

Cite as:

Richard Schubert, Marvin Loba, Jasper Sünnemann, Torben Stolte, and Markus Maurer, “Conformal Prediction of Motion Control Performance for an Automated Vehicle in Presence of Actuator Degradations and Failures,” submitted for publication.

BibTeX:

```
@inproceedings{schubert_conformal_2024,  
  author={{Schubert}, Richard and {Loba}, Marvin and {Suenemann}, Jasper and {Stolte}, Torben and {Maurer}, Markus},  
  title={{Conformal} {Prediction} of {Motion} {Control} {Performance} for an {Automated} {Vehicle}  
    in {Presence} of {Actuator} {Degradations} and {Failures}},  
  year={2024},  
  publisher={submitted for publication}  
}
```

Conformal Prediction of Motion Control Performance for an Automated Vehicle in Presence of Actuator Degradations and Failures*

Richard Schubert¹, Marvin Loba¹, Jasper Sünnemann¹, Torben Stolte¹ and Markus Maurer¹

Abstract—Automated driving systems require monitoring mechanisms to ensure safe operation, especially if system components degrade or fail. Their runtime self-representation plays a key role as it provides a-priori knowledge about the system’s capabilities and limitations. In this paper, we propose a data-driven approach for deriving such a self-representation model for the motion controller of an automated vehicle. A conformalized prediction model is learned and allows estimating how operational conditions as well as potential degradations and failures of the vehicle’s actuators impact motion control performance. During runtime behavior generation, our predictor can provide a heuristic for determining the admissible action space.




Fig. 1. The MOBILE research vehicle in IPG CarMaker, figure from [1].

I. INTRODUCTION

The development of automated driving technology has the potential of (temporarily) relieving drivers of the task of driving and monitoring their vehicle, even in challenging scenarios. However, in order to ensure that these systems operate as desired, it is necessary to consider safety-related constraints, both during design and operation. One key consideration when designing an automated driving system is ensuring that the vehicle can continue to operate safely – among others – if hardware failures arise. Hence, functional boundaries must be carefully evaluated at design time regarding both normal operation and in the event of a degradation or failure. In the sense of the ISO 26262 standard [2], both failure and degradation are the result of a fault, though with different outcome. A *degradation* is the state of a system element with reduced performance, yet with given functionality¹. In contrast, a *failure* denotes the termination of the system’s intended behavior, i.e., the system

is not providing its functionality. In addition to robust design, it is necessary to monitor the system’s performance at runtime in order to detect, isolate, and compensate for failures [4] before they lead to hazardous situations.

Furthermore, to adapt behavioral decisions at runtime accordingly, the system relies on the explicit representation of knowledge about itself, which is referred to as *self-representation* [4]. This involves creating and deploying models that contain a-priori knowledge about the system’s operational limits. The impact thereof lies in the ability to, e.g., trace the impact of degraded or failed components on the overall system behavior. By estimating functional boundaries with respect to performance constraints of the overall system, it is possible to determine the admissible actions.

In this paper, we acknowledge the importance of self-representation and provide an example of how to derive a self-representation model of an automated vehicle’s motion controller in a data-driven manner. Given previous work from our group [1], we use a complex fault-tolerant model-predictive trajectory tracking controller that is able to control a highly over-actuated vehicle as a case study. Modelling the dynamics of the controlled vehicle under various actuator degradations and failures is complex as simple physical models might not be sufficient to capture the system’s behavior – especially when considering over-actuated vehicles. Based on the results of a large-scale numerical simulation, we are able to derive a prediction model that estimates how operational conditions, dynamic constraints, and potential degradations and failures of the vehicle’s actuators impact the controller’s performance at runtime in terms of its deviation from a planned trajectory. In an example scenario, we show how the prediction model can be used to estimate the controller’s performance in order to evaluate the feasibility of actions under given conditions, including evasive maneuvers. This estimation can then be utilized at the behavior generation level in order to constrain the admissible action space.

The remainder of this paper is structured as follows: The following Section II provides an overview of related work and our contributions. In Section III, we describe the motion controller’s role in the functional architecture and present the fault-tolerant controller by [1]. Thereafter, we describe our simulation framework and design a predictor using the generated data in Section IV. Finally, we illustrate the application of the predictor in Section V.

II. RELATED WORK

Previous publications discuss the need to supervise the behavior generation and execution of automated vehicles:

*This research is accomplished within the project AUTOfTech.agil (FKZ 01IS22088R). We acknowledge the financial support by the Federal Ministry of Education and Research of Germany (BMBF).

¹Institute for Control Engineering, TU Braunschweig, Germany {richard.schubert, marvin.loba, j.suennemann, t.stolte, markus.maurer}@tu-braunschweig.de

¹Please note, that the understanding of *degradation* in ISO 26262 is inconclusive, as neither functionality nor performance are defined in the standard. Hence, we use the understanding developed in a previous work [3].

Maurer [5] argues that any system component should provide quality measures that must be aggregated by a performance monitoring function and assessed when behavioral decisions are made. Nolte et al. [4] point out that a system requires a self-representation in order to determine its capabilities [6] and select actions to respond appropriately to the current situation. The set of appropriate actions can be referred to as *admissible action space* [4]. In the context of automated driving, we equate the terms *action* and *driving maneuver*, i.e., an abstraction of possible state progressions for the vehicle's motion [7]. The constraints of the action space can be both *internal* and *external* in nature, i.e., related to the system's internal state and the current situation [4].

In [8], Nolte et al. propose to apply the concept of *skill graphs* as introduced by Reschka [6] to identify the capabilities of an automated vehicle and refine an ISO 26262 compliant design process. This allows to structure system capabilities, assign functional components that contribute to them, and derive safety requirements. Comparable work is presented in the context of the UNICARagil project by Stolte et al. [9]. Both publications propose eliciting requirements related to the vehicle's behavior during system design. These requirements are then refined when considering the vehicle's architecture. Behavioral requirements address quantities related to the vehicle's (externally observable) motion, e.g., the required maximum lateral deviation from a reference path [8]. Technical determinants – e.g., the accuracy of the vehicle's state estimation or the available performance of the vehicle's actuators – influence the quality of the vehicle's behavior execution. In particular, Nolte et al. [8] consider an automated vehicle's objective to follow a specific lane marking on a German highway: The authors require that a certain lateral deviation threshold is not exceeded by the vehicle (with respect to the lane marking). They use a capability graph to successively refine and decompose into further requirements. Focusing on capabilities required for behavior execution, dependencies with respect to the vehicle's localization and motion control capabilities are identified. Inaccurate localization and/or insufficient performance of motion control impact behavior execution and may lead to a violation of requirements. Hence, when considering motion control as an example, it is crucial to ensure that the controller's performance stays within defined bounds.

In this paper, a particular focus is set on highly over-actuated vehicles and their motion control. In this context, Da Silva et al. [10] investigate the impact of actuator fault combinations on an over-actuated vehicle's ability to perform an evasive lane change maneuver. Their offline analysis employs a simulation of the controlled vehicle to acquire several performance metrics, including the maximum lateral deviation from the reference as well as “leave road” and “collision” indicators. In a statistical analysis, they trace the impact of actuator degradations and failures on these metrics – however, only for offline analysis.

At runtime, the vehicle should be able to adapt its planned actions and respect its current control capabilities. Degradations and failures thereof may lead to sudden




Fig. 2. Simplified functional architecture of an automated vehicle, figure taken from Stolte et al. [1], originally based on [15]. Our conformed control performance prediction model relies on self-perception information and is located next to the behavior generation block.

inconsistencies between the planned and actual motion. Ratner et al. [11] further elaborate on the challenge of ensuring consistency between (state-based) motion planning and control models for robots. They propose to bias the planner away from inaccuracies by increasing the cost of states and actions where discrepancies are observed. The authors use a stochastic discrepancy model that is learned online from real-world feedback. A similar approach is presented in [12]. In safety-critical applications, however, using online learning to fit such a model is usually not feasible since model discrepancies must be experienced first before they can be learned. Therefore, Castaneda et al. [13] pair a Gaussian Process model with a Control Barrier Function to ensure that the system stays within defined bounds under uncertainty. Moreover, they propose to acquire new measurements to update their model only in situations where the underlying dynamics are expected to only slightly deviate from the model. To evaluate the space of all actions that can be executed by the system given its actual dynamics, McConachie et al. [14] propose the design of a neural network classifier that predicts the likelihood of a system's action being *reliable* or *unreliable*. An unreliable action is one where the system is likely to run into constraint boundaries in the real-world that are not considered in a simplified model used for planning. McConachie et al. exploit this concept to be able to plan actions in a simplified state space and sort out unreliable actions before executing them.

Ensuring that a controlled system stays within defined bounds is also the objective of *reachability analysis*. Reachability analysis can be used as a design tool offline as well as for *online verification* of dynamic systems [16, 17]. Given model assumptions about the controlled system, reachability analysis is used to explore which states can be reached (or avoided) for given inputs and disturbances within a certain time frame or at a specific time point. *Forward* reachable states can be found by propagating the set of inputs through a dynamics model analytically – e.g., using set- or distribution-based representations [16] – or numerically. *Backward* reachability analysis, on the other hand, aims to identify the set of states from which the system cannot avoid reaching an unsafe state at a later time point [18], regardless of whether a control algorithm tries to prevent this. Recent approaches often adopt *Hamilton-Jacobi*

reachability analysis [19]. However, it is computationally expensive and scales exponentially with the number of states of the underlying system [19, 20]. Systems with more than five states are often considered intractable, especially for online verification [18]. To reduce computational effort, for instance, linearized models with small numbers of states are used [17]. If disturbances are considered, they are often assumed to be bounded – either by fixed boundaries or confidence intervals – but methods for *stochastic reachability analysis* exist as well [16, 21, 22]. Recent approaches for reducing the computational effort at runtime include the learning-based online classification of reachable states [23], their efficient sampling-based approximation [24] as well as decomposition in less complex subsystems [25].

While reachability analysis methods often aim to *guarantee* that certain state regions are reached/avoided, the results are always based on assumptions about the dynamics model, input space, and additional disturbances. Therefore, any model-based guarantee should be considered with care since model-based and additional assumptions may not hold in the real world. For instance, given sudden degradations of the systems, the dynamics representation must be updated accordingly. This requires additional research on the derivation of dynamics models under the influence of degradations and failures, e.g., as presented in [26, 27] for aviation systems. Publications addressing land vehicles are however rare.

Contributions: In this paper, we investigate the dynamics of a highly over-actuated, automated vehicle – which can be described as a nonlinear system with up to 10 states (e.g., [28]) – under the influence of actuator degradations and failures. Inspired by [14], we propose to use a lightweight prediction model that is able to predict whether an automated vehicle can be controlled with sufficiently low lateral deviation from the planned motion under nominal, degraded, and failed actuator conditions. Complementary to the online learning approaches in [11–13], our approach exploits prior knowledge about the system’s dynamics under nominal, degraded, and failed conditions instead of online learning. Combining both approaches in the future could be useful in safety-critical applications. In particular, we use a neural network and capture the complexity and uncertainty of the prediction by predicting *regression quantiles* instead of deterministic values. To ensure statistical rigor, we apply *conformalized quantile regression* [29], which does not require distributional or boundary assumptions. Our data-driven approach is simple to conduct and uses simulation data acquired from a commercial-grade vehicle dynamics simulation model.

III. SIMULATION OF THE MOTION CONTROLLER

In this work, we consider the performance of a motion controller under various dynamic conditions and under the influence of degradations and failures of the vehicle’s actuators. As a case study, we examine the fault-tolerant motion controller by [1]: Stolte et al. point out the advantages of fault-tolerant control as a strategy to widen the motion controller’s functional boundaries with respect to the possibility of actuator degradations and failures. The

nonlinear model-predictive controller is particularly designed for highly over-actuated vehicles, i.e., with four drivable, brakable and steerable wheels. Stolte et al. deploy and assess their controller in a simulation of *MOBILE*, which is an overactuated electric research vehicle [30]. Over-actuation is, for instance, also investigated in the *AUTOtech.agil* [31] and *UNICARagil* projects [32]: Each of the four *UNICARagil* vehicles yields the same actuator topology as *MOBILE*.

In the functional architecture of an automated vehicle according to Ulbrich et al. [15] (as shown in Figure 2), the motion controller is located below the behavior generation block and the subsequent generation of a target trajectory. To inform the behavior generation process, the vehicle relies on scene information derived by its environmental perception functions – for instance, incorporating an estimation of the road/lane geometry. Also, the scene contains information about the ego-vehicle’s internal state. The internal state may contain the dynamic state as well as degradation and failure states that are identified through failure detection and isolation functionalities (FDI). While tracking the target trajectory, the controller receives feedback about the vehicle’s current state through the vehicle’s state estimation functions, e.g., its current position, orientation, and speed. However, the measured information in every time step is subject to uncertainty, e.g., stochastic noise and/or a systematic offset [33], leading to a reduced accuracy.

At runtime, assessing the controller’s performance is important to ensure that safety requirements in terms of behavior execution are met. We consider the vehicle’s deviation from a planned motion for which the deviation introduced by the controller is a crucial determinant. However, simply monitoring a deviation metric itself means that a deviation is only measured if the vehicle is already deviating from its planned motion – potentially without being able to recover. Therefore, using predictive approaches is crucial. Hereafter, we present a data-based prediction model designed for this purpose. This prediction model can act as a heuristic during the behavior generation phase and is therefore located next to the behavior generation block in Figure 2. In order to train the prediction model, we present our simulation framework hereafter: Using an offline reference trajectory generator, we generate a large dataset of trajectories which meet domain-specific and dynamic constraints. In a commercial-grade simulation, the sampled trajectories can be used as the input for the controller and the controller’s performance is assessed with respect to a deviation metric.

A. Reference Trajectory Generation

To generate meaningful input data for the assessment of the motion controller, we aim to mimic the interface between online behavior planning and control in a numerical offline simulation. We use the same trajectory generator as in [34] which is used to generate reference trajectories for a *lane change* maneuver. The generator utilizes a simple point mass model and respects derived *maneuver parameters* as its mathematical constraints. Given these constraints, the optimization algorithm generates a time-minimal trajectory

TABLE I
SIMULATION/PREDICTION INPUTS AND OUTPUTS

Input Variables	Description	Unit	Value Range
ODD and Maneuver Parameters	r_q	direction	[1]
	$w_q^{\min}$	min. lane width	[m]
	$w_q^{\max}$	max. lane width	[m]
	$k_q^{\min}$	min. curvature	[m ⁻¹]
	$k_q^{\max}$	max. curvature	[m ⁻¹]
	μ_q	friction indicator	[1]
	v_{init}	initial speed	[km/h]
	v_{end}	final speed	[km/h]
	a_{max}	max. acceleration	[m/s ²]
Degrations, Failures	$\tilde{\delta}_{D,w}$	steering angle factor	[1]
	$\dot{\tilde{\delta}}_{D,w}$	steering rate factor	[1]
	$\tilde{\tau}_{D,w}$	wheel torque factor	[1]
Output	$\varepsilon_{\text{lat,max}}$	max. lateral deviation	m

for connecting the current and adjacent lane. We refer to [34] for a detailed description. All inputs/constraints of the generation are listed in Table I. The generator outputs a sequence of optimal states

$$\mathbf{x}^*(t) = \left(s^*(t), v^*(t), a_x^*(t), a_y^*(t), \psi^*(t), \dot{\psi}^*(t) \right)^T \quad (1)$$

that the trajectory tracking controller shall follow, with time t , Frenet arc length s , speed v , longitudinal acceleration a_x , lateral acceleration a_y , yaw angle ψ , and yaw rate $\dot{\psi}$. Superscript $*$ denotes reference values. The target and end pose of the maneuver are assumed to be located on the current and adjacent lane's center. Other trajectory generation models could be used here, e.g., explicitly specifying a target pose instead of minimizing time given dynamic limitations [35]. In the remainder of this paper, we only focus on the example of a *lane change*. Note that the optimization-based generation is not limited to this.

With respect to our previous work [34], we extend the trajectory generation approach as follows: To increase realism, we use a dataset of real road segment geometries, collected in the inner city ring road of *Braunschweig, Germany*. Conducting simulations based on these road segments allows us to respect realistic distributions of road widths and curvatures within the operational design domain, i.e., Braunschweig's ring road in this example. For each road segment with two lanes, we receive the lane center lines for the current and adjacent lines. We assume the distance between the two lane center lines to be equal to the width of each lane. While the lane's curvature and width are assumed to be always constant in [34], here, we store the current road segment's minimum and maximum curvature and width $k_q^{\min}, k_q^{\max}, w_q^{\min}, w_q^{\max}$ for a more detailed representation. Given a minimum segment length of 45 m, we receive a set of 403 road segments from the dataset. The value ranges of the considered indicators are given in Table I.

B. Simulation Setup

We apply each generated reference trajectory to the trajectory tracking controller by [1] that is implemented in *MATLAB* and uses a *nonlinear model predictive control*

(NMPC) approach. The controller relies on a nonlinear double-track model paired with an adapted version of the *Fiala* tire model [36]. The MPC uses control inputs $\mathbf{u}_{\text{MPC}} = (\tau_w, \dot{\delta}_w)$ to control the states $\mathbf{x}_{\text{MPC}} = (s, d, \psi, v_x^V, v_y^V, \dot{\psi}^V, \delta_w)$ of the prediction model and outputs optimal values of δ_w^* and τ_w^* to control the plant. Here, τ_w is the wheel torque and $\delta_w, \dot{\delta}_w$ are the steering angle and rate at each wheel $w \in \{\text{FL}, \text{FR}, \text{RL}, \text{RR}\}$. s is the Frenet arc length, d is the lateral deviation, ψ is the yaw angle, and v_x^V, v_y^V are the longitudinal and lateral velocities in the vehicle frame, respectively. In *IPG CarMaker*, the dynamics of the controlled vehicle given τ_w^* and δ_w^* are simulated, the resulting states are recorded and fed back into the control loop. For example, a resulting trajectory of the simulated vehicle is shown in Figure 4 (blue). While not considered in this work, adding noise, delays, and other disturbances could be done in future work. The deviation with respect to the reference trajectory can be measured through

$$\varepsilon_{\text{long}} = |s - s^*|, \varepsilon_{\text{lat}} = |d|, \varepsilon_{\text{yaw}} = |\psi - \psi^*| \quad (2)$$

for the deviation in the longitudinal, lateral, and yaw direction, respectively. All variables above depend on time t . By computing $\varepsilon_{(\cdot),\text{max}} := \max_t \varepsilon_{(\cdot)}(t)$, we can check whether the controller violates a specific deviation limit.

C. Definition of Degradation and Failure Parameters

We particularly focus on the impact of degradations and failures of the vehicle's actuators on the controller's performance. To cluster different degradation and failures, Stolte et al. [1] introduce several failure and degradation categories. They concern actuators that are relevant for motion control, i.e., steering, braking, and powertrain (and tire damage which is not in the focus of the present paper). In a vehicle dynamics model, a degradation or failure can be expressed in the easiest case through a limit or enforced value setting (e.g., to zero) of a specific model variable. In this work, we focus on degradations and failures of *MOBILE*'s steering, powertrain and braking systems, and hence, we represent a degradation or failure as a limited value (range) of the steering angle δ_w , steering rate $\dot{\delta}_w$, and/or wheel torque τ_w , which lays within (or is a subset of) the nominal value range. For the sake of brevity, we use $z \in \{\delta, \dot{\delta}, \tau\}$. We then express such a degradation or failure of an actuator as

$$z_w^{\min} \leq z_{D,w}^{\min} \leq z_w \leq z_{D,w}^{\max} \leq z_w^{\max}, \quad (3)$$

with $z_w^{\min}$ and $z_w^{\max}$ indicating the minimum and maximum values of a variable due to (nominal) technical limitations, and $z_{D,w}^{\min}$ and $z_{D,w}^{\max}$ indicating degraded value ranges.

We assume that z_w can take arbitrary values within the specified range. A *failure* is hence represented, e.g., through one parameter pair being set to zero. In the following, we only consider this type of failure, although other fixed values would also fulfill the criterion of a loss of functionality. In the following, we do not further distinguish between degradations and failures, and only use the term *degradation parameter*.

Hereafter, we consider $|z|_{D,w} := \max\{|z_{D,w}^{\min}|, |z_{D,w}^{\max}|\}$. This is an overestimation of the effects of degradations that

TABLE II
NEURAL NETWORK ARCHITECTURE AND TRAINING PARAMETERS

Layer	Neurons	Activation	Parameter	Value
Linear 1	19	ReLU	Batch Size	64
Linear 2	209	ReLU	Optimizer	RMSprop
Linear 3	209	ReLU	Learning Rate	0.0005
Linear 4	2	ReLU	Num. of Epochs	1200
			Early Stopping	True
			Normalization	Batch Norm

would only affect the positive or negative direction. Then, we can define the normalized degradation parameter

$$\tilde{z}_{D,w} := \frac{|z|_{D,w}}{z_w^{\max}} \in [0, 1], \quad (4)$$

that is normalized by its nominal (i.e., due to technical limitations) limit $z_w^{\max}$, with $\delta_w^{\max} = 30^\circ$, $\dot{\delta}_w^{\max} = 120^\circ/\text{s}$ and $\tau_w^{\max} = 2000 \text{ Nm}$, see [1].

IV. CONFORMALIZED PREDICTION MODEL

Using the data gained during the simulation runs, we develop a data-based model that forecasts the controller's deviation during a specific maneuver under current conditions – including potential actuator degradations and failures. At runtime, we aim to identify those actions that are likely to result in a sufficiently low deviation by exploring a finite set of possible maneuver parameters, given the availability of environmental and FDI information. Hereafter, we consider a lane change maneuver and focus on a single deviation metric as an example. We aim to identify

$$y = f(w_q^{\max}, w_q^{\min}, k_q^{\max}, \dots, \dot{\delta}_{D,w}, \tilde{z}_{D,w}) =: f(\xi), \quad (5)$$

with the input vector ξ and the lateral deviation $y = \varepsilon_{\text{lat,max}}$ as output, as described in Table I. We neglect r_q and μ_q as we limit our experiments to $r_q = \text{left}$ and $\mu_q = 1$ in the following. In this paper, we choose a *neural network* to approximate f with $\hat{y} = \hat{f}(\xi) \approx f(\xi)$. If considering multiple metrics would be desired, either training multiple models or using a multi-output model is possible. Alternatively, similar to [14], a classification approach could be adopted, i.e., using a binary state indicating whether a deviation threshold is exceeded. However, the regression model is more flexible as classification could be added as an additional layer.

Using a deterministic model suggests that any prediction could be made with certainty. However, several sources of uncertainty exist: For example, we only use selected indicators to express the road segment's geometry, e.g., min./max. curvature and width. Additionally, we find that the controller's numerical optimization procedure sometimes tends to find worse solutions for slightly less restricted problems than for more restricted ones, i.e., the deviation sometimes declines when vehicle dynamics demand and/or degradation or failure severity rise. Especially in dynamically challenging scenarios in the presence of degradations and failures, the controller has its difficulties to keep the vehicle on the desired trajectory, resulting in very large and fluctuating (across trajectories) deviation values. In such cases, it is also possible that even the commercial-grade simulation model reaches its limits. Therefore, we find a *stochastic* approach to be more suitable

for our problem. Also, other sources of uncertainty – such as measurement and actuator noise – should be covered by the selected approach if investigated in future work.

A. Conformalized Quantile Regression

We adopt a *quantile regression* [37] approach, which, instead of predicting a single value $\hat{y} \approx y$, results in an interval $\mathcal{I}(\xi_{\text{Test}}) := [\hat{y}_{\text{Lo}}, \hat{y}_{\text{Hi}}]$ for a test sample $\xi_{\text{Test}} \in \mathcal{X}_{\text{Test}}$ in which the true target value is expected to lie with a certain probability. The distance between $\hat{y}_{\text{Lo}}$ and $\hat{y}_{\text{Hi}}$ is referred to as the *interval length* $|\mathcal{I}|$. As ground truth, we consider the results of our simulation. In particular, we use the *conformalized quantile regression* (CQR) method [29, 38], which is a method for constructing prediction intervals for a given significance level α that are guaranteed to have the correct *coverage (probability)* C_α ,

$$P\{y_{\text{Test}} \in \mathcal{I}(\xi_{\text{Test}}) = [\hat{y}_{\text{Lo}}, \hat{y}_{\text{Hi}}]\} \geq 1 - \alpha =: C_\alpha. \quad (6)$$

The calibration of the prediction intervals is achieved using a calibration dataset $\mathcal{X}_{\text{Cal}}$ and holds on unseen test data $\mathcal{X}_{\text{Test}}$ if it is independent and identically distributed with respect to $\mathcal{X}_{\text{Cal}}$. For uncertainty estimation of neural network predictions, CQR is particularly appealing as it is distribution-free, adapts well to heteroscedastic data [29] and yields the aforementioned guarantee. Of course, this guarantee is only useful if the data used for calibration and evaluation is representative with respect to the system's real dynamics – which can only be achieved by using real-world data.

B. Dataset Creation




Fig. 3. Distribution of $\varepsilon_{\text{lat,max}}$ in the dataset

To generate training and evaluation datasets, we use previous simulation results. Given the set of real road segments described in Subsection III-A, we randomly pick N_S segments and generate lane change trajectories on each of them. Next, we employ the offline reference trajectory generator and randomly sample maneuver parameters from the value ranges in Table I. In terms of the maneuver parameters, it would be feasible to use data from real-world test drives within the desired ODD to generate a representative distribution. However, we note that dynamically challenging conditions are rarely encountered in real-world scenarios and therefore choose to draw the maneuver parameters from a uniform distribution, limited by the value ranges in Table I. Any conclusions drawn from this dataset are hence rather conservative.

We receive N_M different sets of maneuver parameters and generate $N_S \cdot N_M$ target trajectories in total. Finally, we additionally generate random sets of degradations and failures and simulate the trajectory tracking behavior for the nominal case as well as for each degradation and failure. We aim to design a predictor that works well even in rarely occurring degradation events and choose uniform distributions accordingly. We generate N_D different sets of degradation parameters, including the nominal case. In the remainder of this paper, we set $N_S = 50$, $N_M = 20$, $N_D = 20$, and receive $N_S \cdot N_M \cdot N_D = 20000$ samples. For each sample, $\varepsilon_{\text{lat,max}}$ is measured and stored. An extract of the distribution is shown in Figure 3. The frequency of large deviations is low due to the fault tolerance of the controller, the over-actuation of the vehicle, and the fact that we include the nominal case for each reference trajectory. In line with [10], we find that larger values of $\varepsilon_{\text{lat,max}}$ are more likely to occur when the maneuver is dynamically challenging or the vehicle's actuators are degraded.

C. Model Training and Evaluation

The neural network and CQR procedure are implemented using the *PyTorch* library, relying on the implementation of Romano et al. [29].² The data is split into a training and a testing set during training with a ratio of 80% to 20%, i.e., $|\mathcal{X}_{\text{Train}}| = 16000$ and $|\mathcal{X}_{\text{Test}}| = 4000$. Of the latter, $|\mathcal{X}_{\text{Cal}}| = 2000$ samples are held out for calibration. Standard scaling is applied to all inputs. While the guarantee in Equation 6 holds for any prediction model to which the calibration in [29] is applied, better models will result in shorter prediction intervals [38]. We therefore conduct a grid search to find optimal hyperparameters for the neural network, i.e., those that minimize the used “pinball loss” [29, 37]. For $C_\alpha = 99\%$, the network architecture that minimizes training loss is shown in Table II. For different values of α , we use the same network architecture but re-train the model.

TABLE III
EMPIRICAL COVERAGE AND INTERVAL LENGTHS ON $\mathcal{X}_{\text{Test}}$

Desired C_α	[%]	90.0	92.5	95.0	97.5	99.0	99.5
Empirical $\hat{C}_\alpha$	[%]	88.40	91.20	93.72	96.67	99.14	99.52
$Q_{10\%}(\mathcal{I})$	[cm]	2.30	1.95	2.91	1.73	5.09	6.01
Median($ \mathcal{I} $)	[cm]	5.20	4.76	5.73	4.84	8.21	9.56
Mean($ \mathcal{I} $)	[cm]	9.90	8.25	10.24	10.68	15.98	19.32
$Q_{90\%}(\mathcal{I})$	[cm]	25.25	19.20	25.23	31.50	41.98	53.49
$Q_{10\%}(\Delta\mathcal{I}_{\text{Hi}})$	[cm]	0.15	0.32	0.37	0.54	1.54	3.43
Median($\Delta\mathcal{I}_{\text{Hi}})$	[cm]	0.58	0.72	0.82	1.16	2.85	5.09
Mean($\Delta\mathcal{I}_{\text{Hi}})$	[cm]	2.93	1.86	3.21	4.51	7.80	12.16
$Q_{90\%}(\Delta\mathcal{I}_{\text{Hi}})$	[cm]	8.12	4.10	6.76	13.21	22.77	35.21

The results on the test set $\mathcal{X}_{\text{Test}}$ are given in Table III. We see that the empirically observed coverage $\hat{C}_\alpha$ is close to the desired coverage C_α (with a maximum deviation of 1.60%) across all desired values of C_α . As expected, the mean of the interval length $|\mathcal{I}|$ as well as its median grow when increasing the value of C_α . With respect to the distribution shown in Figure 3, we can conclude that the intervals are very tight for large portions of the data, i.e., for small values of $\varepsilon_{\text{lat,max}}$.

²<https://github.com/yromano/cqr>, accessed 2024-03-25.




Fig. 4. Nominal path (blue) for a lane change maneuver at $a_{\text{max}} = 3 \text{ m/s}^2$ compared to the resulting path (red) in case of D3 (see Table IV). Blue area indicates lateral spaced covered by prediction $\pm \hat{\varepsilon}_{\text{lat,Hi}}$ with $C_{\alpha,2} = 99\%$. Light red area indicates the additional space covered by the ego-vehicle $\pm w_{\text{veh}}/2$. Red shaded area indicates occupancy by an obstacle.

Given in-depth analyses, we find that larger intervals appear for high values of $\varepsilon_{\text{lat,max}}$. For later applications, we will focus on $\hat{\varepsilon}_{\text{lat,Hi}}$ as it is the upper limit of the interval in which we assume $\varepsilon_{\text{lat,max}}$ to lie with a probability of $\leq \alpha$. To evaluate the degree of overestimation of the CQR method, we compute the difference $\Delta I_{\text{Hi}} := \hat{\varepsilon}_{\text{lat,Hi}} - \varepsilon_{\text{lat,max}}$ for those cases where I actually covers the true value of $\varepsilon_{\text{lat,max}}$. As for $|\mathcal{I}|$, we find that the values of ΔI_{Hi} increase for an increasing value of C_α . We once again find that for higher true values of $\varepsilon_{\text{lat,max}}$, the predictor tends to overestimate.

V. APPLICATION OF THE PREDICTION MODEL

To illustrate the application of the trained prediction model, we consider an example scenario in which the automated vehicle enters a new road segment with two lanes and unexpectedly encounters an obstacle within the lane, i.e., a stopped vehicle. The automated vehicle may choose between two evasive maneuvers: (a) changing lanes to the left or (b) to stop in front of the obstacle. Assuming that the vehicle is able to determine its current dynamic state, degradation and failure information, and relevant operating conditions – such as the current road segment’s geometry – the vehicle shall make an informed decision.

TABLE IV
EXAMPLE DEGRADATIONS D0 . . . D5 AS RANDOM FIXED SETTINGS OF THE DEGRADATION PARAMETERS $\tilde{z}_{D,w}$ FOR THE EXAMPLE SCENARIO. BLUE COLOR INDICATES HIGHER AND RED COLOR LOWER VALUES.

Degradation Parameter Values	D0	D1	D2	D3	D4	D5
$\tilde{\delta}_{D,FL}$	100 %	100 %	30.28 %	100 %	4.72 %	44.55 %
$\tilde{\delta}_{D,FR}$	100 %	24.49 %	100 %	2.24 %	27.75 %	8.63 %
$\tilde{\delta}_{D,RL}$	100 %	33.13 %	100 %	100 %	21.94 %	12.8 %
$\tilde{\delta}_{D,RR}$	100 %	100 %	48.38 %	100 %	100 %	100 %
$\tilde{\delta}_{D,FL}^*$	100 %	100 %	3.92 %	100 %	93.45 %	37.59 %
$\tilde{\delta}_{D,FR}^*$	100 %	85.24 %	100 %	36.29 %	76.94 %	33.5 %
$\tilde{\delta}_{D,RL}^*$	100 %	34.88 %	100 %	100 %	21.62 %	60.5 %
$\tilde{\delta}_{D,RR}^*$	100 %	100 %	74.07 %	100 %	100 %	100 %

We assume the vehicle’s current speed to be $v_{\text{init}} = 50 \text{ km/h}$, which meets the German inner-city speed limit. While many parameters are set due to the initial scene (such as geometrical characteristics of the road segment and the vehicle’s degradation and failure state), the vehicle may still choose different maneuver options, apply different acceleration limits and target speeds. For the sake of brevity,

we only consider degradations of the steering actuators and their impact on the vehicle's lateral dynamics during a possible lane change maneuver at different acceleration limits. We further assume $v_{\text{end}} = 50 \text{ km/h}$ and $\mu_q = 1$.

In order to inform the behavior generation, we make the following considerations: We consider the current value of $w_q^{\text{min}} = 2.78 \text{ m}$ for the road segment's minimum width and the width $w_{\text{Veh}} = 1.96 \text{ m}$ of our test vehicle *MOBILE*. Given the worst case assumption that the maximum lateral deviation is reached where the road segment is narrowest, the free lateral space is $\varepsilon_{\text{lat,free}} = 41 \text{ cm}$ both on the left and right side. Accordingly, the vehicle should not execute maneuvers during which a lateral deviation of $\varepsilon_{\text{lat,free}}$ is exceeded. Hereafter, we only consider the upper prediction bound $\hat{\varepsilon}_{\text{lat,Hi}}$ as described. We visualize our approach in Figure 4 – where it becomes clear that the application of CQR leads to an expected “corridor” of lateral control deviation. However, in addition to the controller's performance, it is crucial to note that the planned trajectory may not ideally connect the current and target lane center, e.g., due to additional uncertainty in the environmental perception and/or map-based information. Furthermore, disturbances in the vehicle's state estimation may arise. Both will require additional lateral margins that we do not address in our framework.

TABLE V

PREDICTION $\hat{\varepsilon}_{\text{lat,Hi}}$ VS. TRUE LATERAL DEVIATION $\varepsilon_{\text{lat,max}}$ IN [m].
BLUE COLOR INDICATES LOW AND RED COLOR INDICATES HIGH TRUE
VALUES. FOR $C_{\alpha,1} = 95\%$ (TOP) AND $C_{\alpha,2} = 99\%$ (BOTTOM).

	D0	D1	D2	D3	D4	D5
Acceleration Limit in [m/s ²]						
1.0	0.03/0.02	0.03/0.02	0.04/0.02	0.03/0.10	0.03/0.02	0.23/0.02
2.0	0.04/0.03	0.04/0.03	0.06/0.03	0.12/0.19	0.26/0.27	0.34/0.08
3.0	0.06/0.05	0.06/0.05	0.08/0.06	0.24/0.19	0.59/0.39	0.40/0.25
4.0	0.08/0.07	0.13/0.07	0.19/0.10	0.43/0.25	0.69/0.38	0.78/0.55
5.0	0.10/0.09	0.17/0.14	0.20/0.13	0.55/0.33	0.63/0.39	0.67/0.46

	D0	D1	D2	D3	D4	D5
Acceleration Limit in [m/s ²]						
1.0	0.08/0.02	0.08/0.02	0.09/0.02	0.12/0.10	0.13/0.02	0.17/0.02
2.0	0.09/0.03	0.11/0.03	0.11/0.03	0.27/0.19	0.18/0.27	0.40/0.08
3.0	0.11/0.05	0.13/0.05	0.15/0.06	0.53/0.19	0.44/0.39	0.62/0.25
4.0	0.13/0.07	0.16/0.07	0.30/0.10	0.75/0.25	0.68/0.38	0.77/0.55
5.0	0.15/0.09	0.26/0.14	0.42/0.13	0.89/0.33	0.87/0.39	0.71/0.46

At runtime, the automated vehicle can use the prediction model and is able to feed possible maneuver (parameter) options into the model, in order to predict the expected bounds of the lateral deviation. Also, we assess the impact of different degradations that are defined in Table IV. In Table V, we present the results of the prediction model: The predicted values of $\hat{\varepsilon}_{\text{lat,Hi}}$ are visualized along with ground truth values of $\varepsilon_{\text{lat,max}}$ from the simulation. As an example, we set $C_{\alpha,1} = 95\%$ (top) and $C_{\alpha,2} = 99\%$ (bottom). In Table V, given the severity of the respective degradation D0...D5 as well as the acceleration limit a_{max} , we see strong changes in both the actual and predicted deviation.

For some cases – even for high accelerations – the expected deviations is much smaller than the value of $\varepsilon_{\text{lat,free}}$, resulting in a remaining margin of $\varepsilon_{\text{lat,free}} - \hat{\varepsilon}_{\text{lat,Hi}}$, which might be required to take account for additional uncertainty in the closed loop. For D4, D5, we see great deviations from the reference trajectory in terms of $\varepsilon_{\text{lat,max}}$.

Especially for high accelerations, the vehicle's lateral dynamics are significantly impacted by the degradation. Furthermore, the predicted values of $\hat{\varepsilon}_{\text{lat,Hi}}$ are very high in these cases and greatly exceeds the true value of $\varepsilon_{\text{lat,max}}$. This is necessary to take account of the uncertainty as described in previous sections. Given Table III, it becomes clear that the degree of overestimation is related to the value of α . To illustrate that, we consider the example of performing a lane change with $a_{\text{max}} = 3 \text{ m/s}^2$ given degradation D3: The model trained for $C_{\alpha,1} = 95\%$ predicts a maximum lateral deviation of up to 0.24m, thus exceeding the true value of $\varepsilon_{\text{lat,max}} = 0.19 \text{ cm}$. However, the prediction $\hat{\varepsilon}_{\text{lat,Hi}}$ is still smaller than $\varepsilon_{\text{lat,free}}$ and hence, the vehicle's action space is not necessarily limited due to an overly conservative prediction (however, it might be when considering additional uncertainty in other system parts). In terms of $C_{\alpha,2}$ we see that the predicted value is even higher, $\hat{\varepsilon}_{\text{lat,Hi}} = 0.53 \text{ m}$, and therefore larger than $\varepsilon_{\text{lat,free}}$, i.e., the corresponding lane change should be rated as infeasible given the prediction. Finetuning the value for α is therefore crucial for the behavior of the prediction model.

VI. CONCLUSION AND FUTURE WORK

In this work, we aim to create a self-representation model of an automated vehicle's motion controller in a data-driven manner. Using an offline simulation, we generate a set of trajectories and train a neural network to predict the maximum lateral deviation introduced by the controller given the maneuver to be executed, operational conditions, and potential actuator degradations or failures. Using conformalized quantile regression, we acknowledge the uncertainty in the prediction and provide prediction intervals for the deviation metric that meet a given significance level. We demonstrate that our predictor is able to estimate the controller's performance with the expected coverage and with low interval length for most cases. Our results show that the feasibility of executing a specific lateral maneuver is a function of the vehicle's operational conditions (i.e., the available free lateral space), its dynamic state, and the presence of an actuator degradation or failure. Our method allows to assess this intuitive relation quantitatively and with the required statistical rigor. Overall, while empirical and desired coverages align well, we find that our method is rather conservative. While conservatism is often desired in safety-critical applications, future work should focus on reducing the average interval length in order to not unnecessarily restrict the vehicle's action space.

While focusing on a lane change maneuver and the vehicle's lateral deviation as an example, we state that with a more complex predictor that covers multiple maneuvers as well as additional performance/deviation metrics, the behavior

generation process can be informed more comprehensively. However, we cannot rule out the possibility that when applying threshold values for different performance measures, the controller's performance could potentially be insufficient for any maneuver option under certain conditions. The behavior generation in such cases can be informed by the predictor but not resolved.

Furthermore, we once again emphasize that further considerations are required to set an appropriate value for C_α . While the statistical representation allows us to inform a stochastic decision-making framework, defining such a value requires a decision with respect to the acceptable level of uncertainty, which is out of the scope of this paper. While the prediction intervals exemplified in this work maintain a feasible length, large values of C_α lead to longer prediction intervals which might render the predictor useless in most scenarios, requiring an enhanced design thereof, e.g., involving more input parameters or a sequential model instead of a sparse "one-shot" feed-forward model. The same might apply to the impact of additional disturbances.

Finally, conducting our analysis and model training using real-world data is a crucial step to ensure that the predictor's output is reliable in real-world experiments, e.g., using a closed test track and high fidelity state estimation to safely acquire ground truth data. Real-world data is expected to introduce even more uncertainty into the prediction, leading to new challenges in the future.

ACKNOWLEDGEMENT

We thank Jens Rieken for providing the road segment data. We thank Xiao Yun for his support with the implementation.

REFERENCES

- [1] T. Stolte et al., "Toward Fault-Tolerant Vehicle Motion Control for Over-Actuated Automated Vehicles: A Non-Linear Model Predictive Approach," *IEEE Access*, vol. 11, pp. 10 499–10 519, 2023.
- [2] International Organization for Standardization, *ISO 26262 Road Vehicles - Functional Safety*, 2018.
- [3] T. Stolte et al., "A Taxonomy to Unify Fault Tolerance Regimes for Automotive Systems: Defining Fail-Operational, Fail-Degraded, and Fail-Safe," *IEEE Transactions on Intelligent Vehicles*, 2021.
- [4] M. Nolte, I. Jatzkowski, S. Ernst, and M. Maurer, "Supporting Safe Decision Making Through Holistic System-Level Representations & Monitoring - A Summary and Taxonomy of Self-Representation Concepts for Automated Vehicles," *arXiv:2007.13807*, 2020.
- [5] M. Maurer, "Flexible Automatisierung von Straßenfahrzeugen mit Rechnersehen," PhD Th., Uni. der Bundw., 2000.
- [6] A. Reschka, "Fertigkeiten- und Fähigkeitsgraphen als Grundlage des sicheren Betriebs von automatisierten Fahrzeugen im öffentlichen Straßenverkehr," PhD Th., TU Braunschweig, 2017.
- [7] I. Jatzkowski et al., "Zum Fahrmanöverbegriff im Kontext automatisierter Straßenfahrzeuge," Tech. Rep., 2021.
- [8] M. Nolte et al., "Towards a skill- and ability-based development process for self-aware automated road vehicles," in *Proc. of ITSC*, 2017, pp. 1–6.
- [9] T. Stolte et al., "Towards Safety Concepts for Automated Vehicles by the Example of the Project UNICARagil," in *Proc. of 29th Aachen Colloquium Sustainable Mobility*, 2020.
- [10] A. Da Silva, C. Birkner, R. N. Jazar, and H. Marzbani, "Crash-Prone Fault Combination Identification for Over-Actuated Vehicles During Evasive Maneuvers," *IEEE Access*, vol. 12, pp. 37 256–37 275, 2024.
- [11] E. Ratner, C. J. Tomlin, and M. Likhachev, "Operating with Inaccurate Models by Integrating Control-Level Discrepancy Information into Planning," in *Proc. of International Conference on Robotics and Automation (ICRA)*, 2023, pp. 7823–7829.
- [12] M. Noseworthy et al., "Active Learning of Abstract Plan Feasibility," in *Proc. of Robotics Science and Systems XVII*, 2021.
- [13] F. Castaneda et al., *Recursively Feasible Probabilistic Safe Online Learning with Control Barrier Functions*, arXiv:2208.10733, 2023.
- [14] D. McConachie, T. Power, P. Mitrano, and D. Berenson, "Learning When to Trust a Dynamics Model for Planning in Reduced State Spaces," *Robotics and Automation Letters*, vol. 5, no. 2, pp. 3540–3547, 2020.
- [15] S. Ulbrich et al., "Towards a functional system architecture for automated vehicles," *arXiv:1703.08557*, 2017.
- [16] M. Althoff, "Reachability analysis and its application to the safety assessment of autonomous cars," PhD Th., TU München, 2010.
- [17] M. Althoff and J. M. Dolan, "Online Verification of Automated Road Vehicles Using Reachability Analysis," *IEEE Transactions on Robotics*, vol. 30, no. 4, pp. 903–918, 2014.
- [18] K. Leung et al., "On infusing reachability-based safety assurance within probabilistic planning frameworks for human-robot vehicle interactions," in *Proc. of Int. Symposium on Experimental Robotics*, 2020, pp. 561–574.
- [19] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-Jacobi reachability: A brief overview and recent advances," in *Proc. of 56th Conference on Decision and Control (CDC)*, 2017, pp. 2242–2253.
- [20] I. M. Mitchell and C. J. Tomlin, "Overapproximating Reachable Sets by Hamilton-Jacobi Projections," 2003.
- [21] A. P. Vinod and M. M. Oishi, "Stochastic reachability of a target tube: Theory and computation," *Automatica*, vol. 125, p. 109 458, 2021.
- [22] A. Shetty and G. X. Gao, "Predicting State Uncertainty Bounds Using Non-Linear Stochastic Reachability Analysis for Urban GNSS-Based UAS Navigation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 9, pp. 5952–5961, 2021.
- [23] V. Rubies-Royo, D. Fridovich-Keil, S. Herbert, and C. J. Tomlin, "A Classification-based Approach for Approximate Reachability," in *Proc. of Int. Conference on Robotics and Automation (ICRA)*, 2019, pp. 7697–7704.
- [24] T. Lew and M. Pavone, "Sampling-based reachability analysis: A random set theory approach with adversarial sampling," in *Proc. of Conference on robot learning*, 2021, pp. 2055–2070.
- [25] M. Chen et al., "Decomposition of Reachable Sets and Tubes for a Class of Nonlinear Systems," *IEEE Transactions on Automatic Control*, vol. 63, no. 11, pp. 3675–3688, 2018.
- [26] T. Lombaerts, P. Chu, J. A. B. Mulder, and O. Stroosma, "Fault tolerant flight control, a physical model approach," *Advances in Flight Control Systems*, 2011, Publisher: IntechOpen.
- [27] D. Asadi, M. Sabzehparvar, E. M. Atkins, and H. A. Talebi, "Damaged Airplane Trajectory Planning Based on Flight Envelope and Motion Primitives," *Journal of Aircraft*, vol. 51, no. 6, pp. 1740–1757, 2014.
- [28] M. Nolte, R. Schubert, C. Reisch, and M. Maurer, "Sensitivity Analysis for Vehicle Dynamics Models – An Approach to Model Quality Assessment for Automated Vehicles," in *Proc. of Intelligent Vehicles Symposium (IV)*, 2020, pp. 1162–1169.
- [29] Y. Romano, E. Patterson, and E. Candes, "Conformalized Quantile Regression," in *Proc. of Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [30] P. Bergmiller, "Towards Functional Safety in Drive-by-Wire Vehicles," PhD Thesis, TU Braunschweig, 2014.
- [31] R. van Kempen et al., "AUTOTech.agil: Architecture and Technologies for Orchestrating Automotive Agility," in *Proc. of 32. Aachen Colloquium Sustainable Mobility*, 2023.
- [32] T. Woopen et al., "UNICARagil - Disruptive modular architectures for agile, automated vehicle concepts," in *27th Aachen Colloq.*, 2018.
- [33] M. Werling, "Ein neues Konzept für die Trajektoriengenerierung und -stabilisierung in zeitkritischen Verkehrsszenarien," PhD Th., KIT, 2011.
- [34] R. Schubert, M. Nolte, A. de La Fortelle, and M. Maurer, "ODD-Centric Contextual Sensitivity Analysis Applied To A Non-Linear Vehicle Dynamics Model," in *Proc. of IV*, 2023, pp. 1–8.
- [35] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal trajectory generation for dynamic street scenarios in a Frenet Frame," in *Proc. of ICRA*, 2010, pp. 987–993.
- [36] R. Y. Hindiyeh and J. Christian Gerdes, "A Controller Framework for Autonomous Drifting: Design, Stability, and Experimental Validation," *Journal of Dynamic Systems, Measurement, and Control*, vol. 136, no. 051015, 2014.
- [37] R. Koenker and G. Bassett Jr, "Regression quantiles," *Econometrica: journal of the Econometric Society*, pp. 33–50, 1978, Publisher: JSTOR.
- [38] A. N. Angelopoulos and S. Bates, *A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification*, arXiv:2107.07511 [cs, math, stat], 2022.