

Learning dynamics models for velocity estimation in autonomous racing

Jan Węgrzynowski, Grzegorz Czechmanowski, Piotr Kicki and Krzysztof Walas

Abstract—Velocity estimation is of great importance in autonomous racing. Still, existing solutions are characterized by limited accuracy, especially in the case of aggressive driving or poor generalization to unseen road conditions. To address these issues, we propose to utilize Unscented Kalman Filter (UKF) with a learned dynamics model that is optimized directly for the state estimation task. Moreover, we propose to aid this model with the online-estimated friction coefficient, which increases the estimation accuracy and enables zero-shot adaptation to the new road conditions. To evaluate the UKF-based velocity estimator with the proposed dynamics model, we introduced a publicly available dataset of aggressive maneuvers performed by an F1TENTH car, with side-slip angles reaching 40° . Using this dataset, we show that learning the dynamics model through UKF leads to improved estimation performance and that the proposed solution outperforms state-of-the-art learning-based state estimators by 17% in the nominal scenario. Moreover, we present unseen zero-shot adaptation abilities of the proposed method to the new road surface thanks to the proposed learning-based tire dynamics model with online friction estimation.

I. INTRODUCTION

In recent years, we have observed a huge interest in self-driving technologies, starting from driver assistance systems and ending with fully autonomous vehicles. Implementing these technologies reduces transportation costs and air pollution, and in the long run, gives a promise of reduced traffic and much safer and more accessible transportation [1]. Most of the research in this area is focused on calm maneuvers that cover most of the driving time. However, from a safety perspective, it is especially important to develop systems prepared for both everyday scenarios and emergencies that may require aggressive maneuvers to avoid a collision.

An application that exploits aggressive maneuvers on a daily basis is autonomous racing (see Fig. 1). Thus, we want to focus on it in this paper. In particular, we tackle the problem of velocity estimation, which is an important part of both perception and control pipelines, as it allows us to undistort the LiDAR measurements and provide state feedback. Although it is possible to use exteroceptive sensors, like GNSS or ground speed sensors, to measure the velocity, they increase the cost of the platform and are prone to failure or bias in challenging environments. Therefore, we focus on the problem of velocity estimation using only proprioception, namely Inertial Measurement Unit (IMU) and motor speed.

All authors are with IDEAS NCBR, Warsaw, Poland and with Institute of Robotics and Machine Intelligence, Poznan University of Technology, Poznan, Poland `name.surname@ideas-ncbr.pl`

Both code of the proposed method and dataset are available at <https://github.com/SzalonyJohnny/Learning-dynamics-models-for-velocity-estimation>



Fig. 1. The main goal of this paper is to learn a dynamics model of an F1/10 car for accurate velocity estimation during aggressive maneuvers.

One of the most popular approaches to this problem relies on using Kalman Filtering, especially Extended Kalman Filter (EKF) [2], [3] and UKF [4] [5]. These methods strongly rely on the accuracy of the car dynamics model, which is crucial in parts of the state space in which the considered system is hard to observe. This can be addressed by using the hybrid car dynamics models, which consist of the physical car model aided with a neural network to enable accurate dynamics modelling [6], [7]. However, these models are trained to predict the state evolution equally accurately among the whole state space, which may not be optimal from the state estimation point of view. It is also possible to utilize fully data-driven state estimators like Recursive Neural Networks (RNNs) [8], [9] as they allow for highly accurate and flexible approximation of the car and measurement dynamics. However, pure machine learning-based methods are typically limited in terms of generalization abilities, like adjusting to changing tire-road friction [10].

To address the aforementioned issues, we propose to exploit the concept of differentiable Kalman filters [11], [12] in the process of learning the dynamics model of a racing car. Although this approach is not new, it is mostly used to train measurement models [12], [13], [14] and its impact on learning data-driven dynamics models of the complex robotics systems for state estimation is not well studied [15]. Surprisingly, learning a dynamics model through a differentiable Kalman filter is said to have no impact on the state estimation accuracy [12]. Still, this paper argues the opposite based on the thorough experimental analysis. In fact, we claim that learning the dynamics model in this way follows a better optimization criterion, as opposed to training the model to be an accurate one-step predictor, and allows the neural network-based model to focus on the accurate prediction along the more important state space dimensions from the estimation point of view. Moreover, we show that in

the considered case learning an accurate car model is more important for the prediction step than for the observation.

Moreover, we exploit the structure of the UKF to estimate not only the car velocities but also the tire-road friction, which allows the learned state estimator to adapt to new road conditions without the need to collect new data and retrain. This zero-shot test time adaptability is crucial in deploying reliable state estimators in the real world, which is typically the main drawback of machine learning-based solutions. Finally, we analyze the sources of the improved velocity estimation performance while learning the dynamics model through UKF.

To compare the proposed approach with state-of-the-art state estimation approaches and validate the aforementioned claims, we introduce a dataset of very dynamic driving of a 1:10 scale F1 car model [16]. This allows us to safely examine the state estimation in severe conditions, including huge side-slip, slip ratios, lateral accelerations, and decreased tire-road friction.

Our contributions can be summarized as follows:

- We show that learning the car dynamics model through UKF increases the state estimation accuracy and allows for outperforming state-of-the-art estimators based on RNNs, while maintaining an interpretability of UKF.
- We propose to aid a hybrid of learning-based and analytical car dynamics models with an online estimated tire-road friction, which increases the estimation accuracy and enables adaptation to the new conditions during test time in a zero-shot manner.
- We introduce a publicly available 52 minutes long dataset of the aggressive maneuvers performed by a FITENTH car [16] (see Fig. 1) with side-slip angles reaching 40° , which is far beyond what one can find in the publicly available datasets.

II. RELATED WORK

In this section, we discuss related papers both from the perspective of the considered application, i.e., velocity estimation of a car during aggressive maneuvers, and the approach we propose, i.e., machine learning-based car dynamics modeling and online tire-road friction estimation.

One of the most classic yet effective approaches to car velocity estimation utilizes the EKF [2], [3] that fuses the data from several sensors and uses the car dynamics model to estimate the lateral velocity. In turn, advances in the field of machine learning inspired many researchers to incorporate neural networks into the Kalman filtering pipeline. The authors of [4] and [5] used a neural network to perform as a virtual side-slip sensor, which measurements are then filtered out using UKF. In [4] the neural network is trained supervisedly to predict the side-slip based on data from the simulation, while in [5] the real data is used and the loss is computed based on the output of the UKF, similarly like in this paper. A different approach was introduced in [17], where the RNN is utilized to predict the Kalman gain, while the rest of the pipeline is not taking advantage of the representational power of neural networks.

The racing car velocity estimation task can also be solved without using Kalman filtering but by exploiting the RNNs trained in an end-to-end manner to maintain some latent representation of the robot state and estimate the velocities out of it [8], [9]. A Gated Recurrent Units (GRUs) based RNN has been utilized in [8] and allowed for an accurate estimation of the vehicle velocity during maneuvers with high side-slip (10° at the rear axle) and slip ratio ($\approx 20\%$), without using external velocity sensors. Authors of [9] furthered this research by using Long Short-Term Memory (LSTM) units to process the sensory information and fully connected layers to estimate the velocity based on the previous estimates and the hidden state of the LSTM. This approach led to superior accuracy on a real-world, full-scale car dataset compared to traditional model-based and other learning-based approaches. Nonetheless, a significant limitation of RNN-based estimators is their lack of built-in uncertainty estimation mechanisms, a critical component ensuring safety during aggressive driving maneuvers. Moreover, due to the end-to-end machine learning-based solutions, these models are potentially more susceptible to the *out-of-distribution* data, which is essential for safety across diverse road conditions.

To make the machine learning-based state estimators more robust, some recent works focused on enhancing the learning-based models with some physical structure [6], [18]. In [18], the neural network used to predict the state estimate is fitted into a kinematic structure of the vehicle to satisfy kinematic laws. However, it lacks insights into the vehicle's dynamics. In turn, in [6], like in this paper, an UKF with a hybrid vehicle dynamics model is considered. They enhance the classical model with a residual neural ordinary differential equation and train them to predict the state derivative that best matches the next observation. Our paper shows that learning residual dynamics model along the classical one lead to similar results as for fully-learned dynamics, both having problems with generalization. Instead, we propose to aid the classical dynamics model with a neural network modeling the tire-forces, like in [7] or [19]. Authors of [7] follow the structure of the Pacejka tire model and use a neural network to predict its parameters, while authors of [19] introduce an *ExpTanh* tire model. In turn, we do not impose this kind of structure on the tire model but use the neural network to model the tire forces almost directly.

One of the main determinants of the forces generated by tires is the tire-road friction coefficient [20]. There were several attempts to learn how to predict it in real-world scenarios using neural networks [10] trained *end-to-end* or hybrid models [20], [21] adapted online based on prediction errors. We follow the second approach, but instead of adapting some friction predictor, we estimate the friction coefficient online using UKF, and we use it to scale the neural network tire model output. Moreover, we show that training the dynamics model through the UKF leads to a notably better state estimation performance than simply training the dynamics model to be an accurate one-step predictor as is typically proposed in literature.

III. MATERIALS AND METHODS

A. Problem definition

In this paper, we consider a problem of the estimation of the vehicle state x defined by

$$x = [v_x, v_y, r, \omega_s], \quad (1)$$

where v_x and v_y are longitudinal and lateral velocities in the local vehicle frame, r is a yaw rate, and ω_s is the motor speed scaled by wheel radius and gear ratio. We assume that the estimation should be performed using (i) the proprioceptive sensors that almost every robot is equipped with, i.e. IMU, which gives us linear accelerations a^{IMU} and angular velocities ω^{IMU} , and velocity measurements ω_s from the electric motor driver, as well as (ii) control signals consisting of steering angle δ and motor current I_q .

The considered task is not trivial due to the aggressive nature of the racing driving style that includes both longitudinal and lateral slips and, therefore, limited observability of the system state. Because the system state is not measurable directly, a reliable and precise system model is necessary to obtain the state estimates from time series of measurements and control inputs. Moreover, we would like to consider the case in which the parameters of the environment are non-stationary and non-uniform, so we are not making popular assumptions about the uniform tire-road friction coefficient over the whole driving area and driving time. In fact, we want to challenge the state estimators to generalize to different friction conditions, preferably without tedious data collection and retraining.

Note that a similar problem to the one presented above is typically solved by incorporating the IMU measurements directly into the simplified dynamics model [2], [3]. However, this approach is not feasible in the problem considered in this paper due to the lack of an external velocity sensor (line GNSS or optical velocity sensor) that can correct the drift caused by the integration of measured acceleration.

B. Proposed solution

To address the problem introduced in the previous section, we propose to use an UKF framework [22], as it is a standard approach to state estimation of nonlinear systems. In particular, we want to exploit the differentiable implementation of UKF [12] to learn an accurate hybrid (understood as a combination of classical and learning-based) model of the system dynamics that will allow us to estimate the robot state during aggressive driving precisely. The general scheme of the proposed approach is presented in Fig. 2. It consists of the classical UKF prediction and update steps, with tunable vehicle and noise models based on the gradient of the estimation loss function. In the next subsections, we describe these elements in detail.

1) *UKF update*: In the update step, following the standard UKF algorithm [22], we compute the difference between the sensor measurements y and the output of the measurement function $h(\hat{x}, u)$, and scale it by the Kalman gain to obtain a state estimate update. The yaw rate r and wheels speed ω_s




Fig. 2. The general scheme of the differentiable UKF-based state estimation, with learnable vehicle and noise models. We propose to train the car dynamics and noise models through UKF to minimize the state estimation error.

are directly measured and comparable with their estimates. However, to compare with longitudinal a_x^{IMU} and lateral a_y^{IMU} accelerations from IMU we need to compute the estimated accelerations in the inertial frame by $\hat{a}_x = \hat{v}_x - \hat{r}\hat{v}_y$ and $\hat{a}_y = \hat{v}_y + \hat{r}\hat{v}_x$, where \hat{v}_x, \hat{v}_y are the local acceleration estimates based on the dynamics model $\dot{\hat{x}} = f(\hat{x}, u)$ and $\hat{r}, \hat{v}_x, \hat{v}_y$ are the estimates of the r, v_x and v_y .

2) *UKF prediction*: In the prediction step, the system dynamics model estimates the future state based on the current state and control inputs. This smooths out the information obtained from update steps and adjusts it to comply with the dynamics of the observed system. The accuracy of the dynamics model is particularly important in those parts of the state space where access to reliable measurements is limited. Still, in our specific case, its importance is even greater because it is also a key part of the measurement model.

In this paper, we analyze the whole family of different vehicle dynamics models, ranging from single track car with Pacejka tire model to the fully-learned vehicle model, which are inspired by the approaches from literature [2], [23], [24], and extend them in a more structured way of learning the vehicle dynamics.

Single track model We generally consider the vehicle dynamics as a continuous function $\dot{x} = f(x, u)$, which can be integrated to obtain the next discrete state. In our considerations, we omit simplistic dynamics models, like a point mass or kinematic bicycle, as they are not feasible in the case of dynamic driving. Instead, we employ a single-track model for its accurate representation of the basic vehicle dynamics, which can be defined by

$$\begin{aligned} \dot{v}_x &= \frac{1}{m} (F_{x_r} + F_{x_f} \cos(\delta) - F_{y_f} \sin(\delta) - F_{drag} + mv_y r), \\ \dot{v}_y &= \frac{1}{m} (F_{x_f} \sin(\delta) + F_{y_r} + F_{y_f} \cos(\delta) - mv_x r), \\ \dot{r} &= \frac{1}{I_z} ((F_{x_f} \sin(\delta) + F_{y_f} \cos(\delta)) l_f - F_{y_r} l_r), \\ \dot{\omega}_s &= \frac{1}{I_e} (k_\phi I_q - \rho \cdot F_{x_f} - \rho \cdot F_{x_r} - \tau_t), \end{aligned} \quad (2)$$

where m is the vehicle mass, F_{drag} is the aerodynamic drag force acting of the vehicle, assumed to scale quadratically with v_x , δ is the steering angle, I_z is the vehicle moment of inertia around the z -axis, l_r and l_f represents distances from center of mass to rear and front axis respectively, ρ is the wheel radius, I_e is powertrain inertia, k_ϕ is the motor's torque constant, while $F_{x_r}, F_{x_f}, F_{y_r}, F_{y_f}$ represents the x

(longitudinal) and y (lateral) components of the lumped tire forces generated by the rear and front virtual wheels. We assume a DC motor model with torque proportional to motor current I_q . This torque is reduced by the transmission system losses, which are modeled as $\tau_t = k_{tC} \text{sgn}(\omega_s) + k_{tv}\omega_s$ acting on the motor shaft, where k_{tC} , k_{tv} represents Coulomb and viscous friction coefficient respectively.

Pacejka tire model The most challenging part of (2) is how to model the lumped forces that model the complex phenomenon of the tire-road interaction. One of the most popular approaches to this matter is to express these forces using Pacejka Magic Formula [25]

$$\begin{aligned} F_x &= \mu D_x \sin(C_x \arctan(B_x \kappa \\ &\quad - E_x (B_x \kappa - \arctan(B_x \kappa))))), \\ F_y &= \mu D_y \sin(C_y \arctan(B_y \alpha \\ &\quad - E_y (B_y \alpha - \arctan(B_y \alpha)))). \end{aligned} \quad (3)$$

where $B_x, C_x, D_x, E_x, B_y, C_y, D_y, E_y$ represents the parameters of the model for longitudinal and lateral lumped tire forces, μ is the friction coefficient between the tire and driven surface, κ is the slip ratio and α is the slip angle. Note that we use the same parameters to obtain the longitudinal forces acting on the front or rear tire, but in the case of lateral force, the B_y, C_y, D_y, E_y parameters can be different for the front and rear virtual wheels. Similarly, the slip angle α is computed separately for specific wheels.

Neural car model Tire forces, commonly represented with (3), offer good generalization but may not adequately capture complex tire behaviors due to their limited flexibility. This issue led to the development of data-driven car models that typically utilize neural networks to model vehicle dynamics. One of the most straightforward approaches in this spirit [23] is to let the neural network approximate the whole dynamics

$$\dot{x} = f_{NN}(x, u). \quad (4)$$

However, despite neural networks' great capacity to surpass classical models in making predictions on data similar to their training environment, their applicability to new scenarios is limited, as they offer limited interpretability and generalization.

Neural tire model To address these issues, we propose a hybrid approach that integrates analytically derived models with neural networks, aiming to achieve both generalizability and accuracy. Analytical models address simpler behaviors, while neural networks tackle the more complex aspects, such as tire forces [19], [24]. In our solution, similarly to [24], a neural network is applied to represent all tire forces

$$[F_{x_r}, F_{x_f}, F_{y_r}, F_{y_f}] = F_{NN}(x, u), \quad (5)$$

while the vehicle dynamics is expressed by (2).

Friction-scaled neural tire model To enhance the adaptability and interpretability of the neural tire model (5), we employ the friction coefficient μ to scale tire forces generated by the neural network such that

$$[F_{x_r}, F_{x_f}, F_{y_r}, F_{y_f}] = \mu \cdot F_{NN}(x, u). \quad (6)$$

This allows us to adjust the predictions of the neural tire model to the driving conditions that may not be present during training just by knowing the estimate of the friction coefficient $\hat{\mu}$ during the test. In particular, we exploit the UKF framework and we use it to estimate the friction coefficient online, which is described in more detail in Section III-B.5.

3) *Noise model*: A crucial component of Bayesian estimation methods is the noise model [26]. It quantifies the uncertainty of the process and measurement models using covariance matrices R and Q . In this paper, we omit the diagonal noise model due to its oversimplification. Instead, we incorporate the off-diagonal elements into the covariance matrices and define the noise model by

$$R = L_R L_R^T + I \cdot \epsilon, \quad Q = L_Q L_Q^T + I \cdot \epsilon, \quad (7)$$

where L_R and L_Q are lower triangular matrices with trainable elements, I is the identity matrix, and ϵ is small number (e.g. $1e-7$) which enhances the numerical stability. This formulation guarantees that both R and Q matrices remain positive definite throughout the training process.

The noise model introduced above assumes that the covariance matrices are constant, which may not be true, e.g., the noise on the IMU may increase when the car is moving very fast on uneven surfaces. To account for these kinds of effects, we utilize a heteroscedastic noise model, which has been shown to notably impact UKF-based state estimation [12]. In this noise model, we predict the elements of the L_R and L_Q matrices based on the current state estimate \hat{x} using linear regression and compute R and Q using (7).

4) *Model training procedure*: All the dynamics models described in Section III-B.2, whether classically derived from first principles or exploiting neural networks, are parameterized with θ , which include, e.g. tire parameters, drag coefficient, neural network weights, etc.. Some of these parameters have physical meaning and must remain positive, which we ensure by learning their logarithms. The values of θ have a tremendous impact on the dynamics model accuracy, so they need to be identified or learned prior to the use of the model in the state estimator. Thanks to the differentiability of the considered models, one can optimize their parameters θ using gradients $\nabla_{\theta} \mathcal{L}$ of some loss function \mathcal{L} w.r.t θ .

To learn an accurate dynamics model $f(x, u)$, one must decide on the loss function \mathcal{L} formulation. The most popular way is to train the model to minimize the square of the one-step prediction error $e_k = (x_{k+1} - \text{ODESolve}(f(x_k, u_k), T_s))^2$, where T_s is the step duration and ODESolve is a numerical integrator, which in our case is the 4-th order Runge-Kutta method. In this way, the model is optimized to accurately predict the local state derivative \dot{x} such that its integral fits the measured evolution of the system state $x_{k+1} - x_k$.

Although training dynamics models to be accurate one-step predictors is a well-established approach, it does not consider the context of their application. This would make no difference for perfect models, but in case of their limited

capacity, it may be suboptimal from the specific task point of view. Therefore, we propose to utilize the differentiable implementation of UKF [12] and optimize the parameters of the dynamics model so that the resulting UKF achieves the smallest possible state estimation errors instead of the minimal one-step prediction errors. This way, we optimize a better criterion from the final task point of view.

Besides learning the dynamics model through UKF, one can also use the same procedure to update the parameters of the noise models, i.e., L_R and L_Q for the homoscedastic process and observation noise models or parameters of the heteroscedastic noise model (see Section III-B.3). The importance of this step is thoroughly described in [26].

Moreover, we observe that simultaneous learning of the dynamics model and noise parameters is also an important factor in the final estimation performance, as it allows for the deliberate adaptation of the noise models to the changing dynamics model. In particular, in the case of the state estimation, we may not need a dynamics model that makes small prediction errors on average for the whole state space. Instead, it may prefer to focus on some key components of the state space that are hard to be measured, and compensate for the resulting prediction errors of the remaining components by reducing their corresponding measurement covariances.

Following the training methods described above, the dynamics model training can be separated into two parts:

- pretraining based on the one-step prediction loss minimized on the batch of samples,
- estimation fine-tuning, where noise model and prediction model parameters are tuned through the differentiable implementation of UKF on sequences of samples by minimizing the Mean Squared Error (MSE) between ground-true states and state estimates (see Fig. 2).

In this way, we can quickly pretrain an accurate dynamics model and then adapt it to the desired task of velocity estimation. Although it is possible to train the model through UKF from scratch, we have found this approach to be inefficient and prone to getting stuck in local minima.

5) *Friction coefficient estimation*: In Section III-B.2, in particular in (6), we proposed to enhance the tire dynamics model by scaling the output of the neural network-based model by the friction coefficient μ . In many papers regarding the racing car velocity estimation, the tire-road friction coefficient is assumed to be constant, e.g. [5], [19], which seems quite impractical taking into account possible changes of the weather conditions or wear out of the tires. Alternatively, it may be treated as a disturbance [8], [9], to which the prediction model needs to become robust, which may lead to suboptimal estimation performance and require collecting data with a wide spectrum of μ . In turn, in [10], only several specific friction coefficients are considered and assumed to be constant and homogeneous on large track areas.

To overcome these simplifying assumptions we propose to estimate μ online exploiting the UKF framework, by adding

μ to the state vector

$$x = [v_x, v_y, r, \omega_s, \mu]. \quad (8)$$

As we cannot model the dynamics of the friction coefficient, we assume it to be constant $\dot{\mu} = 0$, but at the same time, we allow the adjustments of μ during the update step of the UKF, based on the observed discrepancy between the observed and predicted measurements.

C. Dataset

The development of state estimation algorithms requires the availability of high-quality, diverse data. Unfortunately, previous studies [8], [9], [17] often did not publish datasets, limiting the ability to benchmark and improve estimation methods. The most relevant publicly available dataset for developing agile estimation algorithms is Targa Sixty-Six [27], which contains data from a Ferrari 250 LM driven at Palm Beach International Raceway at a speed reaching 200 km/h. However, even though they feature vehicles traveling at speeds exceeding 220 km/h, they do not extensively cover high side-slip maneuvers that occur while operating vehicles at the limits of handling. This notable gap in the availability of challenging, real-world datasets has propelled us to develop and publicly release our dataset.

The dataset introduced in this paper is created from real-world data obtained using an FITENTH vehicle [16] and an external motion capture system. The experimental vehicle, an Xray GTXE'22, is a four-wheel drive, 1/8th scale RC car. It weighs 4.5 kg and is powered by a single motor controlled by a motor driver. It captures raw sensor data at varying frequencies: the IMU at 400Hz and the motor RPM, motor q-axis current, and steering angle at 100Hz. High-precision positional and orientation data were captured using an OptiTrack motion capture system at 100 Hz. Data collection utilized two computers (one onboard for sensor data and another for motion capture) synchronized with Chrony to minimize time discrepancies. To facilitate accurate state estimation, these raw measurements were linearly interpolated and sampled at 100 Hz. Ground truth vehicle velocities were calculated by differentiating OptiTrack position data with a Savitzky-Golay filter [28], using a second-order polynomial and an 8 ms filter window. The velocities were then sampled at the same timestamps as the raw measurements.

This way, we collected 52 minutes of driving time with a remotely controlled car. Thanks to this, we were able to collect a diverse set of high slip and high acceleration maneuvers, for which the 99th percentile (to ignore outliers) of the slip angle at the rear axle exceeds maximum values reported in previous works [8], [9], [27] by a substantial margin (see Fig. 3). In fact, our dataset features dynamic maneuvers with a slip angle at the rear axle surpassing 40° (99th percentile), which presents a much more demanding estimation challenge than existing datasets.

Moreover, typically, the racing datasets contain data for specific tire-road conditions and neglect the complexity of real-life racing, e.g., tire wear during the race or wet asphalt because of the rain. To take such circumstances into account




Fig. 3. Distribution of the rear axle slip angle in the introduced dataset. Maximum value reported in [8] (red) and 99th Percentile from Targa Sixty-Six dataset [27] (purple) are added as reference.

and facilitate the research in the direction of robust data-driven state estimators we increased the diversity of the introduced dataset by using 4 different sets of tires, each offering different coefficients of friction - tire A (20.1% of the dataset, $\mu_A \approx 0.65$), tire B (18.7%, $\mu_B \approx 0.58$), tire C (50.2%, $\mu_C \approx 0.43$) and tire D (11%, $\mu_D \approx 0.45$). Tires C and D are meant to simulate the conditions of the decreased friction, e.g., due to the wet road surface, which we obtained by covering tires A and B with two different types of tape.

IV. EXPERIMENTS

In this section, we examine the performance of the UKF with dynamics models introduced in Section III-B.2 and relate it to the state-of-the-art state estimation methods based on RNNs [8], [9]. To do so, we use the dataset introduced in Section III-C.

A. Experimental setup

We consider the following dynamics models:

- 1) PC – a classical single track vehicle model (2) with Pacejka tire model (3),
- 2) PCR – PC model with learned residual dynamics [6],
- 3) NN – a fully neural network-based vehicle model (4) as proposed in [23],
- 4) NNT – a classical single track vehicle model (2) with neural network-based tire model (5) inspired by [24],
- 5) NNTF – a classical single track vehicle model (2) with the proposed neural network-based tire model scaled by online-estimated friction coefficient (6).

All of these models are used in the UKF framework and are trained to minimize the one-step prediction loss for 1000 epochs. To evaluate the impact of learning through the UKF, as described in Section III-B.4, we fine-tune some models for another 1000 epochs on the 500 steps-long state estimation problems (see Section IV-E), and denote it with subscript $_U$. Validation and testing were conducted using 10s sequences. Moreover, in all cases, we optimize the parameters of the noise models through UKF [26]. All trainings were performed with a learning rate of $5 \cdot 10^{-4}$. For all models, the default noise model is homoscedastic with a full cross-covariance matrix (7), while models with

heteroscedastic noise are denoted by the subscript $_H$ (see Section III-B.3). Neural network tire models employ Multilayer Perceptrons (MLP) that consist of three layers, with each layer comprising 64 neurons, while PCR and NN are using 3 layer MLP with 256 neurons each.

To further extend the set of the baseline methods, we implemented RNN-based state estimators presented in [8] – GRU and [9] – LSTM. We used the same neural networks and training procedure as detailed in these papers. Since GRU and Robust LSTM models require time to stabilize the hidden states for accurate estimations, we adopted the recommendation from [8], and we neglect the first 200 samples (2s, marked with gray in Fig. 4) in both training and evaluation. This technique ensures that the hidden states and estimates of the models are well-formed, allowing a fair comparison.

To obtain a single estimation loss value from all state variables, we weigh the states by the inverse of their variability, calculated as the mean absolute difference of subsequent samples, which, for our dataset, after normalizing to 1 are equal to 0.223, 0.506, 0.157, and 0.114 for v_x, v_y, r , and ω_s respectively. Note, that RNN-based approaches estimate only the lateral, longitudinal, and rotational velocity, which ground true values are provided by a motion capture system. On the other hand, the dynamics of wheel speed is an important aspect in the whole vehicle motion description and must be modeled in our UKF estimator. Thus for evaluation purposes, we set the weight of ω_s to 0.

B. Velocity estimation in known conditions

First, we evaluate the considered methods using data from a single tire type (A), dividing it into training, validation, and test sets with proportions of 70%, 20%, and 10%, respectively. In Tab. I, we present the results obtained on the test set by reporting MSE, Mean Absolute Error (MAE), and 99th percentile of Absolute Error (99%-AE).

Our proposed method NNTF $_{UH}$ achieved the smallest RMS velocity estimation errors (0.062 m/s in the longitudinal and 0.07 m/s in the lateral directions) and significantly outperformed both baseline models [8], [9] in terms of all considered metrics, improving results of GRU by 17%, 6% and 33% for MSE, MAE, and 99%-AE respectively. The results show that each of the introduced components, i.e., online estimated friction coefficient, learning through UKF and heteroscedastic noise model, play their role in achieving state-of-the-art performance. Moreover, one can see that in terms of all considered learning-based dynamics models, learning through UKF improves the estimation accuracy, which shows the importance of adapting the model to the target task. Interestingly, even when tested under the same conditions as in training, RNN-based models exhibit poor worst-case performance - higher 99%-AE than UKF-based state estimators with learned dynamics models.

TABLE I
COMPARISON OF ESTIMATION PERFORMANCE ON DATA FROM THE TIRE USED DURING TRAINING

Method	PC	PCR [6]	NN [23]	NN $_U$	NNT	NNT $_U$	NNTF	NNTF $_U$	NNTF $_{UH}$	GRU [8]	LSTM [9]
MSE	0.0654	0.0089	0.0078	0.0058	0.0088	0.0087	0.0085	0.0072	0.0063	0.0076	0.0338
MAE	0.1459	0.0543	0.0558	0.0464	0.0578	0.0564	0.0574	0.0539	0.0493	0.0524	0.1277
99%-AE	0.5893	0.2512	0.2385	0.1681	0.24	0.2444	0.2184	0.1777	0.1743	0.2608	0.3825

TABLE II

COMPARISON OF ESTIMATION PERFORMANCE FOR THE TIRE UNSEEN DURING TRAINING

Method	PC	PCR[6]	NN[23]	NN _U	NNT	NNT _U	NNTF	NNTF _U	NNTF _{UH}	NNTF _{UH} *	GRU [8]	LSTM [9]
MSE	0.489	0.0769	0.0770	0.0781	-	-	0.0311	0.0259	0.0237	0.0188	0.0967	0.1328
MAE	0.255	0.159	0.154	0.162	-	-	0.105	0.0991	0.0939	0.0810	0.1811	0.2224
99%-AE	2.621	0.758	0.763	0.734	-	-	0.596	0.553	0.504	0.377	0.8735	1.0165

It is worth noting that, in general, the best performance was achieved by a fully neural network-based vehicle model trained through UKF. However, as we will show next, this kind of end-to-end dynamics model generalizes poorly to the data outside of the training distribution.

C. Zero-shot generalization to new road conditions

The velocity estimation for an autonomous vehicle is a safety-critical task that must remain reliable under changing road conditions. Thus, we evaluate the method’s capability for zero-shot generalization in new tire/road friction scenarios. To do so, we use models trained on tire A and evaluate them using the tire C dataset (33% smaller μ). The results of this experiment are presented in Tab. II. Due to the difference in between the learned and actual friction coefficients, the UKF with NNT model becomes unstable, which we denote by ‘-’. For the remaining methods, we observe a notable drop in estimation performance. However, it is the smallest for the proposed NNTF models. These results prove its unparalleled zero-shot generalization abilities to the new road conditions, as NNTF_{UH} achieves 3 and 4 times smaller MAE than NN and GRU respectively. Notably, even better results may be achieved when prior knowledge of μ is available. If we assume that initial friction coefficient estimate $\hat{\mu}$ is equal to 0.4, instead of 0.6, the tracking performance of the proposed method further increases (see NNTF_{UH}* in Tab. II).

In Fig. 4, we present a sample challenging test sequence, which shows the differences in the velocity estimation between the considered solutions and the evolution of the friction coefficient estimate $\hat{\mu}$. One can see that the proposed solution tracks the car velocities very accurately and loses the track only in extremely hard conditions, i.e., high sideslip (5.7s) and simultaneous hard braking and turning (7.5s). Moreover, we observe that despite the wrong initialization of $\hat{\mu}$ it rapidly converges to a region close to the ground truth value of the tire/road friction coefficient μ . The most accurate estimator NNTF_{UH}* achieved RMS velocity errors equal to 0.176m/s and 0.128m/s in longitudinal and lateral velocity, respectively.

D. Impact of the dynamics model on the prediction accuracy

To analyze the sources of improved estimation performance while using dynamics models trained through UKF, we compared a NNTF_U and NNTF on a one-step prediction and UKF state estimation tasks. The results shown in Fig. 5 indicate that the performance in the one-step prediction task achieved by NNTF_U improved slightly for v_x , remained unchanged for v_y , and dropped significantly for r and ω_s . This outcome may be attributed to the critical role of the v_x in the vehicle dynamics and the fact that r and ω_s are directly measurable. Interestingly, in the case of UKF estimation





Fig. 4. Estimation performance on tire C - unseen during training.

Fig. 5. Difference of MSE between NNTF_U and NNTF models in one-step prediction task (blue) and estimation task (red). Positive numbers indicates lower errors of NNTF_U model.Fig. 6. Impact of the dynamics model used in the prediction step (border) and observation model (hatching) on the squared estimation error. Red denotes the NNTF_U, while blue NNTF.

task, using NNTF_U instead of NNTF leads to significantly better estimation of all states. We attribute this result to the simultaneous optimization of dynamics and noise models while training through UKF, which enables the dynamics model to focus on hard to measure states and allows for adequate noise model adaptation.

Moreover, due to the fact that the vehicle dynamic model is a part of both the prediction and update step, we tested which of them is responsible for the gain in estimation performance between NNTF and NNTF_U. To do so, we evaluated UKFs for which the dynamics models used for prediction and




Fig. 7. MSE and average training epoch duration (for batch size 256 using AMD Ryzen 9 5900X) as a function of training sequences length.

update steps may be different and were either NNTF or NNTF_U (see Fig. 6). One can see that from the estimation accuracy perspective, the use of the dynamics model trained through UKF is particularly important in the prediction step and leads to notably smaller errors (18% decrease in MSE).

E. Impact of the training sequence length

One of the main parameters when training dynamics models through UKF is the length of the training sequence. We analyze its impact on the estimation accuracy and training time in Fig. 7. Similarly to [12] we observe that increasing training sequence length leads to decrease in estimation loss but at the same time lengthens the training. In all previous experiments, we choose a sequence length of 500 samples as a compromise between training time and accuracy.

V. CONCLUSIONS

In this paper, we analyzed the impact of learning the dynamics model using UKF on the velocity estimation performance in autonomous racing. We show, in contrast to previous studies [12], that learning the dynamics model through a differentiable filter leads to significantly improved estimation performance. Moreover, we propose to extend the tire forces modeling approach presented in [24] with an online estimated tire-road friction coefficient. This approach not only outperforms the RNN-based state estimators in the nominal conditions but also enables a zero-shot generalization to the significantly lower friction coefficient, which is not achieved by any considered state-of-the-art method. We extended these results with an analysis of the sources of the improved prediction accuracy due to learning through UKF, which we attribute to the possibility of simultaneous adaptation of the dynamics and noise models. Finally, we introduced a publicly available dataset of aggressive maneuvers that covers a significantly larger range of conditions than existing datasets.

REFERENCES

- [1] T. Litman, "Autonomous vehicle implementation predictions: Implications for transport planning," Victoria, BC, CA, 2023.
- [2] A. Wischniewski, *et al.*, "Vehicle dynamics state estimation and localization for high performance race cars," in *IFAC-PapersOnLine*, vol. 52, 07 2019, pp. 154–161.
- [3] N. B. Gosala, *et al.*, "Redundant Perception and State Estimation for Reliable Autonomous Racing," in *2019 Int. Conf. on Robot. and Autom. (ICRA)*, May 2019, pp. 6561–6567.
- [4] D. Kim, *et al.*, "Vehicle sideslip angle estimation using deep ensemble-based adaptive Kalman filter," *Mechanical Syst. and Signal Process.*, vol. 144, p. 106862, Oct. 2020.
- [5] A. Bertipaglia, *et al.*, "An Unscented Kalman Filter-Informed Neural Network for Vehicle Sideslip Angle Estimation," Mar. 2023.
- [6] K. Y. Chee and M. A. Hsieh, "LEARNEST: LEARNing Enhanced Model-based State ESTimation for Robots using Knowledge-based Neural Ordinary Differential Equations," Feb. 2023.
- [7] T. Kim, H. Lee, and W. Lee, "Physics embedded neural network vehicle model and applications in risk-aware autonomous driving using latent features," in *2022 IEEE/RSJ Int. Conf. on Intel. Robots and Syst. (IROS)*, 2022, pp. 4182–4189.
- [8] S. Srinivasan, *et al.*, "End-to-End Velocity Estimation for Autonomous Racing," *IEEE Robot. and Autom. Letters*, vol. 5, no. 4, pp. 6869–6875, Oct. 2020, conf. Name: IEEE Robot. and Autom. Letters.
- [9] A. B. Ghosn, *et al.*, "Robust LSTM-based Vehicle Velocity Observer for Regular and Near-limits Applications," in *2023 IEEE Intel. Vehicles Symposium (IV)*. Anchorage, AK, USA: IEEE, June 2023, pp. 1–6.
- [10] N. Lampe, K.-P. Kortmann, and C. Westerkamp, "Neural Network based Tire-Road Friction Estimation Using Experimental Data," *IFAC-PapersOnLine*, vol. 56, no. 3, pp. 397–402, Jan. 2023.
- [11] T. Haarnoja, *et al.*, "Backprop kf: Learning discriminative deterministic state estimators," in *Adv. in Neural Inf. Process. Syst.*, vol. 29, 2016.
- [12] A. Kloss, G. Martius, and J. Bohg, "How to train your differentiable filter," *Autonomous Robots*, vol. 45, no. 4, pp. 561–578, May 2021.
- [13] X. Liu, *et al.*, "Enhancing State Estimation in Robots: A Data-Driven Approach with Differentiable Ensemble Kalman Filters," Aug. 2023.
- [14] I. Buchnik, *et al.*, "Learned Kalman Filtering in Latent Space with High-Dimensional Data," in *2023 IEEE Int. Conf. on Acoustics, Speech and Signal Process. (ICASSP)*, June 2023, pp. 1–5.
- [15] Y. Bai, *et al.*, "State of art on state estimation: Kalman filter driven by machine learning," *Annual Reviews in Control*, vol. 56, p. 100909, 2023.
- [16] M. O'Kelly, *et al.*, "Fltenth: An open-source evaluation environment for continuous control and reinforcement learning," in *Proceedings of the NeurIPS 2019 Competition and Demonstration Track*, ser. Proceedings of Machine Learning Research, H. J. Escalante and R. Hadsell, Eds., vol. 123. PMLR, 08–14 Dec 2020, pp. 77–89.
- [17] A. L. Escoriza, *et al.*, "Data-Driven Kalman-Based Velocity Estimation for Autonomous Racing," in *2021 IEEE Int. Conf. on Autonomous Syst. (ICAS)*. Montreal, QC, Canada: IEEE, Aug. 2021, pp. 1–5.
- [18] M. D. Lio, M. Piccinini, and F. Biral, "Robust and Sample-Efficient Estimation of Vehicle Lateral Velocity Using Neural Networks With Explainable Structure Informed by Kinematic Principles," *IEEE Trans. on Intel. Transport. Syst.*, vol. 24, no. 12, pp. 13 670–13 684, Dec. 2023.
- [19] F. Djeumou, *et al.*, "Autonomous Drifting with 3 Minutes of Data via Learned Tire Models," in *2023 IEEE Int. Conf. on Robot. and Autom. (ICRA)*, May 2023, pp. 968–974.
- [20] J. Matuško, I. Petrović, and N. Perić, "Neural network based tire/road friction force estimation," *Eng. Applications of Artificial Intelligence*, vol. 21, no. 3, pp. 442–456, Apr. 2008.
- [21] K. F. Löwenstein, *et al.*, "Physics-informed online learning of gray-box models by moving horizon estimation," *European Journal of Control*, vol. 74, p. 100861, Nov. 2023.
- [22] E. Wan and R. Van Der Merwe, "The unscented kalman filter for nonlinear estimation," in *Proceedings of the IEEE 2000 Adaptive Syst. for Signal Process., Communications, and Control Symposium (Cat. No.00EX373)*, 2000, pp. 153–158.
- [23] A. Nagabandi, *et al.*, "Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning," in *2018 IEEE Int. Conf. on Robot. and Autom. (ICRA)*, 2018, p. 7559–7566.
- [24] N. Xu, *et al.*, "Tire force estimation in intelligent tires using machine learning," *IEEE Trans. on Intel. Transportation Syst.*, vol. 23, no. 4, pp. 3565–3574, 2022.
- [25] H. B. Pacejka, *Tire and Vehicle Dynamics*. Butterworth-Heinemann, 2012.
- [26] I. Greenberg, N. Yannay, and S. Mannor, "Optimization or architecture: How to hack kalman filtering," in *Thirty-seventh Conf. on Neural Inf. Process. Syst.*, 2023.
- [27] J. C. Kegelmann, L. K. Harbott, and J. C. Gerdes, "Targa Sixty-Six," 2016, Stanford Digital Repository.
- [28] A. Savitzky and M. J. E. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.