

An Unscented Kalman Filter-Informed Neural Network for Vehicle Sideslip Angle Estimation

Bertipaglia, Alberto; Alirezaei, Mohsen; Happee, Riender; Shyrokau, Barys

DOI

[10.1109/TVT.2024.3389493](https://doi.org/10.1109/TVT.2024.3389493)

Publication date

2024

Document Version

Final published version

Published in

IEEE Transactions on Vehicular Technology

Citation (APA)

Bertipaglia, A., Alirezaei, M., Happee, R., & Shyrokau, B. (2024). An Unscented Kalman Filter-Informed Neural Network for Vehicle Sideslip Angle Estimation. *IEEE Transactions on Vehicular Technology*, 73(9), 12731-12746. <https://doi.org/10.1109/TVT.2024.3389493>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

An Unscented Kalman Filter-Informed Neural Network for Vehicle Sideslip Angle Estimation

Alberto Bertipaglia^{ID}, Graduate Student Member, IEEE, Mohsen Alirezai^{ID}, Riender Happee^{ID},
and Barys Shyrokau^{ID}

Abstract—This paper proposes a novel vehicle sideslip angle estimator, which uses the physical knowledge from an Unscented Kalman Filter (UKF) based on a non-linear single-track vehicle model to enhance the estimation accuracy of a Convolutional Neural Network (CNN). The model-based and data-driven approaches interact mutually, and both use the standard inertial measurement unit and the tyre forces measured by load sensing technology. CNN benefits from the UKF the capacity to leverage the laws of physics. Concurrently, the UKF uses the CNN outputs as sideslip angle pseudo-measurement and adaptive process noise parameters. The back-propagation through time algorithm is applied end-to-end to the CNN and the UKF to employ the mutualistic property. Using a large-scale experimental dataset of 216 manoeuvres containing a great diversity of vehicle behaviours, we demonstrate a significant improvement in the accuracy of the proposed architecture over the current state-of-art hybrid approach combined with model-based and data-driven techniques. In the case that a limited dataset is provided for the training phase, the proposed hybrid approach still guarantees estimation robustness.

Index Terms—State estimation, sideslip angle, physics-informed neural network, unscented Kalman filter, machine learning.

I. INTRODUCTION

ACTIVE vehicle control systems rely on the sideslip angle and yaw rate information to ensure stability and controllability [1], [2]. Whereas low-cost gyro sensors measure the yaw rate, the vehicle sideslip angle must be estimated. Its direct measurement is possible via optical speed sensors or real-time kinematic positioning-global navigation satellite system (RTK-GNSS), but they are expensive to be installed in passenger vehicles [3]. Hence, the development of filter architectures is required to estimate the sideslip angle in real-time and with the desired accuracy error below

Manuscript received 8 March 2023; revised 15 January 2024; accepted 8 April 2024. Date of publication 16 April 2024; date of current version 19 September 2024. This work was supported by The Dutch Science Foundation NWO-TTW research within the EVOLVE under Project nr. 18484. The review of this article was coordinated by Prof. Yan Chen. (Corresponding author: Alberto Bertipaglia.)

Alberto Bertipaglia, Riender Happee, and Barys Shyrokau are with the Department of Cognitive Robotics, Delft University of Technology, 2628 CD Delft, The Netherlands (e-mail: a.bertipaglia@tudelft.nl; r.happee@tudelft.nl; b.shyrokau@tudelft.nl).

Mohsen Alirezai is with the Department of Mechanical Engineering, University of Eindhoven, 5612 AZ Eindhoven, The Netherlands, and also with the Digital Industry-Software-Simulation and Testing Services, 5708 JZ Helmond, The Netherlands (e-mail: m.alirezai@tue.nl).

Digital Object Identifier 10.1109/TVT.2024.3389493

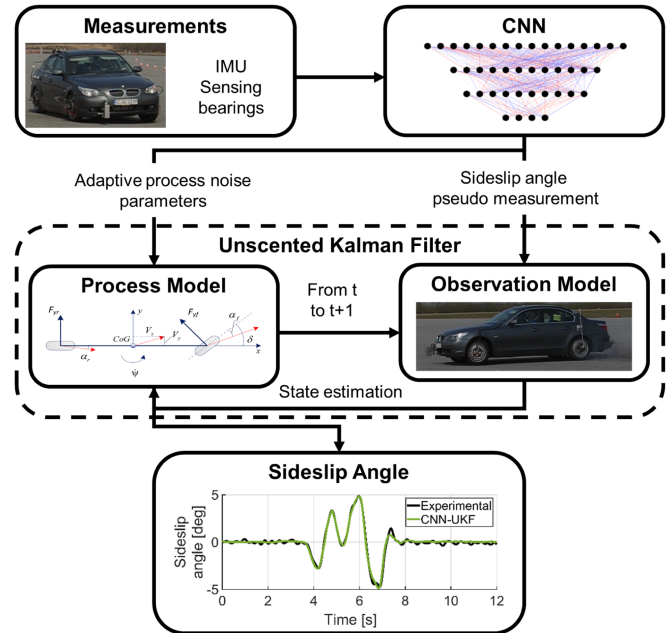


Fig. 1. Framework overview of the CNN-UKF approach. A CNN provides a sideslip angle pseudo-measurement and the process noise parameters to a UKF based on a single-track vehicle model. The UKF monitors and weights the CNN's estimation through physical laws.

one degree in high excitation driving conditions [4]. Sideslip angle estimation is particularly challenging for the following aspects:

- A large diversity of vehicle manoeuvres, e.g. steady-state, transient, low, and high excitation.
- The highly non-linear behaviour of tyres leads to a substantial limitation due to tyre model accuracy.
- Data collection requires expensive and high calibration sensitive instruments.
- Numerous external disturbances, e.g. bank angle, road slope, and road friction coefficient.

Several approaches have been proposed for vehicle sideslip angle estimation [5], [6]. They are split into three main groups, i.e. model-based, data-driven and hybrid approaches. The model-based approach relies on the physical knowledge of a vehicle model for state estimation. Open-loop deterministic models are insufficient to provide an accurate estimation, so stochastic closed-loop observers, e.g. extended Kalman filter

(EKF), unscented Kalman filter (UKF), and particle filters are currently applied to estimate unmeasurable states. EKF and UKF are the industrial state-of-the-art for vehicle sideslip angle estimation because their accuracy can be guaranteed in a specific operating region, and their properties are easily assessed [7]. However, they both struggle in transient and high excitation driving conditions due to the growing non-linearities in the vehicle model [8]. Another branch of model-based approaches is based on deterministic observers adapted to deal with stochastic noise, e.g. Luenberger observer, the sliding-mode observer and H-infinity methods. Despite their advantage in computational complexities, their need for a higher model fidelity makes them less utilised in industry [9], [10]. Nevertheless, these model-based approaches require extensive system knowledge. The data-driven approach has higher accuracy than the model-based approach when enough quality data are provided in the training phase [8]. Different neural network (NN) architectures have been proposed, e.g. feed forward neural network (FFNN) and recurrent neural network (RNN). However, they all lack interpretability and generalisation capabilities. Thus, a purely data-driven approach is hardly applicable for safety applications in the automotive domain [7]. The third approach, named hybrid, combines the pros of the model-based and data-driven approaches. It improves the model-based accuracy thanks to the NN outputs and, simultaneously, gives the data-driven approach an interpretability thanks to the vehicle model. In the proposed hybrid architectures [7], [11], [12], the model and the neural network work in a unidirectional way. Thus, the model in the hybrid architecture relies on the NN knowledge without backward communication, reducing the approach's potential.

This paper proposes a new hybrid approach for vehicle sideslip angle estimation. Its novelty is in the mutualistic relationship between the model-based approach, characterised by a UKF based on a single-track model, and the data-driven approach, represented by a Convolutional Neural Network (CNN). The proposed approach consists of a sequential hybrid architecture in which the CNN passes the pseudo-measurement of the sideslip angle, the level of distrust of its estimation and the process noise parameters of the vehicle model to the UKF, see Fig. 1. A key aspect of the proposed hybrid approach is the training process, which allows the development of a physics-informed NN [13]. Considering that the non-linear vehicle dynamics are described in a UKF, the physics-informed NN will also be referenced as a UKF-informed NN. The training is end-to-end, so the Back-Propagation Through Time (BPTT) algorithm moves through the CNN, the UKF and backwards. Thus, the CNN is constrained to respect the physical laws of vehicle dynamics. Furthermore, it allows the CNN to estimate variables for which a reference is unavailable, i.e. the process noise parameters and pseudo-measurement level of distrust. This will lead to a high estimation accuracy compared to the state-of-the-art hybrid approaches, which always separate the data-driven component from the model-based one during the training [11], [12], [14], [15], [16]. The split proposed in the literature makes the hybrid approach training lighter from a memory and computational point of view. However, it does not

TABLE I
OVERVIEW OF THE VEHICLE SIDESLIP ANGLE ESTIMATION APPROACHES

Approach	Features & Authors
Model-based	Kalman filter based on a kinematic model [17]–[19]
	Kalman filter based on IMU & GNSS [20]–[26]
	EKF based on a dynamic model [27]–[33]
	UKF based on a dynamic model [34]–[39]
	Sliding mode observer [40]–[42]
	H-infinity observer [43], [44]
	Luenberger observer [45]
Data-driven	Hybrid - dynamic & kinematic models [3], [46]–[48]
	Online gradient descent [49]
	Modular estimation scheme [50], [51]
	FFNN [52], [53]
	RNN [54]–[56]
Hybrid	ANFIS [14]
	Kernel-based LPV [57]
	FFNN, ANFIS & UKF [15], [58], [59]
	RNN (GRU) & UKF [16]
	Differentiable EKF [60]
	Kalman filter & FFNN [61]
	Piecewise Affine & Takagy-Sugeno [62]
	FFNN & Kalman in the back-propagation [63]
	KalmanNet [64]
	Kinematic model & RNN (GRU) [7]
	Deep Ensemble Network (LSTM) & EKF [11], UKF [12]

allow the model-based approach to understand when it can rely on the CNN and at the same time, it does not allow the CNN to learn the physical laws of vehicle dynamics. The performance is evaluated using a large-scale, real-world experimental dataset. The dataset contains a great diversity of driving situations, recorded with a constant high friction coefficient.

The paper is organised as follows. Section II contains a summary of the previous works and the main paper contributions. Section III describes the CNN and UKF used in the proposed hybrid approach. Section IV describes how the experiments are conducted and evaluated. Section V reports the obtained results, and Section VI summarises the conclusions and future research paths.

II. RELATED WORKS

A summary of the three categories, i.e. *model-based*, *data-driven* and *hybrid*, is presented in Table I.

The first approach is called *model-based* and relies on the laws of physics. The vehicle behaviour can be described using the geometric constraints, i.e. kinematic model, or considering the forces and moments acting on the vehicle, i.e. dynamic model. The kinematic model requires only geometrical parameters and does not need extensive vehicle parametrisation because its reliability depends mainly on sensing capabilities. The state-of-art kinematic observer [19] is based on a linear parameter varying system, where the states are the vehicle velocities, and the accelerations are the inputs. This approach leads to high accuracy in transient manoeuvres, but the model is not observable in nearly steady-state conditions [3]. Hence to avoid unobservability, a heuristic function is applied to lead the lateral velocity to zero when the vehicle is moving straight or nearly straight [19]. The downside is the amount of data necessary to define the heuristic function. Moreover, despite the performance improvement, it is still susceptible to integration

error and sensor drifting. Thus, in recent publications [20], [21], [22], [23], the measurements from the Inertial Measurement Unit (IMU) are coupled with those from a Global Navigation Satellite System (GNSS) to increase the amount of information available for the estimator. The velocities measured by the GNSS are integrated into an estimation-prediction framework, which estimates the sideslip angle and partially compensates for the error induced by the low GNSS sampling rate [20]. However, GNSS/IMU fusion kinematic approach still suffers from the low GNSS sampling rate. Furthermore, a high-precision GNSS is too expensive as the standard sensor in passenger vehicles, and signal reception cannot always be assured. Therefore, it is mainly applied to racing [24]. Thus, a solution is to consider dynamic models to rely less on the sensor signal quality. Dynamic models allow a more robust noise computation of the accelerations than kinematic models [3]. However, dynamic models require a more profound knowledge of vehicle parameters and the presence of a tyre model, which is a critical source of uncertainty [39]. EKF and UKF are the state-of-art estimation techniques for the model-based approach, and the process and the observation noises are commonly assumed to be Gaussian and uncorrelated. The EKF uses a first-order Taylor series expansion to linearise around the current mean and covariance. It has excellent accuracy in nearly steady-state conditions and when the vehicle behaves close to linearity, i.e. up to a lateral acceleration of approximately 4 m/s^2 [8]. When the vehicle behaves with strong non-linearities, UKF assures a better estimation accuracy because it linearises up to the third order of the Taylor series expansion [8]. However, both observers suffer from the mismatches between the physical and modelled tyre behaviour. A possible solution is to combine the pros of dynamic and kinematic models to develop a hybrid kinematic-dynamic observer [46], [47]. This family combines the accuracy in transient manoeuvres of the kinematic models and the better robustness to sensor noise of the dynamic models. The kinematic and the dynamic filters work simultaneously, and the final sideslip angle estimation is a weighted average of the two approaches. The weights are chosen according to the lateral acceleration signal [47]. However, the weighting coefficients' tuning is complex, and the optimum solution varies according to the considered manoeuvres. Another solution to combine dynamic and kinematic models is the development of a modular scheme to estimate in sequential steps tyre forces, longitudinal and lateral velocities [50]. It consists of monitoring the wheel capacities under combined slip at each vehicle corner to estimate the individual forces and velocities. The approach is experimentally validated in different road conditions, but the results do not show its performance when the vehicle is driven at the limit of handling. Thus, the approach applicability to evasive manoeuvres is limited.

Despite EKF and UKF being the most implemented filters for vehicle sideslip angle estimation [10], also different kinds of observers are proposed in the literature for their advantages in computational complexity and theoretical guarantees of convergence [9], e.g. sliding mode observer [40], [41], [42], H-infinity [43], [44], state-dependent Riccati equations [32] and Luenberger observer [45]. Particularly interesting is the combination of an adaptive sliding mode observer to estimate

the lateral tyre force with an adaptive compensation algorithm to estimate the sideslip angle [44]. Despite the improved performance compared with EKF in an experimental scenario, there is no comparison with UKF, which is the state-of-the-art in extreme driving conditions. Moreover, these observers fail to perform well in evasive manoeuvres due to increased model mismatches, especially in tyre forces [65].

A solution to enhance the state estimation robustness to tyre model inaccuracies of dynamic model is the introduction of adaptive tyre models [27], [58] or new proprioceptive load-sensing technology, e.g. intelligent bearings or smart tyres [39], [66]. The Kalman filter can use tyre force measurements as an additional feedback to improve the estimation and magnify the Kalman gain, especially in the case of non-linear vehicle behaviour. The enhanced vehicle safety and the sensor's cost efficiency (lower than 1000 € per vehicle) make the innovative load-sensing technologies candidate to become part of the standard sensor setup for passenger vehicles [39].

A *data-driven* approach reduces extensive requirements of system knowledge compared to the model-based approach. A deep NN with eight hidden layers, each having a different number of long short-term memory (LSTM) cells, is proposed [55]. Despite the increased training time of such a deep NN, the authors state that a smaller network was incapable of reaching the level of accuracy of deeper NN. The issue is that deep NNs are prone to overfit, and their performance strongly lacks generalisation capabilities. To overtake this issue, an NN classification is applied to choose which available NN is most suitable for specific road conditions [54]. Each of the three FFNNs is built with a single hidden layer, and they are trained with three different datasets corresponding to three different road friction conditions, i.e. dry, wet and icy. Moreover, the performance of data-driven approach can be enhanced by the availability of tyre force measurements [8]. In this case, a FFNN with two hidden layers outperforms the accuracy of a more complex RNN architecture based on LSTM cells. A FFNN also exceeds the performance of various model-based approaches, even if it tends to sporadic higher maximum error. However, the data-driven performance is highly dependent on the amount of representative data, and the data-driven approach will lack performance as soon as the dataset contains a lower amount of data in a particular range of the sideslip angle.

Although the data-driven approach generally has a better estimation performance than the model-based approach, it is impossible to guarantee robust performance over vehicle operating conditions. Conversely, a model-based approach based on a dynamic model with tyre force measurements has lower accuracy, but its performance is consistent over the working region [8]. Thus, a *hybrid* model-based and data-driven approach is proposed. We employ two leading typologies, model-to-NN and NN-to-model, as explained below.

The model-to-NN family aims to augment the number of the NN's inputs using the output of a vehicle model. This will transfer some immeasurable physical states to a NN. A kinematic vehicle model can compute the derivative of the sideslip angle, which is used as extra input for the following RNN based on a Gated Recurrent Unit (GRU) cell [7]. The kinematic model

provides the NN with a pseudo-measurement that contains significant errors, biases and drift. However, with the extra vehicle model information, the NN reduces the sideslip angle's Mean Squared Error (MSE) of the non-informed NN by 2.7%, 5.6%, and 1.2%, respectively, for dry, wet and snow conditions [7]. The slight improvement shows the benefits of developing a hybrid approach and highlights the importance of providing a more accurate pseudo-measurement.

The NN-to-model family aims to provide a sideslip angle pseudo-measurement to the following EKF/UKF. In this case, the NN output is post-processed by a Kalman filter to improve the sideslip angle estimation. One of the first approaches combines an Adaptive Neuro-Fuzzy Inference System (ANFIS) [14] with a UKF to estimate sideslip angle. The model-based component is employed as a filter to minimise the noise of the NN output and the variance of the estimation mean square error [15]. The ANFIS is trained using synthetic data generated through a high-fidelity simulation environment (CarSim). The ANFIS-UKF improves the performance of only an ANFIS [14] by 21% on average for five manoeuvres with high friction conditions. However, the presented figures show a maximum value of sideslip angle of only 3deg in absolute value, which makes the estimation performance easier than in extreme driving conditions. Furthermore, there is no explanation of how the observation noise parameter related to the pseudo-sideslip angle is tuned. This value is essential because it defines the level of distrust the UKF can give to the NN output. A similar approach involves integrating an FFNN with a UKF based on a kinematic vehicle model [59]. Contrary to the ANFIS-UKF approach, the model-based component of the hybrid approach is responsible for filtering the estimation noise and correcting the NN output. This is possible thanks to a proportional feedback correction, which improves the performance of the pseudo-sideslip angle. Although the NN is trained using synthetic measurements, the approach is validated using experimental data and the presented results. Unfortunately, all the results are normalised, so it is impossible to understand if the vehicle was driven at the limit of handling. The presented approach improves the data-driven approach estimation accuracy of 73.3%. However, there needs to be an analysis of how to decide the distrust level of the pseudo-sideslip angle. Otherwise, when the NN is uncertain due to a lack of data, it will negatively influence the UKF's performance. Furthermore, the kinematic vehicle model is highly susceptible to measurement noise and is not observable in steady-state driving. For this reason, the FFNN is substituted by a deep ensemble NN in more recent publications [11], [12].

Deep Ensemble (DE) of RNN, based on LSTM cells, estimates a sideslip angle pseudo-measurement and its level of distrust, which are then provided as extra measurements to a UKF [11], [12]. The level of distrust is modified through a user-defined linear function before being used by the UKF. This further step is mandatory to scale the NN's distrust level to a meaningful value for the UKF. This hybrid architecture reduces the Root Mean Squared Error (RMSE) by, on average, 8% vs the RNNs [11]. The extra tuning of the level of distrust can easily lead the approach to overfit. Moreover, the level of distrust is

computed through the standard deviation of the sideslip angle pseudo-measurements estimated by the RNNs. This does not lead to a physics-informed NN, so it is still complex to assess the properties of this hybrid approach. The reason is that the DE-RNN is not aware of the performance of the UKF, so the estimated level of distrust is not scaled according to the UKF's accuracy. Vice-versa, a physics-informed NN learns the Kalman filter's precision during the training, providing the best level of distrust to maximise the hybrid sideslip angle estimation.

This work proposes a hybrid approach employing a mutual relationship between the model-based and data-driven approaches for vehicle sideslip angle estimation. The inputs and outputs of the NN are, respectively, inputs and measurements of the UKF. The end-to-end training enforces the mutual relationship, meaning that the back-propagation algorithm passes through the NN, UKF and vice-versa. The UKF benefits from the CNN when the mismatch in the UKF process model is particularly significant, for instance, due to the high non-linearity in the tire model or the non-modelled phenomena. Thus, the CNN provides a vehicle sideslip angle pseudo-measurement to the UKF, which guides it towards an accurate estimation. At the same time, a purely model-based approach, unable to rely on any extra information, would only face a decrease in the estimation accuracy. On the other hand, the CNN gains from the UKF a physics domain knowledge, which helps CNN improve its robustness and accuracy in all the real-life situations that were not fully described by the data used during the training. This is particularly relevant in emergencies that cannot be adequately represented in a vehicle sideslip angle estimation dataset.

The main contribution of this paper is threefold.¹ The first is a mutual hybrid approach in which the CNN is trained end-to-end with the UKF, developing an innovative stochastic, deterministic state estimation method [63] for vehicle sideslip angle estimation. Thus, the UKF observation model has access to CNN's deterministic estimation of a pseudo-measurement of the sideslip angle and its level of distrust. On the other hand, CNN is informed by the following stochastic state estimation filter during the training, so it learns how to estimate the level of distrust independently without requiring extra tuning after its training, as previously proposed in the literature [11], [12]. Hence, the proposed hybrid approach enhances the accuracy (MSE) of the state-of-the-art model-based, data-driven, and hybrid approach for vehicle sideslip angle estimation.

The second contribution is that the proposed hybrid architecture considers the hetero-stochasticity of the model dynamics [63], [67], [68]. Thus, the proposed approach estimates the UKF process model uncertainties online thanks to the CNN and UKF end-to-end training, which helps improve the estimation accuracy. Moreover, it provides higher robustness than the state-of-art, even when a limited dataset is used for the training.

The third contribution is that the proposed hybrid architecture is a UKF-informed NN, which means that the NN incorporates the domain knowledge described in the UKF, and it complies with the vehicle dynamics laws of physics. Thus, the proposed

¹The code for our method will be released upon paper acceptance.

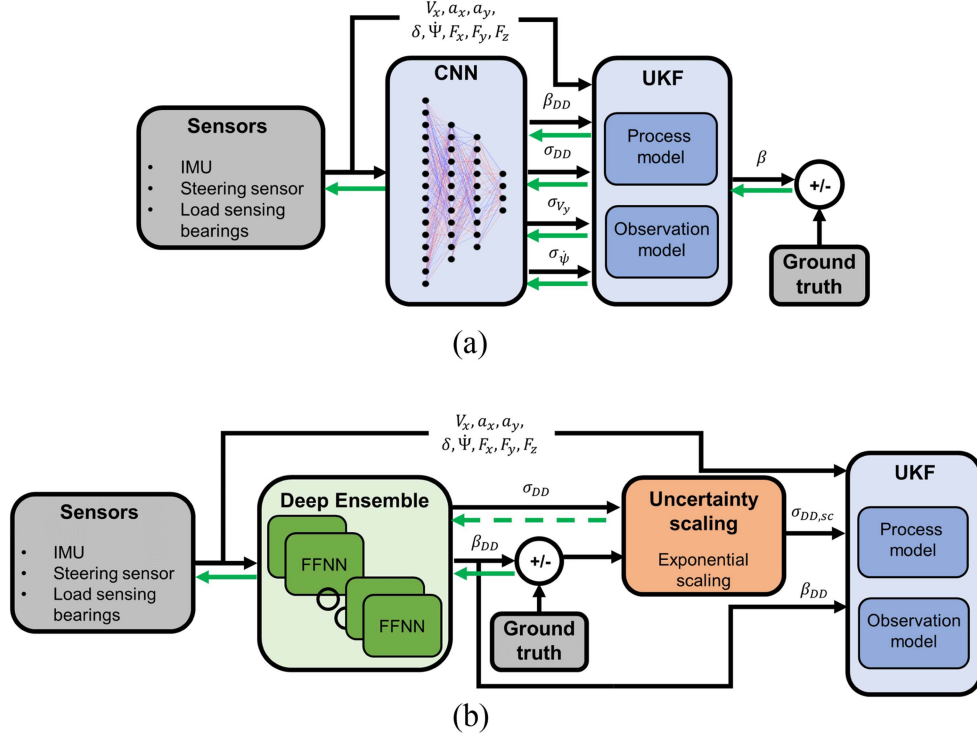


Fig. 2. (a) Shows the proposed hybrid approach architecture (CNN-UKF). (b) Shows the baseline hybrid architecture (DE-UKF) proposed by [11]. The black arrows show the flow of information during the online estimation, while the green arrows show the flow of information during the back-propagation. The dashed green arrow represents a term used in the cost function computation but not used by the optimiser to update the NN weights.

approach has a lower Maximum Error (ME) than a state-of-art model-based [39] and data-driven [8] approach, as well as the state-of-art hybrid approach [11].

III. UKF-INFORMED NEURAL NETWORK

This section describes the proposed hybrid approach based on a CNN end-to-end trained with a UKF (CNN-UKF). A comparison between the proposed mutualistic hybrid approach and the hybrid unidirectional baseline [11] is represented in Fig. 2(a) and (b), respectively. The proposed approach develops a UKF-informed NN, where the NN is constrained to respect the vehicle dynamics. At the same time, the baseline (DE-UKF) corresponds to a UKF augmented by the DE outputs. The approach's discretisation is performed through a zero-order hold method [50] due to its good trade-off between simplicity and accuracy. The discretisation works at 100 Hz, the standard frequency for vehicle state estimation.

A. Data-Driven Component

A straightforward CNN can cope with the complexity of the task because the approach's strength is inside the hybrid architecture. It consists of an input layer, two hidden layers and an output layer.

Seventeen measurements form the input layer (x): longitudinal and lateral accelerations a_x and a_y respectively, longitudinal

velocity V_x , road wheel angle δ , yaw rate $\dot{\psi}$, and longitudinal, lateral and vertical tyre forces for each of the four wheels, respectively F_x , F_y and F_z . Before being used, the input measurements are normalised because each input has a different physical meaning and order of magnitude. Thus, all the inputs are mapped onto the interval $[0, 1]$ to speed up and stabilise the training process [69]. A different normalisation method which scales the data to a mean of zero and a standard deviation of one has been tested. Still, the mapping onto the interval $[0, 1]$ produced the best results after the training.

The two hidden layers consist of 200 and 100 neurons and Rectified Linear Unit (ReLU) activation functions. The hidden layers are 2D convolutions with kernel sizes 1×1 , 0 padding, stride equal to 1 and active bias. The CNN uses a dropout regularisation technique equal to 0.2 and a Xavier weight initialisation to avoid overfitting.

The output layer is formed by four neurons corresponding to the pseudo-measurement of the sideslip angle β_{DD} , the level of distrust in the pseudo-measurement σ_{DD} , the uncertainty of the UKF process model lateral velocity σ_{V_y} and the uncertainty of the UKF process model $\sigma_{\dot{\psi}}$. A reference is available only for β_{DD} , but the other three outputs strongly affect the estimation of the model-based component, which is used in the training loss function; see Section III-C for further details. Thus, all four CNN outputs are correctly trained during the end-to-end training. σ_{V_y} , $\sigma_{\dot{\psi}}$ and σ_{DD} are further processed with a sigmoid function to constrain their values inside the meaningful interval $[0, 1]$. This




Fig. 3. Single-track vehicle model.

last step assures that the CNN can produce uncertainties which not lead to UKF failure.

B. Model-Based Component

A UKF based on a non-linear single-track model with tyre axle forces computed by the Dugoff tyre is chosen as the model component of this study [8], see Fig. 3. The Dugoff tyre is adapted to only assume pure lateral slip condition, reducing the computational complexity [8], [28], [31]. The tyre model parameters (p_t), i.e. tyre cornering stiffness, peak friction coefficient and velocity reduction friction coefficient, are optimised offline using experimental skidpad measurements [34], [36], [39]. The implemented optimisation is a genetic algorithm due to its efficiency with a non-linear and non-convex cost function. The vehicle's symmetry is exploited to merge the left and right wheels into a single central axle, which emulates the entire vehicle's behaviour. This model considers only the in-plane dynamics, so the lateral weight transfer, roll and pitch dynamics are ignored. The static weight distribution is considered together with the effect of steady-state longitudinal weight transfer concerning the normal forces on the front and rear axle. A UKF is implemented for its superior estimation accuracy when the vehicle behaves strongly non-linearly. The vehicle states (x_s) are the V_y and the $\dot{\psi}$, while the vehicle inputs (u_v) are the V_x and the δ . The stochastic process model is responsible for predicting the next time steps of the states according to the following equation:

$$\dot{x}_s(t) = f(x_s(t), u_v(t), p_t) + \omega(t) \quad (1)$$

where $f(x_s(t), u_v(t), p_t)$ is the non-linear single track vehicle model, (2), and ω is the vector containing the process noise parameters $[\sigma_{V_y}, \sigma_{\dot{\psi}}]$.

$$f(x_s, u_v, p_t)$$

$$= \begin{cases} \dot{V}_y = \frac{1}{m} (F_{yf}(x_s, u_v, p_t) \cos(\delta) + F_{yr}(x_s, u_v, p_t)) + \\ -V_x \dot{\psi} \\ \ddot{\psi} = \frac{1}{I_{zz}} (l_f F_{yf}(x_s, u_v, p_t) \cos(\delta) - l_r F_{yr}(x_s, u_v, p_t)) \end{cases} \quad (2)$$

where m (1970 kg) is the vehicle mass, I_{zz} (3498 kg m²) is the vehicle moment of inertia about the vertical axis, l_f (1.47 m) and l_r (1.41 m) are, respectively, the distance of front and rear axles from the vehicle CoG. F_{yf} and F_{yr} are, respectively, the lateral tyre forces at the front and rear axles. The process noise

parameters, σ_{V_y} and $\sigma_{\dot{\psi}}$, are assumed Gaussian and uncorrelated and they capture the uncertainties due to:

- The mismatch between the physical and modelled vehicle behaviour.
- The discretisation error.
- The various operational environments in which the sensors operate.

The filter performance is strongly connected with the process noise parameters, so these are initially tuned using a two-stage Bayesian optimisation (TSBO) [70]. During the estimation, they are computed online by CNN. This is only possible thanks to the mutualistic relationship between CNN and the UKF.

The observation model is responsible for comparing the process model predictions with the available measurements, according to the following equation.

$$y_m(t) = g(x_s(t), u_v(t), p_t) + v(t) \quad (3)$$

where $g(x_s(t), u_v(t), p_t)$ is the measurement vehicle model, (4), and v is the vector containing the observation noise parameters $[\sigma_{a_{yme}}, \sigma_{\dot{\psi}_{me}}, \sigma_{F_{yfme}}, \sigma_{F_{yrme}}, \sigma_{DD}]$.

$$g(x_s, u_v, p_t)$$

$$= \begin{cases} a_{yme} = \frac{1}{m} (F_{yf}(x_s, u_v, p_t) \cos(\delta) + F_{yr}(x_s, u_v, p_t)) \\ \dot{\psi}_{me} = \dot{\psi} \\ F_{yfme} = F_{yf}(x_s, u_v, p_t) \\ F_{yrme} = F_{yr}(x_s, u_v, p_t) \\ \beta_{DD} = \text{atan}\left(\frac{V_y}{V_x}\right) \end{cases} \quad (4)$$

where a_{yme} , $\dot{\psi}_{me}$, F_{yfme} and F_{yrme} are the vehicle measurements, and β_{DD} is the pseudo-measurement, corresponding to the CNN's output. The observation noise parameters $\sigma_{a_{yme}}$ (0.033 m/s²), $\sigma_{\dot{\psi}_{me}}$ (0.001 rad/s), $\sigma_{F_{yfme}}$ (26N) and $\sigma_{F_{yrme}}$ (56N) are the uncertainties of the vehicle measurements and they compensate for the sensor noises. They are tuned by a statistical analysis of the vehicle sensor measurements, which consists of computing the standard deviation of the low-pass measured signal when the steering angle is null and the longitudinal velocity is constant [70]. The variable σ_{DD} is the level of distrust assigned to the pseudo-measurement β_{DD} provided by CNN. The level of distrust computed by the CNN differs from a classic uncertainty measurement because it corresponds to the uncertainty of the pseudo-measurement scaled to match the weight of the noise parameters.

The observability analysis is performed to assess under which conditions it is possible to infer the internal states given the vehicle inputs and measurement. Being the mode highly non-linear, only the local observability around an operating point can be computed, performing a system linearisation [34]. The observation matrix is built on the Jacobian matrices of the process and observation models, and it is full rank, equal to two, for all the operating regions in which $\delta \neq k\pi$, $\forall k \in \mathbf{Z}$ and $V_x \neq 0$, where \mathbf{Z} is the integers set. The second condition is always respected because the measurement is considered when V_x is higher than 5 m/s. The steering angle is always inside the range $|\delta| \leq \pi/2$, so the only realistic unobservability happens

when $\delta = 0$. However, the vehicle sideslip angle is relevant for lateral dynamics, so it only happens when $\delta \neq 0$.

C. Training Phase

The UKF-informed CNN is trained in a supervised way using a labelled dataset. The training is split into two phases: pre-training and end-to-end learning.

1) *Pre-Training*: It consists of the back-propagation algorithm applied only to the CNN to speed up and stabilise the following end-to-end training phase. The pre-training loss function is constituted by the sum of σ_{V_y} , $\sigma_{\dot{\psi}}$ MSE losses and by the β_{DD} , σ_{DD} Gaussian negative log-likelihood loss. The MSE loss functions ($MSE_{L, \sigma_{V_y}}$ and $MSE_{L, \sigma_{\dot{\psi}}}$) are represented as:

$$\begin{aligned} MSE_{L, \sigma_{V_y}} &= \frac{1}{N} \sum_{i=1}^N (\hat{\sigma}_{V_y, i} - \sigma_{V_y, i})^2 \\ MSE_{L, \sigma_{\dot{\psi}}} &= \frac{1}{N} \sum_{i=1}^N (\hat{\sigma}_{\dot{\psi}, i} - \sigma_{\dot{\psi}, i})^2 \end{aligned} \quad (5)$$

where N is the mini-batch size (256), $\hat{\sigma}_{V_y}$ (0.0007 m/s) and $\hat{\sigma}_{\dot{\psi}}$ (0.002 rad/s) are the initial process model uncertainties tuned by the TSBO for the model-based approach. These losses steer the CNN to predict the process model uncertainties with a meaningful order of magnitude. For σ_{DD} , the Gaussian negative log-likelihood loss function ($NLL_{L, \beta_{DD}}$) represented in the following equation is chosen:

$$\begin{aligned} NLL_{L, \beta_{DD}}(\beta_{DD}, \sigma_{DD}^2 | \beta_{me}) \\ = \frac{1}{2} \sum_{i=1}^N \left(\log(\max(\sigma_{DD, i}, \epsilon)) + \frac{(\beta_{me, i} - \beta_{DD, i})^2}{\max(\sigma_{DD, i}, \epsilon)} \right) \end{aligned} \quad (6)$$

where ϵ (10^{-6}) is a constant term for stability, and β_{me} is the sideslip angle ground truth. Thus, this loss function adjusts CNN's weight to maximise the likelihood of the observed data, performing negative log-likelihood minimisation. This process naturally leads to the CNN learning to predict both the mean (β_{DD}) and the variance (σ_{DD}^2). By predicting the variance alongside the mean, the pre-training leads the CNN to provide a point estimate and a measure of confidence or uncertainty associated with each prediction. Thus, it assures a meaningful σ_{DD} , even without its ground truth.

The sideslip angle ground truth is measured through the Corrsys-Datron optical speed sensor installed in the vehicle's front bumper. The sensor reference system is moved to correspond with the vehicle CoG. The measurement is filtered using a zero-phase low-pass filter (bandwidth 5 Hz) because the training phase is sensitive to extreme outliers or noisy references [6]. The cost function is minimised by a mini-batch stochastic gradient descent algorithm based on a standard ADAM optimiser with a learning rate (0.0008). The training procedures' user-defined parameters are optimised through a Bayesian optimisation.

2) *End-to-End Learning*: The back-propagation through time (BPTT) algorithm is applied end-to-end to the CNN and

UKF. This step creates a mutualistic relationship between the model-based and data-driven approaches. Thus, the UKF is treated as a computation graph unrolled through time, so the CNN-UKF is discriminatively trained over the entire mini-batch length and not on a single step. The procedure to compute the loss function gradient is close to [63], but in the proposed study, a UKF is implemented rather than a linear Kalman filter. The first step is the computation of a loss of a function ($L(\theta)$) that connects the output of the UKF-CNN (β) structure with the available ground truth (β_{me}), given the CNN weights (θ). The training phase minimises the loss function error between the estimated and the measured sideslip angle, allowing to correctly estimate β and all the variables influencing it, so σ_{DD} , $\sigma_{\dot{\psi}}$ and σ_{V_y} . The loss function depends on the CNN's weights (θ) and is based on the following equation:

$$L(\theta) = \frac{1}{N} \sum_{i=1}^N (\beta_{me, i} - \beta_{DD, i})^2 + \frac{1}{N} \sum_{i=1}^N (\beta_{me, i} - \beta_i)^2 \quad (7)$$

where β_{DD} is the output of the CNN, and β_{me} is the sideslip angle ground truth. The first loss function part, $\frac{1}{N} \sum_{i=1}^N (\beta_{me, i} - \beta_{DD, i})^2$, helps the CNN to estimate the correct pseudo-measurement β_{DD} . The second part of the loss function, $\frac{1}{N} \sum_{i=1}^N (\beta_{me, i} - \beta_i)^2$, is affected by all four CNN outputs and the UKF. The following step to train the proposed UKF-CNN is the computation of the loss function's gradient with respect to the CNN weights ($\nabla_{\theta} L(\theta)$). This is performed following the typical BPTT algorithm. Moving backwards from the loss function, the $\nabla_{\theta} L(\theta)$ is computed by a recursive computation of the loss function gradient with respect to the vehicle states from $t-1$ to t according to:

$$\frac{\partial L}{\partial x_{s, t-1}} = \frac{\partial ukf_{t-1}}{\partial x_{s, t-1}} \frac{\partial L}{\partial ukf_{t-1}} + \frac{\partial x_{s, t}}{\partial x_{s, t-1}} \frac{\partial L}{\partial x_{s, t}} \quad (8)$$

where ukf_{t-1} represents all the functions that describe the UKF algorithm, i.e. process model, observation model and Kalman gain computation. The UKF process and observations model, see (2) and (7) respectively, are non-linear models with a non-differentiable point only in the tyre force model. Thus, the Dugoff tyre model, described in (9), has been modified to be fully differentiable according to (10).

$$\begin{aligned} \mu &= \mu_0 \left(1 - e_r V_x \sqrt{\kappa^2 + \tan(\alpha)^2} \right) \\ \lambda &= \frac{\mu F_z (1 - \kappa)}{2 \sqrt{(C_x \kappa)^2 + (C_y \tan(\alpha))^2}} \\ f_{\lambda} &= \begin{cases} \lambda (2 - \lambda), & \text{if } \lambda \leq 1 \\ 1 & \end{cases} \\ F_y &= \frac{C_y \tan(\alpha) f_{\lambda}}{1 - \kappa} \end{aligned} \quad (9)$$

$$f_{\lambda} = \frac{2}{e^{-4\lambda} + 1} - 1 \quad (10)$$

where μ_0 is the peak friction coefficient, e_r is the friction reduction coefficient, V_x is the longitudinal velocity, κ and α are the longitudinal and lateral slip, C_x is the longitudinal slip stiffness, and C_y is the cornering stiffness. Regarding all the UKF operations, they are differentiable and available in the open-source machine learning platform TensorFlow or PyTorch. The gradient computation continues applying the chain rule to (8) and moving backwards, computing the derivative to each CNN weight as for a normal NN. This step is performed automatically by the chosen machine learning platform, PyTorch. The training is based on a mini-batch stochastic gradient descent algorithm (mini-batch size of 256) based on a standard ADAM optimiser with a learning rate equal to 0.0008.

IV. EXPERIMENT SETUP

This section describes how the experiments have been conducted and how the proposed approach has been compared to the baseline methods.

A. Experimental Setup and Dataset

The experiments have been conducted at the Automotive Testing Papenburg GmbH with the test platform based on a BMW Series 545i. The test vehicle was instrumented with the standard IMU, Kistler wheel force transducers and SKF intelligent bearings for each wheel, a dual antenna GNSS from Oxford Technical Solutions and a Corrsys-Datron non-contact optical sensor to measure the sideslip angle (measurement accuracy of $\pm 0.2^\circ$). The high-end optical speed sensor was used to measure the ground truth of the vehicle sideslip angle. The vehicle was equipped with a dSPACE 1007 AutoBox as a real-time control platform. All the equipment was interconnected through the Controller Area Network (CAN) interface, and the sampling rate was set up at 100 Hz. The intelligent bearings demonstrate a similar accuracy to the wheel force transducer [71], the most common sensor technique in research for tyre force measurement. Thus, the tyre forces in the training dataset are taken from the wheel force transducers, making the paper easier to reproduce. The dataset contains 216 manoeuvres corresponding to two hours of driving and consists of standard vehicle dynamics manoeuvres, e.g. double lane change, slalom, random steer, J-turn, spiral, braking in the turn, and steady-state circular tests, together with recorded laps at the handling track. All manoeuvres were driven on dry asphalt with tyres inflated according to the manufacturer's specifications. The bank angle and the road slope were negligible, and the friction coefficient was approximately constant. Two different electronic stability control settings (On, Off) were used. All the measurements were recorded at 100 Hz, the standard frequency for vehicle state estimation. A statistical outlier removal has been applied to remove extreme outliers. However, particular attention is paid to not deleting edge case measurements, which are the most valuable data. Furthermore, all the manoeuvres were manually inspected to check the outlier removal efficacy. The measurements are considered when V_x is higher than 5 m/s and are filtered using a low-pass zero-phase




Fig. 4. Log distribution of sideslip angle and lateral acceleration. Each bin corresponds to 1deg and 1 m/s².

filter with a cut-off frequency of 5 Hz based on a finite impulse response technique [6].

The log distribution of the sideslip angle and lateral acceleration is represented in Fig. 4. The lateral acceleration is almost spread equally in the range $[-10, 10]$ m/s². In contrast, the sideslip angle measurements mainly distribute in the range $[-3, 3]$ deg. The latter is a common phenomenon because it is challenging to perform manoeuvres with a high sideslip angle, even when the vehicle has a very high lateral acceleration. Especially in dry road conditions, only a professional driver can induce a high sideslip angle.

A second dataset is selected from the same measurements. It will be referenced as limited dataset because it only contains measurements of when the vehicle has lateral acceleration $|a_y| \leq 7$ m/s². This simulates the cost and complexity of recording a large number of manoeuvres in which the vehicle is driven at the extreme vehicle behaviour, but not at the handling limits. Such a situation is common in the automotive field because, at the handling limits, the driver can easily lose the vehicle's control. Thus, the limited dataset will be used to analyse the proposed hybrid approach regarding its robustness and generalisation capabilities.

Both datasets are split into three sub-sets: training (75%), validation (15%) and test (10%). The test set contains the same manoeuvres for both the full and limited datasets. It consists of manoeuvres representing the entire driving behaviour, but more focus is paid to highly non-linear situations. It includes 23 manoeuvres: two braking in the turn, two skidpad, five J-turn, four slalom, four lane change, two random steers, three spiral and one lap track.

B. Key Performance Indicators

The performance of the different approaches is assessed through four key performance indicators (KPIs), which are commonly used in sideslip angle estimation [8], [34], [39].

- The MSE assesses the overall estimation performance.
- The non-linear MSE (MSE_{nl}) corresponds to the MSE computed only when $|a_y| \geq 4$ m/s². It measures the estimation performance when the vehicle behaves non-linearly.
- The absolute maximum error (ME) measures the worst estimation performance.
- The non-linear ME (ME_{nl}) measures the worst estimation performance in the case of non-linear vehicle behaviour.

The non-linear KPIs analyse the hybrid approach performance in the most critical scenarios. The MSE and MSE_{nl} are used to evaluate the estimation accuracy, while ME and ME_{nl} are used to assess temporary high errors in the estimation. The latter is relevant to assess whether the estimation is always coherent with the physical vehicle behaviour.

C. Baseline Methods

The proposed hybrid approach is compared with the state-of-art model-based, data-driven and hybrid approaches. All the considered baselines are adapted and optimised to use the same sensor setup and dataset, ensuring an objective and fair comparison.

The model-based approach is a UKF-based on a single-track model with tyre force measurements, as presented in [8]. The process noise parameters are tuned with the TSBO, and the observation noise parameters associated with the tyre force measurements are adapted online to enhance the observer's performance. The adaptability is related to the reduction of the level of noise coupled with tyre force measurements, this increases the Kalman gain when the vehicle behaves non-linearly. Thus, the effect of the Kalman gain is magnified during manoeuvres at the handling limit. Otherwise, a magnified Kalman gain when the vehicle behaves linearly could influence the vehicle states to follow the measurement sensor noises. The adaptability is triggered with a hysteresis loop to avoid the chattering phenomenon.

The data-driven approach is a FFNN that uses IMU and tyre force measurements as inputs, as evaluated in [8]. A simple FFNN reaches a better performance than a RNN when the tyre force measurements are included in the input set because the RNN prediction power is insufficient to compensate for the higher numbers of parameters to be trained. The NN is formed by two hidden layers with respectively 250 and 125 neurons each and ReLU activation functions. It uses a dropout regularisation technique (0.2) and a Xavier initialisation to avoid overfitting. An early stopping method with patience equal to 20 is applied for the same reason. The MSE is the loss function minimised by a mini-batch stochastic gradient descent algorithm based on a standard ADAM optimiser with a learning rate (0.001). The mini-batch size is 1024. For the training procedures, user-defined parameters are optimised through a Bayesian optimisation.

The hybrid approach is a deep ensemble-UKF (DE-UKF) [11] adapted to maximise the estimation performance on a dataset with tyre force measurements. The DE is formed by 20 FFNNs trained independently on the same dataset. The FFNNs different estimations are combined in a model averaging. Hence, the final β_{DD} is the mean of the FFNNs estimations, and σ_{DD} is the variance of the different model estimations. Each FFNN is trained using a Gaussian negative log-likelihood cost function optimised through mini-batch stochastic gradient descent based on an ADAM optimiser with a learning rate (0.0008). The epoch's number for each FFNN is 30. DE relies on the stochasticity of neural network training, which allows every FFNN to converge to a different set of parameters. However, the estimation accuracy is low when all models predict incorrectly,

TABLE II
SIDESLIP ANGLE ESTIMATION COMPARISON USING THE FULL DATASET

Approaches	MSE [deg ²]	MSE _{nl} [deg ²]	ME [deg]	ME _{nl} [deg]
Model-based	0.161	0.277	1.111	0.991
Data-driven	0.096	0.157	1.293	1.123
DE-UKF	0.087	0.157	0.981	0.822
CNN-UKF	0.086	0.118	0.979	0.776

and there is no guarantee that the σ_{DD} will be high. This especially happens when the error is in the low sideslip angle range because the NNs estimations tend to be closer. A high level of distrust suggests that the UKF does not rely on the data-driven pseudo-measurement but trusts the estimation of the UKF process model. Vice-versa, when the level of distrust is low, the UKF considers the neural network estimation reliable. σ_{DD} must be scaled before being used by the UKF because the output of the DE does not match the weight of the other noise parameters. Otherwise, the UKF puts too much trust in β_{DD} . The scaling is based on an exponential function ((11)) which differentiates approximately similar σ_{DD} .

$$\sigma_{DD,sc} = 10^{p_1} \sigma_{DD}^{p_2} \quad (11)$$

where p_1 (−4.2690 for the full dataset and −1.4353 for the limited one) and p_2 (0.7901 for the full dataset and 1.465 for the limited one) are two scaling parameters tuned using a Bayesian Optimisation. The values of p_1 and p_2 change according to the dataset because they strongly influence the DE's estimation performance. If p_1 and p_2 are not re-tuned for the limited dataset, the UKF will put too much trust in the DE, even if it lacks performance.

V. RESULTS

This section demonstrates the performance of the proposed approach. Section V-A analyses how accurate the proposed approach is with respect to the baselines when it is trained using a full dataset. Section V-B shows the results of the robustness analysis when only a limited dataset is available. This demonstrates that the data-driven approach is highly influenced by the amount and quality of the data. Section V-C proves the approaches' robustness to different tyre model parameters.

A. Full Dataset Results

The CNN-UKF, the DE-UKF and the data-driven approach have been trained using the full dataset.

The overall comparison is presented in Table II. Both hybrid approaches perform better than the model-based and data-driven approaches considering all four KPIs. This highlights the importance of the hybrid architecture for vehicle sideslip angle estimation. For instance, the model-based approach has a higher MSE and MSE_{nl} than the data-driven approach but a lower average ME and ME_{nl} . The hybrid approaches have the same estimation accuracy (MSE and MSE_{nl}) as the data-driven approach without the average higher ME. The reason is that in a hybrid approach, data-driven estimation is always validated through the model-based approach.




Fig. 5. Distribution of the sideslip angle error when the vehicle $|a_y| > 4 \text{ m/s}^2$ for every approach in the test set. Each bin is 0.25deg wide. The x represents the mean and the line between the vertical symbols ($| - |$) is the standard deviation of the sideslip angle error.

It can be seen that CNN-UKF outperforms the three other approaches for all the proposed KPIs. However, it does not have the same benefits in magnitude for all of them. The overall MSE and ME of DE-UKF and CNN-UKF are comparable. The minor improvements for the linear vehicle behaviour are respectively 1.15% and 0.20% in favour of the CNN-UKF. Anyhow, if the performance is evaluated when the vehicle behaves non-linearly, the CNN-UKF will strongly outperform DE-UKF with an improvement of 24.84% for the MSE_{nl} and 5.60% for the ME_{nl} . A possible explanation is that the end-to-end training informs CNN about the vehicle dynamics compensating for the lower amount of data in this operating condition.

On the contrary, the DE during the training is not aware of the physical vehicle behaviour, so it is subjected to a decay in performance where the dataset has fewer samples. The DE becomes aware of the UKF performance only while tuning the level of distrust scaling parameters. Furthermore, the process model noise parameters are online adapted in the CNN-UKF, allowing the UKF to accommodate better the mismatches between the physical and modelled vehicle behaviour.

Similar conclusions can be stated from the log distribution of the sideslip angle error in the non-linear operating range, see Fig. 5. The data-driven and the hybrid approaches have a similar amount of β error samples in the range $[-1.5, 1.5] \text{ deg}$. In contrast, the model-based approach suffers from the lower accuracy of the vehicle model in the non-linear operating region. However, the data-driven approach and partially the DE-UKF are more prone to high estimation errors ($\geq 1.5 \text{ deg}$) than the model-based and CNN-UKF. The latter outperforms all other approaches and has the β error mean closest to zero and the lowest standard deviation. Hence, a UKF coupled with a data-driven approach has the same performance as a data-driven approach in a low error range, but it reduces the sporadic high errors of a purely data-driven approach. Furthermore, the end-to-end training and the process noise parameters adaption allow the CNN-UKF to maximise the hybrid capability especially when the vehicle $|a_y| > 4 \text{ m/s}^2$. Fig. 6 analyses how the estimation performance change for different manoeuvres. The model-based approach has a weak accuracy, especially in braking-in-the-turn, J-turn and skidpad tests. The braking-in-the-turn involves a coupling between the longitudinal and lateral dynamics, which is not modelled in the used single-track vehicle model. In a J-turn




Fig. 6. Sideslip angle MSE_{nl} comparison for every group of manoeuvres.




Fig. 7. Slalom manoeuvre. Comparison of the sideslip angle estimation between all four approaches.

manoeuvre, the vehicle is driven at the limits of handling, where the mismatches between the physical and modelled vehicle behaviour are higher. Whereas for skidpad tests, the explanation is that it is a quasi steady-state manoeuvre, so the vehicle yaw acceleration is almost null, and the difference between estimated and measured tyre forces becomes essential for the β estimation. The tyre model is one of the most significant uncertainty sources in the model-based approach. The data-driven approach almost constantly behaves better than the model-based but worse than the hybrid approaches for estimation accuracy. However, it outperforms the DE-UKF in a spiral manoeuvre, and a possible explanation is that the DE-UKF puts too much trust in the UKF process model. The CNN-UKF outperforms all the other approaches in five out of seven manoeuvres. Particularly relevant is the improvement in the slalom and spiral manoeuvres. The slalom has the highest number of sideslip angle peaks (Fig. 7), which are the most difficult moments to estimate sideslip. Spiral manoeuvres are particularly challenging because it has an extra turn respect the J-turn. Fig. 7 shows the sideslip angle estimation in a slalom manoeuvre at the handling limits. All four approaches provide a reliable estimation, but the CNN-UKF outperforms the other approaches when the vehicle reaches a β peak of 10deg at around 14 s. This is a typical situation where a correct estimation of β is essential to help the vehicle control system maintain vehicle stability. Thus, an improved estimation in this condition is particularly relevant for safety. The already mentioned high non-linearities reduce the accuracy of the model-based approach. The data-driven approach lacks accuracy at 15 s due to the few data in the training set describing this vehicle's operation point. The DE-UKF improves the estimation performance between 5 s and 13 s combining the pros of the model-based and data-driven approach, but it lacks performance at around 14 s. CNN-UKF improves the estimation accuracy not only in the range of [5, 10]

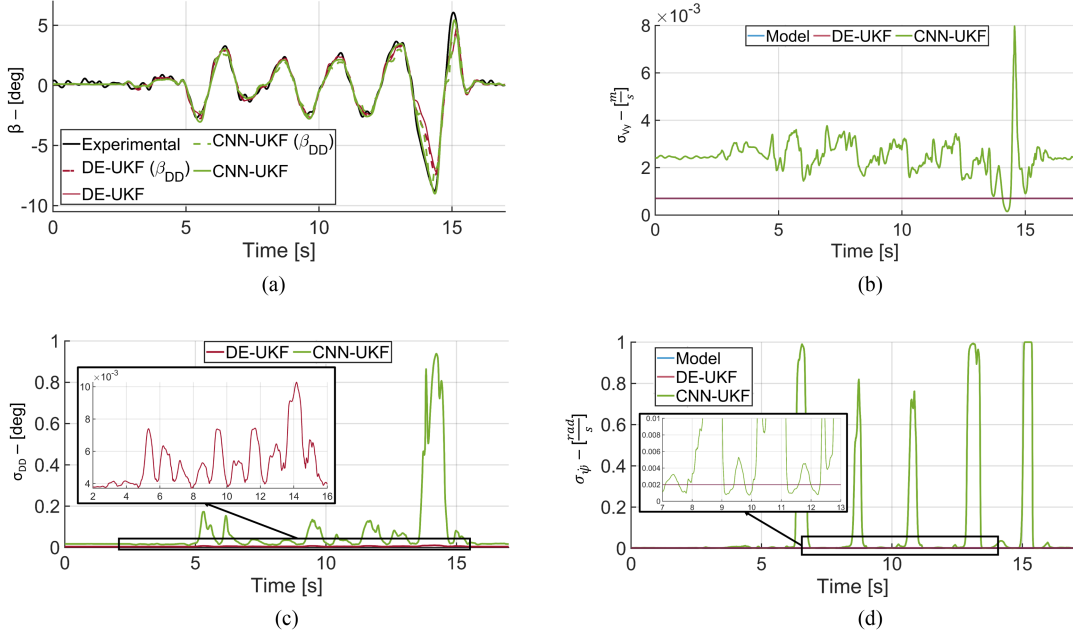


Fig. 8. (a) Shows the estimated and the pseudo-measurement of the sideslip angle. (b) Shows the process noise parameter associated with the V_y . (c) Shows the level of distrust in the NN for the hybrid approaches. (d) Shows the process noise parameter associated with the $\dot{\psi}$.

s but also in the highest peak at 14 s, as can be observed in Fig. 8.

Fig. 8(a) shows the β and β_{DD} for the hybrid approaches. CNN-UKF and DE-UKF β_{DD} s lack accuracy between 12s and 15s, but the CNN-UKF β is accurate because the UKF is correctly weighting the UKF process model's information with the NN's pseudo-measurement. Vice-versa, the UKF of the DE-UKF puts too much trust in β_{DD} . When the β_{DD} error rises, the corresponding level of distrust (Fig. 8(c)) also grows. CNN-UKF and DE-UKF σ_{DD} s have the same order of magnitude in normal driving, but the one related to CNN-UKF rises much more than the DE-UKF. This broader range makes the proposed approach much less confident in the NN when its output is incorrect. This is not possible for the DE-UKF due to its training process. The DE-UKF does not have end-to-end training, so its σ_{DD} cannot match the weight of the other UKF noise parameters. The DE-UKF σ_{DD} non-linear scaling compensates only partially this issue. Fig. 8(c) clearly demonstrates how the CNN-UKF distrust level range is $[10^{-3}, 1]$, while the range for the DE-UKF is only $[10^{-3}, 10^{-2}]$.

Another explanation for the better performance of the CNN-UKF is related to the online adaptation of the process noise parameters. The adaptive parameters allow the UKF to know the current mismatches between the modelled and physical vehicle behaviour. The process noise parameters of the DE-UKF and model-based approach are constant, so they correspond to a trade-off between the different driving conditions. Vice-versa, the CNN-UKF relies on optimal tuned process noise parameters every instant. Fig. 8(b) and (d) show the values of σ_{V_y} and $\sigma_{\dot{\psi}}$, respectively. As expected from the literature [72], both increase with the growth of vehicle non-linearities. This further proves that CNN-UKF behaves according to physical vehicle motion.

σ_{V_y} has a peak at 14s, corresponding to the last vehicle's right turn, where the rear inner tyre is even detached from the ground due to the aggressiveness of the manoeuvre. This extreme condition is created by a transient lateral load transfer (not modelled) which strongly influences lateral tyre force production, resulting in a significant V_y model mismatch. Moreover, the effect of the front axle longitudinal force (F_{xf}) on the lateral velocity is not modelled ($\frac{F_{xf} \sin(\delta)}{m}$). Overall, the constant and the adapted process noise parameter have the same magnitude. Still, the one associated with CNN-UKF is generally bigger (apart from 13s to 14s). The reason is that the constant σ_{V_y} was optimised, considering also less aggressive manoeuvres where the vehicle model is more reliable.

The process noise parameter $\sigma_{\dot{\psi}}$ rises by two orders of magnitude when the vehicle has a high sideslip angle. At the same time, when β is low, the adapted $\sigma_{\dot{\psi}}$ is slightly lower than the constant process noise parameter. A possible explanation is that the mismatches of the modelled $\dot{\psi}$ are higher than that of V_y . The meaningful adaptability of the process parameter noises shows the CNN-UKF has an insight into vehicle dynamics physics and it can online compensate for it. Similar conclusions are obtained from the spiral manoeuvre represented in Fig. 9. Here, the CNN-UKF approach outperforms the accuracy of all other three approaches, particularly from 5s to 13s. The performance of the CNN-UKF is similar to sum of the best estimation between the data-driven approach, from 4s to 6s, and the model-based approach, from 6s to 10s.

The test set also contains a recording of an entire lap in a racing circuit, where the effect of combined slip is maximal. Fig. 10(b) shows the vehicle's lateral and longitudinal acceleration, and it highlights how the driver is pushing the vehicle at




Fig. 9. Spiral manoeuvre. Comparison of the sideslip angle estimation between all four approaches.




Fig. 10. (a) Compares the sideslip angle estimation between four approaches in a portion of a racing track. (b) Shows the recorded lateral and longitudinal acceleration of the vehicle. It highlights the combined slip situation at which the vehicle is driven. (c) Shows the level of distrust in the NN for the hybrid approaches.

the limit of handling in all the corners, see [1, 7] s, [16, 19] s and [23, 29] s. The sideslip angle estimation performance of the four approaches is represented in Fig. 10(a). The model-based approach has the lowest performance, especially in the range [1, 7] s. This result is expected because the implemented Dugoff tyre model works in pure slip conditions. A similar conclusion is also visible in the spiral manoeuvre, Fig. 9. The data-driven approach performs better than a model-based approach. However, it has the maximum absolute error at 23s and 26s, where the vehicle performs a cornering while braking. Both proposed hybrid approaches have higher accuracy than the others, but the CNN-UKF has the best performance. A possible explanation is the physic-informed NN architecture, which allows evaluating a very accurate NN level of distrust. Fig. 10(c) shows the NN level of distrust for the DE-UKF and CNN-UKF. While the DE-UKF level of distrust is almost constant along the manoeuvre, the one associated with CNN-UKF has two peaks in correspondence with the data-driven maximum errors. This allows the

TABLE III
SIDESLIP ANGLE ESTIMATION COMPARISON USING THE LIMITED DATASET

Approaches	MSE [deg ²]	MSE _{nl} [deg ²]	ME [deg]	ME _{nl} [deg]
Model-based	0.161	0.277	1.111	0.991
Data-driven	0.223	0.358	1.445	1.284
DE-UKF	0.157	0.270	1.103	0.983
CNN-UKF	0.156	0.269	1.099	0.975

CNN-UKF to avoid following the high estimation error of the data-driven component. It is further proof of how the CNN-UKF is a physics-informed NN in which the UKF and NN are mutually cooperating to improve the overall estimation of the hybrid approach.

B. Robustness Analysis Using the Limited Dataset

A sideslip angle filter must not only be accurate, but it should be robust to a different amount of qualitative data during the training and tuning phase. Hence, to prove the robustness of the proposed hybrid approach, its estimation performance is compared with the baseline methods when they all have been trained using the limited dataset.

The overall comparison is presented in Table III. Here, the model-based approach shows the same performance as with the full dataset (see Table II), because it is not influenced by the amount of data. As expected the other approaches show a reduced performance with the limited dataset where the MSE is more than doubled while the ME sees a moderate increase. Now the data-driven approach has the worst performance in all four KPIs. The accuracy loss is higher than 30% for all the indicators, without a particular weakness in one of the proposed KPIs. An explanation is that the dataset does not have representative data of the vehicle driven with $|a_y| \geq 7$ m/s², so it must generalise much more than with a full dataset. Significantly, the NN must reconstruct the extreme non-linear vehicle behaviour, the most complex vehicle operating region, without having representative data for these conditions.

The model-based and hybrid approaches' performance is very similar, but DE-UKF and CNN-UKF have the best KPIs. The explanation is that hybrid approaches use the best estimation accuracy of the model and the NN together. Both hybrid approaches strongly rely on the estimation of the UKF process model because they cannot put much trust in the data-driven part. However, the NN still has benefits when the vehicle behaves linearly due to the excellent amount of data in that range. This highlights how the hybrid approaches improve the robustness of both model-based and data-driven approaches. The hybrid approach shows a minor improvement compared with the model-based approach. However, the result is significant because it highlights how the hybrid approach is as robust as a model, even if trained with a limited training dataset. On the contrary, a purely data-driven approach is not robust for using a small training set resulting in poor estimation accuracy.

The CNN-UKF performs slightly better than DE-UKF in all four KPIs. However, the CNN-UKF outperforms the DE-UKF, mainly for the MSE and MSE_{nl}. The main reason is the




Fig. 11. Distribution of the sideslip angle error when the vehicle $|a_y| > 4 \text{ m/s}^2$ for every approach in the test set. Each bin is 0.25 deg wide. The x represents the mean and the line between the vertical symbols (| — |) is the standard deviation of the sideslip angle error. Results based on the limited dataset, see Fig. 12 for best results.




Fig. 12. J-turn manoeuvre. Comparison of the sideslip angle estimation between all four approaches using the limited dataset.

adaptability of the process noise parameters, which cope with the change of vehicle model mismatches in the various vehicle operating points. However, the improvement in accuracy is not enough to be considered significant ($< 5 \%$).

Fig. 11 shows the sideslip angle error log distribution in the non-linear operating range. All the approaches which rely on a model highly outperform the data-driven approach. The latter have β error samples in the range $[-3.8, 4] \text{ deg}$, while the other approaches have β errors between $[-1.8, 1.8] \text{ deg}$. This proves that the data-driven approach is highly prone to high estimation errors when trained with a limited dataset. The performance of the model-based and hybrid approaches is very similar. They also share an equal error distribution. The data-driven approach slightly outperforms the other approaches in the very low error range $[-0.3, 0.6] \text{ deg}$. This explains why the hybrid approaches are more accurate overall than the model-based one, despite mainly relying on it.

Fig. 12 shows the sideslip angle estimation in a J-turn manoeuvre at the handling limits. The model-based and hybrid approaches behave almost identically, and all strongly outperform the purely data-driven approach. The only visible differences are between $[1.5, 3] \text{ s}$, where the CNN-UKF captures slightly better the conclusion of the peak and between $[4, 7] \text{ s}$ where the DE-UKF is closer to the β reference.

The major difference between DE-UKF and CNN-UKF is visible from the comparison of the β_{DD} , see Fig. 13. The β_{DD} computed by the CNN-UKF is highly outperforming the one estimated by the DE-UKF. The explanation is that the CNN-UKF is trained end-to-end, so the output of the CNN




Fig. 13. J-turn manoeuvre. Comparison of the estimated β and β_{DD} between the hybrid approaches using the limited dataset.

has physical information that the NN uses to increase its accuracy. The DE is trained independently, performing similarly to the purely data-driven approach. An higher accuracy of β_{DD} implies that the following UKF can rely on a better sideslip angle pseudo-measurement. This proves the benefits of using a physical informed-NN. Despite this, the β estimation of DE-UKF and CNN-UKF is similar because the model-based approach still outperforms both β_{DD} . Thus, both hybrid approaches mainly rely on the UKF. Due to the high chances of dealing with a limited dataset, the hybrid approach is fundamental to improving vehicle sideslip angle estimation.

However, the performance of the proposed CNN-UKF approach is still influenced by the amount and quality of data in the training set. Thus, it still represents a limitation of the proposed approach that must be addressed in the future. This highlights the importance of defining standards procedure to collect valuable and broad datasets. Regardless, the proposed CNN-UKF allows the introduction of possible solutions for lack of data, e.g., weakly-supervised learning during the end-to-end training, which allows for using data recorded without expensive sensors.

C. Robustness Analysis to Tyre Model Parameters

An essential property of vehicle sideslip angle estimators is the robustness to vehicle parameter variations. The variation of cornering stiffness strongly influences the sideslip angle estimation, because it can vary from its nominal value due to numerous factors, e.g. tyre pressure, temperature, and wear [39]. Thus, to prove the robustness of the proposed approach, Fig. 14 shows how the MSE_{nl} (Fig. 14(b)) and the MSE (Fig. 14(a)) of the vehicle sideslip angle vary at the variation of the axle cornering stiffness. In particular, the front and rear axle cornering stiffnesses are changed by $\pm 10\%$ [39]. Overall, it is visible that the model-based approach is the most sensible to the variation of the inner model, while the data-driven approach, which does not have a physical model, is not affected. Both hybrid approaches are less influenced by the variation in the physical model than the purely model-based approach. However, Fig. 14 shows that the CNN-UKF performance is more influenced by the model mismatch than the DE-UKF. A possible explanation is that the CNN-UKF varies the level of distrust in the NN during the manoeuvre according to how it learned during the training, giving more trust to the UKF at some specific moments. When it faces

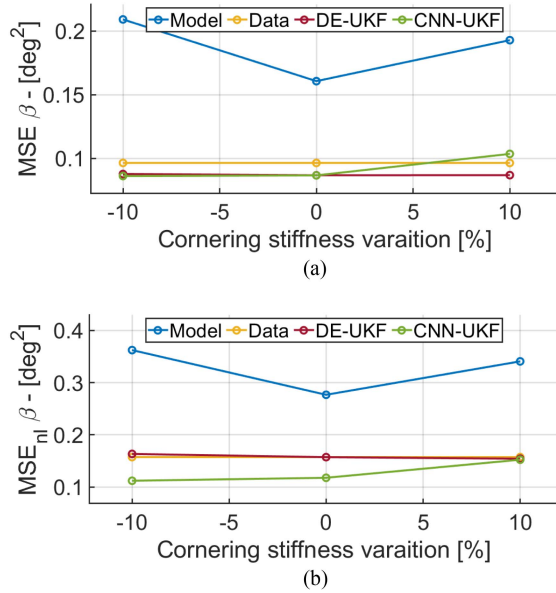


Fig. 14. (a) Shows how the axle cornering stiffness influences the vehicle sideslip angle MSE of all four approaches. (b) Shows how the axle cornering stiffness influences the vehicle sideslip angle MSE_{nl} of all four approaches.

a different model mismatch than previously learned, it is more influenced by it than the DE-UKF, which does not vary the level of distrust to the NN, always prioritising the data-driven side of the approach. Despite this, the CNN-UKF consistently has a lower MSE_{nl} than the DE-UKF, see Fig. 14(b) and only when both axle cornering stiffness of the vehicle model are increased by 10% as, on average, a higher MSE than the DE-UKF. Thus, it can be concluded that hybrid approaches are more robust to parameter uncertainties than purely model approaches and that the DE-UKF performance is less affected by the internal model accuracy than CNN-UKF.

VI. CONCLUSION

The paper presents a novel hybrid approach to vehicle sideslip angle estimation, which involves utilising the physical knowledge from a UKF based on a single-track vehicle model to enhance the estimation accuracy of a CNN. Using a large-scale experimental dataset of 216 manoeuvres, it has been shown that the hybrid approach is more accurate than purely model-based or data-driven approaches. Moreover, the CNN-UKF is slightly reducing the MSE of the DE-UKF. However, when the MSE_{nl} is compared, the CNN-UKF outperforms the DE-UKF by 25%, providing a much higher accuracy in the most critical operating region for active vehicle control systems. The CNN-UKF, thanks to the end-to-end training, is forcing the CNN to comply with the vehicle physics, reducing the ME and ME_{nl} of all other approaches. When a limited dataset is provided, the proposed hybrid approach has a minor improvement in the estimation robustness over the model-based and the DE-UKF approach for all the KPIs. The CNN-UKF is highly outperforming the estimation of a purely data-driven approach. Future works involve testing

the generalisation capability of the CNN-UKF utilising a dataset with different levels of road grip, e.g. wet, snow or icy roads.

REFERENCES

- [1] N. Chowdhri, L. Ferranti, F. S. Iribarren, and B. Shyrokau, "Integrated nonlinear model predictive control for automated driving," *Control Eng. Pract.*, vol. 106, 2021, Art. no. 104654.
- [2] M. T. Peterson, T. Goel, and J. C. Gerdes, "Exploiting linear structure for precision control of highly nonlinear vehicle dynamics," *IEEE Trans. Intell. Veh.*, vol. 8, no. 2, pp. 1852–1862, Feb. 2023.
- [3] Y.-W. Liao and F. Borrelli, "An adaptive approach to real-time estimation of vehicle sideslip, road bank angles, and sensor bias," *IEEE Trans. Veh. Technol.*, vol. 68, no. 8, pp. 7443–7454, Aug. 2019.
- [4] H. F. Grip, L. Imsland, T. A. Johansen, J. C. Kalkkuhl, and A. Suissa, "Vehicle sideslip estimation," *IEEE Control Syst. Mag.*, vol. 29, no. 5, pp. 36–52, Oct. 2009.
- [5] J. Liu, Z. Wang, L. Zhang, and P. Walker, "Sideslip angle estimation of ground vehicles: A comparative study," *IET Control Theory Appl.*, vol. 14, pp. 3490–3505, 2020.
- [6] M. Viehweger et al., "Vehicle state and tyre force estimation: Demonstrations and guidelines," *Veh. Syst. Dyn.*, vol. 59, no. 5, pp. 675–702, 2021.
- [7] T. Gräber, S. Lupberger, M. Unterreiner, and D. Schramm, "A hybrid approach to side-slip angle estimation with recurrent neural networks and kinematic vehicle models," *IEEE Trans. Intell. Veh.*, vol. 4, no. 1, pp. 39–47, Mar. 2019.
- [8] A. Bertipaglia, D. de Mol, M. Alirezai, R. Happee, and B. Shyrokau, "Model-based vs data-driven estimation of vehicle sideslip angle and benefits of tyre force measurements," 2022, *arXiv:2206.15119*.
- [9] J. Qi, A. F. Taha, and J. Wang, "Comparing Kalman filters and observers for power system dynamic state estimation with model uncertainty and malicious cyber attacks," *IEEE Access*, vol. 6, pp. 77155–77168, 2018.
- [10] D. Chindamo, B. Lenzo, and M. Gadola, "On the vehicle sideslip angle estimation: A literature review of methods, models, and innovations," *Appl. Sci.*, vol. 8, no. 3, 2018, Art. no. 355. [Online]. Available: <https://www.mdpi.com/2076-3417/8/3/355>
- [11] D. Kim, K. Min, H. Kim, and K. Huh, "Vehicle sideslip angle estimation using deep ensemble-based adaptive Kalman filter," *Mech. Syst. Signal Process.*, vol. 144, 2020, Art. no. 106862.
- [12] D. Kim, G. Kim, S. Choi, and K. Huh, "An integrated deep ensemble-unscented Kalman filter for sideslip angle estimation with sensor filtering network," *IEEE Access*, vol. 9, pp. 149681–149689, 2021.
- [13] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, 2019.
- [14] B. L. Boada, M. J. L. Boada, A. Gauchía, E. Olmeda, and V. Díaz, "Sideslip angle estimator based on ANFIS for vehicle handling and stability," *J. Mech. Sci. Technol.*, vol. 29, no. 4, pp. 1473–1481, 2015.
- [15] B. L. Boada, M. J. L. Boada, and V. Díaz, "Vehicle sideslip angle measurement based on sensor data fusion using an integrated ANFIS and an unscented Kalman filter algorithm," *Mech. Syst. Signal Process.*, vol. 72, pp. 832–845, 2016.
- [16] P. M. Sieberg, S. Blume, N. Harnack, N. Maas, and D. Schramm, "Hybrid state estimation combining artificial neural network and physical model," in *Proc. Intell. Transp. Syst. Conf.*, 2019, pp. 894–899.
- [17] A. J. Tuononen, "Vehicle lateral state estimation based on measured tyre forces," *Sensors*, vol. 9, no. 11, pp. 8761–8775, 2009.
- [18] A. K. Madhusudhanan, M. Corno, and E. Holweg, "Vehicle sideslip estimator using load sensing bearings," *Control Eng. Pract.*, vol. 54, pp. 46–57, 2016.
- [19] D. Selmanaj, M. Corno, G. Panzani, and S. M. Savaresi, "Vehicle sideslip estimation: A kinematic based approach," *Control Eng. Pract.*, vol. 67, pp. 1–12, 2017.
- [20] W. Liu, X. Xia, L. Xiong, Y. Lu, L. Gao, and Z. Yu, "Automated vehicle sideslip angle estimation considering signal measurement characteristic," *IEEE Sensors J.*, vol. 21, no. 19, pp. 21675–21687, Oct. 2021.
- [21] X. Xia, E. Hashemi, L. Xiong, A. Khajepour, and N. Xu, "Autonomous vehicles sideslip angle estimation: Single antenna GNSS/IMU fusion with observability analysis," *IEEE Internet Things J.*, vol. 8, no. 19, pp. 14845–14859, Oct. 2021.

- [22] W. Li, Z. Xie, P. K. Wong, Y. Hu, G. Guo, and J. Zhao, "Event-triggered asynchronous fuzzy filtering for vehicle sideslip angle estimation with data quantization and dropouts," *IEEE Trans. Fuzzy Syst.*, vol. 30, no. 8, pp. 2822–2836, Aug. 2022.
- [23] X. Ding, Z. Wang, and L. Zhang, "Event-triggered vehicle sideslip angle estimation based on low-cost sensors," *IEEE Trans. Ind. Inform.*, vol. 18, no. 7, pp. 4466–4476, Jul. 2022.
- [24] A. Wischnewski, T. Stahl, J. Betz, and B. Lohmann, "Vehicle dynamics state estimation and localization for high performance race cars," *IFAC-PapersOnLine*, vol. 52, no. 8, pp. 154–161, 2019.
- [25] R. Song and Y. Fang, "Vehicle state estimation for INS/GPS aided by sensors fusion and SCKF-based algorithm," *Mech. Syst. Signal Process.*, vol. 150, 2021, Art. no. 107315.
- [26] J.-H. Yoon and H. Peng, "Robust vehicle sideslip angle estimation through a disturbance rejection filter that integrates a magnetometer with GPS," *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 1, pp. 191–204, Feb. 2014.
- [27] S. van Aalst, F. Naets, B. Boulkroune, W. De Nijs, and W. Desmet, "An adaptive vehicle sideslip estimator for reliable estimation in low and high excitation driving," *IFAC-PapersOnLine*, vol. 51, no. 9, pp. 243–248, 2018.
- [28] M. Gadola, D. Chindamo, M. Romano, and F. Padula, "Development and validation of a Kalman filter-based model for vehicle slip angle estimation," *Veh. Syst. Dyn.*, vol. 52, no. 1, pp. 68–84, 2014.
- [29] C. Vaseur, S. Van Aalst, and W. Desmet, "Robust vehicle state and tire force estimation: Highlights on effects of road angles and sensor performance," in *Proc. Intell. Veh. Symp.*, 2021, pp. 28–33.
- [30] L. Li, G. Jia, X. Ran, J. Song, and K. Wu, "A variable structure extended Kalman filter for vehicle sideslip angle estimation on a low friction road," *Veh. Syst. Dyn.*, vol. 52, no. 2, pp. 280–308, 2014.
- [31] F. Cheli, D. Ivone, and E. Sabbioni, "Smart tyre induced benefits in sideslip angle and friction coefficient estimation," *Sensors Instrum.*, vol. 5, pp. 73–83, 2015.
- [32] R. van Hoek, M. Alirezai, A. Schmeitz, and H. Nijmeijer, "Vehicle state estimation using a state dependent Riccati equation," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 3388–3393, 2017.
- [33] T. Ando et al., "Localization using global magnetic positioning system for automated driving bus and intervals for magnetic markers," *IEEE Trans. Intell. Veh.*, vol. 8, no. 1, pp. 878–887, Jan. 2023.
- [34] M. Doumiati, A. C. Victorino, A. Charara, and D. Lechner, "Onboard real-time estimation of vehicle lateral tire-road forces and sideslip angle," *IEEE/ASME Trans. Mechatron.*, vol. 16, no. 4, pp. 601–614, Aug. 2011.
- [35] S. Antonov, A. Fehn, and A. Kugi, "Unscented Kalman filter for vehicle state estimation," *Veh. Syst. Dyn.*, vol. 49, no. 9, pp. 1497–1520, 2011.
- [36] H. Heidfeld and M. Schünemann, "Optimization-based tuning of a hybrid UKF state estimator with tire model adaption for an all wheel drive electric vehicle," *Energies*, vol. 14, no. 5, 2021, Art. no. 1396. [Online]. Available: <https://www.mdpi.com/1996-1073/14/5/1396>
- [37] K. B. Singh, "Vehicle sideslip angle estimation based on tire model adaptation," *Electronics*, vol. 8, no. 2, 2019, Art. no. 199. [Online]. Available: <https://www.mdpi.com/2079-9292/8/2/199>
- [38] J. Bechthoff and R. Isermann, "Cornering stiffness and sideslip angle estimation for integrated vehicle dynamics control," *IFAC-PapersOnLine*, vol. 49, no. 11, pp. 297–304, 2016.
- [39] V. Mazzilli et al., "On the benefit of smart tyre technology on vehicle state estimation," *Veh. Syst. Dyn.*, vol. 60, pp. 3694–3719, 2022.
- [40] J. Stephant, A. Charara, and D. Meizel, "Virtual sensor: Application to vehicle sideslip angle and transversal forces," *IEEE Trans. Ind. Electron.*, vol. 51, no. 2, pp. 278–289, Apr. 2004.
- [41] J. Stéphant, A. Charara, and D. Meizel, "Evaluation of a sliding mode observer for vehicle sideslip angle," *Control Eng. Pract.*, vol. 15, no. 7, pp. 803–812, 2007.
- [42] J. Zhang et al., "Adaptive sliding mode-based lateral stability control of steer-by-wire vehicles with experimental validations," *IEEE Trans. Veh. Technol.*, vol. 69, no. 9, pp. 9589–9600, Sep. 2020.
- [43] H. Zhang, G. Zhang, and J. Wang, "Sideslip angle estimation of an electric ground vehicle via finite-frequency H_∞ approach," *IEEE Trans. Transp. Electrification*, vol. 2, no. 2, pp. 200–209, Jun. 2016.
- [44] S. Cheng, L. Li, B. Yan, C. Liu, X. Wang, and J. Fang, "Simultaneous estimation of tire side-slip angle and lateral tire force for vehicle lateral stability control," *Mech. Syst. Signal Process.*, vol. 132, pp. 168–182, 2019.
- [45] N. Ding, W. Chen, Y. Zhang, G. Xu, and F. Gao, "An extended Luenberger observer for estimation of vehicle sideslip angle and road friction," *Int. J. Veh. Des.*, vol. 66, no. 4, pp. 385–414, 2014.
- [46] O. Galluppi, M. Corno, and S. M. Savaresi, "Mixed-kinematic body sideslip angle estimator for high performance cars," in *Proc. Eur. Control Conf.*, 2018, pp. 941–946.
- [47] E. Villano, B. Lenzo, and A. Sakhnevych, "Cross-combined UKF for vehicle sideslip angle estimation with a modified duffoff tire model: Design and experimental results," *Meccanica*, vol. 56, no. 11, pp. 2653–2668, 2021.
- [48] G. Park, "Vehicle sideslip angle estimation based on interacting multiple model Kalman filter using low-cost sensor fusion," *IEEE Trans. Veh. Technol.*, vol. 71, no. 6, pp. 6088–6099, Jun. 2022.
- [49] W. Chen, D. Tan, and L. Zhao, "Vehicle sideslip angle and road friction estimation using online gradient descent algorithm," *IEEE Trans. Veh. Technol.*, vol. 67, no. 12, pp. 11475–11485, Dec. 2018.
- [50] E. Hashemi, M. Pirani, A. Khajepour, A. Kasaiezadeh, S.-K. Chen, and B. Litkouhi, "Corner-based estimation of tire forces and vehicle velocities robust to road conditions," *Control Eng. Pract.*, vol. 61, pp. 28–40, 2017.
- [51] M. Jalali, E. Hashemi, A. Khajepour, S.-K. Chen, and B. Litkouhi, "Integrated model predictive control and velocity estimation of electric vehicles," *Mechatronics*, vol. 46, pp. 84–100, 2017.
- [52] H. Sasaki and T. Nishimaki, "A side-slip angle estimation using neural network for a wheeled vehicle," *SAE Trans.*, vol. 109, pp. 1026–1031, 2000.
- [53] S. Melzi and E. Sabbioni, "On the vehicle sideslip angle estimation through neural networks: Numerical and experimental results," *Mech. Syst. Signal Process.*, vol. 25, no. 6, pp. 2005–2019, 2011.
- [54] A. Bonfitto, S. Feraco, A. Tonoli, and N. Amati, "Combined regression and classification artificial neural networks for sideslip angle estimation and road condition identification," *Veh. Syst. Dyn.*, vol. 58, no. 11, pp. 1766–1787, 2020.
- [55] A. Tonoli, N. Amati, and J. Ghosh, "A deep learning based virtual sensor for vehicle sideslip angle estimation: Experimental results," in *WCX World Congr. Experience*, Warrendale, PA, USA: SAE International, Apr. 2018.
- [56] J. Liu, Z. Wang, and L. Zhang, "A time-delay neural network of sideslip angle estimation for in-wheel motor drive electric vehicles," in *Proc. IEEE 91st Veh. Technol. Conf.*, 2020, pp. 1–5.
- [57] V. Breschi, S. Formentin, G. Rallo, M. Corno, and S. M. Savaresi, "Vehicle sideslip estimation via kernel-based LPV identification: Theory and experiments," *Automatica*, vol. 122, 2020, Art. no. 109237.
- [58] M. Acosta, S. Kanarachos, and M. E. Fitzpatrick, "Robust virtual sensing for vehicle agile manoeuvring: A tyre-model-less approach," *IEEE Trans. Veh. Technol.*, vol. 67, no. 3, pp. 1894–1908, Mar. 2018.
- [59] T. Novi, R. Capitani, and C. Annicchiarico, "An integrated artificial neural network–unscented Kalman filter vehicle sideslip angle estimation based on inertial measurement unit measurements," *IMEch, Part D: J. Automobile Eng.*, vol. 233, no. 7, pp. 1864–1878, 2019.
- [60] N. A. Piga, U. Pattacini, and L. Natale, "A differentiable extended Kalman filter for object tracking under sliding regime," *Front. Robot. AI*, Frontiers Media SA, vol. 8, 2021, Art. no. 686447.
- [61] W. Jeon, A. Chakrabarty, A. Zemouche, and R. Rajamani, "Simultaneous state estimation and tire model learning for autonomous vehicle applications," *IEEE/ASME Trans. Mechatron.*, vol. 26, no. 4, pp. 1941–1950, Aug. 2021.
- [62] Q. Zhang, H. Jing, Z. Liu, Y. Jiang, and M. Gu, "A novel PWA lateral dynamics modeling method and switched T-S observer design for vehicle sideslip angle estimation," *IEEE Trans. Ind. Electron.*, vol. 69, no. 2, pp. 1847–1857, Feb. 2022.
- [63] T. Haarnoja, A. Ajay, S. Levine, and P. Abbeel, "Backprop KF: Learning discriminative deterministic state estimators," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 4383–4391.
- [64] G. Revach, N. Shlezinger, X. Ni, A. L. Escoriza, R. J. G. Van Sloun, and Y. C. Eldar, "KalmanNet: Neural network aided Kalman filtering for partially known dynamics," *IEEE Trans. Signal Process.*, vol. 70, pp. 1532–1547, 2022.
- [65] N. Xu, Y. Huang, H. Askari, and Z. Tang, "Tire slip angle estimation based on the intelligent tire technology," *IEEE Trans. Veh. Technol.*, vol. 70, no. 3, pp. 2239–2249, Mar. 2021.
- [66] S. Kerst, B. Shyrokau, and E. Holweg, "A model-based approach for the estimation of bearing forces and moments using outer ring deformation," *IEEE Trans. Ind. Electron.*, vol. 67, no. 1, pp. 461–470, Jan. 2020.
- [67] K. Kersting, C. Plagemann, P. Pfaff, and W. Burgard, "Most likely heteroscedastic Gaussian process regression," in *Proc. 24th Int. Conf. Mach. Learn.*, 2007, pp. 393–400.
- [68] M. Bauza and A. Rodriguez, "A probabilistic data-driven model for planar pushing," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2017, pp. 3008–3015.
- [69] J. Brownlee, *Better Deep Learning: Train Faster, Reduce Overfitting, and Make Better Predictions*. Melbourne, VIC, Australia: Machine Learning Mastery, 2018.

- [70] A. Bertipaglia, B. Shyrokau, M. Alirezaei, and R. Happee, "A two-stage Bayesian optimisation for automatic tuning of an unscented Kalman filter for vehicle sideslip angle estimation," in *Proc. Intell. Veh. Symp.*, 2022, pp. 670–677.
- [71] S. Kerst, B. Shyrokau, and E. Holweg, "Reconstruction of wheel forces using an intelligent bearing," *SAE Int. J. Passenger Cars Electron. Elect. Syst.*, vol. 9, pp. 196–203, 2016.
- [72] M. Acosta, S. Kanarachos, and M. E. Fitzpatrick, "Optimized tire force estimation using extended Kalman filter and fruit fly optimization," in *Proc. IEEE Ind. Electron. Soc.*, 2017, pp. 4074–4079.



Riender Happee received the Ph.D. degree from the Delft University of Technology, Delft, The Netherlands, in 1992. From 1992 to 2007, he investigated road safety and introduced biomechanical human models for impact and comfort with TNO Automotive. He is currently a Full Professor with the Delft University of Technology, where he, investigates the human interaction with automated vehicles focusing on safety, motion comfort, and acceptance.



Alberto Bertipaglia (Graduate Student Member, IEEE) received the M.Sc. degree, (*cum laude*), in automotive engineering from Politecnico di Torino, Turin, Italia, in 2020. He is working toward the Ph.D. degree with the Department of Cognitive Robotics, the Delft University of Technology, Delft, Netherlands. His work is part of a national research project aiming to develop innovative active vehicle controls to improve vehicle stability and increase the safety of automated vehicles performing evasive manoeuvres.



Barys Shyrokau received the Dipl. Eng. degree (*cum laude*) in mechanical engineering from Belarusian National Technical University, Minsk, Belarus, in 2004, and the joint Ph.D. degree in control engineering from Nanyang Technological University, Singapore, and Technical University Munich, Munich, Germany, in 2015. He is currently an Associate Professor with the Section of Intelligent Vehicles, Delft University of Technology, Delft, The Netherlands. His research interests include vehicle dynamics and control, motion comfort, and driving simulator technology.

He was the recipient of Scholarship's and awards, such as SAE, FISITA, DAAD, SINGA, ISTVS, and CADLM.



Mohsen Alirezaei received the Ph.D. degree in mechanical engineering, robotics, and control from the Delft University of Technology, Delft, The Netherlands, in 2011. He was a Postdoc Researcher with Delft University of Technology, in 2012. He was a Senior Scientist with Integrated Vehicle Safety Department, TNO automotive during 2012–2019, and a part time Assistant Professor with the Delft University of Technology during 2015–2019. He is currently a Fellow Scientist with Siemens Industry Software and Services, Helmond, The Netherlands and a part time

Assistant Professor with the Eindhoven University of Technology, Eindhoven, The Netherlands. His research interests include verification and validation of automated and cooperative automated driving, and advance driver assistance systems.