darmstadt, 64289 Darmstadt, Germany, and also with Dr. Ing. h.c. F. Porsche AG, 71287 Weissach, Germany (e-mail: loeckel@ias.tu-darmstadt.de). AQ3

Maximilian Schaller is with ETH Zürich, 8092 Zürich, Switzerland.

Peter van Vliet is with Dr. Ing. h.c. F. Porsche AG, 71287 Weissach, Germany.

Jan Peters is with the Computer Science Department, Technische Universität Darmstadt, 64289 Darmstadt, Germany.

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TSMC.2023.3285588.

Digital Object Identifier 10.1109/TSMC.2023.3285588

[1] A driver model represents a vehicle control policy aiming to mimic the behavior of the human race driver in order to support full vehicle simulations.

Fig. 1. Race car simulator at Porsche Motorsport [5]: Realistic visualization, a vehicle cockpit mounted on an actuated platform, and a high-fidelity vehicle dynamics model facilitate rapid testing of new vehicle configurations with the human driver in the loop. The vehicle model is developed in-house, has 14 degrees of freedom, and is validated using real-track data. It is accompanied by a high-resolution, laser-scanned track model. Details about the simulator can be found in [6]. This simulator is used to generate demonstration data from professional race drivers for our adaptive human driver model. Consequently, the simulator's vehicle model is taken to evaluate the human driver model, intending to support the future vehicle development process.

actuation b. This policy should be able to robustly maneuver a race car at the handling limits while being similar to the unknown internal driving policy π^E of human experts. At the same time, this expert policy is nondeterministic due to natural human imprecision and intentional adaptation, and able to generalize to new situations as, for example, new race tracks. In this work, we aim to approach the problem of modeling this behavior by:

1) identifying and understanding certain aspects of the most important adaptation and learning mechanisms through related work and expert interviews;
2) using these findings to considerably extend a data-based driver modeling approach;
3) evaluating the developed methods using data from professional race drivers and a state-of-the-art motorsports simulation environment.

Consequently, the resulting driver-specific control policy π^M should be able to generalize to unseen tracks and exhibit certain adaptation characteristics of the human driver. We thereby focus on the adaptation *result*, finishing laps with sufficient performance.

B. Related Work

This section discusses related work in all relevant fields, from methods to analyze or achieve optimal racing performance, to past work on the analysis, modeling, and imitation of human driver behavior, and research on the analysis of human adaptation behavior.

Optimal Racing Performance: To model the physics of a car in different driving situations, a variety of approaches with different complexity is available [1]. In classical control-based approaches, such vehicle models can be used to predict the driving behavior in standard maneuvers or to estimate the vehicle performance on a particular race track using lap time simulation approaches [7], [8], [9]. In the field of autonomous driving or racing, more recent research aims to achieve optimal performance with (data-driven) model predictive control (MPC) [10], [11], [12]. Furthermore, reinforcement learning can be used to train an agent that outperforms human drivers in simulated race environments [2], [13].

HDiL Simulation and Analysis: However, individual human driver behavior, being an important component of the vehicle-driver-entity, is often not sufficiently considered by these methods. This fact encourages motorsport teams to utilize HDiL simulation approaches, where the real driver operates the vehicle within a realistic simulation environment, facilitating faster prototyping and more realistic predictions of the true vehicle performance [6]. Furthermore, HDiL simulators enable the study of human driver behavior, for instance, perceptual and cognition skills of professional and nonprofessional race car drivers [14].

Modeling of Human Driver Behavior: Accordingly, a variety of related work describes car racing from the driver's perspective, analyzes racing techniques, driving lines, and the complex decision-making processes in greater detail, and contributes to a better understanding of the human driver in general [15], [16], [17]. Nevertheless, the task of modeling this behavior remains highly challenging. A number of approaches for building a driver model for different use cases mainly rely on conventional control architectures in partial driving scenarios [18], [19]. Using a cognitive architecture based on adaptive control, the driving behavior is modeled in a highway environment [20]. Some recently developed methods utilize imitation learning techniques to imitate human drivers: using supervised learning, random forests were trained to predict car control inputs from basic vehicle states [21] and it was shown that a feedforward neural network is able to track a driving line generated by a human [22]. Furthermore, methods based on (inverse) reinforcement learning were used to mimic drivers in highway driving scenarios [23], [24], [25], and were extended to imitate human behavior in a short-term race driving setting based on visual features [26]. By imitating a coach, reinforcement learning also enables end-to-end urban driving [27]. Besides that, research also targets specific human individuals [28], [29], [30] and hierarchical modeling [31]. These studies give insights into autonomous driving and driver modeling, but most of them are designed for urban driving and lack the ability to adapt when used for race car driving.

Probabilistic Modeling of Driver Behavior (ProMoD): Among the research on the modeling of human driver behavior, the ProMoD framework was demonstrated to be capable of completing full laps with a competitive performance by mimicking professional race drivers [3], [4]. The data-based and modular approach learns distributions of driving lines represented by probabilistic movement primitives (ProMPs) [32], [33] and trains a recurrent neural network on human race driver data in a supervised fashion. Furthermore, the driver identification and metrics ranking algorithm (DIMRA) was developed to classify individual driving styles using clustering algorithms and was later used as an evaluation method for the learned driver model [4].

Human Adaptation Behavior: Related to this topic, there seems to be a shift from linear and time-invariant models of human manual control to nonlinear and time-varying approaches that are apparent in current research trends [34]. In particular, adaptation over time is identified as a key aspect of human behavior that should and can be modeled by moving toward time-varying models. While the ProMoD framework is shown to work well in many situations, it is still lacking the functionality of a time-varying model, i.e., the ability to learn to drive on unknown tracks and to adapt and learn from gathered experience from driven laps. As such learning and adaptation aspects play fundamental roles in competitive motorsports, any robust and accurate driver modeling approach should be able to reflect them.

Human adaptation behavior w.r.t. adaption times for changing road types in a driving simulator is analyzed, yet not modeled in the work of [35]. Past research on modeling driver adaptation to sudden changes in the vehicle dynamics takes into account limb impedance modulation and updating of the driver's internal representation of the vehicle dynamics [36]. However, the latter work focuses exclusively on lateral dynamics with a first-principles approach without a superordinate objective such as lap time.

Among these approaches, ProMoD offers a solid foundation for this work, as the modeling approach is able to dynamically control a car in a race driving setting, mimicking individual driver behavior without achieving super-human performance. In this work, we considerably modify and extend ProMoD to model human driving adaptation—to the best of our knowledge, for the first time in the racing context. With the modular architecture, the driving policy adaptation remains interpretable. We considerably enhance the quality of a modern driver modeling approach, contribute to a better understanding of human race driver behavior, and aim to pave the way for more accurate vehicle simulations and, potentially, future autonomous racing.

II. Methodology

As a proper understanding of the human race driver is fundamental for modeling its learning techniques, we ground our methodology on key insights from literature, supplemented by findings of an expert interview with a professional race engineer[2] for LMP1[3] race cars. To derive modeling principles, we summarize these literature and expert insights into adaptation principles and select the most distinct of them with the help of a simple, questionnaire-based survey conducted with professional race drivers and expert motorsport engineers, as an extra layer of expertise. The adaptation principles identified in Section II-A are followed by a short summary of the recently presented ProMoD driver modeling framework in Section II-B. In Section II-C, we present a novel way to generalize the driver model to new tracks. Finally, Section II-D introduces a new method to optimize driving similar to a race driver based on previous laps.

[2]A race engineer works at the interface between the driver and the vehicle, trying to help the driver work with the vehicle and to find a vehicle setup tailored to the driver's needs.

[3]Le-Mans-Prototypes represent a top class of race cars used in different endurance racing series with races lasting up to 24 h.

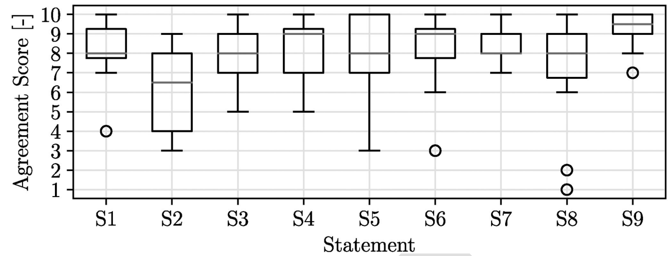

Fig. 2. Agreement levels of 12 experts with statements **S1**–**S9**. The experts were asked to choose an agreement level between 1 and 10 with step size 1. Red lines indicate the median. Boxes represent the interquartile range. The whiskers measure 1.5 times the interquartile range.

A. Adaptation Principles

Race drivers constantly pursue better racing performance in the presence of new tracks and modified vehicle setups. In this section, we aim to understand the most important principles for their adaptation behavior. We gather the following key insights from literature, extended with an expert interview[4] of a race engineer in the Appendix. We aggregate these two sources of insights into summarizing statements **S1**–**S9** detailed in the following, and finally conduct a simple, questionnaire-based survey to directly ask additional experts for their agreement with these statements. In Fig. 2, we measure the agreement of 12 additional experts (including drivers, race engineers, vehicle engineers, and tire engineers) with the 9 statements.

Objective (Delta) Lap Time: In order to (iteratively) optimize lap time [37], race drivers pay attention to the delta lap time, which is the difference between the current and the last (or best) lap time[9] (**S1**). Modifications to the vehicle setup and environmental changes are only considered a posteriori, which means that race drivers usually do not plan with them, but only react after experiencing them[9] (**S2**).

Risk Awareness: Race drivers are particularly risk-aware and constantly test for the vehicle limits [17], starting from a safe region and improving their driving incrementally[9] (**S3**).

Hierarchy: The choice of brake points heavily influences the speed profile of the entire corner [16], [38]. Subsequently, the speed profile heavily influences the driving line. Race drivers control brake points, speed profile, and driving line hierarchically, in this order[9] (which means that brake points are the main tuning knob) (**S4**).

Initialization—Driving on New Tracks: When starting on a new track, drivers tend to compare all new situations and corners to their experience from other tracks [16], [38], to get an initial guess of reasonable brake points and driving lines, which are subsequently refined[9] (**S5**). The initialization of brake points begins already before starting to drive, while the speed profile and driving line are initialized during the first few laps[9] (**S6**). After the first few laps, drivers are able to complete the lap with a close to competitive lap time[9] (**S7**).

Iteration—Adaptation Rules and Quantities: The general adaptation strategy seems to be similar for all drivers, where adaptation of the braking (brake points and peak brake pressure) is particularly important[9] (**S8**). By fine-tuning brake

[4]Findings from the expert interview are marked with this footnote. A summary of the interview is given in the Appendix.

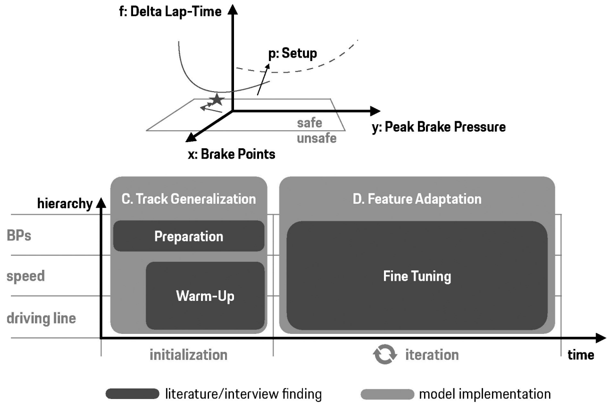

Fig. 3. Top: Iterative adaptation process visualized as an optimization problem. Bottom: Three phases of driver adaptation to solve the above optimization problem, arranged in hierarchy-time plane. Dark color denotes findings from the expert interview and related work, whereas light color signifies how the respective findings are implemented in the adaptive model.

points and peak brake pressure, drivers manage to achieve better performances[9] (**S9**).

Overall, the agreement level of the additional experts with the above statements is high. The lowest median agreement is 7 for **S2** on environmental changes viewed as a disturbance. The corresponding lower end of the interquartile range is 4, much less than 7 (or higher) for all other statements. Hence, we do *not* base our subsequent design choices on **S2**. Further, we observe outliers that might be connected to the diverse backgrounds of the 12 experts. Two tire engineers strongly disagree with **S8** on braking being particularly important for adaptation, applicative as a rule for all drivers. In contrast, both asked drivers strongly agree with this statement. Since drivers are the modeling target themselves, we decide to approve the main expert's statement that brake points are the key control variables. To summarize and simplify the problem, we set up the following qualitative model: Race drivers optimize delta lap time as a function of brake points, peak brake pressure, and other variables as visualized in Fig. 3. This function is parameterized through the vehicle setup. To solve this problem, the brake point variables are initialized in the *Preparation* phase in a safe region, i.e., such that the lap can be completed. Speed and driving line are initialized in hierarchical order during the *Warm-Up* phase. Afterward, drivers iteratively adapt and try out changes on all three hierarchical levels during *Fine-Tuning*.[5] Eventually, they arrive close to the optimizer shown as a star on the top of Fig. 3. This point usually lies close to the boundary of the safe set, as the driver will be operating the vehicle at the handling limits. As these generalization and adaptation capabilities are fundamental for professional race drivers, a driver model used for full vehicle simulations is required to have them as well. In the following, the basic ProMoD framework will be derived and subsequently extended with these skills.

B. ProMoD

The recently presented ProMoD framework combines knowledge and ideas from both race driver behavior and autonomous driving architecture. It consists of multiple modules as visualized in Fig. 4, where each of these modules represents fundamental steps in the decision-making process of a human race driver [3], [4].

[5]In the following, heuristically defined *control points* will be introduced for different vehicle states to directly adapt all three hierarchy levels.

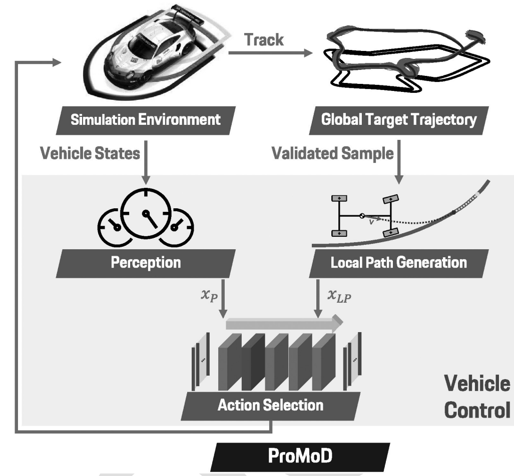

Fig. 4. Original ProMoD framework to imitate human race drivers in simulation [4]: *Global Target Trajectory* holds a distribution of potential target driving lines, relating to the driver's mental image of a driving corridor. *Local Path Generation* and *Perception* calculate a feature vector based on the current situation on track and a sampled target driving line. *Action Selection* maps the features to driver actions. Feeding back the predicted actions to the simulation environment closes the loop.

Our novel generalization and adaptation methods are based on this architecture, which is summarized in the following.

Global Target Trajectory: Every driver keeps a mental image of the whole race track in their head, knowing approximately where to brake, to turn in, and to accelerate again in each corner. However, this imagined driving corridor is not precise, i.e., it incorporates variance, and additionally changes over time with gathered experience. Hence, we model the global target trajectory with a distribution over potential driving lines, which could be interpreted as a driving corridor, using ProMPs [32], [33]. For this purpose, both the spatial and the temporal information of every demonstrated driving line on a particular track is projected to a lower-dimensional weight space. We define a series of equally distributed radial basis functions (RBFs)

$$b_j(s) = \exp\left(-\frac{\left(s - c_j\right)^2}{2h}\right) \tag{1}$$

with function index $j \in \{1, 2, \ldots, N_{\mathrm{BF}}\}$, track distance s, constant width h, and c_j being the equally distributed centers of the functions. All basis functions are assembled into the basis function matrix $\boldsymbol{\Phi}_s \in \mathbb{R}^{N_s \times N_{\mathrm{BF}}}$, where the jth column contains $b_j(s)$ evaluated at N_s points, equidistant in terms of track distance. Subsequently, $\boldsymbol{\Phi}_s$ is aggregated into

$$\boldsymbol{\Psi}_s = \mathrm{diag}(\boldsymbol{\Phi}_s, \boldsymbol{\Phi}_s, \ldots, \boldsymbol{\Phi}_s) \in \mathbb{R}^{nN_s \times nN_{\mathrm{BF}}} \tag{2}$$

for n variables that the trajectory consists of. The weight vector

$$\boldsymbol{w}_i = \left(\boldsymbol{\Psi}_s^T \boldsymbol{\Psi}_s + \epsilon \boldsymbol{I}\right)^{-1} \boldsymbol{\Psi}_s^T \boldsymbol{\tau}_{s,i} \in \mathbb{R}^{nN_{\mathrm{BF}}} \tag{3}$$

is derived using ridge regression for each demonstration trajectory $\boldsymbol{\tau}_{s,i} \in \mathbb{R}^{nN_s}$ and regularization factor ϵ. By fitting a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{w}}, \boldsymbol{\Sigma}_{\boldsymbol{w}})$ over the N demonstration weights with mean $\boldsymbol{\mu}_{\boldsymbol{w}}$ and variance $\boldsymbol{\Sigma}_{\boldsymbol{w}}$

$$\boldsymbol{\mu}_{\boldsymbol{w}} = \frac{1}{N}\sum_{i=1}^{N} \boldsymbol{w}_i \in \mathbb{R}^{nN_{\mathrm{BF}}}, \tag{4}$$

$$\boldsymbol{\Sigma}_{\boldsymbol{w}} = \frac{1}{N}\sum_{i=1}^{N} \big(\boldsymbol{w}_i - \boldsymbol{\mu}_{\boldsymbol{w}}\big)\big(\boldsymbol{w}_i - \boldsymbol{\mu}_{\boldsymbol{w}}\big)^T \in \mathbb{R}^{nN_{\mathrm{BF}} \times nN_{\mathrm{BF}}} \tag{5}$$

we are able to describe the distribution of driving lines for a driver on a particular track efficiently. Subsequently, an arbitrary number of new driving lines which are similar to all demonstrations can be generated by sampling a weight vector from this distribution, $\boldsymbol{w}^* \sim \mathcal{N}\big(\boldsymbol{\mu}_{\boldsymbol{w}}, \boldsymbol{\Sigma}_{\boldsymbol{w}}\big)$, and using

$$\boldsymbol{\tau}^* = \boldsymbol{\Psi}_s \boldsymbol{w}^* \tag{6}$$

to retrieve a new driving line in the original formulation which could be subsequently used as *target trajectory*. While sampling this target trajectory all at once models the human driver's ahead-of-time plan based on experience and knowledge of the whole track, real-time planning based on the current state on the track is performed by ProMoD's *Local Path Generation* module.

Local Path Generation: For any situation on track, a human driver continuously plans the upcoming path a few seconds ahead. We use this module to mimic the path planning by calculating constrained polynomials and multiple preview features[6] based on the current vehicle state and the *target trajectory*. These local path features are denoted as $\mathbf{x}_{\mathrm{LP}}$.

Perception: In addition to the path planning features, each driver relies on additional information about their surroundings, such as visual information or experienced accelerations. These perception features, which mostly relate to basic vehicle states, are gathered inside this module and are denoted as $\mathbf{x}_P$.

Action Selection: The action selection process, i.e., the mapping from the current state (as described by the feature vector $\mathbf{x} = [\mathbf{x}_{\mathrm{LP}}\ \mathbf{x}_P]$) to human-like control actions $\mathbf{a}$, is learned using a recurrent neural network. It is trained on all available demonstration data for a particular driver, aiming to imitate its individual driving style and incorporating the dynamics of the action selection process.

This modular and hierarchical structure, compared with end-to-end learning such as in [13], increases interpretability when tuning the driving behavior. After the recurrent neural network, i.e., the action selector, is trained, it serves as a controller that drives the car by following the global reference trajectory. Subsequently, by modifying the global reference trajectory, the driver model can be adjusted for performance or generalization. Compared with the direct adaptation of the action selection policy (parameters of the recurrent neural network), the adaptation of the global reference trajectory has the following advantages: 1) fewer parameters to update; 2) an interpretable adaptation process; and 3) predictable and understandable results. In the following, we present methods to generalize and adapt this driver model in two different phases. Section II-C introduces *Track Generalization*, addressing the *Preparation* and *Warm-Up* steps identified from the literature and interview (see Fig. 3). Section II-D describes *Feature Adaptation*, modeling the iterative *Fine-Tuning*.

[6]Examples are a predicted lateral offset or a predicted speed difference from the target driving line. More details are given in [4].

Algorithm 1 Estimating a Driving Line Distribution + Sampling

$\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa,dy}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa,dy} \leftarrow$ BuildProMP($\mathcal{D}$)
$x'(s), y'(s), \kappa'(s) \leftarrow$ BuildDrivingLine($\mathcal{B}_{\mathrm{left}}, \mathcal{B}_{\mathrm{right}}$)
$\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa'} \leftarrow$ RidgeRegression($\kappa'(s)$)
$\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'} \leftarrow \mathbf{0}$
$\boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'} \leftarrow$ EstimateVariance($\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa'}, \boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa,dy}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa,dy}$)

for $i \leftarrow 1, N_{\mathrm{samples}}$ **do**
  $\boldsymbol{w}_i^{*dy} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'}\right)$
  $x_i^*(s), y_i^*(s) \leftarrow$ Reconstruct($x'(s), y'(s), \boldsymbol{w}_i^{*dy}$)
  $\Delta t_i^*(s) \leftarrow$ EstimateSpeed($x_i^*(s), y_i^*(s), \mathcal{P}$)
end for

C. Track Generalization: Generate Driving Line Distributions

In order to generate first laps on a new, yet unknown track, it is required to learn a reasonable driving line distribution for the *Global Target Trajectory* module. All other modules of ProMoD are track-independent by definition and remain unmodified. Hence, we construct a driving line distribution for a new track based on its borders (assumed to be known) and prior knowledge from other tracks. Inspired by the results from Section II-A, we propose the methodology described in Algorithm 1. We utilize a novel ProMP description, conventional methods to fit driving lines based on geometric boundaries, and a method to estimate the variance of the driving line around the track based on experience from other tracks.

ProMPs on Demonstration Data: To encode prior knowledge from other tracks, we use all available driving line data from all known tracks $\mathcal{D}$ and calculate ProMPs with a modified representation as driving line distributions for each track separately. In particular, we take the vehicle positions in the Cartesian space for all laps on a given track and map them to a curvilinear description $x(s),\ y(s) \mapsto dy(s),\ \kappa(s)$ for each track. Thereby, dy represents the lateral deviation from a reference line and κ the line curvature, both based on the reference line distance s. While there is an overlap between the information in dy and κ, both representations are needed for subsequent calculations. Similar to the computation of RBF weights via ridge regression in the Cartesian space, driving lines are now represented by weight vectors $\boldsymbol{w}^{dy}$ and $\boldsymbol{w}^{\kappa}$ for dy, κ, and RBFs in the curvilinear space with equidistant discretization. Assuming a Gaussian distribution, we retrieve mean weight vectors $\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa}$, $\boldsymbol{\mu}_{\boldsymbol{w}}^{dy}$ and variances $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa}$, $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy}$ to describe the distribution of all available driving lines on a particular track. By iterating this process for all available tracks, we can aggregate all driving line information into $\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa,dy}$, $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa,dy}$. In the following, we estimate a driving line distribution for an unknown track by combining this stochastic information with a conventional path planning method.

Generate Mean Driving Trajectory: We start by estimating a mean driving trajectory which is only based on the given track boundaries $\mathcal{B}_{\text{left}}$ and $\mathcal{B}_{\text{right}}$. As the generation of a reasonable and collision-free path around the track is required, we decide to use Elastic Bands [39], [40]. While being computationally efficient and easy to interpret, this method exhibited reasonable driving line estimates with sufficient accuracy. The resulting trajectory is now taken as the reference and mean driving line for the new track. Similarly to the ProMP calculation on the available demonstration data, the curvature $\kappa'(s)$ of this Elastic Band driving line is projected to the lower-dimensional weight space and set as the mean curvature $\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa'}$ with $\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'} = \mathbf{0}$ by definition.

Variance Estimation: Using this mean trajectory and the existing corner information from other tracks, we estimate the variance with a sliding window approach. For this purpose, we are moving along the estimated mean driving line's curvature $\kappa'(s)$ and compare the current situation, described by a sequence of curvatures, to all situations on all known tracks as encoded in $\boldsymbol{\mu}_{\boldsymbol{w}}^{\kappa,dy}$, $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{\kappa,dy}$. By finding the most similar corner measured by the absolute difference between curvatures, we are now able to iteratively build $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'}$, which describes the variance of driving lines on the new track.[7]

Sampling and Reconstruction: Using the Elastic Band estimate $x'(s)$, $y'(s)$ for the mean driving line and the modified ProMP $\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'} = \mathbf{0}$, $\boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'}$ describing the lateral deviation from the mean, we are now able to sample new driving lines for the new track. In particular, we draw a sample weight vector $\boldsymbol{w}_i^{*dy} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\boldsymbol{w}}^{dy'}, \boldsymbol{\Sigma}_{\boldsymbol{w}}^{dy'}\right)$ and retrieve the lateral deviation $dy_i^*(s)$ as $\boldsymbol{\Phi}_s \boldsymbol{w}_i^{*dy}$. Now, it is possible to construct a sample driving trajectory in the Cartesian space using

$$x_i^*(s) = x'(s) - \sin\bigl(\phi_i^*(s)\bigr) dy_i^*(s) \tag{7}$$
$$y_i^*(s) = y'(s) + \cos\bigl(\phi_i^*(s)\bigr) dy_i^*(s) \tag{8}$$

where ϕ_i^* is the mean heading angle of the vehicle and equals 0 when the vehicle drives purely into x-direction.

Speed Profile: In addition to the trajectory of the vehicle, ProMoD requires a speed profile for the *Local Path Generation* module. Since this velocity profile depends on the vehicle and its setup and is hard to estimate using the available demonstration data, we follow a more robust approach based on vehicle dynamics. For each sampled vehicle trajectory $x_i^*(s)$, $y_i^*(s)$, we utilize a conventional lap time estimation approach based on the vehicle performance envelope $\mathcal{P}$ to retrieve an approximate speed profile [7], [41].

Simulation: The sampled driving lines with corresponding speed profiles can now be used to reconstruct the original ProMP formulation within the previously presented ProMoD framework. Initializing with a reduced performance envelope $\mathcal{P}$ represents the *Preparation* phase on a new track and allows for safely simulating first laps. By iteratively expanding $\mathcal{P}$ and simulating the resulting driving lines and speed profiles, ProMoD is able to cautiously approach the vehicle limitations, aiming to mimic the *Warm-Up* phase. The complete process facilitates simulations on new tracks for which no demonstration data exists, enhancing our driver modeling framework with track familiarization abilities to generate first fast laps. After becoming familiar with a track, human drivers continuously optimize their performance, as shown in Section II-A. Hence, ProMoD should also be adaptable and learn from experience, which necessitates adaptation techniques.

[7]We use the curvature κ to find similar corners since it naturally describes the corner shape. The lateral deviation dy is used for sampling, as it allows for a more robust reconstruction.

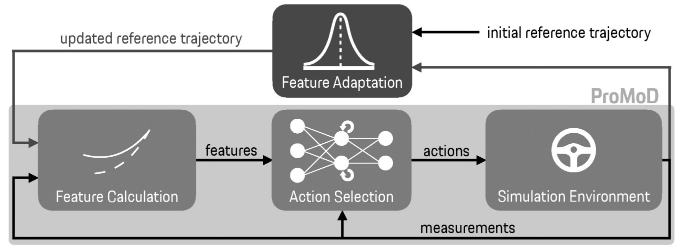

Fig. 5. Feature adaptation (red) extending the original ProMoD framework (gray), consisting of feature calculation (summarizes *Local Path Generation* and *Perception*), *Action Selection*, and the simulation environment. Every finished lap is analyzed and the reference trajectory is adapted correspondingly.

D. Feature Adaptation

Professional race drivers master the skill of continuously optimizing their performance by analyzing past laps and adapting accordingly. With an additional feedback loop as shown in Fig. 5, ProMoD is enabled to mimic this learning process to a certain extent. By only adapting the global target trajectory, which is used to compute local path planning features $\mathbf{x}_{\text{LP}}$, the behavior of ProMoD can be influenced. At the same time, ProMoD maintains its ability to imitate human drivers as the action selection module remains unchanged. In the following, we use *Conditioning* and *Scaling* to modify the global target trajectory while keeping it human-like:

Conditioning: Recall that the ProMPs for the global target trajectory are represented by a Gaussian weight distribution $p(\boldsymbol{w}) = \mathcal{N}\bigl(\boldsymbol{w} \mid \boldsymbol{\mu}_{\boldsymbol{w}}, \boldsymbol{\Sigma}_{\boldsymbol{w}}\bigr)$ with mean weight vector $\boldsymbol{\mu}_{\boldsymbol{w}}$ and covariance matrix $\boldsymbol{\Sigma}_{\boldsymbol{w}}$. We are now able to alter this distribution by conditioning the prior distribution to a new (algorithmically chosen) observation $\boldsymbol{x}_{s'}^* = \{\boldsymbol{y}_{s'}^*, \boldsymbol{\Sigma}_y^*\}$ at a specific location $s = s'$, as presented in [33]. Here, the control point $\boldsymbol{y}_{s'}^* \in \mathbb{R}^n$ is an algorithmically chosen target state (see Paragraph *Adaptation Process* for details) of the vehicle position and velocity to be reached at distance s', and variance $\boldsymbol{\Sigma}_y^* \in \mathbb{R}^{n \times n}$ is the confidence of this choice. The conditional distribution $p\bigl(\boldsymbol{w} \mid \boldsymbol{x}_{s'}^*\bigr)$ remains Gaussian with updated parameters

$$\boldsymbol{\mu}_{\boldsymbol{w}}^{[\text{new}]} = \boldsymbol{\mu}_{\boldsymbol{w}} + \boldsymbol{L}\bigl(\boldsymbol{y}_{s'}^* - \boldsymbol{\Psi}_s^T \boldsymbol{\mu}_{\boldsymbol{w}}\bigr), \tag{9}$$
$$\boldsymbol{\Sigma}_{\boldsymbol{w}}^{[\text{new}]} = \boldsymbol{\Sigma}_{\boldsymbol{w}} - \boldsymbol{L} \boldsymbol{\Psi}_{s'}^T \boldsymbol{\Sigma}_{\boldsymbol{w}} \tag{10}$$

where

$$\boldsymbol{L} = \boldsymbol{\Sigma}_{\boldsymbol{w}} \boldsymbol{\Psi}_{s'} \Bigl(\boldsymbol{\Sigma}_y^* + \boldsymbol{\Psi}_{s'}^T \boldsymbol{\Sigma}_{\boldsymbol{w}} \boldsymbol{\Psi}_{s'}\Bigr)^{-1} \tag{11}$$

relates the variances of the prior distribution and the new observation with $\boldsymbol{\Psi}_{s'} \in \mathbb{R}^{nN_{\text{BF}} \times n}$ representing the value of all basis functions at $s = s'$ [33].

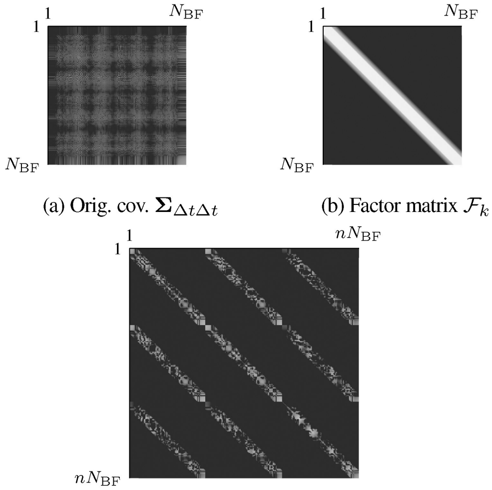

(a) Orig. cov. $\mathbf{\Sigma}_{\Delta t \Delta t}$ (b) Factor matrix $\mathcal{F}_k$

(c) Pre-processed cov. $\mathbf{\Sigma}_{\boldsymbol{w}}^{\text{masked}}$

Fig. 6. Masking the covariance matrix. (a) Part of the covariance matrix for a single variable ($\mathbf{\Sigma}_{\Delta t \Delta t} \in \mathbb{R}^{N_{\mathrm{BF}} \times N_{\mathrm{BF}}}$), where brighter colors indicate higher covariances. Far-off-diagonal correlations in the data potentially result from different vehicle setups in the demonstration data but are difficult to consider during conditioning. (b) Factor matrix for a single variable, where the elements on the diagonal are one, and off-diagonal entries are fading out to zero using bandwidth k. Here, k is selected such that distant and nonconsecutive turns cannot mutually influence each other. (c) Resulting matrix $\mathbf{\Sigma}_{\boldsymbol{w}}^{\text{masked}} \in \mathbb{R}^{nN_{\mathrm{BF}} \times nN_{\mathrm{BF}}}$ for three variables after masking, filtering out correlations over larger distances.

This procedure allows to move brake points or to shift apexes[8] by conditioning the prior distribution utilizing a set of rules derived from Section II-A. In the meantime, the correlations between different locations are taken into consideration by the covariance matrix which is learned from the data so that the whole trajectory is modified correspondingly. However, when using the prior variance without further consideration, conditioning at a specific turn potentially affects distant turns due to nonzero covariances in the data, as shown for $\mathbf{\Sigma}_{\Delta t \Delta t}$ in Fig. 6(a). As such a large effect across multiple turns is not considered to be human-like, we aim to reduce it by masking the original matrix using a factor matrix $\mathcal{F}_k \in \mathbb{R}^{N_{\mathrm{BF}} \times N_{\mathrm{BF}}}$ shown in Fig. 6(b). By multiplying $\mathcal{F}_k$ element-wise with each submatrix of $\mathbf{\Sigma}_{\boldsymbol{w}}$, we retrieve a masked matrix for conditioning

$$\mathbf{\Sigma}_{\boldsymbol{w}}^{\text{masked}} = \begin{bmatrix} \mathcal{F}_k \circ \mathbf{\Sigma}_{xx} & \mathcal{F}_k \circ \mathbf{\Sigma}_{xy} & \mathcal{F}_k \circ \mathbf{\Sigma}_{x\Delta t} \\ \mathcal{F}_k \circ \mathbf{\Sigma}_{yx} & \mathcal{F}_k \circ \mathbf{\Sigma}_{yy} & \mathcal{F}_k \circ \mathbf{\Sigma}_{y\Delta t} \\ \mathcal{F}_k \circ \mathbf{\Sigma}_{\Delta tx} & \mathcal{F}_k \circ \mathbf{\Sigma}_{\Delta ty} & \mathcal{F}_k \circ \mathbf{\Sigma}_{\Delta t\Delta t} \end{bmatrix} \tag{12}$$

which effectively lowers the influence of conditioning on distant regions as shown in Fig. 6(c).[9] This matrix can then replace $\mathbf{\Sigma}_{\boldsymbol{w}}$ for effective local *Conditioning*.

Scaling: In order to fully utilize the vehicle's potential on straights, the speed profile can be adapted to influence the throttle actuation and braking behavior of ProMoD. Since the neural network performs trajectory tracking, aiming to minimize the control error between the reference speed and the actual speed, its output signals tend to fluctuate during intervals of full throttle. Therefore, if the actual velocity is larger than the reference velocity, ProMoD tends to accelerate less, even if the virtual driver is on a straight and expected to drive as fast as possible. This problem can be effectively solved by smoothly scaling the reference speed on long straights.

Algorithm 2 Adaptation Process

Input: $\boldsymbol{\mu}_{\boldsymbol{w}}^0$, $\hat{\mathbf{\Sigma}}_{\boldsymbol{w}}^0$, envelope
$\mathbf{\Sigma}_{\boldsymbol{w}}^0 \leftarrow$ PROCESSVARIANCE($\hat{\mathbf{\Sigma}}_{\boldsymbol{w}}^0$)
$\boldsymbol{\tau}_0^{ref} \leftarrow$ CALCMEANTRAJECTORY($\boldsymbol{\mu}_{\boldsymbol{w}}^0$)
$\boldsymbol{\mathcal{I}}^{track} \leftarrow$ ANALYSETRACK($\boldsymbol{\tau}_0^{ref}$)
for $i = 0, 1, 2, \ldots$ **do**
  $\boldsymbol{\tau}_i \leftarrow$ SIMULATE($\boldsymbol{\tau}_i^{ref}$)
  $\boldsymbol{y}_s^* = \emptyset$
  if not ISCOMPLETED($\boldsymbol{\tau}_i$) **then**
    $\boldsymbol{y}_s^* \leftarrow \boldsymbol{y}_s^* \cup$ ANALYSEDL($\boldsymbol{\tau}_i, \boldsymbol{\mathcal{I}}^{track}$, envelope)
    if $\boldsymbol{y}_s^* == \emptyset$ or SLIPCHECK($\boldsymbol{\tau}_i, \boldsymbol{\mathcal{I}}^{track}$) **then**
      $\boldsymbol{y}_s^* \leftarrow \boldsymbol{y}_s^* \cup$ ADAPTSPEED($\boldsymbol{\tau}_i, \boldsymbol{\mathcal{I}}^{track}$)
    end if
  else
    $\boldsymbol{y}_s^* \leftarrow \boldsymbol{y}_s^* \cup$ CHECKINENVELOPE($\boldsymbol{\tau}_i$, envelope)
  end if
  $\boldsymbol{\mu}_{\boldsymbol{w}}^{i+1,0}, \mathbf{\Sigma}_{\boldsymbol{w}}^{i+1,0} = \boldsymbol{\mu}_{\boldsymbol{w}}^{i+1}, \mathbf{\Sigma}_{\boldsymbol{w}}^{i+1}$
  for $j = 1, 2, \ldots$, number of items in $\boldsymbol{y}_s^*$ **do**
    $\boldsymbol{\mu}_{\boldsymbol{w}}^{i+1,j}, \mathbf{\Sigma}_{\boldsymbol{w}}^{i+1,j} \leftarrow$ COND($\boldsymbol{\mu}_{\boldsymbol{w}}^{i+1,j-1}, \mathbf{\Sigma}_{\boldsymbol{w}}^{i+1,j-1}, \boldsymbol{y}_{s,j}^*$)
  end for

  $\hat{\boldsymbol{\tau}}_{i+1}^{ref} \leftarrow$ CALCMEANTRAJECTORY($\boldsymbol{\mu}_{\boldsymbol{w}}^{i+1}$)
  if ISCOMPLETED($\boldsymbol{\tau}_i$) **then**
    $\boldsymbol{\tau}_{i+1}^{ref} \leftarrow$ SPEEDSCALING($\hat{\boldsymbol{\tau}}_{i+1}^{ref}, \boldsymbol{\tau}_i, \boldsymbol{\mathcal{I}}^{track}$)
  else
    $\boldsymbol{\tau}_{i+1}^{ref} \leftarrow \hat{\boldsymbol{\tau}}_{i+1}^{ref}$
  end if
end for

Adaptation Process: The complete adaptation process, shown in Algorithm 2, is inspired by the insights from Section II-A and uses both introduced methods, *Conditioning* and *Scaling*, to continuously adapt ProMoD based on gathered experience. After simulating a lap, an initial check is done whether the lap was completed successfully. If this is not the case, the situation where the vehicle left the track is analyzed and the ProMP is conditioned using two subprocedures.

1) *Driving-Line Check and Adaptation:* As seen in Section II-A, the turn-in is the most important phase during cornering. Hence, the driving line is compared to the permissible driving corridor, represented by track borders or by the envelope of all demonstrations from the human drivers, and the largest deviation before the apex is found. Then, a new control point $\boldsymbol{y}_{s'}^*$ is added for *Conditioning* at this position, shifting the driving line distribution toward the permissible area.
2) *Velocity Adaptation:* If no valid adaptation is found or extreme tire slip occurs, a control point will be added to reduce the target speed shortly before the track was left.

In practice, ProMoD can eventually complete each critical corner when the target speed is low enough. Subsequently, the completed laps can be further adapted to improve the lap time and to keep the driving line in the envelope by:

1) Checking and reducing smaller deviations from the permissible driving corridor: Just like during real racing, ProMoD sometimes slightly exceeds the theoretically allowed driving corridor but still manages to complete

[8] An apex is defined as the closest point to the inner side of a corner, typically coinciding with the locally maximal curvature of the driving line.

[9] While the assumption of a fixed bandwidth k is not entirely human-like, it turned out to be sufficient to introduce the required adaptation characteristics. Future work may focus on finding a variable, distance-dependent masking to further enhance human likeness.

the lap. These situations are checked and additional control points are introduced for *Conditioning*.

2) Checking acceleration intervals and *Scaling* of the speed: As discussed before, sometimes ProMoD does not utilize the full vehicle potential during acceleration phases on straight lines. Hence, speed scaling is used to further increase the performance on already completed laps.

By introducing this process, we are able to encourage ProMoD to learn from the experience of previous laps, to correct mistakes, and to increase performance, matching the requirements illustrated in Fig. 3.

III. Evaluation

In this work, we use data of professional race drivers gathered from the HDiL simulator shown in Fig. 1 to train and evaluate our driver model. All rollouts of our driver model are simulated using the same in-house developed vehicle model of a high-performance race car, guaranteeing realistic vehicle dynamics and facilitating comparability to the human demonstrations. The task of driving the simulated race car is highly challenging as its only driver assistance system is *Traction Control*. In order to safeguard intellectual property, all plots in this article are shown normalized.

A. Track Generalization

We evaluate the presented track generalization method of our ProMoD framework on two race tracks, Motorland Aragón (AGN) and the Yas Marina Circuit in Abu Dhabi (ABD), and exclude demonstration data from these tracks during training. For each track, we initially estimate driving line distributions according to the methodology presented in Section II-C and draw N_{samples} driving lines from these distributions. When using these driving line samples for simulation on the corresponding unknown tracks, ProMoD is capable of completing full laps on the respective race track, as visualized in Fig. 7 for ABD.

For AGN, the track generalization method achieves comparable results considering the similarities of the resulting driving line and driver action distributions with the human driver. Furthermore, we compare the performance of ProMoD and the human driver on both tracks with equal vehicle setups. Fig. 8 visualizes the resulting lap time distributions, normalized to the median lap time of the human driver on each track, respectively.

Here, ProMoD is able to achieve lap times close to those of the human driver, with a slightly increased median due to small deviations in the expected speed profiles as visible in Fig. 7(a) between reference distances 0.1 and 0.2. These deviations result from the herein utilized conventional lap simulation approach [7], [41] that marginally underestimates the available acceleration potential and hence permissible speed of the vehicle in dynamic situations. This is a reasonable limitation, as the track generalization method is mainly intended to safely finish first laps on a new track with a close to competitive performance. In contrast to baseline machine learning models and also to conventional lap simulation approaches [7], [41] that rely on simplified vehicle models and do not consider human characteristics, extensive evaluations of the basic ProMoD framework in earlier research [3], [4] already demonstrated that the framework can robustly mimic human driving styles in a variety of settings. These findings are underlined by an extended evaluation of the adapted ProMoD model in the following section.

B. Feature Adaptation

The feature adaptation process is tested on two different tracks, the Silverstone Circuit (SVT) and Motorland Aragón (AGN), as these tracks turned out to be particularly difficult to finish for the driver model and, hence, are a suitable environment to demonstrate the applicability of our method. We start with an evaluation of the local effects of *Conditioning* and *Scaling* by showing the executed adaptations, the resulting changes in terms of driving line, and the selected actions of the driver model. Subsequently, we test the complete adaptation process on both tracks, showing that the method is able to pass previously unfinished turns and to improve lap time.

Local Effect—Adaptation: The local effects of adaptation are presented in Figs. 9 and 10, visualizing adaptations of the driving line and the speed profile, as well as the resulting action signals and driven lines. Here, ProMoD fails initially at Turn (T) 6/7 of SVT due to considerably exceeding the vehicle potential as shown in Fig. 9(b). In order to adapt the speed profile effectively, three control points are used to set the lower peak speed value, resulting in earlier braking and consequently helping to avoid the mistake and pass the turn. At the same time, with the purpose of reducing the curvature and avoiding corner-cutting, the driving line is pulled outwards around fifty meters before the first apex as shown in Fig. 9(a). After two iterations of simultaneously adapting both the speed profile and the driving line, ProMoD succeeds in this turn. Note that such intermediate iterations are part of our modeling *algorithm* and not part of the adaptation model itself, that is resembled by the final iterate of speed profile and driving line.

Local Effect—Scaling: Scaling is particularly useful on straights if ProMoD initially does not fully utilize the vehicle potential due to a modified vehicle setup and a too conservative prior target speed definition. Its effect becomes apparent when observing the throttle actuation signal. With a higher reference speed, the model tends to utilize full throttle more often on long straights, as shown in Fig. 11. Consequently, the fluctuations of the throttle signal in those intervals are eliminated, and the lap time is improved by about 0.2 s.

Adaptation Process: The developed adaptation process for ProMoD has been successfully tested on SVT and AGN as visualized in Fig. 12. While it requires four iterations to complete SVT, ProMoD needs more iterations for AGN since it fails at more locations. On both tracks, the learning speed is slower compared to a race driver, but ProMoD ultimately succeeds in completing a lap after less than 20 iterations, with at most five iterations for a problematic turn. To indicate the adaptation progress, the lap progress and the portion of the

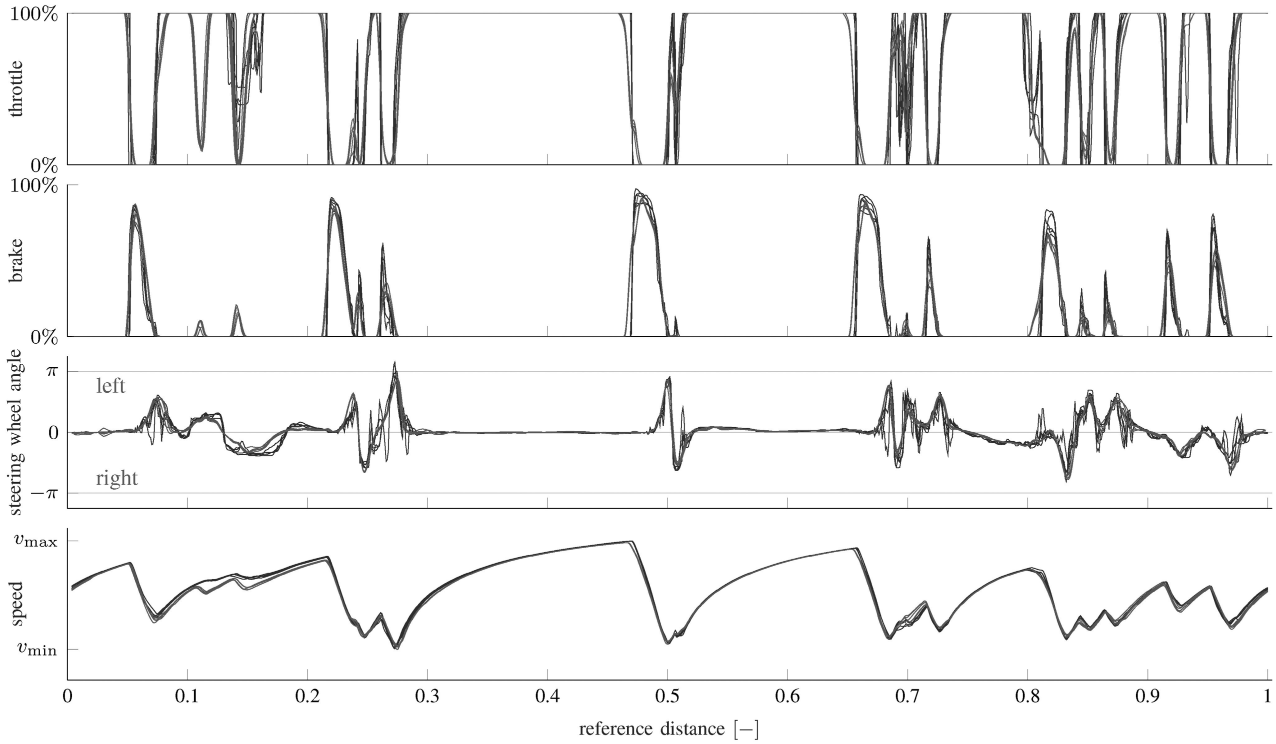

(a)

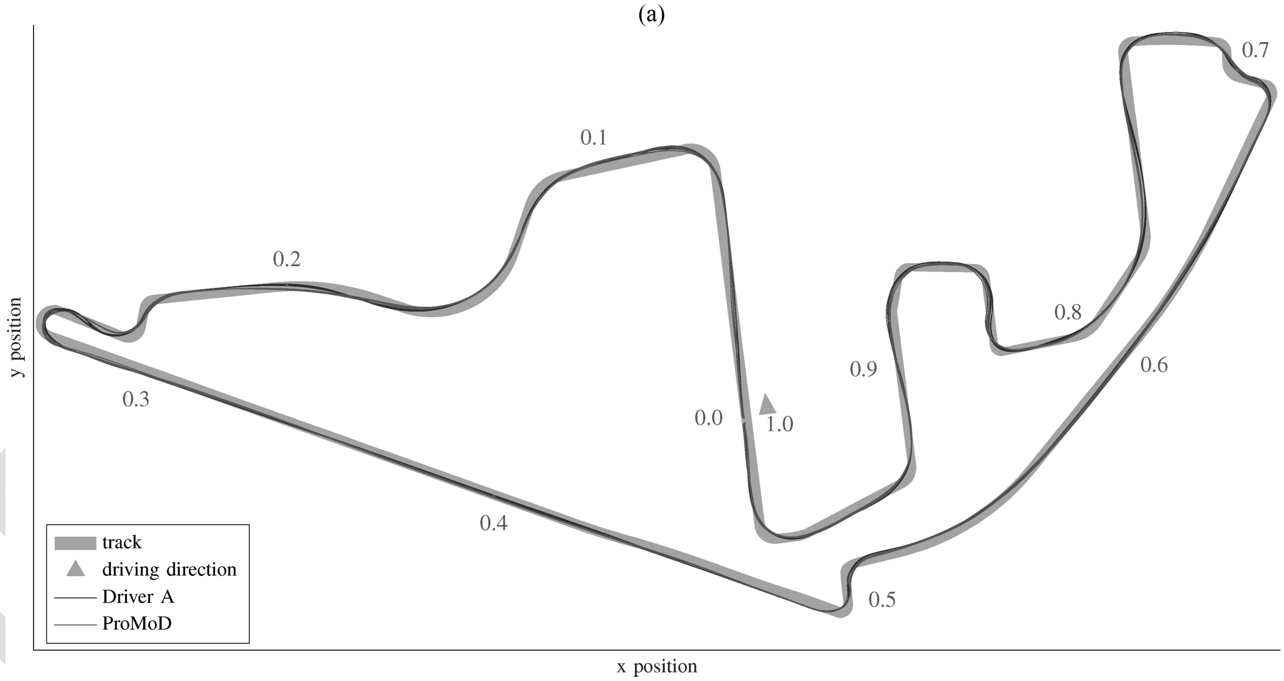

(b)

Fig. 7. Track Generalization results on ABD: We compare five laps of the human driver (dark gray) to five laps of the track-generalized ProMoD framework (red) with an identical vehicle setup. (a) Comparison of the driver actions and the resulting speed profiles over the normalized track reference distance. Here, ProMoD is able to approximately reproduce the throttle, braking, and steering activity of the real driver considering the braking points, actuation speeds, and amplitudes. The velocity profile shows small deviations after the first corner where ProMoD does not fully utilize the vehicle potential due to a slightly over-conservative speed profile estimation in this region. (b) Resulting simulated driving lines around the track (light gray) where numbers indicate the reference distance. The position of the start/finish line and the driving direction is indicated by the bright blue triangle. Here, ProMoD is able to generalize and approximately follows the demonstrations of the human driver even though they were not used during training for this race track. Some deviations are present at particularly challenging locations (e.g., the hairpin corner on the left), which, however, do not prevent ProMoD from finishing the lap with reasonable performance. These deviations may be reduced by using adaptation methods to learn from the gathered experience on the track.

lap with full throttle are plotted over the number of iterations, corresponding to the objective of finishing laps and optimizing the lap time, respectively, while imitating the human drivers.

DIMRA: Finally, we use DIMRA to evaluate the adapted model regarding the similarity of its driving style to that of the target human driver [4]. In Fig. 13, each marker

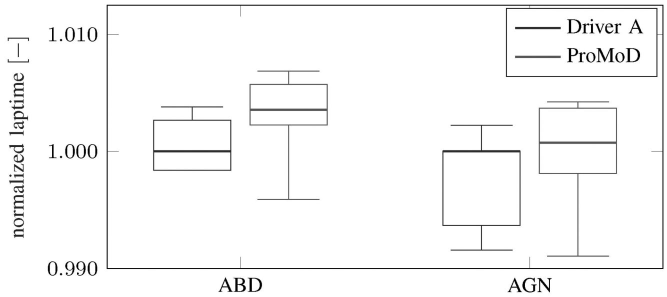

Fig. 8. Lap time comparison for track generalization on race tracks ABD and AGN: Times are normalized to the median demonstration lap time of the corresponding track. The whiskers correspond to the minimum/maximum values, the boxes indicate the upper/lower quartiles, and the thick central line shows the median value. Here, ProMoD is able to finish laps on unknown race tracks, less than 0.5% slower than the human driver in the median and at a competitive pace for its fastest laps. The slightly slower median lap time might be a result of a yet nonoptimal speed profile or driving line distribution.

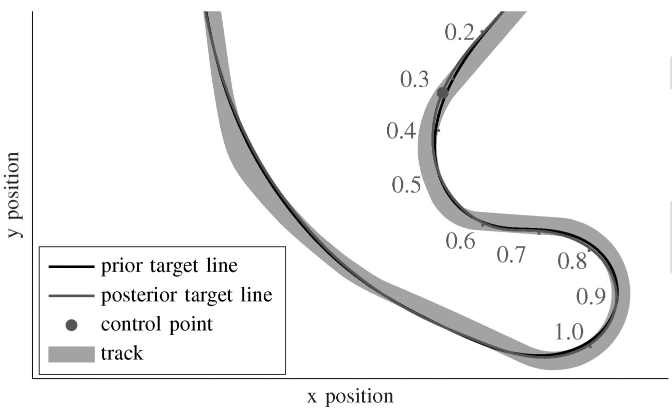

(a)

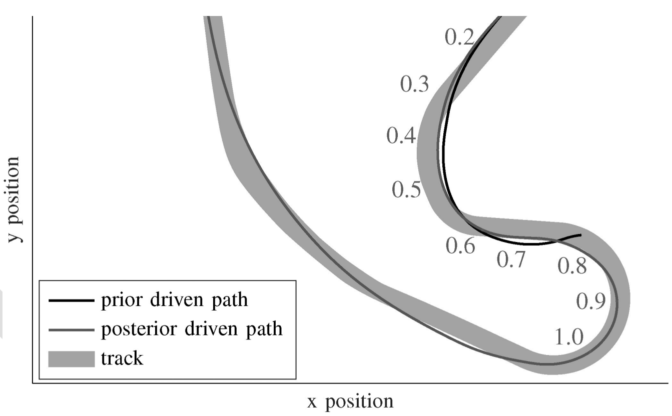

(b)

Fig. 9. Adaptation of the target line for T6/7 on SVT and the resulting driven paths. (a) Prior (black) and posterior (red) target lines. The posterior target line is pulled outwards before the first apex using a control point at corner entry, as ProMoD initially exceeded the vehicle potential and left the track. (b) Resulting lines driven by ProMoD. After simultaneous adaptation of the target line and the velocity profile, ProMoD is able to successfully finish this turn.

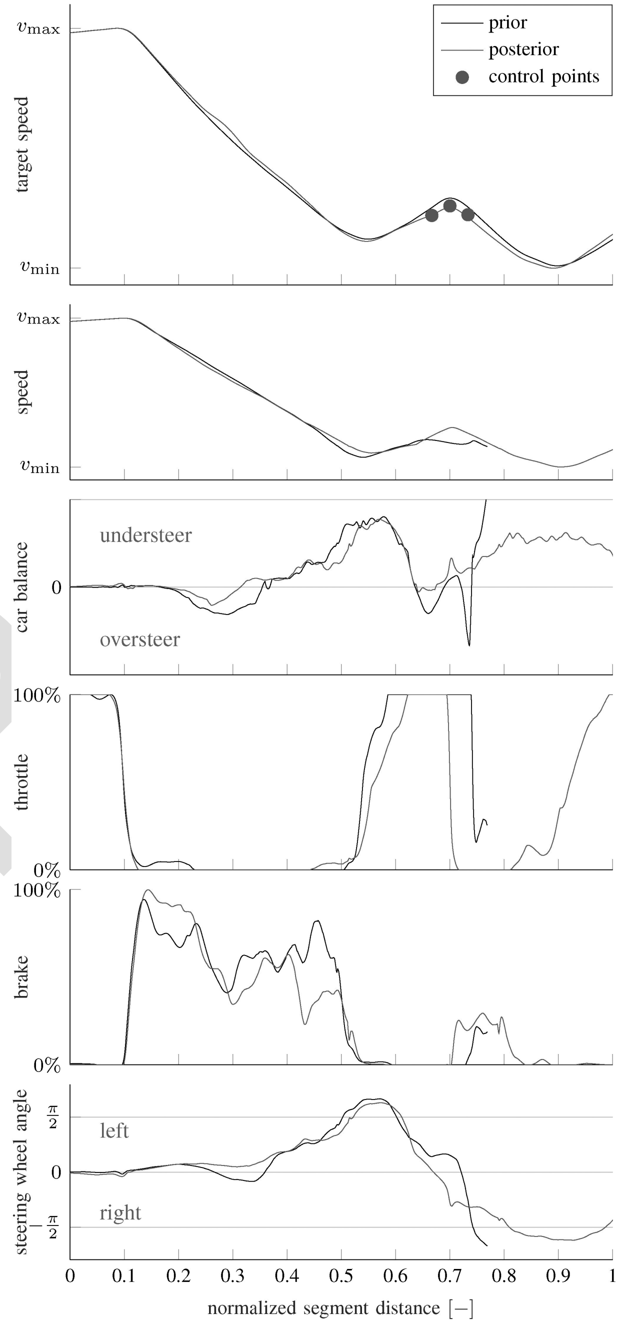

Fig. 10. Target speed and resulting vehicle states and driver actions over the normalized segment distance before and after adaptation (two iterations) of the target speed profile for T6/7 on SVT: Via three control points, the target speed profile is adapted while its general shape is preserved. The car balance refers to the dynamic driving state. When operated close to the friction limit (e.g., while cornering), the car balance typically assumes an oversteer (over-rotating, negative values) or understeer (under-rotating, positive values) state [1]. Before adaptation, at normalized segment distance 0.25, the vehicle oversteers and ProMoD is able to recover the vehicle by countersteering, at the cost of losing speed. However, at distance 0.65, ProMoD largely exceeds the grip potential, sliding over both axles which forces the vehicle off the track [see Fig. 9(b)]. After adapting the speed profile and driving line, ProMoD is able to keep the vehicle safely on track. Via *Action Selection*, ProMoD automatically increases the braking force during the first turn-in, accelerates later, and lifts the throttle and brakes earlier for the following turn.

represents a single lap with three metrics characterizing the individual driving style: throttle speed, brake speed, and the time of simultaneously pressed brake and throttle pedals. This plot indicates that after adaptation, the driver model remains capable of mimicking the individual characteristics of a specific driver while considerably differing from the others.

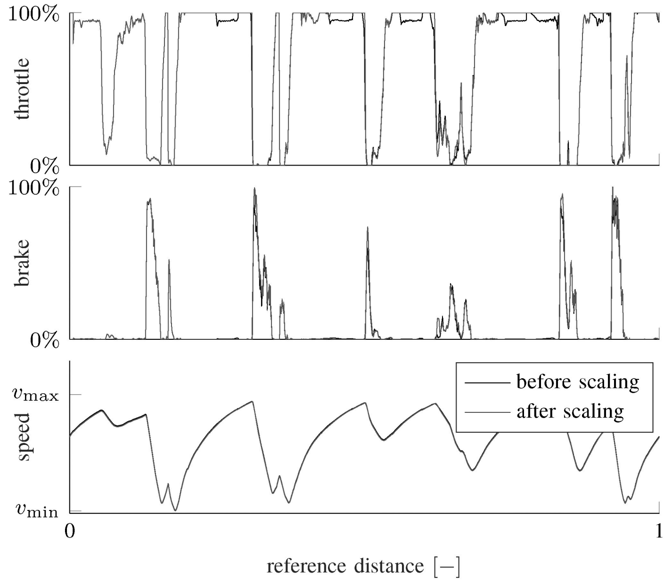

Fig. 11. Effect of *Speed Scaling* on straights: After scaling, ProMoD effectively utilizes the longitudinal potential of the vehicle and uses full throttle on most straights. For intervals where ProMoD would fail in subsequent turns due to the increased speed, scaling is prevented.

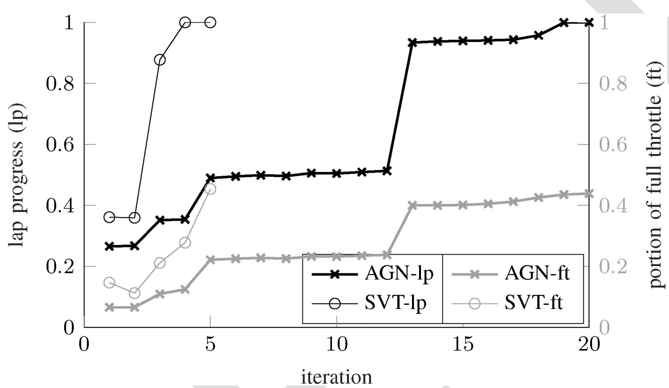

Fig. 12. Adaptation progress of ProMoD on AGN and SVT: For both tracks, ProMoD succeeds in completing a previously unfinished lap within 20 iterations, shown by lap progress (*lp*). The portion of full throttle is denoted by *ft*, where average expert values are 0.6152 and 0.5289 on SVT and AGN, respectively. Additional iterations can be used to further increase performance.

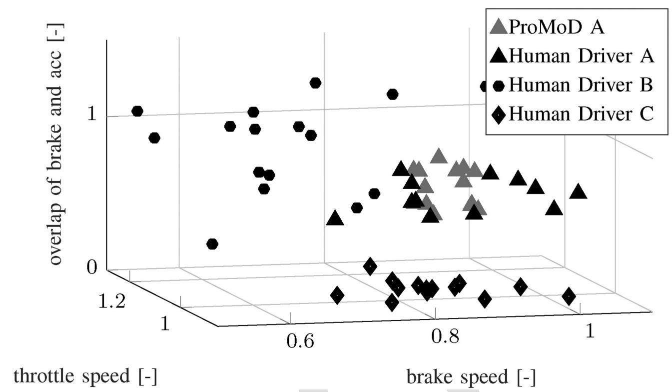

Fig. 13. Top three DIMRA driving style metrics of ProMoD and human drivers on SVT. ProMoD accurately mimics the individual driving style of driver A while still being distinguishable from two other professional race drivers.

IV. Conclusion

In this article, we collect insights into the general adaptation behavior and the learning processes of professional race drivers and derive new methods to extend ProMoD, an advanced modeling method for race driver behavior. With the purpose of understanding driver behavior in general and identifying the most important adaptation processes, this work starts with key insights from related work and experts inside and outside of the cockpit. Based on this acquired knowledge, we develop a novel method that estimates human-like driving line distributions for unknown tracks. These distributions can be used to simulate complete laps with almost competitive performances and human-like driver control inputs in a professional motorsport driving simulator. Subsequently, we present a feature adaptation method that allows ProMoD to learn from the gathered experience of previous laps. We demonstrate the model's ability to continuously learn from mistakes and to improve driving performance in terms of lap completion and time. This work contributes to the modeling and a better understanding of driver behavior, paving the way for advanced full-vehicle simulations with consideration of the human driver and potentially future autonomous racing.

Due to its modular architecture, ProMoD might be extended in various ways in future research. For feature adaptation and optimization, new methods may be introduced such as generating a more human-like masking matrix. Besides that, the neural network of the *Action Selection* module could be adapted to learn from experience using reinforcement learning techniques, or real track data may be used to provide more demonstration data. In order to better understand and model the efficient and complex adaptation process of human race drivers, approaching our modeling problem from the perspective of behavioral science is worth to be explored. On top of the development of the new adaptation methods, additional performance criteria related to the human adaptation process over subsequent laps could be defined for a more holistic assessment of the adaptation methods and improvement of the model. Furthermore, human-like qualitative feedback, which is based on encountered problems during driving, could help to further support the vehicle development process. In addition, our driver model may be extended to a multiagent environment with opponents on the race track, facilitating a more accurate prediction of true racing performance and potentially optimizing full racing strategies. Finally, ProMoD might be applied to similar use cases with the target of modeling human behavior in dynamic environments with small stability margins.

Appendix

Expert Interview

Is there a universal adaption rule that applies to all drivers and tracks?

Indeed, it turns out that adaption strategies are very similar across different drivers, tracks, and vehicles, in spite of the individual driving behavior, the various layouts of the tracks and the continuously modified vehicle setups. The driver's main goal is to "brake as late as possible, and accelerate as early as possible." The resulting driving line, the turn-in, and the on-throttle behavior are seen as a consequence of pursuing that goal.

How do drivers drive their first laps on a new track?

When faced with a new track, what a driver would do can be divided into three phases: 1) preparation; 2) warm-up; and 3) subsequent fine tuning.

1) *Preparation:* Drivers come to a new track with a memorized "database of corner information," collected from their prior experience, simulator sessions, statistical data, etc. First, drivers characterize each new corner by comparing it with those in their memory and assemble a first guess of the driving line. Since every corner is unique, this first guess is usually a rough approximation. At this point, it is helpful to consult other drivers to improve the initial guess. Finally, they set brake points, utilizing signs in the environment such as brake markers. Having concretized all prior information and exchanged opinions with fellow drivers of specific positions for hitting the brake pedal, the drivers start their first laps on a new track.
2) *Warm-Up:* Race drivers are particularly talented in assessing risk. They usually start off with a slow and safe speed profile, which they adapt from lap to lap to higher velocities. This process can take very few iterations. For example, one driver managed to reach a competitive lap time on the Le Mans circuit surprisingly after only five laps.
3) *Fine Tuning:* After warming up, drivers are able to complete the lap with a close to competitive lap time, which they then try to improve incrementally. Usually, drivers do not reach a global optimum but are aware of how to improve. High- and changing-speed corners are the most difficult ones, where spinning should be prevented, as it is extremely difficult to control.

Which quantities do race drivers adapt and how? Do they pay attention to specific metrics?

Although the goal of improving lap time is sound and clear, the real optimization process is indeed very complicated, and many factors have to be taken into consideration. The following three aspects are most critical during optimization.

1) *Delta Lap Time:* The adaption behavior of race drivers is result-oriented. They are not paying much attention to the exact speed values at local points around the track, but rather to the lap time difference to the previous or best lap. The association with the optimization problem is visualized on the top of Fig. 3.
2) *Brake Point:* Hitting the brake is where the corner starts. It is the most crucial tuning knob, not only because it influences the speed profile, but also since it is the source of any issues arising throughout the following corner. I.e., all issues should be traced back to the brake point, and cannot be locally analyzed.
3) *Peak Brake Pressure:* The driver attempts to predict the future state of the car when making decisions. In the presence of slip, however, uncertainty about the vehicle state is introduced, eventually leading to wrong predictions by the driver. Therefore, slip management is crucial during cornering, with the maximum brake pressure helping to anticipate imminent slip.

How do race drivers behave when the vehicle setup is modified? Will they preadapt their strategy according to the setup?

It is extremely complicated to analyze the car and the behavior of the driver simultaneously. Therefore, when new vehicle setups are tested, the drivers do not and are not expected to have much idea of what has been adapted on the car. Sometimes, race engineers would do blind tests in order to isolate the influences of the modified setups from those of the drivers.

References

[1] W. F. Milliken and D. L. Milliken, *Race Car Vehicle Dynamics*. Warrendale, PA, USA: Soc. Autom. Eng., 1995.

[2] P. R. Wurman et al., "Outracing champion Gran Turismo drivers with deep reinforcement learning," *Nature*, vol. 602, no. 7896, pp. 223–228, 2022.

[3] S. Löckel, J. Peters, and P. van Vliet, "A probabilistic framework for imitating human race driver behavior," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 2086–2093, Apr. 2020.

[4] S. Löckel, A. Kretschi, P. van Vliet, and J. Peters, "Identification and modelling of race driving styles," *Veh. Syst. Dyn.*, vol. 60, no. 8, pp. 2890–2918, 2022.

[5] Porsche AG. "Simulator." Nov. 2019. [Online]. Available: https://presse.porsche.de/prod/presse_pag/PressResources.nsf/Content?ReadForm&languageversionid=1043793&hl=christophorus-393-misc

[6] K. Fritzsche, F. Strecker, and M. Huneke, "Simulation-based development process for Motorsport and series production," *ATZ Worldwide*, vol. 119, no. 9, pp. 60–63, 2017.

[7] B. Siegler, "Lap time simulation for racing car design," Ph.D. dissertation, Dept. Comput. Sci., Univ. Leeds, Leeds, U.K., 2002.

[8] D. P. Kelly, "Lap time simulation with transient vehicle and tyre dynamics," Ph.D. dissertation, School Eng., Cranfield Univ., Cranfield, U.K., 2008.

[9] G. Perantoni and D. J. N. Limebeer, "Optimal control for a formula one car with variable parameters," *Veh. Syst. Dyn.*, vol. 52, no. 5, pp. 653–678, 2014.

[10] U. Rosolia, X. Zhang, and F. Borrelli, "Data-driven predictive control for autonomous systems," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 1, no. 1, pp. 259–286, 2018.

[11] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 3363–3370, Nov. 2020.

[12] T. Novi, "A control system framework for autonomous vehicles at the limits of handling," Ph.D. dissertation, School Eng., Univ. Florence, Florence, Ialy, 2019.

[13] F. Fuchs, Y. Song, E. Kaufmann, D. Scaramuzza, and P. Dürr, "Super-human performance in gran turismo sport using deep reinforcement learning," *IEEE Robot. Autom. Lett.*, vol. 6, no. 3, pp. 4257–4264, Jul. 2021.

[14] P. M. van Leeuwen, S. de Groot, R. Happee, and J. C. F. de Winter, "Differences between racing and non-racing drivers: A simulator study using eye-tracking," *PLoS ONE*, vol. 12, no. 11, Nov. 2017, Art. no. e0186871.

[15] R. Bentley, *Ultimate Speed Secrets: The Complete Guide to High-Performance and Race Driving*, Motorbooks Int., Minneapolis, MN, USA, 2011.

[16] M. Krumm, *Driving on the Edge: The Art and Science of Race Driving*. London, U.K.: Icon, 2015.

[17] J. C. Kegelman, "Learning from professional race car drivers to make automated vehicles safer," Ph.D. dissertation, Dyn. Design Lab., Stanford Univ., Stanford, CA, USA, 2018.

[18] R. A. Hess and A. Modjtahedzadeh, "A control theoretic model of driver steering behavior," *IEEE Control Syst. Mag.*, vol. 10, no. 5, pp. 3–8, Aug. 1990.

[19] L. König, *Ein virtueller Testfahrer für den querdynamischen Grenzbereich*, vol. 43. Heidelberg, Germany: expert Verlag, 2009.

[20] D. D. Salvucci, "Modeling driver behavior in a cognitive architecture," *Human Factors J. Human Factors Ergonom. Soc.*, vol. 48, no. 2, pp. 362–380, 2006.

[21] P. Cichosz and Ł. Pawełczak, "Imitation learning of car driving skills with decision trees and random forests," *Int. J. Appl. Math. Comput. Sci.*, vol. 24, no. 3, pp. 579–597, 2014.

[22] H. Wei, W. Ross, S. Varisco, P. Krief, and S. Ferrari, "Modeling of human driver Behavior via receding horizon and artificial neural network controllers," in *Proc. IEEE Conf. Decis. Control*, 2013, pp. 6778–6785.

[23] A. Kuefler, J. Morton, T. Wheeler, and M. J. Kochenderfer, "Imitating driver Behavior with generative adversarial networks," in *Proc. IEEE Intell. Veh. Symp.*, 2017, pp. 204–211.

[24] R. Bhattacharyya et al., "Modeling human driving behavior through generative adversarial imitation learning," 2020, *arXiv:2006.06412*.

[25] Z. Huang, J. Wu, and C. Lv, "Driving behavior modeling using naturalistic human driving data with inverse reinforcement learning," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 8, pp. 10239–10251, Aug. 2022.

[26] Y. Li, J. Song, and S. Ermon, "InfoGAIL: Interpretable imitation learning from visual demonstrations," in *Proc. NIPS*, 2017, pp. 3815–3825.

[27] Z. Zhang, A. Liniger, D. Dai, F. Yu, and L. Van Gool, "End-to-end urban driving by imitating a reinforcement learning coach," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 15222–15232.

[28] Z. Deng et al., "A probabilistic model for driving-style-recognition-enabled driver steering behaviors," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 52, no. 3, pp. 1838–1851, Mar. 2022.

[29] S. Schnelle, J. Wang, H. Su, and R. Jagacinski, "A driver steering model with personalized desired path generation," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 47, no. 1, pp. 111–120, Jan. 2017.

[30] B. Shi et al., "Evaluating driving styles by normalizing driving behavior based on personalized driver modeling," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 45, no. 12, pp. 1502–1508, Dec. 2015.

[31] Y. Wang et al., "Imitation learning of hierarchical driving model: From continuous intention to continuous trajectory," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 2477–2484, Apr. 2021.

[32] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, "Probabilistic movement primitives," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 2616–2624.

[33] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, "Using probabilistic movement primitives in robotics," *Auton. Robots*, vol. 42, no. 3, pp. 529–551, 2018.

[34] M. Mulder et al., "Manual control cybernetics: State-of-the-art and current trends," *IEEE Trans. Human–Mach. Syst.*, vol. 48, no. 5, pp. 468–485, Oct. 2018.

[35] A. Ronen and N. Yair, "The adaptation period to a driving simulator," *Transp. Res. F Traffic Psychol. Behav.*, vol. 18, pp. 94–106, May 2013.

[36] R. Deborne, R. Gilles, and A. Kemeny, "Modeling driver adaptation capabilities in critical driving situations," in *Proc. SAE World Congr. Exhibit.*, 2012, pp. 1–10.

[37] F. Braghin, F. Cheli, S. Melzi, and E. Sabbioni, "Race driver model," *Comput. Struct.*, vol. 86, nos. 13–14, pp. 1503–1516, 2008.

[38] R. Bentley, *Speed Secrets: Professional Race Driving Techniques*. Minneapolis, MN, USA: MotorBooks Int., 1998.

[39] S. Quinlan and O. Khatib, "Elastic bands: Connecting path planning and control," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1993, pp. 802–807.

[40] S. K. Gehrig and F. J. Stein, "Collision avoidance for vehicle-following systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 8, no. 2, pp. 233–244, Jun. 2007.

[41] B. Siegler, A. Deakin, and D. Crolla, "Lap time simulation: Comparison of steady state, quasi-static and transient racing car cornering strategies," *SAE Trans.*, vol. 109, no. 6, pp. 2575–2581, 2000.

AQ4

Stefan Löckel received the B.Eng. degree in mechatronics from DHBW Mannheim, Mannheim, Germany, in 2015, the M.Sc. degree in electrical engineering and control from TU Kaiserslautern, Kaiserslautern, Germany, in 2017, and the Ph.D. degree in computer science from TU Darmstadt, Darmstadt, Germany, in 2022.

He is currently employed with Dr. Ing. h.c. F. Porsche AG, Weissach, Germany, where he works on machine-learning models for the development of future sports cars.

Siwei Ju received the B.Sc. degree in mechanical engineering from the Beijing Institute of Technology, Beijing, China, in 2018, and the M.Sc. degree in automation engineering from RWTH Aachen, Aachen, Germany, in 2021. She is currently pursuing the Ph.D. degree with Dr. Ing. h.c. F. Porsche AG, Weissach, Germany, and Intelligent Autonomous System, Computer Science Department, TU Darmstadt, Darmstadt, Germany. AQ5

Her research focuses on race driver modeling using reinforcement learning and imitation learning.

Maximilian Schaller received the B.Sc. degree in mechanical and process engineering from TU Darmstadt, Darmstadt, Germany, in 2019, and the M.Sc. degree in robotics, systems and control from ETH Zürich, Zürich, Switzerland, in 2022.

He is currently affiliated with the Electrical Engineering Department, Stanford University, Stanford, CA, USA. His research focuses on optimization for control and learning for time-series modeling.

Peter van Vliet received the M.Sc. degree in mechanical engineering from the Delft University of Technology, Delft, The Netherlands, in 2005, specializing in systems and control.

Since 2012, he has been with Dr. Ing. h.c. F. Porsche AG, Weissach, Germany, working on a variety of topics around high-performance road and race cars. His current interests are research activities in the field of machine learning, learning control systems and data driven, and advanced method development.

Jan Peters has been a Full Professor (W3) for Intelligent Autonomous Systems with the Computer Science Department, Technische Universität Darmstadt, Darmstadt, Germany, since 2011, and, at the same time, he has been the Head of the Research Department on Systems AI for Robot Learning, German Research Center for Artificial Intelligence (Deutsches Forschungszentrum für Künstliche Intelligenz, DFKI), Kaiserslautern, Germany, since 2022. He is also a Founding Research Faculty Member of the Hessian Center for Artificial Intelligence, Darmstadt. AQ6

Dr. Peters has received the Dick Volz Best 2007 U.S. Ph.D. Thesis Runner-Up Award, the Robotics: Science and Systems—Early Career Spotlight, the INNS Young Investigator Award, the IEEE Robotics and Automation Society's Early Career Award, as well as numerous best paper awards. In 2015, he received an ERC Starting Grant. He was an ELLIS Fellow and a Program Director in 2020 and an AAIA Fellow in 2021.