Robust Conformal Prediction for Adaptive Motion Planning Among Interactive Agents

RO57035 Thesis

Mayank Prashar

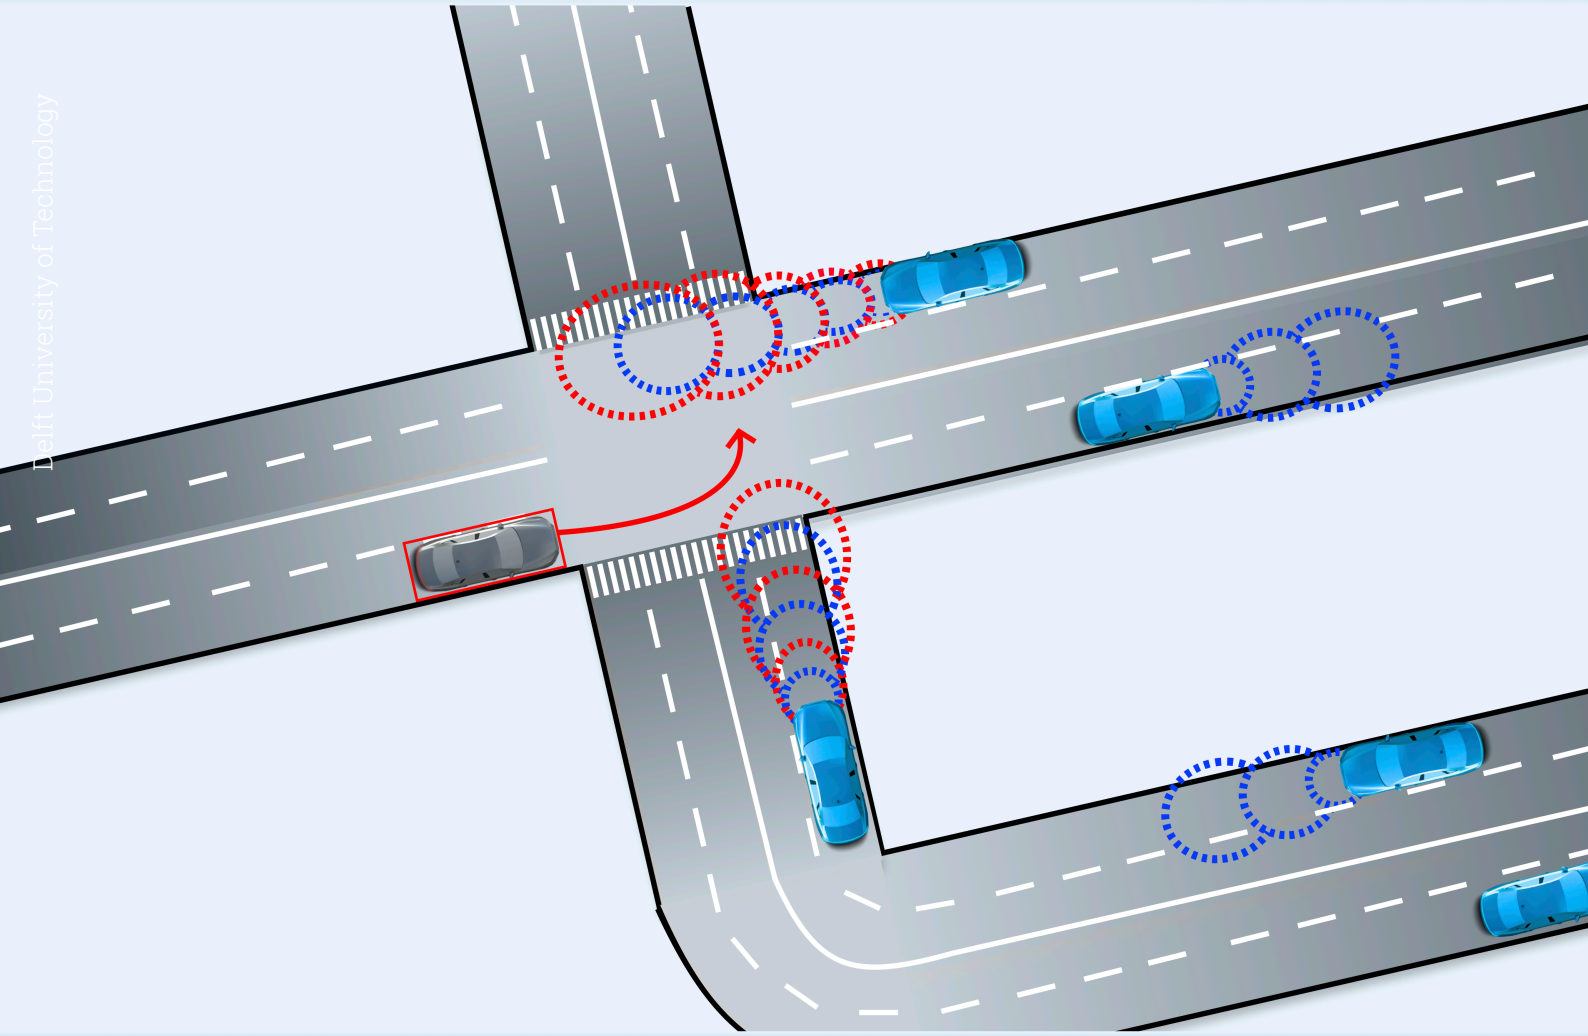

Delft University of Technology

TU Delft

Robust Conformal Prediction for Adaptive Motion Planning Among Interactive Agents

by

Mayank Prashar

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Thursday August 29, 2024 at 11:30 AM.

Student number: 4643526

Thesis committee:		
	Prof. Javier Alonso Mora	TU Delft, Main supervisor & chair
	Prof. Lars Lindemann	University of Southern California, Daily supervisor
	Prof. Cosimo Della Santina	TU Delft
	Prof. Luca Laurenti	TU Delft

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Abstract

Autonomous motion planning requires the ability to safely reason about learned trajectory predictors, particularly in settings where an agent can influence other agents' behavior. These learned predictors are essential for anticipating the future states of uncontrollable agents, whose decision-making process can be difficult to model analytically. Thus, uncertainty quantification of these predictors is crucial for ensuring safe planning and control. In this work, we introduce a framework for interactive motion planning in unknown dynamic environments with probabilistic safety assurances. We adapt a model predictive controller (MPC) to distribution shifts in learned trajectory predictors when other agents react to the ego agent's plan. Our approach leverages tools from conformal prediction (CP) to detect when the other agent's behavior deviates from the training distribution and employs robust CP to quantify the uncertainty in trajectory predictions during these agent interactions. We propose a method for estimating interaction-induced distribution shifts during runtime and the Huber quantile for enhanced outlier detection. Using a KL divergence ambiguity set that upper bounds the distribution shift, our method constructs prediction regions with probabilistic assurances in the presence of distribution shifts caused by interactions with the ego agent. We evaluate our framework in interactive scenarios involving navigation around autonomous vehicles in the BITS simulator, demonstrating enhanced safety and reduced conservatism. The code of this work will be made publicly available [1]

[1] `https://github.com/mayank176/robust_CP_interactive_planning`

Acknowledgments

First and foremost, I would like to express my gratitude to my thesis supervisors Prof. Lars Lindemann and Prof. Javier Alonso Mora for their invaluable guidance, support, and feedback throughout this thesis project. I would like to sincerely thank Prof. Lars Lindemann for providing me with the opportunity to be a visiting student in his lab at the USC Thomas Lord Department of Computer Science, introducing me to a fascinating thesis topic, and providing expertise and feedback that have shaped this thesis to its current form. I would like to deeply thank Prof. Javier Alonso Mora for his guidance and support in doing my thesis abroad and in becoming a more proficient researcher. I would also like to extend my appreciation to Prof. Luca Laurenti, who, along with my supervisors, supported my endeavor to conduct my thesis research abroad. Additionally, I would like to thank the thesis committee for their time, consideration and support in reviewing my thesis report and presentation.

I am deeply grateful to my family for their persistent and unconditional support and for making it possible for me to pursue this education. I would also like to thank my friends in both Delft and Los Angeles for their invaluable support along my academic journey.

Lastly, I would like to express my sincere gratitude to the Delft University of Technology, where I completed both my BSc. in Aerospace Engineering and MSc. in Robotics. The educational experience has been instrumental in shaping my academic and professional journey.

Mayank Prashar
Delft, August 2024

Contents

List of Figures

List of Tables

1

Introduction

1.1. Problem Statement

The increasing adoption of autonomous systems across various domains such as autonomous vehicles (AV's), drones, legged robots, manipulators, and collaborative robots- has sparked a significant push to make these systems safe, capable of reasoning under uncertainty, and able to learn and adapt to unstructured environments. Among these systems, AV's (see Figure 1.1) have garnered substantial investment and attention due to their significant potential to impact mobility and the environment ranging from improved road safety [54], reduced fuel consumption [5] to economic advantages [19] and enhanced on-demand mobility solutions [80]. Recent industry surveys have identified consumer safety as a primary bottleneck to the widespread adoption of AV's [17]. This challenge has prompted extensive research and development efforts in both academia and industry [37], [21], [63], [71] [55], [85], [94], [46], [52], focusing on ensuring safe operations of AV's in unstructured environments. A key aspect of this safety challenge lies in the ability of autonomous systems to reason about and interact safely with other agents in dynamic environments [38], [47], [61].

Traditional approaches to motion planning for autonomous systems have relied on predict-then-plan schemes. This involves first predicting the motion of other agents using either heuristic [77], [32], or learned models [1] and then creating a plan for the ego agent based on these predictions. However, these methods do not consider the influence of the ego agent's plan on other agents' behaviors, leading to overly conservative behaviors that have been shown to lead to unsafe situations [76]. Recent works have addressed this limitation of predict-then-plan schemes through the combination of learning-based and optimization-based methods in planning and control frameworks. These works aim to leverage the expressivity of learning-based methods to capture interactions [60] and the constraint satisfaction guarantees of optimization-based methods [49]. Notable approaches of this methodology in industry include Waymo's use of transformers for behavior prediction, integrated with an optimization-based planner in a receding horizon fashion [84] as well as Tesla's tree-based planner to model intra-object dependencies (i.e., interactions), integrated with a neural planner trained on offline optimized trajectories [70].

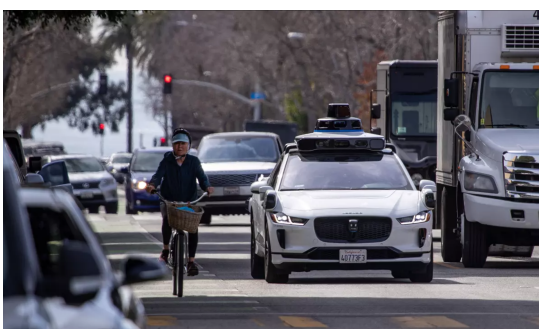

Figure 1.1: Waymo full self-driving car [28]

While the above methods have made significant progress, a critical gap remains in providing probabilistic safety assurances for interactive planning. Existing methods for uncertainty quantification in interactive planning assume specific agent behavior models or predictors. These approaches restrict the type of disturbances and prediction models considered, lack safety assurances, and show overly conservative unsafe behavior. Moreover, current state-of-the-art prediction models capture interactions through ego-conditioning, where the prediction of other (non-ego) agents is conditioned on the ego agent's motion plan. However, integrating ego-conditioned models with optimization-based planners is computationally prohibitive, making it challenging to design planners that can leverage these prediction models and maintain the constraint satisfaction guarantees of optimization-based methods, which are vital for providing safety assurances.

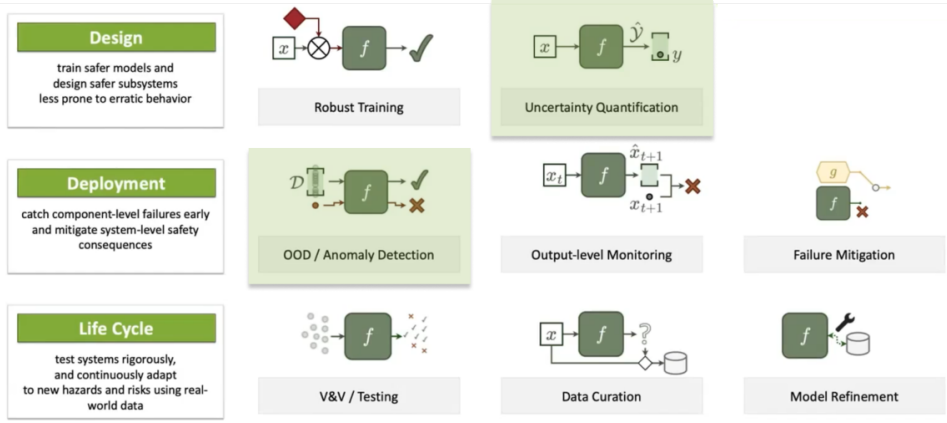

Figure 1.2: Pillars of AI Safety Research for Safety-Critical Robot Autonomy, adapted from [56]. This thesis proposes methods that integrate uncertainty quantification and anomaly detection (highlighted in green) in interactive planning frameworks to provide safety assurances.

To address these challenges, we focus on using methods from uncertainty quantification and anomaly detection to provide safety and adaptability of planners under interactive scenarios (see Figure 1.2). In particular, this thesis proposes to bridge this gap by developing an MPC-based framework for safe interactive planning with valid safety assurances. We use tools from conformal prediction (CP), a statistical tool that rigorously quantifies uncertainty in complex deep-learned trajectory predictors without making assumptions about the data distribution or predictive model [78], [62]. By incorporating trajectory prediction regions to design safety constraints, CP has been established as an effective framework in planning tasks [43], [23]. However, CP assumes exchangeability between train and test distributions i.e., such that there is no distribution shift between the distributions. As current training data predominantly captures well-structured, non-interactive behavior [83], [33], real-world test scenarios often involve interactive behaviors where the ego agent's plan can influence other agents' behaviors. This discrepancy necessitates a method to quantify uncertainty in trajectory predictors under interaction-induced distribution shifts. To achieve this, we use CP to quantify uncertainty in learned trajectory predictors under interactions. To ensure safety amid interaction-induced distribution shifts, we perform outlier detection using CP to detect when the other agent's behavior deviates from the training distribution and employ robust CP methods [48] to capture nearby distributions upper bounded by an KL divergence for these outliers. This approach allows us to maintain safety assurances when encountering behaviors that deviate from the training distribution due to interactions. We summarize our research questions in Section 1.2, contributions in Section 1.3, and provide a formal problem formulation in Section 3.1.

1.2. Research Questions

- *What techniques can **quantify** the uncertainty of complex deep-learned trajectory forecasting methods, and how can we detect outliers reliably?*

  In particular, we focus on using techniques from conformal prediction (CP), an uncertainty quantification tool that makes no assumptions about the data distribution or predictive model. We explore variants of CP, and how they can be used to quantify uncertainty in learned trajectory predictors under interaction-induced distribution shifts.

- *How can learning-based trajectory forecasting methods be **integrated** into interactive motion planning frameworks with probabilistic safety assurances?*

  We focus on making minimal assumptions on the predictor type to create a framework that can be compatible with a wide range of predictive models, including both ontological and phenomenological methods.

- *How can we **model** the interactions between agents when the counterfactual motion of other agents is unknown?*

  Given the challenges of predicting interactions, we focus on how we can develop an explainable model that captures interaction-induced uncertainty at a distributional level rather than at a (single) trajectory level.

1.3. Contributions

- We develop an adaptive MPC-based framework for safe interactive planning that accounts for distribution shifts in learned trajectory predictors resulting from ego-conditioning. Our approach leverages tools from conformal prediction to detect and quantify these shifts, enabling the planner to adapt during runtime.
- We propose a method for estimating interaction-induced distribution shifts during runtime and a Huber quantile for enhanced outlier detection.
- We empirically show that our framework works in interactive scenarios involving navigation around autonomous vehicles in the Bi-level imitation for traffic simulation (BITS) simulator [87].

1.4. Outline

The remainder of this thesis is organized as follows: Chapter 2 provides a comprehensive overview of related works in interactive planning, planning with safety assurances, motion prediction, and outlier detection techniques. Chapter 3 formally introduces the problem addressed in this thesis, it then introduces the reader to conformal prediction (CP), an uncertainty quantification tool, and presents methods for applying CP in settings involving distribution shifts and finite-horizon trajectories. Chapter 4 provides the main contribution of this thesis, proposing a safe interaction-robust model predictive control (MPC)-based planner. This chapter details a framework that encompasses enhanced outlier detection using the Huber quantile, a technique to estimate interaction-induced distribution shifts, and an MPC that incorporates these elements to ensure safety assurances for interactive planning. Chapter 5 presents and analyzes simulation results using the Bi-level Imitation for Traffic Simulation (BITS) simulator. This chapter discusses implementation details and provides both qualitative and quantitative analyses of the results. Chapter 6 concludes the thesis by summarizing the key findings and proposing future research directions and open problems in the field. Appendix A includes the research paper resulting from this thesis work.

2

Related Work

2.1. Interactive Planning

Initial approaches in interactive planning include modeling the behavior of other agents with Gaussian uncertainty without considering the influence of the ego agent. However, ignoring the influence of the ego agent was found to result in conservative plans, as shown in the freezing robot problem [76]. To take into account interaction, approaches employ joint optimization via Gaussian processes [76]. Furthermore, interaction potential terms are incorporated with Gaussian processes to capture interactions [75]. Additionally, other works explore artificial potential fields [86] and reinforcement learning to capture interactions in dense crowds [11].

Reachability analysis is another approach for modeling multi-agent interactions [51]. This method provides a framework for determining possible system states given specific dynamics and time horizons. Forward reachability, or open-loop safety, has been utilized to identify potential states of other agents, allowing the ego agent to plan collision-free trajectories [3]. While this approach offers safety assurances, it often leads to conservative behaviors and can result in planning infeasibility, particularly over extended time horizons. Backward reachable sets, or closed-loop safety, which involve backward propagation of joint agent dynamics, can mitigate some conservatism but are often computationally intractable.

Other model-based approaches use game theory to formulate coupled trajectory optimization problems, assuming the availability of objective models for all agents. Such approaches are well-known for reasoning about bi-directional interactions. However, they assume simplified behavior models such as the rational model [58] or noisy Boltzmann models [27] and are computationally prohibitive when scaled to more agents. Other works have explored partially observable Markov decision processes (POMDP) to model interactions; however, it can result in an intractable partially observable stochastic game (POSG) [59].

2.2. Motion Planning with Probabilistic Safety Assurances

Approaches for providing probabilistic safety assurances in planning frameworks typically attempt to quantify uncertainty in other agents' motions. A common assumption is that an agent's motion is a random process, and its distribution can be determined provided independent and identically distributed (i.i.d.) samples of trajectories.

Various methods have been developed to address this challenge. Scenario optimization [10] computes a predicted set for other agents based on previously observed scenarios, which are then formulated into chance constraints. These chance constraints have been applied using various models, including Gaussian process models [44], Gaussian uncertainty [93], and arbitrary uncertainty distributions [31]. Other approaches have leveraged the access to human-in-the-loop datasets of driving behaviors to solve mixed integer linear programs to find the minimum area set that contains trajectories within a given confidence level [24]. Similarly, quantile regression has been used to learn bounds over trajec-

tory datasets for a specified confidence level, which are then incorporated in a Tube MPC scheme [25]. Additional techniques for generating bounds on trajectory prediction models include post-bloating, support vector machines, and conformal prediction as shown in [14]. Fundamentally, these methods can be viewed as addressing a classification problem, aiming to ensure that 1 - δ of predicted trajectories should be within some learned bounds [16].

Recent works have focused on using conformal prediction (CP) to provide safety assurances. Lindemann et al. [43] use CP to construct prediction regions for learning-based trajectory predictors (see Figure 2.1) and use these to design a model predictive controller that uses these prediction regions as collision avoidance constraints. This framework provides probabilistic safety assurances over the task horizon by using Boole's inequality. However, this method is conservative and results in larger prediction regions for longer prediction horizons that can potentially make the MPC infeasible. Yu et al. [89] extended the work from [43] to further reduce the conservatism of prediction regions by a normalization technique inspired by [18]. In another method to reduce conservatism, Tonkens et al. [74] use copulas to construct prediction regions for longer horizons. Taking into account general distribution shifts, Dixit et al. [23] use adaptive CP to construct the prediction regions but only provide asymptotically valid guarantees [29], [30]. We note that the aforementioned works provide safety assurances over the task horizon. Additionally, these methods do not guarantee planner feasibility, besides recent work that provides feasibility guarantees with a shrinking-horizon framework, where the horizon extends to the end of the mission and is solved at each time step [65].

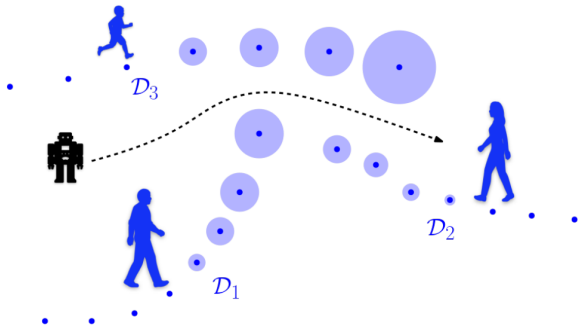

Figure 2.1: Safe planning framework using CP regions (blue circles) to quantify the uncertainty of learned trajectory predictors [43].

Other methods that address distribution shifts [20] propose a method that can detect general distribution shifts using out-of-distribution (OOD) detection with an ensemble of multi-layer perceptrons (MLPs). They then use a conservative reachability-based fallback controller when such shifts are identified. However, their approach uses simplified MLPs that don't represent SOTA trajectory forecasting methods [60], [90], making it difficult to know if this OOD method would translate to SOTA models in a computationally efficient manner as they use an ensemble of predictors. In another line of work, Lekeufack et al. [42] introduce conformal decision theory, an approach that couples prediction uncertainty with system control through a risk notion adjusted by a conformal control variable λ. Rather than using prediction regions for navigation, the method employs λ to inform the controller. However, this approach may still lead to conservative actions and requires careful parameter tuning.

Other works have explored different aspects of the planning problem. Muthali et al. [53] have constructed reachable sets calibrated with CP. Sun et al. [68] focus on uncertainties in the dynamics model and propose using CP with diffusion dynamics models in offline reinforcement learning tasks. In the context of run-time monitors, Zhao et al. [91] propose monitoring algorithms using robust CP [48] under signal temporal logic (STL) tasks for general distribution shift, though the method assumes a worst-case shift at all times without adapting to actual distribution shifts during test time. Other works propose conformal predictive safety filters [67] for reinforcement learning which ensures the learned policy avoids the CP regions, and Luo et al. [45] use CP to detect unsafe situations with a guarantee on the false negative rate. As a result, the following challenges remain: ensuring planner feasibility while providing safety assurances, adapting to real-time interaction-induced distribution shifts, and scaling outlier detection and/or OOD detection methods to more sophisticated trajectory forecasting models.

2.3. Motion Prediction

Trajectory forecasting methods have been extensively studied in literature and can be broadly categorized into ontological and phenomenological methods. Ontological methods capture other agents' interaction dynamics by assuming they strongly tend to avoid collisions (i.e., agents are cooperative). The optimal reciprocal collision avoidance (ORCA) [7] [2], the social force model [32], the intelligent driver model [77], and the interaction potential for Gaussian processes [75] are examples of ontological methods. An advantage of these methods is that they explicitly derive interaction dynamics, making them explainable. Such models work well for short time horizon (<0.5s) trajectory predictions. However, a drawback of these methods is that they do not consider interaction history and scene context and only use the current state to model interactions, making them unsuitable for longer horizon predictions.

On the other hand, phenomenological methods consider deep-learned models to forecast agent trajectories [1], [60], [90]. Methods have been proposed that produce single (unimodal and deterministic) trajectory predictions [1] and also multiple (multimodal and probabilistic) trajectories [60]. These can be further categorized as agent-centric or scene-centric models, which predict trajectories of other agents separately or jointly, respectively. Deep-learning models have been shown to provide better closed-loop motion predictions for longer horizons compared to ontological methods [87]. However, integrating such models into planning and control frameworks presents challenges. These include a lack of explainability, uncertainty in model training (epistemic uncertainty), and complexity in representing the distribution of other agents' motion from finite data (aleatoric uncertainty) [16]. Furthermore, such models make assumptions on the distribution of agents, such as Gaussian Mixture Models [60]. While scene-centric predictions capture more realistic interactions through joint trajectory predictions, they tend to be more computationally expensive as the predictions cannot be parallelized. Moreover, a current limitation of most datasets is that they mostly contain well-structured, goal-directed, and cooperative motion [57]. As these datasets may not encompass all possible behaviors [83], the efficacy of prediction models trained on such data in interactive scenarios is unclear.

2.4. Outlier Detection

Outlier detection methods aim to identify unusual data points or patterns that deviate from normal behavior. Laxhammar et al. [40] propose three explanations for anomalies: rare samples from the same distribution as training data, distribution shifts between training and test data, and true anomalies. There are three main approaches to outlier detection: distance-based, distribution-based, and functional uncertainty-based methods. Distance-based methods, such as conformal anomaly detection proposed by [41], use metrics to identify outliers but can be computationally intensive for real-time applications as they use a nonconformity score that requires all trajectories during test time. Distribution-based methods compare test data against training data distributions, such as the 'p-quantile anomaly' [26] and the use of exchangeability martingales [9] to detect distribution shifts.

Functional uncertainty-based methods address the limitations of distance and distribution-based approaches by considering how changes in input affect the output, which is particularly important for closed-loop control systems. These methods aim to assess how well the training data informs predictions for new inputs and often involve equipping trained models with out-of-distribution (OOD) monitors. One recent approach by [64] uses the Fisher information matrix to determine how a model's parameters influence its output distribution, offering comparable performance to baselines while reducing computational burden for real-time applications.

3

Preliminaries

3.1. Problem Formulation

Consider a discrete-time dynamical system:

$$X_{t+1}^i = f^i(X_t^i, u_t^i) \tag{3.1}$$

where $X_t^i \in \mathcal{X} \subseteq \mathbb{R}^n$ and $u_t^i \in \mathcal{U} \subseteq \mathbb{R}^m$ denote the state and the control input of the agent $i \in \{e, o\}$ at time $t \in \mathbb{N} \cup \{0\}$, respectively. Here, e represents the ego (controllable) agent and o represents the other (uncontrollable) agents. The sets $\mathcal{U}$ and $\mathcal{X}$ denote the set of permissible control inputs and the workspace of the agent, respectively. The measurable function $f^i : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ describes the agents' dynamics. Let $X_t^o := (X_{t,1}^o, ..., X_{t,N}^o) \in \mathbb{R}^{Nn}$ refer to the joint state of N other agents at time t, such that $X_{t,j}^o$ is the state of agent j at time t. The joint other agent trajectories $X_{0:T}^o$ is sampled from an unknown distribution $\mathcal{D}$ that obeys the dynamics f^o:

$$X^o := (X_0^o, X_1^o, \ldots, X_t^o) \sim \mathcal{D}$$

Assumption 1 We have access to a dataset $D := \{X^{o(1)}, \ldots, X^{o(K)}\}$ containing K trajectories where $X^{o(i)} := (X_0^{o(i)}, X_1^{o(i)}, \ldots, X_T^{o(i)})$ are drawn independently and identically distributed (i.i.d.) from a training distribution $\mathcal{D}$, i.e., $X^{o(i)} \sim \mathcal{D}$.

Remark 3.1.1. *Assumption 1 is standard in prior work [43], [23], [89], [24], [14], [31] and not restrictive in practice as current autonomous driving datasets typically contain data from certain geographic regions and mostly contain well-structured, goal orientated, non-interactive trajectories [8], [57].*

Assumption 2 At test time, the ego agent's state X_t^e influences the behavior of the N other agents, shifting their trajectory distribution from the training distribution $\mathcal{D}$ to an unknown, ego-conditioned distribution $\mathcal{D}^e$. The resulting test time trajectories for some time T are denoted as $X_{0,T}^{o,e}$ and are sampled from the unknown distribution $\mathcal{D}^e$:

$$X^{o,e} := (X_0^o, X_1^o(X_0^e), \ldots, X_t^o(X_0^e, X_1^e, \ldots, X_{t-1}^e)) \sim \mathcal{D}^e$$

where these test time trajectories capture the influence of the ego agent's state on the other agents' trajectories.

Remark 3.1.2. *Assumption 2 aligns with established works in both trajectory forecasting [60], [90], [1] and planning [12], [15], [13], where they assume that the ego agent's plan can influence the neighboring agents' plan. This interdependence is particularly evident in trajectory forecasting research, where the distribution of other agents' is learned by conditioning their predicted trajectories on the ego agent's anticipated plan. Moreover, this interdependence is also observed in many real-world interactive driving scenarios such as merging, lane changes, and yielding where the plans of agents are influenced by one another.*

Assumption 3 For any time $t \geq 0$, the Kullback-Leibler (KL) divergence between the test distribution $\mathcal{D}^e$ and training distribution $\mathcal{D}$ can be bounded by an ego-state-dependent function $\epsilon(X_t^e)$, where X_t^e is the ego agent's state. For notational simplicity, we denote the bound as ϵ:

$$D_{KL}(\mathcal{D}^e||\mathcal{D}) \leq \epsilon$$

It is important to note that this bound is not uniform across all ego agent states but varies with X_t^e. We define the ambiguity set $\mathcal{P}_{KL,\epsilon}(\mathcal{D})$ centered at $\mathcal{D}$ with radius ϵ as:

$$\mathcal{P}_{KL,\epsilon}(\mathcal{D}) := \{\mathcal{D}^e \text{ s.t. } D_{KL}(\mathcal{D}^e||\mathcal{D}) \leq \epsilon\}$$

This set captures all distributions $\mathcal{D}^e$ within an ϵ KL divergence radius from the training distribution $\mathcal{D}$. The value ϵ is unknown a priori and variable. It can be estimated during runtime when the ego agent state X_t^e is known. We make no assumptions on the form of the distributions $\mathcal{D}$ and $\mathcal{D}^e$.

Remark 3.1.3. *Assumption 3 is applicable in many practical scenarios involving autonomous vehicles. It models situations where other agents generally follow their expected behavior ($\mathcal{D}$) but may deviate (resulting in $\mathcal{D}^e$) in response to the ego agent's state as agents tend to strongly prefer avoiding collisions (i.e., cooperative). The use of the KL divergence to quantify shifts is also consistent with approaches in related work, where the KL divergence has been used to measure the degree of influence an agent has over another [73] and mutual information as an indicator of forthcoming collisions [50].*

Formally, we present the problem this thesis attempts to solve in Problem 1:

Problem 1 Given the system in (3.1), access to trajectories $X^{o(i)}$ from a distribution $\mathcal{D}$, access to a trajectory $X^{o,e}$ from distribution $\mathcal{D}^e$, the KL divergence $D_{KL}(\mathcal{D}^e||\mathcal{D}) \leq \epsilon$ valid for the current ego agent state X_t^e, and a failure probability $\delta \in (0,1)$, our goal is to compute control inputs u_t^e and trajectory X_t^e such that the collision avoidance constraint $c : \mathbb{R}^n \times \mathbb{R}^{nN} \to \mathbb{R}$ is satisfied with a probability of at least $1-\delta$ for all time t:

$$\mathbb{P}(c(X_t^e, X_t^{o,e}) \geq 0, \forall t \in \{0, \dots, T\}) \geq 1 - \delta \tag{3.2}$$

In this work, we propose an approximate solution to Problem 1. Given that the actual test distribution is unknown, a conservative approach would be to assume a uniform bound on the distribution shift that is valid for all ego agent states and that holds for all time steps. However, this could lead to overly conservative behavior, leading to infeasibility of the problem that is typically seen in robust methods [92].

Therefore to address Problem 1, we propose an approach that dynamically adapts to distribution shifts during test time. Our method combines:

- Real-time outlier detection for distribution shifts using conformal prediction (CP) techniques
- Estimation of the interaction-induced distribution shift with the Kullback-Leibler (KL) divergence
- Robust CP to update collision avoidance constraints to take into account the estimated interaction-induced distribution shift
- An adaptable model predictive control (MPC) framework that incorporates these robust prediction regions

This approach allows us to provide safety assurances while adapting to varying levels of interaction-induced distribution shifts. The following subsections (3.2 and 3.4) in this chapter detail each component of our methodology. In Chapter 4, we propose an interaction robust planner to approximately solve Problem 1.

3.2. Conformal Prediction

Let $\{R^{(i)}\}_{i=0}^k$ be $k+1$ i.i.d. random variables. We denote the nonconformity score, also known as the residual $R^{(i)}$. This score is commonly defined as $R^{(i)} := ||Z^{(i)} - \mu(X^{(i)})||_2$, where the predictor is a mapping $\mu : X^{(i)} \to Z^{(i)}$ from the input space $X^{(i)}$ to the output space $Z^{(i)}$. A large nonconformity score indicates a poor predictive model. Our goal is to construct a prediction region C for $R^{(0)}$ based

on $\{R^{(i)}\}_{i=1}^{k}$, such that the random variable is contained in the prediction region C with a probability $1-\delta$, where $\delta \in (0,1)$ is a user-specified failure probability [78], [62]:

$$\mathbb{P}(R^{(0)} \leq C) \geq 1-\delta \tag{3.3}$$

Using the Quantile Lemma, we construct the prediction region C using the finite-sample corrected $(1-\delta)th$ quantile of the empirical distribution of $\{R^{(i)}\}_{i=0}^{k}$ (Lemma 1 [72]). Firstly, the quantile is defined as:

$$\text{Quantile}(1-\delta, \mathcal{R}) = \inf\{z : \mathbb{P}\{Z \leq z\} \geq 1-\delta\}, Z \sim \mathcal{R} \tag{3.4}$$

where Z is the empirical distribution of $\{R^{(i)}\}_{i=0}^{k}$ constructed using Dirac distributions $\delta_{R^{(i)}}$ centered at $R^{(i)}$. We then determine the quantile, where $\lceil \cdot \rceil$ is the ceiling function:

$$C = \text{Quantile}\left(\frac{\lceil (1-\delta)(K+1) \rceil}{K}; \frac{1}{K}\sum_{i=1}^{K} \delta_{R^{(i)}}\right) \tag{3.5}$$

Furthermore, two requirements must be met before determining C. The quantile must be bounded and the minimum number of data points K is required such that:

$$0 \leq \frac{\lceil (1-\delta)(K+1) \rceil}{K} \leq 1 \tag{3.6}$$

Provided that the requirement (3.6) is satisfied, we sort $\{R^{(1)}, ..., R^{(k)}\}$ in a non-decreasing order and let $R^{(k+1)} := \infty$. We then have $C := R^{(p)}$ to be the pth smallest nonconformity score where $p := \lceil (k+1)(1-\delta) \rceil$. When requirement (3.6) is violated, $C = \infty$.

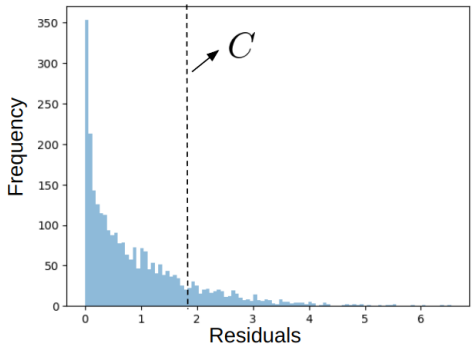

Figure 3.1: Histogram of residuals. C is chosen as the $\lceil (1-\delta)(n+1) \rceil$ smallest of $\{R^{(1)}, \ldots, R^{(n)}, \infty\}$

It is important to note that CP provides marginal coverage guarantees and does not imply conditional coverage. Specifically, the probability bound holds in expectation over the input space and may not hold uniformly across all subsets of the input space. Additionally, CP's validity also holds when the non-conformity scores are exchangeable. Particularly, exchangeability is formally defined as: $(X_i, Y_i)_{i=1}^{n}$ are exchangeable if, for any permutation σ of $[1, n]$:

$$((X_1, Y_1), \ldots, (X_n, Y_n)) \overset{d}{=} \left((X_{\sigma(1)}, Y_{\sigma(1)}), \ldots, (X_{\sigma(n)}, Y_{\sigma(n)})\right). \tag{3.7}$$

For example, Z_1 and Z_2 are exchangeable if $(Z_1, Z_2) \overset{d}{=} (Z_2, Z_1)$. In other words, rearranging the labels does not alter the probability distribution.

Remark 3.2.1. *CP requires exchangeability at minimum, a less stringent assumption than i.i.d. Specifically, i.i.d. refers to conditions where the distribution remains constant, such that all samples are taken from the same probability distribution, and that each sample does not influence other samples (i.e., the sampled random variables are independent). Exchangeability relaxes the second requirement of i.i.d. data as it allows for dependencies between samples that do not affect the overall joint probability distribution when the samples are re-ordered.*

However, in settings of distribution shift where the calibration residuals $\{R^{(i)}\}_{i=1}^{k}$ are drawn from a distribution $\mathcal{R}$ and the test sample $R^{(0)}$ is drawn from a different distribution $\mathcal{R}^e$, the above methodology makes no claims of validity as the exchangeability assumption is not satisfied. This motivates the following robust variant of conformal prediction.

Lemma 1. *Robust Conformal Prediction* [48] provides valid prediction regions under sufficiently small distribution shifts between the calibration and test data. The distribution shift is typically quantified using the f-divergence: Given a closed convex function $f : \mathbb{R} \to \mathbb{R}$ satisfying $f(1) = 0$ and $f(t) = +\infty$ for $t < 0$, the f-divergence between probability distributions $\mathcal{R}^e$ and $\mathcal{R}$ on a set $\mathcal{Z}$ is

$$D_f(\mathcal{R}^e \| \mathcal{R}) := \int_{z \in \mathcal{Z}} f\left(\frac{d\mathcal{R}^e(z)}{d\mathcal{R}(z)}\right) d\mathcal{R}(z) \tag{3.8}$$

where $\mathcal{Z}$ is the support of $\mathcal{R}$, and $\frac{d\mathcal{R}^e(z)}{d\mathcal{R}(z)}$ is the Radon-Nikodyn derivative of $\mathcal{R}^e$ with respect to $\mathcal{R}$. Jensen's inequality guarantees that $D_f(\mathcal{R}^e \| \mathcal{R}) \geq 0$ always. Examples include $f(z) = z \log z$, which is the KL divergence, and $f(t) = \frac{1}{2}(t-1)^2$, which gives the chi^2 divergence.

We denote $\mathcal{R}^e$ and $\mathcal{R}$ as the test and calibration residual distributions, respectively, where $R^{(0)} \sim \mathcal{R}^e$ and $\{R^i\}_{i=1}^{K} \sim \mathcal{R}$, such that their f-divergence $D_f(\mathcal{R}^e, \mathcal{R}) \leq \epsilon$. The adjusted prediction region $\tilde{C}$ in $\mathbb{P}(R^{(0)} \leq \tilde{C}) \geq 1 - \delta$ is computed using an adjusted confidence level $\tilde{\delta}$ that takes into account a set of distributions $\mathcal{P}_{f,\epsilon}(\mathcal{R}) := \{\mathcal{R}^e \text{ s.t. } D_f(\mathcal{R}^e || \mathcal{R}) \leq \epsilon\}$ centered around the calibration residual distribution $\mathcal{R}$ (see Corollary 2.2 [48]):

$$\tilde{C} := \text{Quantile}_{1-\tilde{\delta}}(R^{(1)}, ..., R^{(K)}), \tag{3.9a}$$

$$\tilde{\delta} := 1 - g^{-1}(1 - \delta_n) \tag{3.9b}$$

that is obtained with the following convex optimization problems:

$$\delta_n := 1 - g\left(\left(1 + \frac{1}{K}\right) g^{-1}(1 - \delta)\right) \tag{3.10a}$$

$$g(\beta) := \inf\left\{z \in [0,1] \;\middle|\; \beta f\left(\frac{z}{\beta}\right) + (1-\beta) f\left(\frac{1-z}{1-\beta}\right) \leq \epsilon\right\} \tag{3.10b}$$

$$g^{-1}(\tau) := \sup\left\{\beta \in [0,1] \mid g(\beta) \leq \tau\right\} \tag{3.10c}$$

where the function $g(\beta)$ (3.10b) solves a constrained optimization problem. It seeks the smallest value of z that satisfies the f-divergence constraint between the calibration and test distributions, effectively mapping a nominal confidence level β to a robust confidence level z while considering potential distribution shifts. Complementing this, $g^{-}1(\tau)$ (3.10c) serves as the inverse of $g(\beta)$, identifying the largest β that yields a $g(\beta)$ value not exceeding τ. We also note that robust CP provides prediction regions that take into account distributional uncertainty as they are valid for all distributions within the KL divergence radius. This differs from doing split CP on the test data samples if available, which would only provide guarantees on the test data distribution.

Robust CP assumes a known, fixed amount of shift ϵ. However, our approach addresses scenarios where the shift ϵ is variable and unknown. In our setting, the shift ϵ to be robust against depends on the level of influence the ego agent has on the other agents. Therefore, this necessitates an adaptive approach to determining ϵ during run-time.

In addition, robust CP presents requirements for the amount of calibration data K. From (3.9b) and (3.10a) we require the following holds to compute the adjusted quantile $1 - \tilde{\delta}$:

$$0 \leq \left(1 + \frac{1}{K}\right) g^{-1}(1 - \delta) \leq 1 \tag{3.11}$$

Remark 3.2.2. *Robust CP has been shown to provide the right coverage level for arbitrary distribution shifts. However, the challenge lies in estimating the amount of distribution shift to be robust against. We present an approach to estimate the KL divergence in Lemma 3.*

Remark 3.2.3. *The data requirement for* (3.9)*, as presented in* (3.11)*, implicitly constrains the permissible distribution shift. This constraint arises from the relationship between the original failure probability* δ *and the adjusted failure probability* $\hat{\delta}$*. As* δ *decreases, the upper bound on* $\hat{\delta}$ *(which must remain less than or equal to 1 for quantile computation) becomes more restrictive. Consequently, this limits the extent of distribution shift that can be accommodated while maintaining the desired failure probability (see Figure 3.2).*

Remark 3.2.4. *It has been shown that robust CP can be solved with binary search over feasible* $\beta \in [\tau, 1]$ *with accuracy* ϵ *in time* $\log \frac{1-\tau}{\epsilon}$ *[48]. However, even for low accuracy values, this would not be feasible in an MPC framework. We later show that we interpolate robust CP* $\tilde{\delta}$ *values offline and use them during planning.*

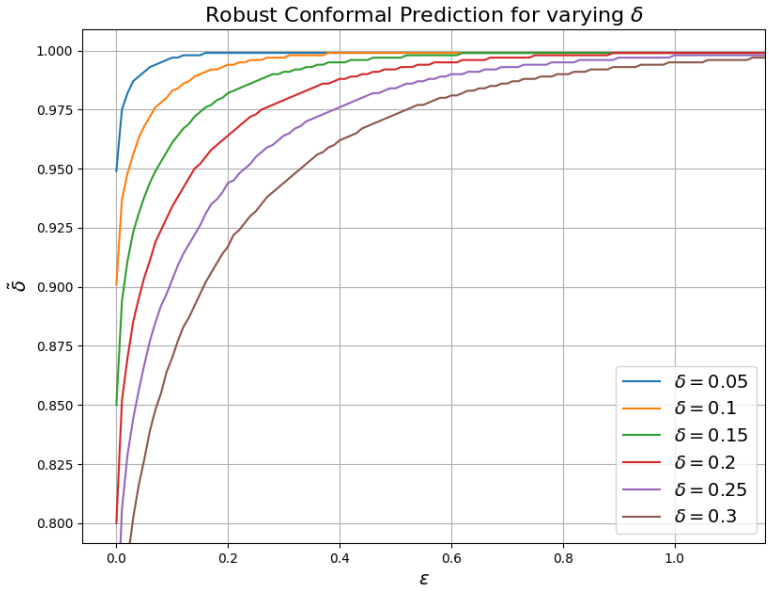

Figure 3.2: Robust CP calculated for varying δ with KL divergence. We note that increasing δ reduces the permissible distribution shift ϵ

While Assumption 3 bounds the trajectory-level distribution shift by ϵ, estimating this shift at the residual level is more efficient due to its one-dimensional nature. The data processing inequality [6] allows us to compute the KL divergence on the nonconformity measure level, which is computationally advantageous and potentially provides tighter bounds as we directly estimate the distribution shift in the nonconformity score metric that we are interested in compared to estimating distribution shifts on multi-dimensional trajectories.

Lemma 2. *Data Processing Inequality* [6]: Given a nonconformity measure $R : X \to \mathbb{R}$, $X^{o,e} \sim \mathcal{D}^e$ and $X^o \sim \mathcal{D}$, let R^e and R follow a distribution $\mathcal{R}^e$ and $\mathcal{R}$ such that they are the push-forward of $\mathcal{D}^e$ and $\mathcal{D}$, respectively. It holds that the f-divergence between two distributions does not increase when pushed through R:

$$D_f(\mathcal{D}^e||\mathcal{D}) \geq D_f(\mathcal{R}^e||\mathcal{R}),$$
$$D_f(\mathcal{D}^e||\mathcal{D}) \leq \epsilon \to D_f(\mathcal{R}^e||\mathcal{R}) \leq \epsilon$$

Therefore, to determine ϵ for solving (3.9) and (3.10), we compute the distribution shift at the residual level rather than at the trajectory level. As we make no assumptions on the data distribution, the approach requires samples from both the calibration data distribution $\mathcal{R}$ and test residual distribution $\mathcal{R}^e$. However, Assumption 3 implies that we lack a priori access to test residuals, necessitating a method to estimate these residuals and, consequently, the current KL divergence for use in robust CP. We elaborate on our proposed framework for determining samples from $\mathcal{R}^e$ in Section 4.3. Given samples from both $\mathcal{R}$ and $\mathcal{R}^e$, we propose using the following KL divergence estimator that can operate efficiently in real-time.

3.3. k-NN-based KL Divergence Estimator

Lemma 3. *k-NN Based KL Divergence Estimator* [81], [82]: We consider the following samples: $\{R_1^e, \ldots, R_L^e\}$ drawn i.i.d. from $\mathcal{R}^e$ and $\{R_1, \ldots, R_M\}$ drawn i.i.d. from $\mathcal{R}$. Let $\eta_i := ||R_i^e - R_{i(k)}^e||_1$ be the absolute difference between R_i^e and its k-nearest neighbor (k-NN) in $\{R_j^e\}_{j \neq i}$. Particularly, the k-NN of R_i^e in $\{R_1^e, \ldots, R_{i-1}^e, R_{i+1}^e, \ldots, R_L^e\}$ is $R_{i(k)}^e$ where the indices $i(1), \ldots i(i-1), i(i+1), \ldots i(L)$ represent an ordering of all other points based on their distance from R_i^e:

$$||R_i^e - R_{i(1)}^e||_1 \leq ||R_i^e - R_{i(2)}^e||_1 \leq \ldots \leq ||R_i^e - R_{i(L)}^e||_1$$

Similarly, $\nu_i := ||R_i^e - R_{i(k)}||_1$ is the absolute difference between R_i^e and its k-NN in $\{R_1, \ldots, R_M\}$. The k-NN of R_i^e in $\{R_1 \ldots R_M\}$ is $R_{i(k)}^e$ where $i(1), \ldots, i(n)$ is such that:

$$||R_i^e - R_{i(1)}||_1 \leq ||R_i^e - R_{i(2)}||_1 \leq \ldots \leq ||R_i^e - R_{i(M)}||_1$$

The k-NN-based KL-divergence estimator is defined as follows [81], [82]:

$$\hat{D}_{KL}(\mathcal{R}^e||\mathcal{R}) = \frac{d}{L}\sum_{i=1}^{L} \log\frac{\nu_k(i)}{\eta_k(i)} + \log\frac{M}{L-1} \tag{3.12}$$

It has been shown under mild regularity conditions that the k-NN divergence estimator (3.12) is asymptotically unbiased (Theorem 1 [82]):

$$\lim_{L,M\to\infty} \mathbb{E}[\hat{D}_{KL}(\mathcal{R}^e||\mathcal{R})] = D_{KL}(\mathcal{R}^e||\mathcal{R})$$

Additionally, the estimator is mean-square consistent (Theorem 2 [82]):

$$\lim_{L,M\to\infty} \mathbb{E}[(\hat{D}_{KL}(\mathcal{R}^e||\mathcal{R}) - D_{KL}(\mathcal{R}^e||\mathcal{R}))^2] = 0$$

We note that the samples must be drawn i.i.d. to prove mean-square consistency. Below we run the KL divergence estimator using pytorch for an example computing the KL divergence for two exponential distributions:

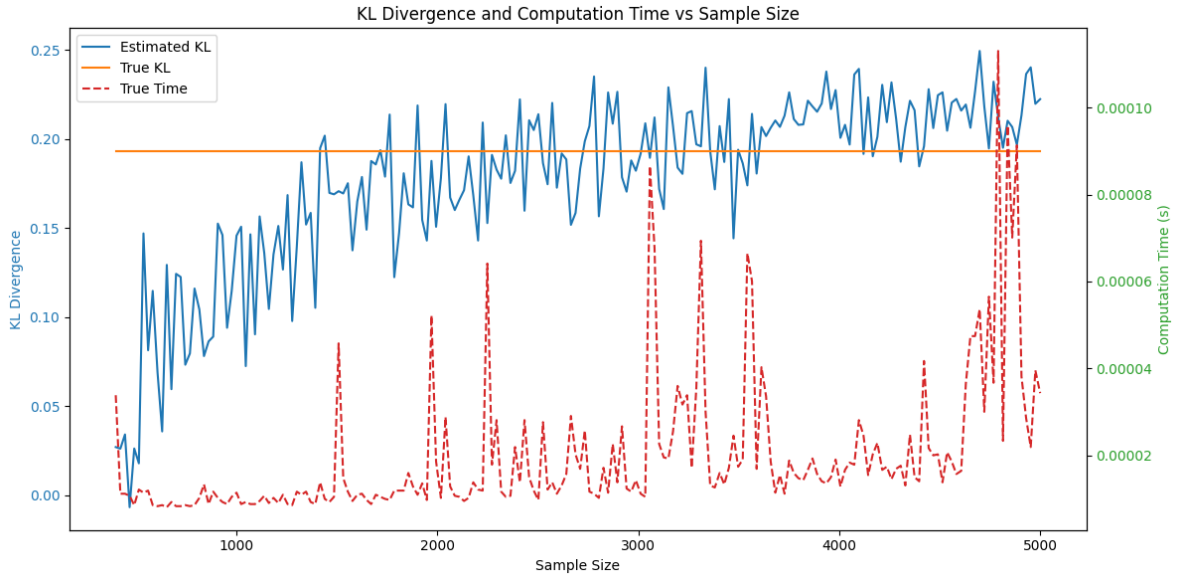

Figure 3.3: Estimator for $D_{KL}(Exp(1)||Exp(2))$ with varying sample size. (Both distributions are sampled with the same number of samples), k=50

3.4. Conformal Prediction for Finite-Horizon Trajectories

We now use CP to obtain prediction regions for predictions denoted $\hat{X}^{o,e}_{\tau|t}$.

Remark 3.4.1. *This methodology treats each trajectory as an independent observation, thereby maintaining the i.i.d. assumption across trajectories while acknowledging the temporal dependence of time steps within each trajectory [69], [43], [18], [89].*

In the multi-step setting, the goal is to construct prediction regions over a task horizon T with $1-\delta$ guarantee. Formally, we want to construct prediction regions defined by $C_{\tau|t}$, where $\tau \in \{t+1,\dots,T\}$ is a future timestep. Let a dataset D contain trajectories of the other agents as stated in Assumption 1. Let $D_{train} \subset D$, $D_{cal} \subset D$, and $D_{train} \cap D_{cal} = \emptyset$. Consider a function PREDICT that is trained on observations $(X^o_0,\dots,X^o_t)$ from D_{train} and outputs predictions $(\hat{X}^o_{t+1|t},\dots,\hat{X}^o_{T|t})$. With the following predictions, we define the nonconformity scores as introduced in Section 3.2 with a normalization term $\sigma_{\tau|t,j} > 0$ determined using a subset of the training data $D_{train,1}$ to conform to the exchangeability assumption in CP. We normalize with the maximum prediction error for an individual timestep. The nonconformity score for all future times $\tau \in \{t+1,\dots,T\}$ and other agents j is defined as [89], [18]:

$$\begin{aligned} R^{(i)}_{\tau|t} &:= \max_{\substack{(\tau,j)\in \\ \{1,\dots,T\}\times\{1,\dots,N\}}} \frac{||X^{o(i)}_{\tau,j} - \hat{X}^{o(i)}_{\tau|t,j}||}{\sigma_{\tau|t,j}}, \\ \sigma_{\tau|t,j} &:= \max_{i\in\{1,\dots,K\}} ||X^{o(i)}_{\tau,j} - \hat{X}^{o(i)}_{\tau|t,j}|| \end{aligned} \tag{3.13}$$

This approach enables us to achieve more efficient prediction regions for each agent compared to using Boole's inequality [43], [23]. Provided the nonconformity scores $R^{(i)}_{\tau|t}$ determined using D_{cal}, we can construct prediction regions for finite-horizon trajectories (i.e., over multiple future predictions) in an open-loop and closed-loop setting. The open-loop ensures the validity of all τ-step ahead prediction regions (i.e., $\tau|0$) while the closed-loop ensures the validity of all one-step ahead prediction regions (i.e., $t+1|t$). Particularly, the following two statements hold for open-loop and closed-loop regions, respectively:

$$\begin{aligned} P(&||X^{o(i)}_{\tau,j} - \hat{X}^{o(i)}_{\tau|0,j}|| \le C_{OL}\sigma_{\tau|0,j}, \\ &(\tau,j) \in \{t+1,\dots,T\}\times\{1,\dots,N\}) \ge 1-\delta \end{aligned} \tag{3.14}$$

$$\begin{aligned} P(&||X^{o(i)}_{t+1|t,j} - \hat{X}^{o(i)}_{t+1|t,j}|| \le C_{CL}\sigma_{t+1|t,j}, \\ &(t,j) \in \{0,\dots,T-1\}\times\{1,\dots,N\}) \ge 1-\delta \end{aligned} \tag{3.15}$$

where $C_{CL} := R^{(p)}_{CL}$ and $C_{OL} := R^{(p)}_{OL}$ are the pth smallest closed loop and open loop nonconformity scores, respectively, with $p := \lceil (k+1)(1-\delta) \rceil$.

However, in settings of interaction-induced distribution shift such that during test time the other agents are under some level of influence from the ego agent, the prediction regions in (3.14) and (3.15) are not valid. To address this, we require an estimate of the distribution shift ϵ which depends on the test nonconformity scores. By solving (3.9) and (3.10), we can then obtain updated prediction regions $\tilde{C}_{CL}$ and $\tilde{C}_{OL}$ that remain valid under the specified distribution shift. However, evaluating the nonconformity score defined in (3.13) at test time t is challenging, as only states $(\hat{X}^{o,e}_0,\dots,\hat{X}^{o,e}_t)$ are observable and not the future states $(\hat{X}^{o,e}_{t+1},\dots,\hat{X}^{o,e}_T)$. To overcome this limitation, we introduce the time-lagged nonconformity score, detailed in Section 4.2.

Remark 3.4.2. *We note that as $\sigma_{\tau|t,j}$ is computed using data from the proper training set, the maximum prediction error could be less than what could be observed in in-distribution data not trained on. This is because the prediction model is trained to minimize the prediction error (i.e., the residuals). Consequently, this approach leads to an underestimation of the normalization factor $\sigma_{\tau|t,j}$, which results in a slightly conservative outcome.*

4

Interaction Robust MPC-based Planner

4.1. Framework Overview

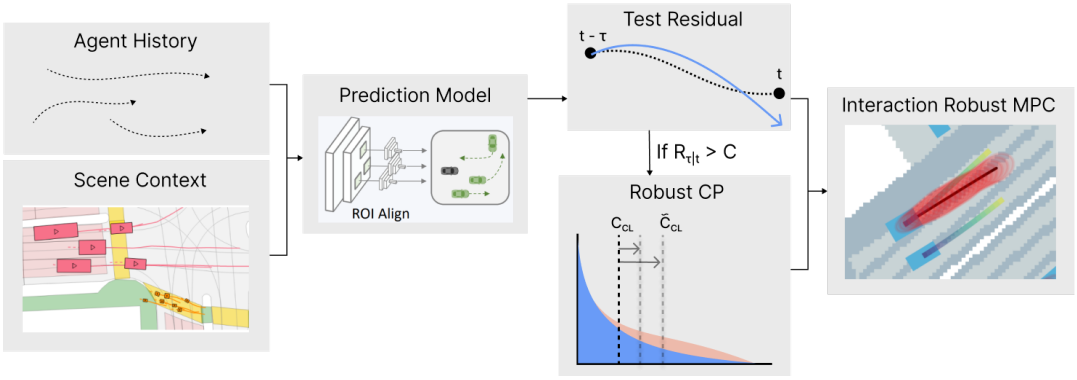

Figure 4.1: Proposed framework that employs robust CP to take into account interaction uncertainty due to the ego agent's influence on other agents when an outlier is detected.

In this thesis, we propose an interactive planning framework that uses robust CP for learned trajectory predictors to address interaction uncertainty when outliers are detected (see Figure 4.1). The approach uses observed data of other agents to compute test-time residuals to detect outliers, defined as scores not contained within the prediction regions constructed with the calibration residuals. When an outlier is detected, the prediction region is updated to adapt to the interaction uncertainty by estimating real-time distribution shifts (i.e., KL divergence) in other agents' behavior caused by the ego agent's presence. The estimated shift is used to update the prediction regions by solving a series of convex optimization problems as defined in equations (3.9) and (3.10) to ensure $1 - \delta$ coverage under distribution shift.

Using outlier detection also addresses the planner's computational efficiency and feasibility under distribution shifts. The framework reduces computational burden in multi-agent settings by estimating shifts only when CP guarantees with calibration residuals under-cover (i.e., when an outlier is detected), rather than for all agents. This also benefits the planner's feasibility by limiting changes in collision avoidance constraints across iterations to only agents under interaction uncertainty.

Lastly, under Assumption 1 in section 3.1 we use an agent-centric prediction model that can predict trajectories of other agents, taking into account their dynamics and semantic information, as detailed in [87]. Our framework is compatible with both deterministic and probabilistic trajectory predictors. The PREDICT function takes as input the observed trajectory of other agents $X^{o,e}_{0:t}$, scene information S and outputs a prediction $\hat{X}^{o,e}_{\tau|t}$ of the other agents.

4.2. Outlier Detection

The proposed framework employs outlier detection using CP to identify when the observed behavior of other agents deviates from the predicted behavior. This requires computing a nonconformity score R^e during test time to determine if it lies in the prediction region $C_{\tau|t}$, computed based on calibration residuals defined in (3.13). An outlier is detected when $R^e > C_{\tau|t}$ However, recall that the nonconformity score defined in (3.13) cannot be evaluated at time t as only the previous states $(\hat{X}_0^{o,e}, \ldots, \hat{X}_t^{o,e})$ are observable, while the future states $(\hat{X}_{t+1}^{o,e}, \ldots, \hat{X}_T^{o,e})$ are unknown. Therefore, we define a time-lagged nonconformity score that evaluates the τ step-ahead prediction error that was made τ time steps ago and is valid when $t \geq \tau$:

$$R^e := \max_{\substack{(\tau,j)\in \\ \{1,\ldots,T\}\times\{1,\ldots,N\}}} \frac{||X_{\tau,j}^{o,e} - \hat{X}_{\tau|t-\tau,j}^{o,e}||}{\sigma_{\tau|t,j}} \tag{4.1}$$

Furthermore, the prediction region $C_{\tau|t}$ is computed using the calibration residuals (3.13) and follows a Beta distribution. With a sufficiently large calibration set, we expect the outlier rate to approach the calibrated failure probability δ. However, it is important to note that the prediction region $C_{\tau|t}$ is computed using calibration residuals from diverse driving scenarios where the frequency of certain scenarios can be much lower than more commonly observed scenarios. Also known as the 'curse of rarity' [22] where the occurrence of safety-critical scenarios (i.e., scenarios in the long tail of the distribution) are observed rarely. This imbalance can lead to epistemic uncertainty in regions of the input space with limited data [36]. Consequently, this may lead to large nonconformity scores that lead to conservatism while determining the prediction region $C_{\tau|t}$ and thus fewer detected outliers. However, we want a prediction region that takes into account epistemic uncertainty due to the curse of rarity. It should be able to detect more outliers such that we leverage our proposed conservative controller that is designed to capture distribution shifts.

Inspired by robust regression [34], we address this challenge by obtaining a less conservative prediction region $C_{\tau|t}$ for outlier detection. We use a quantile that minimizes the Huber loss to reduce epistemic uncertainty. The Huber loss gives less weight to residuals above a threshold parameter λ by using a linear loss and more weight to residuals below λ by using a quadratic loss [35]:

$$L_\lambda(R^{(i)}) = \begin{cases} \frac{1}{2}(R^{(i)})^2 & \text{if } R^{(i)} \leq \lambda, \\ \lambda\delta(|R^{(i)}| - \frac{1}{2}\lambda) & \text{if } R^{(i)} > \lambda. \end{cases} \tag{4.2}$$

where λ is commonly chosen to be $MAR/0.6745$ [35], with MAR denoting the median absolute residual. We note that this is a tunable parameter that can be adjusted to ensure δ outliers are detected during experiments. To compute MAR we use a subset of the training data $D_{train,2}$ to conform to the exchangeability assumption in CP. To compute the Huber quantile, we substitute (4.2) in the quantile loss function:

$$L_{q,\lambda}(R^{(i)}) = \begin{cases} \delta \cdot L_\lambda(R^{(i)} - q), & \text{if } R^{(i)} \geq q, \\ (1-\delta) \cdot L_\lambda(R^{(i)} - q), & \text{otherwise}, \end{cases} \tag{4.3}$$

The quantile is found by minimizing this Huber quantile loss over the calibration set of size k:

$$\hat{q} = \arg\min_q \sum_{i=1}^{k} L_{q,\lambda}(R^{(i)}) \tag{4.4}$$

where $C_{CL,h} := \hat{q}$ in the case the closed-loop residuals are minimized. Formally, we classify the test residual as an outlier if $R^e > C_{CL,h}$. The offline computation of residuals for conformal prediction is summarized below in Algorithm 1.

Algorithm 1 Residuals for Conformal Prediction

Input: Failure probability δ, calibration dataset $\mathcal{D}_{\text{cal}}$, training dataset $\mathcal{D}_{\text{train}}$, prediction and task horizons H and T, threshold parameter λ
Output: $C_{CL}, C_{OL}, C_{CL,h}, C_{OL,h}, \sigma_{\tau|t}$
$p \leftarrow (|\mathcal{D}_{\text{cal}}| + 1)(1 - \delta)$
for $t = 0$ **to** $T - 1$ **do** ▷ Using Training Data
  for $\tau = t + 1$ **to** $t + H$ **do**
    Obtain predictions $\hat{X}^{o(i)}_{\tau|t}$ for each $X^{o(i)} \in \mathcal{D}_{\text{cal}}$
  end for
  $\sigma_{\tau|t} \leftarrow \max_i ||X^{o(i)}_{\tau,j} - \hat{X}^{o(i)}_{\tau|t,j}||$
end for
for $t = 0$ **to** $T - 1$ **do** ▷ Using Calibration Data
  for $\tau = t + 1$ **to** $t + H$ **do**
    Obtain predictions $\hat{X}^{o(i)}_{\tau|t}$ for each $X^{o(i)} \in \mathcal{D}_{\text{cal}}$
    $R^{(i)}_{CL} \leftarrow \max_\tau \frac{\left\|X^{o(i)}_{\tau|t} - \hat{X}^{o(i)}_{\tau|t}\right\|}{\sigma_{\tau|t}}$ for each $X^{o(i)} \in \mathcal{D}_{\text{cal}}$
    $R^{(i)}_{OL} \leftarrow \max_\tau \frac{\left\|X^{o(i)}_{\tau|t} - \hat{X}^{o(i)}_{\tau|t}\right\|}{\sigma_{\tau|0}}$ for each $X^{o(i)} \in \mathcal{D}_{\text{cal}}$
  end for
end for
$R^{|\mathcal{D}_{\text{cal}}+1|}_{CL} \leftarrow \infty, R^{|\mathcal{D}_{\text{cal}}+1|}_{OL} \leftarrow \infty$
Sort $R^{(i)}_{CL}$, and $R^{(i)}_{OL}$ in non-decreasing order
$C_{CL} \leftarrow R^{(p)}_{CL}, C_{OL} \leftarrow R^{(p)}_{OL}$
Compute Huber quantiles $\hat{q}_{OL}$, and $\hat{q}_{CL}$ by solving (4.4)
$C_{CL,h} \leftarrow \hat{q}_{CL}, C_{OL,h} \leftarrow \hat{q}_{OL}$

4.3. Capturing Distribution Shift

Provided an outlier is detected, we propose updating the prediction regions to take into account interaction uncertainty which is inherent with longer prediction horizons [24]. Particularly, we use robust CP to dynamically adapt to the aleatoric uncertainty arising from interactions with the ego agent. This method provides prediction regions that are valid for a KL divergence ambiguity set defined as: $\mathcal{P}_{KL,\epsilon}(\mathcal{R}) := \{\mathcal{R}^e \text{ s.t. } D_{KL}(\mathcal{R}^e||\mathcal{R}) \leq \epsilon\}$. where $\mathcal{R}^e$ represents the distribution of residuals at test time under interaction, $\mathcal{R}$ is the calibration residual distribution, and ϵ is the KL divergence bound. To compute the updated prediction region $\hat{C}$ the KL divergence must be estimated. However, this would require multiple test residual samples to estimate a distribution and thus distribution shift. Furthermore, we have access to only calibration residuals and a single trajectory prediction of the other agent.

To address this limitation, we propose using an interaction potential term to estimate the test distribution. This approach is on the basis that residuals increase in magnitude due to interaction uncertainty leading to larger prediction errors:

Assumption 4. Let the test residuals R^e be a function of the calibration R residuals with the following heuristic interaction potential to model the influence of the ego agent:

$$R^e = \frac{R}{1 - \gamma \exp(-\frac{1}{2h^2} d^2)} \tag{4.5}$$

where d is the L^2 norm distance between the ego and other agents, and γ and h are parameters that determine the strength of the interaction (see Figure 4.4). This interaction potential allows us to ego-condition M samples of the calibration residuals, providing a heuristic estimate of the test residual distribution. We can then use this estimate to determine the KL divergence radius ϵ on the residual level using equation (3.12). Given the estimated ϵ, we use robust CP to find the updated prediction region $\tilde{C}_{CL}$ by solving (3.9) and (3.10). This process allows us to dynamically adapt the predictions to take into account interaction uncertainty that can vary at each time step by computing a new estimate ϵ.

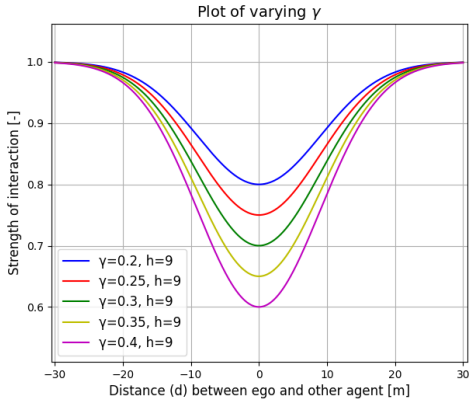

Figure 4.2: Increasing γ increasing the strength of interaction near the agent

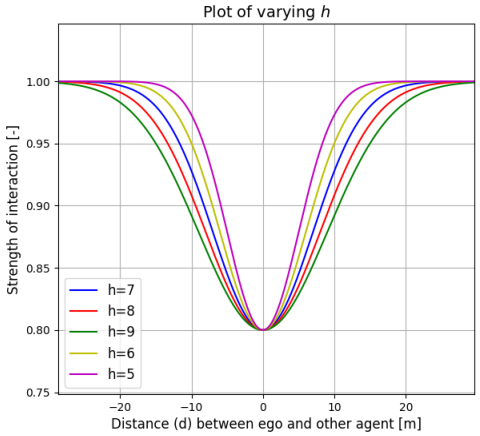

Figure 4.3: Increasing h increasing the influence of interaction at greater distances d

Figure 4.4: Interaction potential parameters γ and h

4.4. Model Predictive Control

With the prior sections, we formulate an optimization problem for the ego agent given a cost function J, the initial condition of the ego and other agents, and predictions of the other agents' trajectories. The objective is to compute control inputs u_t^e such that the collision probability with the other agents is less than or equal to δ over the task horizon T:

$$\min_{X^e, u^e} \sum_{k=t}^{t+H-1} J(X_{k+1}^e, u_k^e) \tag{4.6a}$$

s.t.,

$$X_0^e = \xi^e, \tag{4.6b}$$

$$X_{k+1}^e = f^e(X_k^e, u_k^e), \qquad \forall k \in \{t, \dots, t+H-1\} \tag{4.6c}$$

$$u_k^e \in \mathcal{U},\ X_{k+1}^e \in \mathcal{X}, \qquad \forall k \in \{t, \dots, t+H-1\} \tag{4.6d}$$

$$c(X_k^e, \hat{X}_{k,j}^{o,e}) \geq 0, \forall (k,j) \in \{t, \dots, t+H-1\} \times \{1, \dots, N\} \tag{4.6e}$$

The optimization problem in (4.6) formulates an approach that uses robust CP prediction regions as constraints for a model predictive control (MPC) method to compute control inputs u^e over a prediction horizon H. The objective function J in (4.6a) minimizes a step-wise cost function over the prediction horizon H. The constraints (4.6b) set the initial condition where ξ^e is the current state of the ego agent. Constraint (4.6c) enforces the ego agent dynamical constraints represented by f^e. Constraint (4.6d) ensures the control inputs u^e and states X^e remain within their feasible sets $\mathcal{U}$ and $\mathcal{X}$, respectively. Constraint (4.6e) enforces the probabilistic safety constraint. It utilizes the robust conformal prediction regions for $\hat{X}_k^{o,e}$ when an outlier is detected and prediction regions from calibration residuals when no outlier is detected. The function $c(\cdot, \cdot)$ represents the safety criterion, for example, to ensure a safe distance between the ego agent and the predicted states and regions of the other N agents.

The above optimization problem (4.6) is then solved by linearizing the cost, dynamics, and constraints to solve the quadratic program in a closed-loop (receding horizon) strategy. We now summarize the proposed interaction robust MPC in Algorithm 2:

Algorithm 2 MPC with Robust Conformal Prediction

Input: Failure probability δ, calibration dataset $\mathcal{D}_{\text{cal}}$, prediction and task horizons H and T
Output: Control input u_t^e
for $t = 0$ **to** $T - 1$ **do** ▷ Online perform planning loop
    Sense X_t^e and $X_t^{o,e}$
    Obtain predictions $\hat{X}_{\tau|t}^{o,e}$ for $\tau = t+1, \ldots, t+H$
    if test residual from (4.1) $> C_{CL,h}$ **then**
        Obtain M samples of R^e from (4.5)
        Obtain ϵ using the KL divergence estimator (3.12)
        Obtain $\tilde{C}_{\tau|t} := \tilde{C}_{CL}\sigma_{\tau|t}$ given ϵ
    else
        Obtain $C_{\tau|t} := C_{CL}\sigma_{\tau|t}$
    end if
    Calculate controls $u_t^e, \ldots, u_{T-1}^e$ as the solution of (4.6)
    Apply u_t^e to the ego agent
end for

5

Results

5.1. Implementation Details

We evaluate our approach in interactive scenarios by simulating autonomous vehicle interactions with other human-driven vehicles using the Bi-level Imitation for Traffic Simulation (BITS) environment [87] with the Nuscenes dataset [8].

Remark 5.1.1. *We select the Nuscenes dataset for our experiments, aligning with Assumption 1. This choice is supported by the analysis in [33], which introduces a courtesy metric for quantifying inter-agent interactions. In particular, the courtesy metric computes the KL divergence between an agent's motion with and without the presence of the ego agent. They show (see Figure 2 in [33]) that the Nuscenes dataset contains minimal interactions among agents, thus satisfying Assumption 1.*

This BITS closed-loop simulation offers a comprehensive way to evaluate our proposed approach over replaying dataset scenes. Each simulation scene is initialized with one of the 100 test scenes from the NuScenes dataset and then evolves forward with a policy for each agent. We select one agent as the ego agent and control it using our proposed method. We let the remaining agents run the BITS policy. This policy generates diverse traffic behaviors by leveraging a bi-level imitation learning framework, which captures high-level decision-making and low-level control of human drivers (see Figure 5.1 ego agent is one without red prediction regions and with MPC horizon plotted).

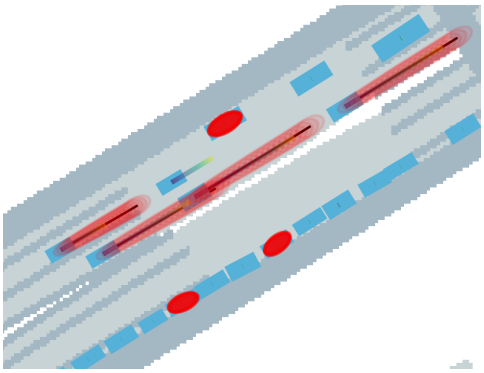

Figure 5.1: BITS simulator environment [87]

It is important to note that the BITS environment generates diverse in-distribution behaviors. This stems from the fact that simulation tools are limited by their training data, which can limit the scenarios they can generate. Therefore, we test our approach by introducing out-of-distribution behaviors. Specifically, we vary other agents' speeds by up to $\pm$ 20% from the BITS policy's planned trajectories. This aids in creating more interactions and tests our proposed method's adaptability to varying behaviors. Furthermore, we also spawn new agents every 4 seconds around the ego agent. This also aids in increasing potential interactions by creating more challenging and dense traffic scenarios.

Remark 5.1.2. *Generating out-of-distribution behaviors: Our approach to generating diverse behaviors in the simulator by varying the speed of the other agents also follows from prior works [20], and [88], where they vary the velocities of pedestrians to create out-of-distribution scenarios.*

The BITS policy that the other agents follow uses three main components (see parameters in Table 5.3 and 5.4). Firstly, a spatial goal network extracts multiple goals for each other agent. Secondly, a goal-conditioned policy generates plans to reach these sampled goals. Finally, a trajectory predictor is used to predict the future trajectories of the neighboring agents. These components are then integrated into a cost-minimization framework to determine the optimal action.

Our closed-loop planner uses three modules: a trajectory predictor, a high-level route planner, and our proposed interaction robust MPC that plans the trajectory. We use the deterministic BITS prediction model as the trajectory predictor without ego-conditioning. We train the model on the NuScenes dataset [8], consisting of 1000 20-second scenes collected in Boston and Singapore with an Nvidia 4070 RTX GPU. Predictions are made by combining local features extracted from RoIAlign, with global scene features extracted from the rasterized input using a ResNet-18 ConvNet backbone (see [87] for a detailed overview). We set the number of history and future frames to be 20 (2 seconds). We then use this prediction model to predict single trajectories for each other agent lasting 20 frames (2 seconds). We similarly train the BITS spatial planner to generate diverse, interactive goal poses for the ego agent. These poses serve as a global plan for the agent to follow. We solve the MPC using Casadi [4] on a CPU and model the ego agent with the extended unicycle dynamics model [39]:

$$\begin{bmatrix} x_{t+1} \\ y_{t+1} \\ v_{t+1} \\ \theta_{t+1} \end{bmatrix} = \begin{bmatrix} x_t + \Delta t \cdot v_t \cos(\theta_t) \\ y_t + \Delta t \cdot v_t \sin(\theta_t) \\ v_t + \Delta t \cdot a_t \\ \theta_t + \Delta t \cdot \omega_t \end{bmatrix} \tag{5.1}$$

where x_t, y_t is the position, v_t is the velocity, θ_t is the heading (orientation), and Δt is the sampling time (0.1 s). The control inputs include the yaw rate ω_t and the longitudinal acceleration a_t. We use a standard quadratic cost that minimizes the control effort and tracking error. Further, we constrain the velocity, acceleration and yaw rate to: $v \in [-5, 50]$, $a \in [-6, 6]$, $\omega \in [-8, 8]$ and the cost values used in the experiments are as follows:

Table 5.1: MPC costs

Cost x	Cost y	Cost v	Cost a	Cost ω
1	5	1	0.5	2

Table 5.2 shows the key simulation parameters. Particularly, we construct CP prediction regions using a failure probability $\delta = 0.2$. Our trajectory predictor predicts other agents' trajectories for a horizon of 2 s and considers agents in a 30 m radius of the ego agent.

Table 5.2: Simulation Parameters

Failure Probability δ	Prediction Horizon	Task Horizon	Ego action timestep	Agent Interaction Distance
0.2	2 s	20 s	0.2 s	30 m

Table 5.3: BITS simulator parameters

Parameter	Value
Mask drivable	True
Number of plan samples	50
Number of action samples	20
Pos to yaw	True
Yaw correction speed	1.0
Diversification clearance	None
Sample	True

Table 5.4: BITS simulator cost parameters

Cost Weights	Value
Collision weight	20.0
Lane weight	20.0
Lane direction weight	20.0
Likelihood weight	0.1
Progress weight	0.005

We evaluate the ego failure rate, which is caused by either collisions or going offroad, the average minimum distance, average jerk, average MPC cost, constraint violations (i.e., use of slack variable that we discuss later below), conformal efficiency, outliers detected, and updated prediction region coverage. The ego failure rate is determined as the percentage of time steps where the ego agent is in a collision or offroad. Conformal efficiency is defined as the average size of the prediction sets over all τ-step ahead predictions [79]. We compare the interaction robust MPC to three baselines:

1. Safe planning in dynamic environments with conformal prediction (SPDE) [43]: in this baseline, distribution shifts are not considered and Boole's inequality is used to construct the finite-horizon prediction regions. However, for a fair comparison, we use the prediction regions constructed in Section 3.4
2. Worst-case $\hat{\delta} = 0.01$: a variant of the interaction robust MPC that uses a constant worst-case distribution shift $\hat{\delta} = 0.01$ without outlier detection
3. Interaction robust with C_{CL}: a variant of the proposed planner that uses the closed-loop prediction region C_{CL} from (3.15) for outlier detection instead of the proposed Huber adjusted prediction region $C_{CL,h}$ from (4.4)

We use the following ellipsoidal collision avoidance constraint that takes into account the ego agent and other agents' extents denoted as (L^e, W^e) and (L^o, W^o), respectively, and prediction regions $\hat{C}_{\tau|t}$ or $C_{\tau|t}$ when an outlier is not detected.

$$c(X_t^e, \hat{X}_t^{o,e}) = \frac{(X_{\tau|t}^e[0] - \hat{X}_{\tau|t}^{o,e}[0])^2}{\frac{L^e}{\sqrt{2}} + \frac{L^o}{\sqrt{2}} + \tilde{C}_{\tau|t}} + \frac{(X_{\tau|t}^e[1] - \hat{X}_{\tau|t}^{o,e}[1])^2}{\frac{W^e}{\sqrt{2}} + \frac{W^o}{\sqrt{2}} + \tilde{C}_{\tau|t}} - 1 \geq 0, \quad \forall t \geq 0, \forall \tau \in \{t+1, \ldots, t+H\} \tag{5.2}$$

The prediction regions $C_{\tau|t}$ are computed using the train-val subset (200 scenes) of the Nuscenes dataset, and we compute 4891 close-loop residuals. Further, $\sigma_{\tau|t,j}$ from (3.13) is trained on a subset of the training data (100 scenes). The histogram of close-loop residuals is as follows:

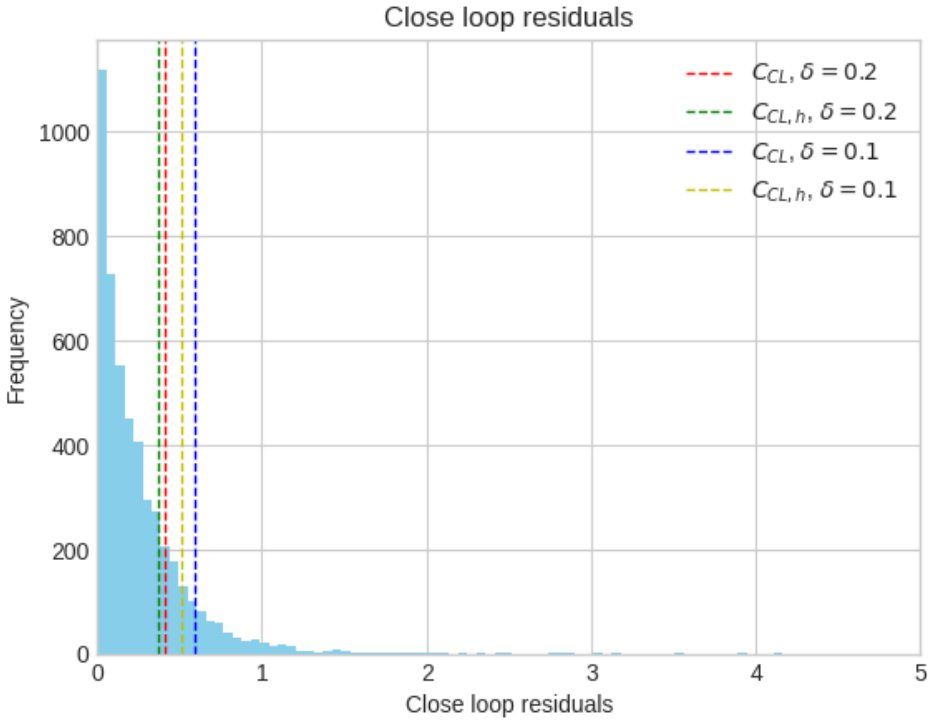

Figure 5.2: Close loop residuals

From observation, we note that the values of $C_{\tau|t}$ (from (3.15)) are similar at different times t. Therefore we define the τ-step ahead prediction region using the minimum value of $C_{\tau|t}$ across all observed t.

Further, we introduce a slack variable for constraint (4.6e) to take into account noise in the orientation of the elliptical constraints due to rasterization (see Figure 5.3). To determine the Huber quantile, we first run simulations using the ground truth data to check how many outliers are detected with no distribution shift and then tune the parameter λ to ensure we detect the expected number of outliers i.e., for $\delta = 0.2$, it is expected that on average 20% of the scene is detected as outliers under no distribution shift (ground truth data). Lastly, we implement robust CP with interpolation to satisfy the run-time requirements of the MPC.

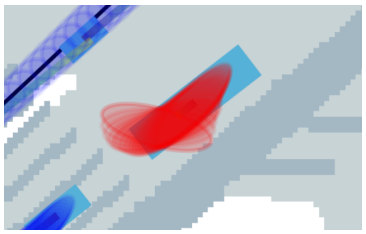

Figure 5.3: Rasterization noise alters the orientation of the elliptical collision avoidance constraint

5.2. Qualitative Results

We present the following key qualitative results that show several insights into the performance of the different planners and prediction model:

- *Prediction Accuracy and Interaction Effects:* We observed frequent detection of outliers near the ego agent, suggesting a decrease in predictor performance during interactions. Interestingly, outliers were also detected at distances where they should have minimal influence on the ego agent's behavior.
- *Impact of Historical Data on Prediction:* A significant factor in reducing conservatism in prediction regions was the number of historical timesteps used. Models trained with 20 history frames (2s) produced notably tighter prediction regions compared to those using 10 frames (1s).
- *Worst-Case MPC Performance:* The worst-case MPC exhibited challenges in scenarios with neighboring agents on both sides, as the conservative prediction regions limited the availability of feasible paths. This resulted in many instances of off-road behavior compared to the other approaches.
- *SPDE Baseline Behavior:* In contrast to the worst-case MPC, the SPDE baseline demonstrated a tendency to navigate much closer to other agents. This behavior occasionally led to collisions in certain scenarios due to insufficient maintenance of safe distances especially because of the added variability in the other agents speed that we added to create out of distribution behaviors.
- *Interaction robust MPC with C_{CL} for outlier detection:* The interaction robust MPC showed improved feasibility compared to the worst-case approach, particularly in cluttered multi-agent scenarios. It identified static vehicles as non-outliers, allowing for planning in lanes with agents on both sides. However, in some instances where other agents demonstrated out-of-distribution behavior, it was not able to reach quickly to the other agents changes.
- *Interaction robust MPC with $C_{CL,h}$ for outlier detection:* An enhanced version of interaction robust MPC with improved outlier detection. capabilities maintained better distances from other agents compared to SPDE, while showing only minor differences from the standard interaction robust MPC. This version detected more outliers that reduced close encounters with other agents in some scenarios. However, the increased outlier detection occasionally led to reduced feasibility in dense scenarios, resulting in jerky or off-road maneuvers.

We note that the proposed interaction robust MPC variant using a Huber quantile for tunable outlier detection achieved a balance between worst-case MPC and standard interaction robust MPC.

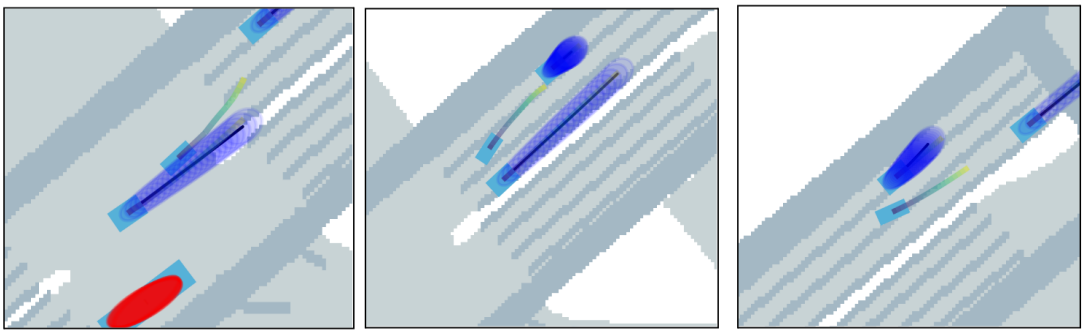

Figure 5.4: Scene 0093, (left to right) shows the ego agent navigating around two other vehicles identified as outliers (blue prediction regions)

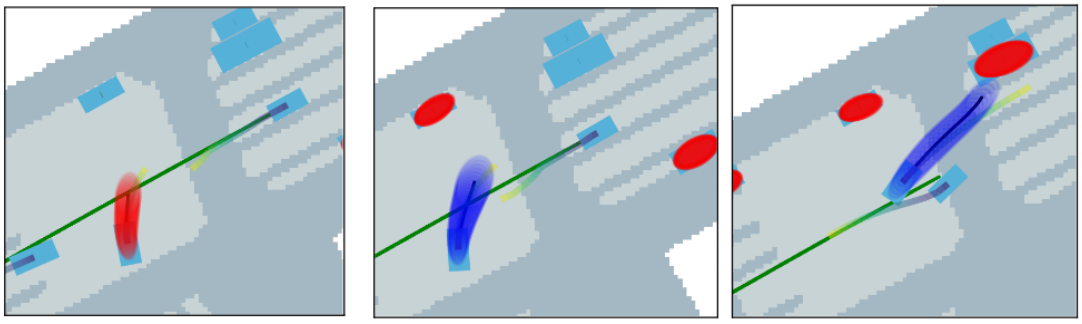

Figure 5.5: Scene 0522, (left to right) shows an agent detected as an outlier and maneuver by the ego agent to avoid collision

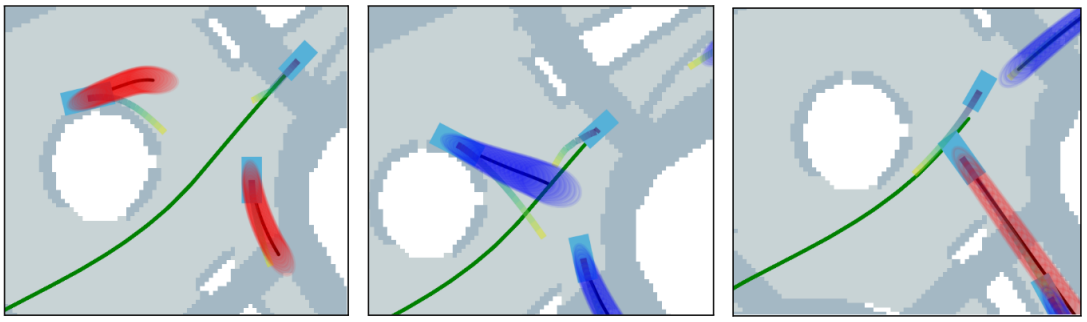

Figure 5.6: Scene 0523, (left to right) Show a roundabout scenario, where an outlier is detected and the ego agent deviates from its plan to prevent a collision

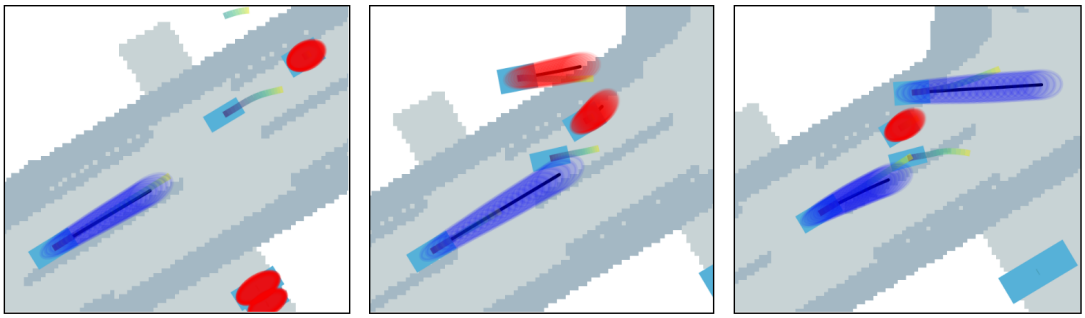

Figure 5.7: Scene 0330, (left to right) shows a lane merge scenario, where the other agent slows down to let the ego agent merge safely

5.3. Quantitative Results

To evaluate our approach, we analyze both conformal prediction metrics and simulation metrics, focusing on safety, outlier detection, and controller performance. We begin by examining results from ground truth data, which we use to calibrate the tuning parameter λ in the Huber loss.

Experiments where the other agents follow ground truth behavior show that the outliers detected using C_{CL} for $\delta = 0.2$ and $\delta = 0.1$ are fewer than expected (see Tables 5.5 and 5.6). We note that we expect to detect the calibrated number of outliers as there is no distribution shift between the calibration and ground truth data used. This can be due to the limitations of using the time-lagged nonconformity score as an estimate for the actual test residual, and can also stem from our argument in Section 4.2 that in nominal (in-distribution) scenarios, out-of-distribution events are rare. This rarity of samples leads to epistemic uncertainty in the model due to limited training, leading to inflated prediction regions to ensure CP coverage. To address this we adjust λ in subsection 4.2 to be $MAR/0.6745$ and re-evaluate using the ground truth policy for the other agent behaviors. Tables 5.5 and 5.6 demonstrate that this λ value leads to a higher number of detected outliers. We note that λ can be further tuned to reach a desired outlier detection rate, in our experiments we adopt the above value for the results that are presented below.

Table 5.5: Outliers detected in ground truth data (no distribution shift), $\delta = 0.2$

	Outliers Detected
Interaction Robust w/$C_{CL,h}$	21%
Interaction Robust w/C_{CL}	17%

Table 5.6: Outliers detected in ground truth data (no distribution shift), $\delta = 0.1$

	Outliers Detected
Interaction Robust w/$C_{CL,h}$	12%
Interaction Robust w/C_{CL}	8%

Table 5.7: Conformal Prediction Metrics, $\delta = 0.2$

	Conformal Efficiency (m) (↓)	Mean Empirical Coverage Level With Outlier	Detected Outliers	Empirical Coverage
Interaction Robust w/$C_{CL,h}$	0.52	0.87	31%	77%
Interaction Robust w/C_{CL}	0.52	0.88	28.9%	78%
Worst-case $\hat{\delta} = 0.01$	0.73	0.99	28.3%	93%
SPDE [43]	0.5	0.8	28.7%	70%

Table 5.8: Simulation Metrics, $\delta = 0.2$

	Average minimum distance (m) (↑)	Constraint Violation (↓)	Average Jerk (ms^{-3})(↓)	Average MPC cost (↓)	Ego failure rate (↓)
Interaction Robust w/$C_{CL,h}$	0.39	20%	1.20	15.52	16%
Interaction Robust w/C_{CL}	0.39	21%	1.22	17.2	19.3%
Worst-case $\hat{\delta} = 0.01$	0.43	28	1.45	28.2	24%
SPDE [43]	0.27	22%	1.23	17.79	22%

Table 5.9: Mean Runtime of Interaction Robust MPC

Predictor	Build and Solve Time	KL divergence Estimator
28Hz	20 Hz	2000 Hz

We now compare the four approaches we conducted experiments on. Table 5.7 and 5.8 summarize the main experimental results. Overall, our findings show that our proposed interaction robust planner using $C_{CL,h}$ for outlier detection demonstrated the best balance between conservatism and almost nearly achieve conformal coverage (i.e., we would like 80% and our adaptive method increases coverage from 69% to 77%), see Figure 5.16 and 5.15. It achieved the lowest use of slack with the constraint violation being 20% and the lowest ego failure rate of 16% (see Figures 5.12 and 5.14) We note that a slack cost of 1000 is used. However, we note that the constraint violations may not all be attributed to the planner but also due to limitations in the rasterization as illustrated in Figure 5.3. The proposed planner has an average minimum distance of 0.39 m that aligns close to the minimum distance achieved with the worst-case planner demonstrating the planner's robustness (see Figure 5.9). We note that the methods have about equal jerk values (see Figure 5.8), instead of the worst-case planner. We highlight that the proposed planner has the lowest MPC cost demonstrating the efficiency of using outlier detection to induce interaction robustness when needed compared to at all times like in the worst-case planner.

Even though the worst-case approach provided the highest conformal coverage of 93%, it resulted in the highest average MPC cost of 28.20 (see Figure 5.10), and the highest constraint violations of 28% which suggests the planner is overly conservative. On the other hand, not being robust to the interaction-induced distribution shifts led to the SPDE baseline having the lowest conformal coverage of 70% and the lowest average minimum distance of 0.27 m. We note that the detected outliers in Table 5.7 represent the drop in coverage under distribution shift (see also Figure 5.13). They are all calculated using the time-lagged nonconformity score (4.1) in a close-loop manner.

By providing near-expected coverage while maintaining robustness behavior as seen with the average minimum distance, the proposed planner achieves results in out-of-distribution settings that are comparable to those obtained in scenarios without distribution shift, demonstrating its robustness to interaction-induced shifts.

Lastly, Table 5.9 shows the runtime for the interaction robust MPC. The predictor and KL-divergence estimator were run on a Nvidia 4070 RTX GPU, and the MPC was run on a CPU using Casadi. The runtime of the KL divergence estimator is shown when using 3500 residual samples and $k = 50$. The number of agents varied across scenes from 5 to 20 agents. As demonstrated, estimating the KL divergence can be done very efficiently allowing us to update the prediction regions individually in the prediction horizon. Another data-efficient approach would be to only take the maximum shift and update all regions with a uniform shift.

Continuing, we highlight that the permissible distribution shift is also influenced by the interaction potential parameters γ and h that were set to 0.2 and 8, respectively. By tuning these parameters the distribution shift could increase leading to more conservatism as shown in the worst-case-shift planner. Thus an appropriate choice of the interaction potential is vital to balancing conservatism and achieving the conformal coverage under distribution shift. Along the same line, the Huber loss tuning parameter λ also has a direct effect on the conservatism of the solution. By having a smaller $C_{CL,h}$ more outliers are detected which would result in more prediction regions with higher coverage leading to conservatism.

In conclusion, experimental results demonstrate the effectiveness of the proposed interaction robust planner using $C_{CL,h}$ for outlier detection. This approach aims to provide an optimal balance between conservatism and conformal coverage and outperforms both the overly conservative worst-case approach and the non-robust SPDE baseline. By achieving near-expected coverage under distribution shift, our planner shows it can adapt to interaction-induced distribution shifts. The method's efficiency is evident in its low MPC cost and minimal use of slack, while still maintaining safe distances comparable to the worst-case planner. Furthermore, the runtime analysis reveals the computational feasibility of our approach, even with a variable number of agents.

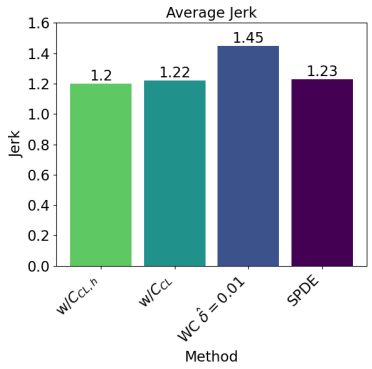

Figure 5.8: Average Jerk (ms^{-3})

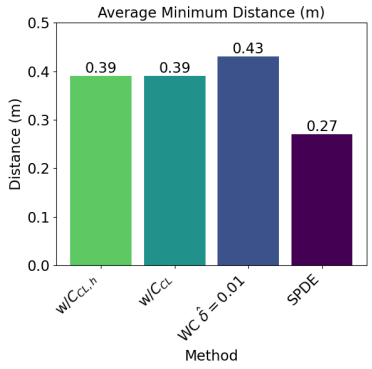

Figure 5.9: Average minimum distance (m)

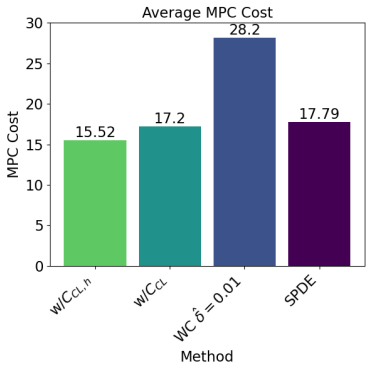

Figure 5.10: Average MPC cost (-)

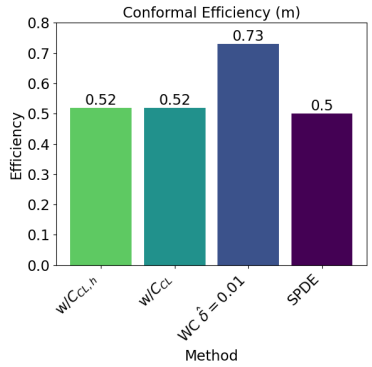

Figure 5.11: Conformal efficiency (m)

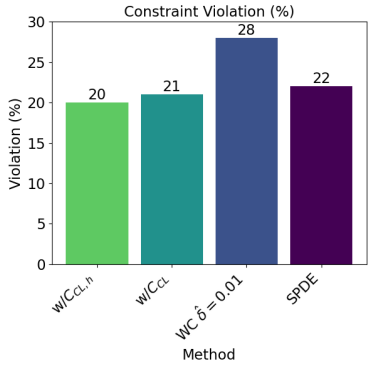

Figure 5.12: Constraint Violation (use of slack variable)

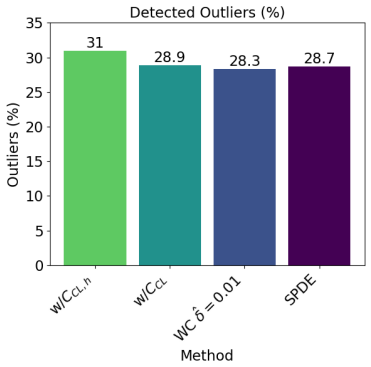

Figure 5.13: Detected outliers

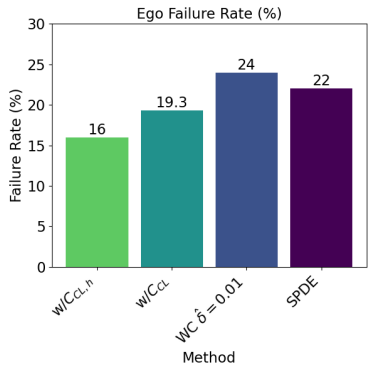

Figure 5.14: Ego failure rate

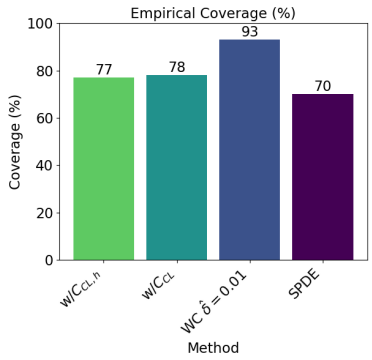

Figure 5.15: Conformal empirical coverage level

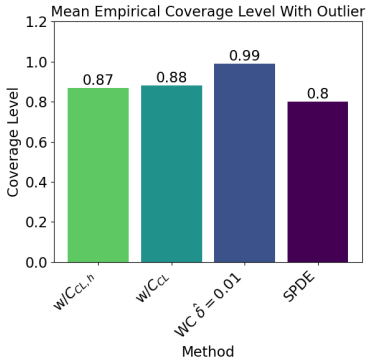

Figure 5.16: Mean empirical coverage level with outlier

6

Conclusion

6.1. Summary

Autonomous motion planning requires the ability to safely reason about learned trajectory predictors, particularly in settings where an agent can influence other agents' behavior. These learned predictors are essential for anticipating the future states of uncontrollable agents, whose decision-making process can be difficult to model analytically. Ensuring safe planning and control requires robust uncertainty quantification of these predictors, especially in scenarios where the ego agent's plan induces a distribution shift in other agents' behaviors. Existing methods often assume specific agent behavior models, limiting their applicability and potentially leading to overly conservative or unsafe behavior. Moreover, while state-of-the-art prediction models capture interactions through ego-conditioning, integrating these models with optimization-based planners is computationally prohibitive.

In this thesis we present an adaptive framework for safe interactive motion planning in unknown dynamic environments, addressing the challenge of uncertainty quantification in learned trajectory predictors when the ego agent interacts with other agents. The approach combines model predictive control (MPC) with tools from conformal prediction (CP) to provide safety assurances in the presence of interaction-induced distribution shifts. The thesis introduces an adaptive MPC-based framework that accounts for distribution shifts in learned trajectory predictors resulting from ego-conditioning, along with a method for estimating these shifts during runtime using a heuristic interaction potential term.

The framework first uses CP to detect outliers in other agents' behaviors. It employs a time-lagged nonconformity score to estimate the test residual and a Huber quantile to tune the outlier detector. When outliers are detected, the system estimates a distribution shift using the interaction potential term and a k-NN based KL divergence estimator. The framework then updates prediction regions using robust CP to account for the estimated distribution shift. These updated regions are then incorporated into an MPC formulation as collision avoidance constraints.

Our interaction robust MPC-based planner is evaluated in diverse autonomous driving scenarios using the Bi-level imitation for traffic simulation (BITS) closed loop simulator. It achieved near-expected conformal coverage (77% vs 80% desired) with the lowest constraint violations (20%) and ego failure rate (16%). Maintaining an average minimum distance of 0.39 m, comparable to the worst-case (conservative) planner, it showed robustness while achieving the lowest MPC cost. This balanced approach outperformed the over-conservative worst-case method (93% coverage but highest MPC cost and violations) and the less robust SPDE baseline (70% coverage, 0.27 m average minimum distance). Notably, our method demonstrated comparable performance in out-of-distribution settings to scenarios without distribution shift, indicating robustness to interaction-induced shifts.

In light of our findings and proposed method, we revisit the research questions that guided this thesis:

- *What techniques can **quantify** the uncertainty of complex deep-learned trajectory forecasting methods, and how can we detect outliers reliably?*

We demonstrated the effectiveness of CP techniques in quantifying uncertainty without making assumptions about data distribution or the predictive model. Our approach, using a time-lagged nonconformity score and a Huber quantile, proved effective in detecting outliers, achieving near-expected conformal coverage (77% vs 80% desired) in simulations.

- *How can learning-based trajectory forecasting methods be **integrated** into interactive motion planning frameworks with probabilistic safety assurances?*

  Our adaptive MPC-based framework successfully integrated robust CP to quantify uncertainty in trajectory predictors under agent interactions. This integration provided probabilistic safety assurances by dynamically adjusting the distribution shift that robust CP accounts for.

- *How can we **model** the interactions between agents when the counterfactual motion of other agents is unknown?*

  We developed an interaction potential term that estimates the test residual distribution given samples from the calibration data, and we propose using a k-NN-based KL divergence estimator to estimate distribution shifts during runtime.

6.2. Limitations and Future Work

The result of this thesis work opens many interesting research directions to address limitations and improve the existing method:

- A limitation of our approach emerges when the ego agent and other agents come into proximity. There could be scenarios where other agents are not cooperative, and as the distance between the agents decreases, the interaction potential term would increase, leading to larger estimated distribution shifts. This would result in larger prediction regions for the other agents and could result in feasibility issues in the ego agent plan due to their proximity. Such scenarios would require slack in the constraints as the distribution shift increases to ensure that the ego agent does not result in taking unsafe or overly conservative behaviors impeding its objective.
- The accuracy of our prediction method could be significantly enhanced by considering the diversity of driving scenarios. Currently, we compute residuals from trajectories across various situations such as turning, lane merging, and highway driving. However, the prediction difficulty, and thus the distribution of residuals, can vary greatly between these scenarios. For instance, predicting the trajectory of a vehicle maintaining a constant speed on a highway is typically easier than anticipating how a vehicle might navigate a turn. To address this, it would be worth labeling scenarios in autonomous driving datasets. By categorizing residuals based on specific driving scenarios, we could create more targeted prediction regions. When an agent is identified to be in a particular scenario (e.g., driving straight on a highway), we would use the residuals specific to that scenario to compute the prediction region. This method could help reduce conservatism and provide context awareness by constructing prediction regions precisely for specific scenarios.
- To address the 'curse of rarity' as introduced in section 4.2, it would be worth exploring how to develop robust prediction models by generating safety-critical scenarios and training prediction models to minimize the quantile and Huber quantile losses. There is existing work to generate safety-critical scenarios to develop robust prediction models [66], and it would be interesting to train these models using the quantile and/or Huber quantile loss to reduce epistemic uncertainty [36].
- Given that our current method works for KL divergences, a major limitation is that the ego-conditioned distribution must have an overlapping domain with the training distribution. This could limit the types of interactive behaviors that can be considered, making it challenging to measure larger interaction uncertainties. Another metric that does not require overlapping domains is the Wasserstein metric, which outputs finite values when the domains do not overlap. Therefore, it would be worth considering to extend the robust CP framework to Wasserstein distances to take into account larger distribution shifts which may not necessarily have overlapping domains.
- Under assumptions on the trajectory predictor type, it would be worth exploring the use of probabilistic trajectory forecasting methods to generate i.i.d. test trajectory samples to compute distribution shift estimates online. In particular, the predictor could be evaluated with and without

ego-conditioning, and the samples generated from both distributions could be used to compute a distribution shift estimate. Moreover, with current probabilistic predictors making assumptions on the distribution (i.e., Gaussian mixture models [60]), closed-form solutions to the shift could be used to then update prediction regions using robust CP. Along the same lines, by predicting the evaluator only once with ego-conditioning, the test samples could be used in an exchangeability martingales framework to test if the samples are in distribution with respect to the calibration residuals.

A

Scientific Paper

This Appendix includes the research paper that has resulted from this thesis work.

Robust Conformal Prediction for Adaptive Motion Planning Among Interactive Agents

Mayank Prashar[1], Javier Alonso-Mora[1], Lars Lindemann[2]

***Abstract*— Autonomous motion planning requires the ability to safely reason about learned trajectory predictors, particularly in settings where an agent can influence other agents' behavior. These learned predictors are essential for anticipating the future states of uncontrollable agents, whose decision-making process can be difficult to model analytically. Thus, uncertainty quantification of these predictors is crucial for ensuring safe planning and control. In this work, we introduce a framework for interactive motion planning in unknown dynamic environments with probabilistic safety assurances. We adapt a model predictive controller (MPC) to distribution shifts in learned trajectory predictors when other agents react to the ego agent's plan. Our approach leverages tools from conformal prediction (CP) to detect when the other agent's behavior deviates from the training distribution and employ robust CP to quantify the uncertainty in trajectory predictors during these agent interactions. We propose a method for estimating interaction-induced distribution shifts during runtime and the Huber quantile for enhanced outlier detection. Using a KL divergence ambiguity set that upper bounds the distribution shift, this method constructs prediction regions with probabilistic assurances in the presence of distribution shifts caused by interactions with the ego agent. We evaluate our framework in interactive scenarios involving navigation around autonomous vehicles in the BITS simulator, demonstrating enhanced safety and reduced conservatism.**

I. INTRODUCTION

A major challenge in motion planning for autonomous systems is the ability to safely reason about interactions with other agents [1]. A typical approach involves predicting the motion of other agents through heuristic [2], [3], or learned models [4], followed by the creation of an ego (controllable) agent plan considering these predictions. Such *predict-then-plan* schemes often overlook the influence of the ego agents' plan on other agents, resulting in overly conservative behavior that has been shown to lead to unsafe situations [5]. State-of-the-art (SOTA) prediction models [6], [7] attempt to address this limitation through ego-conditioning, where the prediction of other agents is conditioned on the ego agent's motion plan. However, integrating ego-conditioned prediction models with iterative planners is computationally prohibitive, and existing methods lack uncertainty quantification tools to provide safety assurances [8], [9], [10]. While iterative planners provide constraint satisfaction guarantees

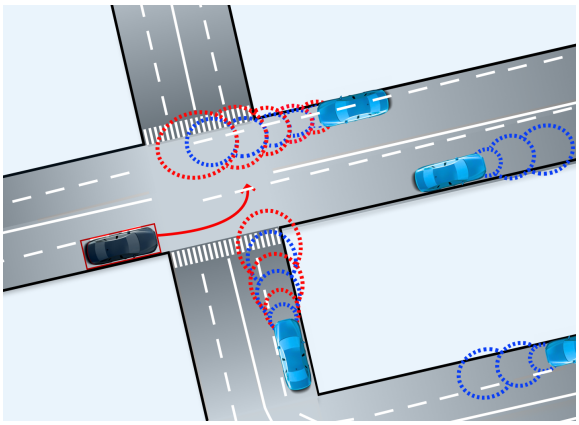

Fig. 1. Using robust conformal prediction, we adjust the prediction regions of other agents (shown in light blue) in response to the ego agent's plan (shown in red) at test time

[11], quantifying the uncertainty of prediction models allows us to design safety constraints with probabilistic assurances [12], [13].

Existing uncertainty quantification methods in interactive motion planning typically assume specific agent behavior models [14], [8], [15], [16] or predictors [5], [17], [18], [19], [20], [21]. These approaches restrict the types of disturbances and prediction models considered, lack safety assurances, or exhibit overly conservative unsafe behavior when attempting to do so [22]. Conformal prediction (CP) is a statistical tool that rigorously provides valid prediction sets to quantify model uncertainty without assumptions on the data distribution or predictive model [23], [24]. It has been proven as an effective framework for uncertainty quantification in complex deep-learned trajectory predictors, offering safety assurances over finite-horizon trajectories [12] [13]. However, CP assumes exchangeability between train and test distributions. While current autonomous driving training data mostly captures well-structured, non-interactive behavior [25], [26], real-world (test) scenarios often involve interactive behaviors where the ego agent's plan can influence other agents' behaviors. Such interactive scenarios necessitate a method to quantify uncertainty in trajectory predictors under interaction-induced distribution shifts.

To address this challenge, this paper proposes an adaptive model predictive control (MPC)-based framework for interactive planning with safety assurances. We use CP to quantify uncertainty in learned trajectory predictors under

[1]Authors are with the Department of Cognitive Robotics (CoR), Delft University of Technology, 2628 CD Delft, Netherlands.

[2]Author is with the Thomas ord Department of Computer Science, University of Southern California, Los Angeles, CA, USA.

interactions (see Figure 1). To ensure safety amid interaction-induced distribution shifts, we detect these shifts using CP and employ robust CP methods [27] to capture nearby distributions upper bounded by a KL divergence. We then formulate an MPC with robust prediction regions as collision avoidance constraints. Our contributions are as follows:

- We develop an adaptive MPC-based framework for safe interactive planning that accounts for distribution shifts in learned trajectory predictors resulting from ego-conditioning. Our approach leverages tools from conformal prediction to detect and quantify these shifts, enabling the planner to adapt during runtime.
- We propose a method for estimating interaction-induced distribution shifts during runtime and a Huber quantile for enhanced outlier detection.
- We empirically show that our framework works in interactive scenarios involving navigation around autonomous vehicles in the Bi-level Imitation for Traffic Simulation (BITS) simulator [28].

II. RELATED WORK

A. Interactive Planning

Early approaches have modeled the behavior of agents with Gaussian uncertainty while not considering the influence of the ego agent, leading to conservative plans as seen in the freezing robot problem [5]. To take into account interaction, initial approaches employed joint optimization via Gaussian processes [5], [29], artificial potential fields [30], and reinforcement learning to capture interactions [31].

Multi-agent interactions have also been modeled via reachability sets [32] that help determine the set of possible states a system can reach given the system's dynamics and time horizon. Even though this provides safety assurances, it results in overly conservative behaviors and infeasibility of the planning problem with longer time horizons [33]. Furthermore, when the exact dynamics of other agents are unknown, providing safety assurances can be challenging.

Other model-based approaches use game theory to formulate a coupled trajectory optimization problem under the assumption that an objective model of all players is available. Such approaches address bi-directional interactions, yet assume simple behavior models such as the rational model [16] or noisy Boltzmann models [15] and are difficult to scale. Partially observable Markov decision processes (POMDP) formulation also results in intractable partially observable stochastic game (POSG) [34].

B. Motion Planning with Probabilistic Safety Assurances

Approaches toward probabilistic safety assurances in planning frameworks typically attempt to quantify uncertainty in other agents' motions. A common assumption is that the other agent's motion is a random process and a distribution can be determined provided independent and identically distributed (i.i.d.) samples of trajectories. Scenario optimization [35] computes a predicted set over other agents based on samples from previously observed scenarios. These are formulated to create chance constraints that contain the given samples with a confidence level. Previous works have used chance constraints with Gaussian process models [36], Gaussian uncertainty [37], and arbitrary uncertainty distributions [38]. Other approaches have assumed access to a human-in-the-loop dataset of driving behaviors, which is used to find the minimum area set that contains trajectories within a given confidence level [39]. Similarly, quantile regression has been used to learn bounds over a dataset, which are then used in a tube MPC [40]. Other techniques to generate bounds on the prediction model include post-bloating, support vector machines, and conformal prediction as shown in [41]. Overall, methods that offer confidence bounds attempt to ensure that 1 - δ of the predicted trajectories should be within some learned bounds [42].

Recent works utilized CP to quantify uncertainty in trajectory predictors and integrated them into an MPC framework, providing safety assurances over finite-horizon trajectories using Boole's inequality [12]. Recent approaches have relaxed the assumption that the training and test trajectories are i.i.d. Using adaptive CP [43], the confidence level is adjusted online to take into account distribution shifts; however, it provides asymptotic guarantees [13]. Other approaches have used robust CP [27] for online verification under general distribution shift [44] and for out-of-distribution detection under general distribution shift, which is then used to adapt a controller [45].

C. Motion Prediction

Behavior models are broadly categorized into ontological and phenomenological methods. Ontological frameworks capture other agents' interaction dynamics by assuming that agents strongly tend to avoid collisions. Examples include the Optimal Reciprocal Collision Avoidance (ORCA) [46], Social Force Model [3], and Intelligent Driver Model [2]. These methods explicitly derive interaction dynamics that make them explainable and work well for short time horizons. However, these methods solely use the current state, making them unsuitable for longer predictions [39]. Phenomenological methods use deep-learned models to predict trajectories [4], [6], [7]. Initial methods produced single (unimodal and deterministic) predictions [4]. Recent works account for uncertainty by generating multiple (multimodal and probabilistic) trajectories [6]. These models are also categorized as agent-centric and scene-centric which predict trajectories of other agents separately and jointly, respectively. However, these methods lack explainability and have uncertainty in training the model (epistemic uncertainty) and learning distributions from finite data (aleatoric uncertainty) [42]. Furthermore, assumptions are made on the distribution of the agents (e.g., Gaussian mixture models) [6]. Scene-centric models, though more realistic, can be computationally expensive due to a lack of parallelization. Moreover, most datasets use well-structured, and cooperative motion [47]. This may not be representative of the diverse behaviors and interactions possible [25]. As prediction models are trained on such data it is unclear how effective they are in interaction-aware scenarios.

III. PRELIMINARIES

A. Problem Formulation

Consider a discrete-time dynamical system $X_{t+1}^i = f^i(X_t^i, u_t^i)$ where $X_t^i \in \mathcal{X} \subseteq \mathbb{R}^n$ and $u_t^i \in \mathcal{U} \subseteq \mathbb{R}^m$ denote the state and the control input of the agent $i \in \{e, o\}$ at time $t \in \mathbb{N} \cup \{0\}$, respectively. Here, e represents the ego (controllable) agent and o represents the other (uncontrollable) agents. The sets $\mathcal{U}$ and $\mathcal{X}$ denote the set of permissible control inputs and the workspace of the agent. The measurable function $f^i : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ describes the agents' dynamics. Let $X_t^o := (X_{t,1}^o, ..., X_{t,N}^o) \in \mathbb{R}^{Nn}$ refer to the joint state of N other agents at time t, such that $X_{t,j}^o$ is the state of agent j at time t. The joint other agent trajectories $X_{0:T}^o$ is sampled from an unknown distribution $\mathcal{D}$ that obeys the dynamics f^o:

$$X^o := (X_0^o, X_1^o, \dots, X_t^o) \sim \mathcal{D}$$

Assumption 1 We have access to a dataset $D := \{X^{o(1)}, \dots, X^{o(K)}\}$ containing K trajectories where $X^{o(i)} := (X_0^{o(i)}, X_1^{o(i)}, \dots, X_T^{o(i)})$ are drawn i.i.d. from a training distribution $\mathcal{D}$.

Assumption 2 At test time, the ego agent's state X_t^e influences the behavior of the N other agents, shifting their trajectory distribution from the training distribution $\mathcal{D}$ to an unknown, ego-conditioned distribution $\mathcal{D}^e$. The resulting test time trajectories are denoted as $X_{0,T}^{o,e}$ and are sampled from the unknown distribution $\mathcal{D}^e$:

$$X^{o,e} := (X_0^o, X_1^o(X_0^e), \dots, X_t^o(X_0^e, X_1^e, \dots, X_{t-1}^e)) \sim \mathcal{D}^e$$

where these test time trajectories capture the influence of the ego agent's state on the other agents' trajectories.

Assumption 3 For any time $t \geq 0$, the Kullback-Leibler (KL) divergence between the test distribution $\mathcal{D}^e$ and training distribution $\mathcal{D}$, is bounded by an ego-state-dependent function $\epsilon(X_t^e)$, where X_t^e is the ego agent's state. For notational simplicity, we denote the bound as ϵ:

$$D_{KL}(\mathcal{D}^e || \mathcal{D}) \leq \epsilon$$

It is important to note that this bound is not uniform across all ego agent states but varies with X_t^e. The ambiguity set $\mathcal{P}_{KL,\epsilon}(\mathcal{D})$ with radius ϵ is defined as:

$$\mathcal{P}_{KL,\epsilon}(\mathcal{D}) := \{\mathcal{D}^e \text{ s.t. } D_{KL}(\mathcal{D}^e || \mathcal{D}) \leq \epsilon\}$$

This set captures all distributions $\mathcal{D}^e$ within an ϵ KL divergence from the training distribution $\mathcal{D}$. The value ϵ is unknown a priori and variable. It can be estimated during runtime when the ego agent state X_t^e is known. We make no assumptions on the form of the distributions $\mathcal{D}$ and $\mathcal{D}^e$.

Assumption 3 is applicable in many practical scenarios involving autonomous vehicles. It models situations where other agents generally follow their expected behavior ($\mathcal{D}$), but may deviate (resulting in $\mathcal{D}^e$) in response to the ego agent's state as agents tend to strongly prefer avoiding collisions. The use of the KL divergence to quantify shifts is also consistent with approaches in related work. Where the KL divergence has been used to measure the degree of influence an agent has over another [48], and mutual information as an indicator of forthcoming collisions [49].

Problem 1 Given the system, access to trajectories $X^{o(i)}$ from a distribution $\mathcal{D}$, access to a trajectory $X^{o,e}$ from distribution $\mathcal{D}^e$, the KL divergence estimate $D_{KL}(\mathcal{D}^e || \mathcal{D}) \leq \epsilon$ valid for the current ego agent state X_t^e, and a failure probability $\delta \in (0, 1)$, our goal is to compute control inputs u_t^e and trajectory X_t^e such that the collision avoidance constraint $c : \mathbb{R}^n \times \mathbb{R}^{nN} \to \mathbb{R}$ is satisfied with a probability of at least $1 - \delta$ for all time t:

$$\mathbb{P}(c(X_t^e, X_t^{o,e}) \geq 0, \forall t \in \{0, \dots, T\}) \geq 1 - \delta \tag{1}$$

B. Conformal Prediction

Let $\{R^{(i)}\}_{i=0}^k$ be $k+1$ i.i.d. random variables. The variable $R^{(i)}$ is referred to as the nonconformity score, and is commonly defined as $R^{(i)} := ||Z^{(i)} - \mu(X^{(i)})||_2$, where the predictor is a mapping $\mu : X^{(i)} \to Z^{(i)}$. A large nonconformity score indicates a poor predictive model. Our goal is to construct a prediction region C for $R^{(0)}$ based on $\{R^{(i)}\}_{i=1}^k$ such that the random variable is contained in the prediction region C with a probability $1-\delta$, where $\delta \in (0, 1)$ is a user-specified failure probability [23], [24]:

$$\mathbb{P}(R^{(0)} \leq C) \geq 1 - \delta \tag{2}$$

Using the Quantile Lemma, we construct the prediction region C using the finite-sample corrected $(1-\delta)th$ quantile of the empirical distribution of $\{R^{(i)}\}_{i=0}^k$ (Lemma 1 [50]). Firstly, the quantile is defined as:

$$\text{Quantile}(1 - \delta, \mathcal{R}) = \inf\{z : \mathbb{P}\{Z \leq z\} \geq 1 - \delta\}, Z \sim \mathcal{R} \tag{3}$$

where Z is the empirical distribution of $\{R^{(i)}\}_{i=0}^k$ constructed using Dirac distributions $\delta_{R^{(i)}}$ centered at $R^{(i)}$. We then determine the quantile, where $\lceil \cdot \rceil$ is the ceiling function:

$$C = \text{Quantile}\left(\frac{\lceil (1-\delta)(K+1) \rceil}{K}; \frac{1}{K}\sum_{i=1}^{K} \delta_{R^{(i)}}\right) \tag{4}$$

Furthermore, two requirements must be met before determining C. The quantile must be bounded and the minimum number of data points K is required such that:

$$0 \leq \frac{\lceil (1-\delta)(K+1) \rceil}{K} \leq 1 \tag{5}$$

Provided that the requirement (5) is satisfied, we sort $\{R^{(1)}, ..., R^{(k)}\}$ in a non-decreasing order and let $R^{(k+1)} := \infty$. We then have $C := R^{(p)}$ to be the pth smallest nonconformity score where $p := \lceil (k+1)(1-\delta) \rceil$. When requirement (5) is violated, $C = \infty$. It is important to note that CP provides marginal coverage guarantees and does not imply conditional coverage. Specifically, the probability bound holds in expectation over the input space, and may not hold uniformly across all subsets of the input space. Additionally, CP's validity also holds when the nonconformity scores are exchangeable such that their joint distribution is unchanged for all permutations σ: $\{R^{\sigma(i)}\}_{i=0}^k$. However, in settings of distribution shift where the calibration residuals $\{R^{(i)}\}_{i=1}^k$ are drawn from a distribution $\mathcal{R}$ and the test

sample $R^{(0)}$ is drawn from a different distribution $\mathcal{R}^e$, the above methodology makes no claims of validity as the exchangeability assumption is not satisfied. This motivates the following robust variant of conformal prediction.

Lemma 1 *Robust Conformal Prediction* [27] provides valid prediction regions under sufficiently small distribution shifts between the calibration and test data. Denote $\mathcal{R}^e$ and $\mathcal{R}$ as the test and calibration residual distributions, respectively, where $R^{(0)} \sim \mathcal{R}^e$ and $\{R^i\}_{i=1}^K \sim \mathcal{R}$, such that their f-divergence $D_f(\mathcal{R}^e, \mathcal{R}) \leq \epsilon$. The adjusted prediction region $\tilde{C}$ in $\mathbb{P}(R^{(0)} \leq \tilde{C}) \geq 1-\delta$ is computed using an adjusted confidence level $\tilde{\delta}$ that takes into account a set of distributions $\mathcal{P}_{f,\epsilon}(\mathcal{R}) := \{\mathcal{R}^e \text{ s.t. } D_f(\mathcal{R}^e||\mathcal{R}) \leq \epsilon\}$ (see Corollary 2.2 [27]):

$$\begin{aligned} \tilde{C} &:= \text{Quantile}_{1-\tilde{\delta}}(R^{(1)}, ..., R^{(K)}), \\ \tilde{\delta} &:= 1 - g^{-1}(1-\delta_n) \end{aligned} \tag{6}$$

that is obtained with the following convex optimization problems:

$$\begin{aligned} \delta_n &:= 1 - g\left(\left(1+\frac{1}{K}\right) g^{-1}(1-\delta)\right) \\ g(\beta) &:= \inf\Big\{z \in [0,1] \;\Big|\; \beta f\left(\frac{z}{\beta}\right) + (1-\beta) f\left(\frac{1-z}{1-\beta}\right) \leq \epsilon\Big\} \\ g^{-1}(\tau) &:= \sup\big\{\beta \in [0,1] \mid g(\beta) \leq \tau\big\} \end{aligned} \tag{7}$$

Robust CP assumes a known, fixed amount of shift ϵ. However, our approach addresses scenarios where the shift ϵ is variable and unknown. In our setting, the shift ϵ to be robust against depends on the level of influence the ego agent has on the other agents. Therefore, this necessitates an adaptive approach to determining ϵ during run time.

While Assumption 3 bounds the trajectory-level distribution shift by ϵ, estimating this shift at the residual level is more efficient due to its one-dimensional nature. The data processing inequality allows us to compute the KL divergence on the nonconformity measure level, which is computationally advantageous and potentially provides tighter bounds compared to estimating distribution shifts on multi-dimensional trajectories.

Lemma 2 *Data Processing Inequality* [51]: Given a nonconformity measure $R : X \to \mathbb{R}$, $X^{o,e} \sim \mathcal{D}^e$ and $X^o \sim \mathcal{D}$, let R^e and R follow a distribution $\mathcal{R}^e$ and $\mathcal{R}$ such that they are the push-forward of $\mathcal{D}^e$ and $\mathcal{D}$, respectively. It holds that the f-divergence between two distributions does not increase when pushed through R:

$$D_f(\mathcal{D}^e||\mathcal{D}) \leq \epsilon \to D_f(\mathcal{R}^e||\mathcal{R}) \leq \epsilon$$

Therefore, to determine ϵ to solve (6), (7), we can estimate the KL divergence using M samples from $\mathcal{R}$ and L samples from $\mathcal{R}^e$. We further elaborate in Section IV-B of our proposed framework how we determine samples from $\mathcal{R}^e$.

Lemma 3 *k-NN Based KL Divergence Estimator* [52], [53]: We consider 1-dimensional samples $\{R_1^e, \ldots, R_L^e\}$ drawn i.i.d. from $\mathcal{R}^e$, and $\{R_1, \ldots, R_M\}$ drawn i.i.d. from $\mathcal{R}$. $\eta_i := ||R_i^e - R_{i(k)}^e||_1$ is the absolute difference between R_i^e and its k-nearest neighbor (k-NN) in $\{R_j^e\}_{j \neq i}$. Similarly, $\nu_i := ||R_i^e - R_{i(k)}||_1$ is the absolute difference between R_i^e and its k-NN in $\{R_1, \ldots, R_M\}$. The k-NN-based KL-divergence estimator is defined as follows [52], [53]:

$$\hat{D}_{KL}(\mathcal{R}^e||\mathcal{R}) = \frac{d}{L}\sum_{i=1}^{L} \log\frac{\nu_k(i)}{\eta_k(i)} + \log\frac{M}{L-1} \tag{8}$$

It has been shown under mild regularity conditions that the k-NN divergence estimator (8) is asymptotically unbiased (Theorem 1 [53]):

$$\lim_{L,M\to\infty} \mathbb{E}[\hat{D}_{KL}(\mathcal{R}^e||\mathcal{R})] = D_{KL}(\mathcal{R}^e||\mathcal{R})$$

and mean-square consistent (Theorem 2 [53]):

$$\lim_{L,M\to\infty} \mathbb{E}[(\hat{D}_{KL}(\mathcal{R}^e||\mathcal{R}) - D_{KL}(\mathcal{R}^e||\mathcal{R}))^2] = 0$$

We note that the samples must be drawn i.i.d. to prove mean-square consistency.

C. Conformal Prediction for Finite-Horizon Trajectories

Let a dataset D contain trajectories of the other agents as stated in Assumption 1. Let $D_{train} \subset D$, $D_{cal} \subset D$, and $D_{train} \cap D_{cal} = \emptyset$. Consider a function PREDICT that is trained on observations $(X_0^o, \ldots, X_t^o)$ from D_{train} and outputs predictions $(\hat{X}^o_{t+1|t}, \ldots, \hat{X}^o_{T|t})$. With the following predictions we define the nonconformity scores as introduced in Section III-B with a normalization term $\sigma_{\tau|t,j} > 0$ determined using a subset of the training data $D_{train,1}$ to conform to the exchangeability assumption in CP. The nonconformity score for all future times $\tau \in \{t+1, \ldots, T\}$ and other agents j is defined as [54], [55]:

$$\begin{aligned} R^{(i)}_{\tau|t} &:= \max_{\substack{(\tau,j)\in \\ \{1,\ldots,T\}\times\{1,\ldots,N\}}} \frac{||X^{o(i)}_{\tau,j} - \hat{X}^{o(i)}_{\tau|t,j}||}{\sigma_{\tau|t,j}}, \\ \sigma_{\tau|t,j} &:= \max_{i\in\{1,\ldots,K\}} ||X^{o(i)}_{\tau,j} - \hat{X}^{o(i)}_{\tau|t,j}|| \end{aligned} \tag{9}$$

This approach enables us to achieve more efficient prediction regions for each agent compared to using Boole's inequality [12], [13]. Provided the nonconformity scores $R^{(i)}_{\tau|t}$ determined using D_{cal}, we can construct prediction regions for finite-horizon trajectories (i.e., over multiple future predictions) in an open-loop and closed-loop setting. The open-loop ensures the validity of all τ-step ahead prediction regions (i.e., $\tau|0$), while the closed-loop ensures the validity of all one-step ahead prediction regions (i.e., $t+1|t$). Particularly, the following two statements hold for open-loop and closed-loop regions, respectively:

$$\begin{aligned} P(||X^{o(i)}_{\tau,j} - \hat{X}^{o(i)}_{\tau|0,j}|| \leq C_{OL}\sigma_{\tau|0,j}, \\ (\tau,j) \in \{t+1,\ldots,T\}\times\{1,\ldots,N\}) \geq 1-\delta \end{aligned} \tag{10}$$

$$\begin{aligned} P(||X^{o(i)}_{t+1|t,j} - \hat{X}^{o(i)}_{t+1|t,j}|| \leq C_{CL}\sigma_{t+1|t,j}, \\ (t,j) \in \{0,\ldots,T-1\}\times\{1,\ldots,N\}) \geq 1-\delta \end{aligned} \tag{11}$$

where $C_{CL} := R^{(p)}_{CL}$ and $C_{OL} := R^{(p)}_{OL}$, are the pth smallest closed loop and open loop nonconformity scores, respectively, with $p := \lceil (k+1)(1-\delta) \rceil$.

However, in settings of interaction-induced distribution shift such that during test time the other agents are under some level of influence from the ego agent, the prediction regions in (10) and (11) are not valid. To address this, we require an estimate of the distribution shift ϵ which depends on the test nonconformity scores. By solving (6) and (7) we can then obtain updated prediction regions $\tilde{C}_{CL}$ and $\tilde{C}_{OL}$ that remain valid under the specified distribution shift. However, evaluating the nonconformity score defined in (9) at test time t is challenging, as only states $(\hat{X}^{o,e}_0, \dots, \hat{X}^{o,e}_t)$ are observable, and not the future states $(\hat{X}^{o,e}_{t+1}, \dots, \hat{X}^{o,e}_T)$. To overcome this limitation, we introduce the time-lagged nonconformity score, detailed in Section IV-A.

IV. INTERACTION ROBUST MODEL PREDICTIVE PLANNER

This section proposes an interactive planning framework that uses robust CP for learned trajectory predictors to address interaction uncertainty. The approach uses observed data of other agents to compute test-time residuals to detect outliers, defined as scores not contained within the prediction regions constructed with the calibration residuals. When an outlier is detected, the prediction region is updated to adapt to the interaction uncertainty by estimating real-time distribution shifts (i.e., KL divergence) in other agents' behavior caused by the ego agent's presence. The estimated shift is used to update the prediction regions by solving a series of convex optimization problems as defined in equations (6) and (7) to ensure $1-\delta$ coverage under distribution shift.

Using outlier detection also addresses the planner's computational efficiency and feasibility under distribution shifts. The framework reduces computational burden in multi-agent settings by estimating shifts only when CP guarantees with calibration residuals under-cover (i.e., when an outlier is detected) rather than for all agents. This also benefits the planner's feasibility by limiting changes in collision avoidance constraints across iterations to only agents under interaction uncertainty.

Lastly, under Assumption 1 in section III-A, we use an agent-centric prediction model that can predict trajectories of other agents taking into account their dynamics and semantic information, as detailed in [28]. Our framework is compatible with both deterministic and probabilistic trajectory predictors. The PREDICT function takes as input the observed trajectory of other agents $X^{o,e}_{0:t}$, scene information S, and outputs a prediction $\hat{X}^{o,e}_{\tau|t}$ of the other agents.

A. Outlier Detection

The proposed framework employs outlier detection using CP to identify when the observed behavior of other agents deviates from the predicted behavior. This requires computing a nonconformity score R^e during test time to determine if it lies in the prediction region $C_{\tau|t}$ computed based on calibration residuals defined in (9). An outlier is detected when $R^e > C_{\tau|t}$. However, recall that the nonconformity score defined in (9) cannot be evaluated at time t as only the previous states $(\hat{X}^{o,e}_0, \dots, \hat{X}^{o,e}_t)$ are observable, while the future states $(\hat{X}^{o,e}_{t+1}, \dots, \hat{X}^{o,e}_T)$ are unknown. Therefore, we define a time-lagged nonconformity score that evaluates the τ step-ahead prediction error that was made τ time steps ago:

$$R^e := \max_{\substack{(\tau,j)\in \\ \{1,\dots,T\}\times\{1,\dots,N\}}} \frac{||X^{o,e}_{\tau,j} - \hat{X}^{o,e}_{\tau|t-\tau,j}||}{\sigma_{\tau|t,j}} \tag{12}$$

That is valid when $t \geq \tau$. Furthermore, the prediction region $C_{\tau|t}$ is computed using the calibration residuals (9) and follows a Beta distribution. With a sufficiently large calibration set, we expect the outlier rate to approach the calibrated failure probability δ. However, it is important to note that the prediction region $C_{\tau|t}$ is computed using calibration residuals from diverse driving scenarios where the frequency of certain scenarios can be much lower than more commonly observed scenarios. Also known as the 'curse of rarity' [56] where the occurrence of safety-critical scenarios (i.e., scenarios in the long tail of the distribution) are observed rarely. This imbalance can lead to epistemic uncertainty in regions of the input space with limited data [57]. Consequently, this may lead to large nonconformity scores that lead to conservatism while determining the prediction region $C_{\tau|t}$ and thus fewer detected outliers. However, we want a prediction region that takes into account epistemic uncertainty due to the curse of rarity. It should be able to detect more outliers such that we leverage our proposed conservative controller that is designed to capture distribution shifts.

To address this challenge and obtain a less conservative prediction region $C_{\tau|t}$ for outlier detection, we use a quantile that minimizes the Huber loss to reduce epistemic uncertainty. The Huber loss gives less weight to residuals above a threshold parameter λ by using a linear loss and more weight to residuals below λ by using a quadratic loss [58]:

$$L_\lambda(R^{(i)}) = \begin{cases} \frac{1}{2}(R^{(i)})^2 & \text{if } R^{(i)} \leq \lambda, \\ \lambda\delta(|R^{(i)}| - \frac{1}{2}\lambda) & \text{if } R^{(i)} > \lambda. \end{cases} \tag{13}$$

where λ is commonly chosen to be $MAR/0.6745$ [58], with MAR denoting the median absolute residual. We note that this is a tunable parameter that can be adjusted to ensure δ outliers are detected during experiments. To compute MAR, we use a subset of the training data $D_{train,2}$, to conform to the exchangeability assumption in CP. To compute the Huber quantile, we substitute (13) in the quantile loss function:

$$L_{q,\lambda}(R^{(i)}) = \begin{cases} \delta \cdot L_\lambda(R^{(i)} - q), & \text{if } R^{(i)} \geq q, \\ (1-\delta) \cdot L_\lambda(R^{(i)} - q), & \text{otherwise,} \end{cases} \tag{14}$$

and the quantile is found by minimizing this Huber quantile loss over the calibration set of size k:

$$\hat{q} = \arg\min_q \sum_{i=1}^{k} L_{q,\lambda}(R^{(i)}) \tag{15}$$

where $C_{CL,h} := \hat{q}$ in the case the closed-loop residuals are minimized. Formally, we classify the test residual as an outlier if $R^e > C_{CL,h}$.

Algorithm 1 Residuals for Conformal Prediction

Input: Failure probability δ, calibration dataset $\mathcal{D}_{\text{cal}}$, training dataset $\mathcal{D}_{\text{train}}$, prediction and task horizons H and T, threshold parameter λ
Output: $C_{CL}, C_{OL}, C_{CL,h}, C_{OL,h}, \sigma_{\tau|t}$
$p \leftarrow (|\mathcal{D}_{\text{cal}}| + 1)(1 - \delta)$
for $t = 0$ **to** $T - 1$ **do** ▷ Using Training Data
  for $\tau = t + 1$ **to** $t + H$ **do**
    Obtain predictions $\hat{X}^{o(i)}_{\tau|t}$ for each $X^{o(i)} \in \mathcal{D}_{\text{cal}}$
  end for
  $\sigma_{\tau|t} \leftarrow \max_i ||X^{o(i)}_{\tau,j} - \hat{X}^{o(i)}_{\tau|t,j}||$
end for
for $t = 0$ **to** $T - 1$ **do** ▷ Using Calibration Data
  for $\tau = t + 1$ **to** $t + H$ **do**
    Obtain predictions $\hat{X}^{o(i)}_{\tau|t}$ for each $X^{o(i)} \in \mathcal{D}_{\text{cal}}$
    $R^{(i)}_{CL} \leftarrow \max_\tau \frac{\left\|X^{o(i)}_{\tau|t} - \hat{X}^{o(i)}_{\tau|t}\right\|}{\sigma_{\tau|t}}$ for each $X^{o(i)} \in \mathcal{D}_{\text{cal}}$
    $R^{(i)}_{OL} \leftarrow \max_\tau \frac{\left\|X^{o(i)}_{\tau|t} - \hat{X}^{o(i)}_{\tau|t}\right\|}{\sigma_{\tau|0}}$ for each $X^{o(i)} \in \mathcal{D}_{\text{cal}}$
  end for
end for
$R^{|\mathcal{D}_{\text{cal}}+1|}_{CL} \leftarrow \infty, R^{|\mathcal{D}_{\text{cal}}+1|}_{OL} \leftarrow \infty$
Sort $R^{(i)}_{CL}$, and $R^{(i)}_{OL}$ in non-decreasing order
$C_{CL} \leftarrow R^{(p)}_{CL}, C_{OL} \leftarrow R^{(p)}_{OL}$
Compute Huber quantiles $\hat{q}_{OL}$, and $\hat{q}_{CL}$ by solving (15)
$C_{CL,h} \leftarrow \hat{q}_{CL}, C_{OL,h} \leftarrow \hat{q}_{OL}$

B. Capturing Distribution Shift

Provided an outlier is detected, we propose updating the prediction regions to take into account interaction uncertainty which is inherent with longer prediction horizons [39]. Particularly, we use robust CP to dynamically adapt to the aleatoric uncertainty arising from interactions with the ego agent. This method provides prediction regions that are valid for a KL divergence ambiguity set defined as $\mathcal{P}_{KL,\epsilon}(\mathcal{R}) := \{\mathcal{R}^e \text{ s.t. } D_{KL}(\mathcal{R}^e||\mathcal{R}) \leq \epsilon\}$, where $\mathcal{R}^e$ represents the distribution of residuals at test time under interaction, $\mathcal{R}$ is the calibration residual distribution, and ϵ is the KL divergence bound. To compute the updated prediction region $\tilde{C}$, the KL divergence must be estimated. However, this would require multiple test residual samples to estimate a distribution and thus the distribution shift. However, we have access to only calibration residuals and a single trajectory prediction of the other agent.

To address this limitation, we propose using an interaction potential term to estimate the test distribution. This approach is on the basis that residuals increase in magnitude due to interaction uncertainty.

Assumption 4. Let the test residuals R^e be a function of the calibration R residuals with the following heuristic interaction potential to model the influence of the ego agent:

$$R^e = \frac{R}{1 - \gamma \exp(-\frac{1}{2h^2} d^2)} \tag{16}$$

where d is the L^2 norm distance between the ego and other agents, γ, and h are parameters that determine the strength of the interaction. This interaction potential allows us to ego-condition M samples of the calibration residuals, providing a heuristic estimate of the test residual distribution. We can then use this estimate to determine the KL divergence ϵ on the residual level using equation (8). Given the estimated ϵ, we use robust CP to find the updated prediction region $\tilde{C}_{CL}$ by solving (6) and (7). This process allows us to dynamically adapt the predictions to take into account interaction uncertainty that can vary at each time step by computing a new estimate ϵ.

C. Model Predictive Control

With the above sections we formulate an optimization problem for the ego agent given a cost function J, the initial condition of the ego and other agents, and predictions of the other agents' trajectories. The objective is to compute control inputs u^e_t such that the collision probability with the other agents is less than or equal to δ over the task horizon T:

$$\min_{X^e, u^e} \sum_{k=t}^{t+H-1} J(X^e_{k+1}, u^e_k) \tag{17a}$$

s.t.,

$$X^e_0 = \xi^e, \tag{17b}$$

$$X^e_{k+1} = f^e(X^e_k, u^e_k), \qquad \forall k \in \{t, \dots, t+H-1\} \tag{17c}$$

$$u^e_k \in \mathcal{U},\ X^e_{k+1} \in \mathcal{X}, \qquad \forall k \in \{t, \dots, t+H-1\} \tag{17d}$$

$$c(X^e_k, \hat{X}^{o,e}_{k,j}) \geq 0, \forall (k,j) \in \{t, \dots, t+H-1\} \times \{1, \dots, N\} \tag{17e}$$

The optimization problem in (17) formulates an approach that uses robust CP prediction regions as constraints for a model predictive control (MPC) method to compute control inputs u^e over a prediction horizon H. The objective function J in (17a) minimizes a step-wise cost function over the prediction horizon H. The constraints (17b) set the initial condition where ξ^e is the current state of the ego agent. Constraint (17c) enforces the ego agent dynamical constraints represented by f^e. Constraint (17d) ensures the control inputs u^e and states X^e remain within their feasible sets $\mathcal{U}$ and $\mathcal{X}$, respectively. Constraint (17e) enforces the probabilistic safety constraint. It utilizes the robust conformal prediction regions for $\hat{X}^{o,e}_k$ when an outlier is detected, and prediction regions from calibration residuals when no outlier is detected. The function $c(\cdot, \cdot)$ represents the safety criterion, for example, to ensure a safe distance between the ego agent and the predicted states and regions of the other N agents.

The above optimization problem (17) is then solved by linearizing the cost, dynamics, and constraints to solve the quadratic program in a closed-loop (receding horizon) strategy.

V. EXPERIMENTS

To evaluate our approach in interactive scenarios, we perform experiments in driving scenarios between an autonomous vehicle and other human-driven vehicles. As we want to evaluate our approach in interactive settings, using

Algorithm 2 MPC with Robust Conformal Prediction

Input: Failure probability δ, calibration dataset $\mathcal{D}_{\text{cal}}$, prediction and task horizons H and T
Output: Control input u_t^e
for $t = 0$ **to** $T-1$ **do** ▷ Online perform planning loop
  Sense X_t^e and $X_t^{o,e}$
  Obtain predictions $\hat{X}_{\tau|t}^{o,e}$ for $\tau = t+1, \ldots, t+H$
  if test residual from (12) $> C_{CL,h}$ **then**
    Obtain M samples of R^e from (16)
    Obtain ϵ using the KL divergence estimator (8)
    Obtain $\tilde{C}_{\tau|t} := \tilde{C}_{CL}\sigma_{\tau|t}$ given ϵ,
  else
    Obtain $C_{\tau|t} := C_{CL}\sigma_{\tau|t}$
  end if
  Calculate controls $u_t^e, \ldots, u_{T-1}^e$ as the solution of (17)
  Apply u_t^e to the ego agent
end for

replayed scenes from the NuScenes dataset would be insufficient to evaluate the planner. Therefore, we conduct closed-loop simulations in the BITS simulation environment [28]. Each simulation scene is initialized with one of the 100 test scenes from the NuScenes dataset and then evolves forward with a policy for each agent. We choose one agent as the ego agent that our proposed planner controls. All the other agents run the BITS policy that generates realistic and diverse behaviors of road vehicles. Additionally, every 4 seconds, we spawn agents around the ego agent to increase the number of interactions in each scene. Our closed-loop planner uses three modules: a trajectory predictor, a high-level route planner, and the interaction robust MPC that plans the trajectory.

We use the deterministic BITS prediction model as the trajectory predictor without ego-conditioning, which generates a single trajectory lasting 20 timesteps (2 seconds). The model is trained on the NuScenes dataset [59], consisting of 1000 20-second scenes collected in Boston and Singapore. The BITS spatial planner is used to generate diverse global plans for the ego agent that provides a goal pose (position and heading). Table I shows the key simulation parameters:

TABLE I
SIMULATION PARAMETERS

Failure Probability δ	Prediction Horizon	Task Horizon	Ego Action Timestep	Agent Interaction Distance
0.2	2 s	20 s	0.2 s	30 m

We evaluate the ego failure rate which is caused by either collisions or going offroad, the average minimum distance, average jerk, average MPC cost, constraint violations (i.e., use of slack variable that we discuss later below), conformal efficiency, outliers detected, and updated prediction region coverage. The ego failure rate is determined as the percentage of time steps where the ego agent is in a collision or offroad. Conformal efficiency is defined as the average size of the prediction sets over all τ-step ahead predictions [60]. We compare the interaction robust MPC to three baselines:

1) Safe planning in dynamic environments with conformal prediction (SPDE) [12]: in which distribution shifts are not considered and Boole's inequality is used to construct the prediction regions. For a fair comparison, we use the prediction regions constructed in Section III-C with this method
2) Worst-case $\hat{\delta} = 0.01$: A variant of the interaction robust MPC that uses a constant worst-case distribution shift $\hat{\delta} = 0.01$ without outlier detection
3) Interaction robust with C_{CL}: A variant of the proposed planner that uses the closed-loop prediction region C_{CL} from (11) for outlier detection instead of the proposed Huber adjusted prediction region $C_{CL,h}$ from (15)

We use the following ellipsoidal collision avoidance constraint that takes into account the ego agent and other agents' extents denoted as (L^e, W^e) and (L^o, W^o), respectively, and prediction regions $\hat{C}_{\tau|t}$ or $C_{\tau|t}$ when an outlier is not detected.

$$c(X_t^e, \hat{X}_t^{o,e}) = \frac{(X_{\tau|t}^e[0] - \hat{X}_{\tau|t}^{o,e}[0])^2}{\frac{L^e}{\sqrt{2}} + \frac{L^o}{\sqrt{2}} + \tilde{C}_{\tau|t}} + \frac{(X_{\tau|t}^e[1] - \hat{X}_{\tau|t}^{o,e}[1])^2}{\frac{W^e}{\sqrt{2}} + \frac{W^o}{\sqrt{2}} + \tilde{C}_{\tau|t}} - 1 \geq 0, \quad \forall t \geq 0, \forall \tau \in \{t+1, \ldots, t+H\} \tag{18}$$

From observation, we note that the values of $C_{\tau|t}$ are similar at different times t. Therefore, we define the τ-step ahead prediction region using the minimum value of $C_{\tau|t}$ across all observed t. Further, we introduce a slack variable for constraint (17e) to take into account noise in the orientation of the elliptical constraints due to rasterization. To ensure the planner runs in real-time, we interpolate values from robust CP. To determine the Huber quantile, we first run simulations using the ground truth data to check how many outliers are detected with no distribution shift. Using closed-loop test residuals, the detected outliers is 17% using ground truth data (no distribution shift). We then initialize the tuning parameter $\lambda = MAR/0.6745$ in the Huber loss and repeat simulations on ground truth data and detect 21% outliers. We note that λ can be further tuned to reach a desired outlier detection rate, in our experiments we adopt the above value for the results that are presented below.

TABLE II
CONFORMAL PREDICTION METRICS, $\delta = 0.2$

	Conformal Efficiency (m) (↓)	Mean Empirical Coverage Level With Outlier	Detected Outliers	Empirical Coverage
Interaction Robust w/$C_{CL,h}$	0.52	0.87	31%	77%
Interaction Robust w/C_{CL}	0.52	0.88	28.9%	78%
Worst-case $\hat{\delta} = 0.01$	0.73	0.99	28.3%	93%
SPDE [12]	0.5	0.8	28.7%	70%

TABLE III
SIMULATION METRICS, $\delta = 0.2$

	Average minimum distance (m) (↑)	Constraint Violation (↓)	Average Jerk (↓)	Average MPC cost (↓)	Ego failure rate (↓)
Interaction Robust w/$C_{CL,h}$	0.39	20%	1.20	15.52	16%
Interaction Robust w/C_{CL}	0.39	21%	1.22	17.2	19.3%
Worst-case $\hat{\delta} = 0.01$	0.43	28	1.45	28.2	24%
SPDE [12]	0.27	22%	1.23	17.79	22%

Our experimental results, summarized in Tables II and III, compare our method with the three baselines. Simulation videos can be found here. Interaction robust planner using

$C_{CL,h}$ outperforms baselines. Our planner achieved the best balance between conservatism and conformal coverage (i.e., we would like 80% and our adaptive method increases coverage from 69% to 77%). It demonstrated the lowest constraint violations (20%) and lowest ego failure rate (16%). Also, it has an average minimum distance of 0.39 m, comparable to the worst-case-shift planner. Notably, it achieved the lowest MPC cost, indicating efficient use of outlier detection for scenarios requiring interaction robustness. The worst-case-shift approach, while providing the highest conformal coverage (93%), resulted in the highest average MPC cost (28.20) and constraint violations (28%), suggesting over-conservatism. Conversely, the SPDE baseline, not robust to interaction-induced distribution shifts, had the lowest conformal coverage (70%) and average minimum distance (0.27 m) as it does not take into account interactions. Our method demonstrates comparable performance in out-of-distribution settings to scenarios without distribution shift by maintaining its conformal coverage level, indicating robustness to interaction-induced shifts. Detected outliers in Table II represent the coverage drop under distribution shift, calculated using the time-lagged nonconformity score (12) in a closed-loop manner.

Runtime analysis: Simulations showed efficient performance across scenes with agents varying from 5 to 20. The predictor and KL divergence estimator run on an Nvidia 4070 RTX GPU at 28 Hz and 2000Hz respectively. The MPC was run on a CPU using Casadi at 20Hz. The KL divergence estimation was performed using 3500 residual samples and $k = 50$.

VI. CONCLUSION

This paper presented a framework for safe interactive planning in unknown dynamic environments, introducing an adaptive MPC approach that leverages robust conformal prediction to address interaction-induced distribution shifts. Our key contributions include a method for real-time detection and quantification of these shifts using conformal prediction and KL divergence estimation and a Huber quantile approach for enhanced outlier detection. Experimental results in the BITS simulator demonstrated the effectiveness of our approach, achieving the lowest constraint violations (20%) and ego failure rate (16%). The proposed planner maintained comparable performance in out-of-distribution settings, indicating robustness to interaction-induced shifts.

REFERENCES

[1] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core challenges of social robot navigation: A survey," *J. Hum.-Robot Interact.*, vol. 12, no. 3, apr 2023. [Online]. Available: https://doi.org/10.1145/3583741

[2] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E*, vol. 62, no. 2, p. 1805–1824, Aug. 2000. [Online]. Available: http://dx.doi.org/10.1103/PhysRevE.62.1805

[3] D. Helbing and P. Molnár, "Social force model for pedestrian dynamics," *Phys. Rev. E*, vol. 51, pp. 4282–4286, May 1995. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevE.51.4282

[4] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 961–971.

[5] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010, pp. 797–803.

[6] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 683–700.

[7] Y. Yuan, X. Weng, Y. Ou, and K. Kitani, "Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. Los Alamitos, CA, USA: IEEE Computer Society, oct 2021, pp. 9793–9803. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/ICCV48922.2021.00967

[8] Y. Chen, U. Rosolia, W. Ubellacker, N. Csomay-Shanklin, and A. D. Ames, "Interactive multi-modal motion planning with branch model predictive control," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5365–5372, 2022.

[9] Y. Chen, S. Veer, P. Karkus, and M. Pavone, "Interactive joint planning for autonomous vehicles," *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 987–994, 2024.

[10] Y. Chen, P. Karkus, B. Ivanovic, X. Weng, and M. Pavone, "Tree-structured policy planning with learned behavior models," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 7902–7908.

[11] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109899002149

[12] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, "Safe planning in dynamic environments using conformal prediction," *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 5116–5123, 2023.

[13] A. Dixit, L. Lindemann, S. X. Wei, M. Cleaveland, G. J. Pappas, and J. W. Burdick, "Adaptive conformal prediction for motion planning among dynamic agents," in *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*, ser. Proceedings of Machine Learning Research, N. Matni, M. Morari, and G. J. Pappas, Eds., vol. 211. PMLR, 15–16 Jun 2023, pp. 300–314. [Online]. Available: https://proceedings.mlr.press/v211/dixit23a.html

[14] N. E. Du Toit and J. W. Burdick, "Robot motion planning in dynamic, uncertain environments," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 101–115, 2012.

[15] J. F. Fisac, E. Bronstein, E. Stefansson, D. Sadigh, S. S. Sastry, and A. D. Dragan, "Hierarchical game-theoretic planning for autonomous vehicles," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE Press, 2019, p. 9590–9596. [Online]. Available: https://doi.org/10.1109/ICRA.2019.8794007

[16] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for autonomous cars that leverage effects on human actions," in *Robotics: Science and Systems*, 2016. [Online]. Available: https://api.semanticscholar.org/CorpusID:7087988

[17] Y. Yoon, C. Kim, J. Lee, and K. Yi, "Interaction-aware probabilistic trajectory prediction of cut-in vehicles using gaussian process for proactive control of autonomous vehicles," *IEEE Access*, vol. 9, pp. 63 440–63 455, 2021. [Online]. Available: https://api.semanticscholar.org/CorpusID:233465718

[18] R. Toledo-Moreo and M. A. Zamora-Izquierdo, "Imm-based lane-change prediction in highways with low-cost gps/ins," *Trans. Intell. Transport. Syst.*, vol. 10, no. 1, p. 180–185, mar 2009. [Online]. Available: https://doi.org/10.1109/TITS.2008.2011691

[19] T. Gindele, S. Brechtel, and R. Dillmann, "A probabilistic model for estimating driver behaviors and vehicle trajectories in traffic environments," *13th International IEEE Conference on Intelligent Transportation Systems*, pp. 1625–1631, 2010. [Online]. Available: https://api.semanticscholar.org/CorpusID:9386434

[20] J. F. Fisac, A. V. Bajcsy, S. L. Herbert, D. Fridovich-Keil, S. Wang, C. J. Tomlin, and A. D. Dragan, "Probabilistically safe robot planning with confidence-based human predictions," *ArXiv*, vol. abs/1806.00109, 2018. [Online]. Available: https://api.semanticscholar.org/CorpusID:44166364

[21] D. Fridovich-Keil, A. V. Bajcsy, J. F. Fisac, S. L. Herbert, S. Wang, A. D. Dragan, and C. J. Tomlin, "Confidence-aware motion prediction for real-time collision avoidance1," *The International Journal of Robotics Research*, vol. 39, pp. 250 – 265, 2020. [Online]. Available: https://api.semanticscholar.org/CorpusID:198432212

[22] K. Zhou and J. C. Doyle, *Essentials of robust control*. Prentice hall Upper Saddle River, NJ, 1998, vol. 104.

[23] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*. Berlin, Heidelberg: Springer-Verlag, 2005.

[24] G. Shafer and V. Vovk, "A tutorial on conformal prediction," *J. Mach. Learn. Res.*, vol. 9, p. 371–421, jun 2008.

[25] W. Wang, C. Liu, and D. Zhao, "How much data are enough? a statistical approach with case study on longitudinal driving behavior," *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 2, pp. 85–98, 2017.

[26] K.-C. Hsu, K. Leung, Y. Chen, J. F. Fisac, and M. Pavone, "Interpretable trajectory prediction for autonomous vehicles via counterfactual responsibility," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 5918–5925.

[27] A. A. Maxime Cauchois, Suyash Gupta and J. C. Duchi, "Robust validation: Confident predictions even when distributions shift," *Journal of the American Statistical Association*, vol. 0, no. 0, pp. 1–66, 2024. [Online]. Available: https://doi.org/10.1080/01621459.2023.2298037

[28] D. Xu, Y. Chen, B. Ivanovic, and M. Pavone, "Bits: Bi-level imitation for traffic simulation," 2022.

[29] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015. [Online]. Available: https://doi.org/10.1177/0278364914557874

[30] M. T. Wolf and J. W. Burdick, "Artificial potential functions for highway driving with collision avoidance," in *2008 IEEE International Conference on Robotics and Automation*, 2008, pp. 3731–3736.

[31] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE Press, 2019, p. 6015–6022. [Online]. Available: https://doi.org/10.1109/ICRA.2019.8794134

[32] I. Mitchell, A. Bayen, and C. Tomlin, "A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on Automatic Control*, vol. 50, no. 7, pp. 947–957, 2005.

[33] M. Althoff and J. M. Dolan, "Set-based computation of vehicle behaviors for the online verification of autonomous vehicles," in *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, 2011, pp. 1162–1167.

[34] D. Sadigh, N. Landolfi, S. Sastry, S. Seshia, and A. Dragan, "Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state," *Autonomous Robots*, vol. 42, 10 2018.

[35] M. Campi and S. Garatti, "Wait-and-judge scenario optimization," *Mathematical Programming*, vol. 167, 07 2016.

[36] B. Luders, G. Aoude, J. Joseph, N. Roy, and J. How, "Probabilistically safe avoidance of dynamic obstacles with uncertain motion patterns," 07 2011.

[37] H. Zhu and J. Alonso-Mora, "Chance-constrained collision avoidance for mavs in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 776–783, 2019.

[38] O. de Groot, B. Brito, L. Ferranti, D. Gavrila, and J. Alonso-Mora, "Scenario-based trajectory optimization in uncertain dynamic environments," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5389–5396, 2021.

[39] K. Driggs-Campbell, V. Govindarajan, and R. Bajcsy, "Integrating intuitive driver models in autonomous planning for interactive maneuvers," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 12, pp. 3461–3472, 2017.

[40] D. D. Fan, A. akbar Agha-mohammadi, and E. A. Theodorou, "Deep learning tubes for tube mpc," *ArXiv*, vol. abs/2002.01587, 2020. [Online]. Available: https://api.semanticscholar.org/CorpusID:211032257

[41] Y. Chen, U. Rosolia, C. Fan, A. Ames, and R. Murray, "Reactive motion planning with probabilisticsafety guarantees," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 1958–1970. [Online]. Available: https://proceedings.mlr.press/v155/chen21e.html

[42] R. Cheng, R. M. Murray, and J. W. Burdick, "Limits of probabilistic safety guarantees when considering human uncertainty," 2021.

[43] I. Gibbs and E. Candes, "Adaptive conformal inference under distribution shift," in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 1660–1672.

[44] Y. Zhao, B. Hoxha, G. Fainekos, J. V. Deshmukh, and L. Lindemann, "Robust conformal prediction for stl runtime verification under distribution shift," in *2024 ACM/IEEE 15th International Conference on Cyber-Physical Systems (ICCPS)*, 2024, pp. 169–179.

[45] P. Contreras, O. Shorinwa, and M. Schwager, "Out-of-distribution runtime adaptation with conformalized neural network ensembles," 2024. [Online]. Available: https://arxiv.org/abs/2406.02436

[46] J. Alonso-Mora, A. Breitenmoser, M. Rufli, P. Beardsley, and R. Siegwart, *Optimal Reciprocal Collision Avoidance for Multiple Non-Holonomic Robots*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 203–216.

[47] S. Pellegrini, A. Ess, K. Schindler, and L. van Gool, "You'll never walk alone: Modeling social behavior for multi-target tracking," in *2009 IEEE 12th International Conference on Computer Vision*, 2009, pp. 261–268.

[48] E. Tolstaya, R. Mahjourian, C. Downey, B. Vadarajan, B. Sapp, and D. Anguelov, "Identifying driver interactions via conditional behavior prediction," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 3473–3479.

[49] R. Michelmore, M. Kwiatkowska, and Y. Gal, "Evaluating uncertainty quantification in end-to-end autonomous driving control," *arXiv preprint arXiv:1811.06817*, 2018.

[50] R. J. Tibshirani, R. Foygel Barber, E. Candes, and A. Ramdas, "Conformal prediction under covariate shift," in *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32. Curran Associates, Inc., 2019.

[51] N. J. Beaudry and R. Renner, "An intuitive proof of the data processing inequality," *Quantum Info. Comput.*, vol. 12, no. 5–6, p. 432–441, may 2012.

[52] Q. Wang, S. R. Kulkarni, and S. Verdu, "A nearest-neighbor approach to estimating divergence between continuous random vectors," in *2006 IEEE International Symposium on Information Theory*, 2006, pp. 242–246.

[53] ——, "Divergence estimation for multidimensional densities via k-nearest-neighbor distances," *IEEE Transactions on Information Theory*, vol. 55, no. 5, pp. 2392–2405, 2009.

[54] X. Yu, Y. Zhao, X. Yin, and L. Lindemann, "Signal temporal logic control synthesis among uncontrollable dynamic agents with conformal prediction," 12 2023.

[55] M. Cleaveland, I. Lee, G. J. Pappas, and L. Lindemann, "Conformal prediction regions for time series using linear complementarity programming," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 19, pp. 20 984–20 992, Mar. 2024. [Online]. Available: https://ojs.aaai.org/index.php/AAAI/article/view/30089

[56] W. Ding, C. Xu, M. Arief, H. Lin, B. Li, and D. Zhao, "A survey on safety-critical driving scenario generation—a methodological perspective," *Trans. Intell. Transport. Sys.*, vol. 24, no. 7, p. 6971–6988, jul 2023. [Online]. Available: https://doi.org/10.1109/TITS.2023.3259322

[57] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods," *Machine Learning*, vol. 110, 03 2021.

[58] P. J. Huber, "Robust Estimation of a Location Parameter," *The Annals of Mathematical Statistics*, vol. 35, no. 1, pp. 73 – 101, 1964. [Online]. Available: https://doi.org/10.1214/aoms/1177703732

[59] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

[60] V. Vovk, V. Fedorova, I. Nouretdinov, and A. Gammerman, "Criteria of efficiency for conformal prediction," in *Conformal and Probabilistic Prediction with Applications*, A. Gammerman, Z. Luo, J. Vega, and V. Vovk, Eds. Cham: Springer International Publishing, 2016, pp. 23–39.

References

[1] Alexandre Alahi et al. “Social LSTM: Human Trajectory Prediction in Crowded Spaces”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016, pp. 961–971. DOI: 10.1109/CVPR.2016.110.

[2] Javier Alonso-Mora et al. “Optimal Reciprocal Collision Avoidance for Multiple Non-Holonomic Robots”. In: *Distributed Autonomous Robotic Systems: The 10th International Symposium*. Ed. by Alcherio Martinoli et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 203–216. ISBN: 978-3-642-32723-0. DOI: 10.1007/978-3-642-32723-0_15.

[3] Matthias Althoff and John M. Dolan. “Set-based computation of vehicle behaviors for the online verification of autonomous vehicles”. In: *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*. 2011, pp. 1162–1167. DOI: 10.1109/ITSC.2011.6083052.

[4] Joel A. E. Andersson et al. “CasADi: A software framework for nonlinear optimization and optimal control”. In: *Mathematical Programming Computation* 11.1 (2019), pp. 1–36.

[5] Matthew Barth, Kanok Boriboonsomsin, and Guoyuan Wu. “Vehicle Automation and Its Potential Impacts on Energy and Emissions”. In: *Road Vehicle Automation*. Ed. by Gereon Meyer and Sven Beiker. Cham: Springer International Publishing, 2014, pp. 103–112. ISBN: 978-3-319-05990-7. DOI: 10.1007/978-3-319-05990-7_10. URL: https://doi.org/10.1007/978-3-319-05990-7_10.

[6] Normand J. Beaudry and Renato Renner. “An intuitive proof of the data processing inequality”. In: *Quantum Info. Comput.* 12.5–6 (May 2012), pp. 432–441. ISSN: 1533-7146.

[7] Jur van den Berg et al. “Reciprocal n-Body Collision Avoidance”. In: *Robotics Research*. Ed. by Cédric Pradalier, Roland Siegwart, and Gerhard Hirzinger. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 3–19. ISBN: 978-3-642-19457-3.

[8] Holger Caesar et al. “nuScenes: A Multimodal Dataset for Autonomous Driving”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2020.

[9] Feiyang Cai and Xenofon Koutsoukos. “Real-time Out-of-distribution Detection in Learning-Enabled Cyber-Physical Systems”. In: *2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCPS)*. 2020, pp. 174–183. DOI: 10.1109/ICCPS48487.2020.00024.

[10] Marco Campi and Simone Garatti. “Wait-and-judge scenario optimization”. In: *Mathematical Programming* 167 (July 2016). DOI: 10.1007/s10107-016-1056-9.

[11] Changan Chen et al. “Crowd-Robot Interaction: Crowd-Aware Robot Navigation With Attention-Based Deep Reinforcement Learning”. In: *2019 International Conference on Robotics and Automation (ICRA)*. Montreal, QC, Canada: IEEE Press, 2019, pp. 6015–6022. DOI: 10.1109/ICRA.2019.8794134. URL: https://doi.org/10.1109/ICRA.2019.8794134.

[12] Yuxiao Chen et al. “Interactive Joint Planning for Autonomous Vehicles”. In: *IEEE Robotics and Automation Letters* 9.2 (2024), pp. 987–994. DOI: 10.1109/LRA.2023.3332474.

[13] Yuxiao Chen et al. “Interactive Multi-Modal Motion Planning With Branch Model Predictive Control”. In: *IEEE Robotics and Automation Letters* 7.2 (2022), pp. 5365–5372. DOI: 10.1109/LRA.2022.3156648.

[14] Yuxiao Chen et al. “Reactive motion planning with probabilisticsafety guarantees”. In: *Proceedings of the 2020 Conference on Robot Learning*. Ed. by Jens Kober, Fabio Ramos, and Claire Tomlin. Vol. 155. Proceedings of Machine Learning Research. PMLR, 16–18 Nov 2021, pp. 1958–1970. URL: https://proceedings.mlr.press/v155/chen21e.html.

[15] Yuxiao Chen et al. “Tree-structured Policy Planning with Learned Behavior Models”. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. 2023, pp. 7902–7908. DOI: 10.1109/ICRA48891.2023.10161419.

[16] Richard Cheng, Richard M. Murray, and Joel W. Burdick. *Limits of Probabilistic Safety Guarantees when Considering Human Uncertainty*. 2021. arXiv: 2103.03388 [cs.RO].

[17] Derek Chiao et al. *Autonomous vehicles moving forward: Perspectives from industry leaders*. 2024. URL: https://www.mckinsey.com/features/mckinsey-center-for-future-mobility/our-insights/autonomous-vehicles-moving-forward-perspectives-from-industry-leaders#/ (visited on 01/01/2024).

[18] Matthew Cleaveland et al. "Conformal Prediction Regions for Time Series Using Linear Complementarity Programming". In: *Proceedings of the AAAI Conference on Artificial Intelligence* 38.19 (Mar. 2024), pp. 20984–20992. DOI: 10.1609/aaai.v38i19.30089. URL: https://ojs.aaai.org/index.php/AAAI/article/view/30089.

[19] Compass Transportation and Technology. *The Economic and Social Value of Autonomous Vehicles: Implications from Past Network-Scale Investments*. Report. Securing America's Future Energy (SAFE), 2018.

[20] Polo Contreras, Ola Shorinwa, and Mac Schwager. *Out-of-Distribution Runtime Adaptation with Conformalized Neural Network Ensembles*. 2024. arXiv: 2406.02436 [cs.RO]. URL: https://arxiv.org/abs/2406.02436.

[21] John Dahl et al. "Collision Avoidance: A Literature Review on Threat-Assessment Techniques". In: 4 (Dec. 2018), pp. 101–113. DOI: 10.1109/TIV.2018.2886682.

[22] Wenhao Ding et al. "A Survey on Safety-Critical Driving Scenario Generation—A Methodological Perspective". In: *Trans. Intell. Transport. Sys.* 24.7 (July 2023), pp. 6971–6988. ISSN: 1524-9050. DOI: 10.1109/TITS.2023.3259322. URL: https://doi.org/10.1109/TITS.2023.3259322.

[23] Anushri Dixit et al. "Adaptive Conformal Prediction for Motion Planning among Dynamic Agents". In: *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*. Ed. by Nikolai Matni, Manfred Morari, and George J. Pappas. Vol. 211. Proceedings of Machine Learning Research. PMLR, 15–16 Jun 2023, pp. 300–314. URL: https://proceedings.mlr.press/v211/dixit23a.html.

[24] Katherine Driggs-Campbell, Vijay Govindarajan, and Ruzena Bajcsy. "Integrating Intuitive Driver Models in Autonomous Planning for Interactive Maneuvers". In: *IEEE Transactions on Intelligent Transportation Systems* 18.12 (2017), pp. 3461–3472. DOI: 10.1109/TITS.2017.2715836.

[25] David D. Fan, Ali-akbar Agha-mohammadi, and Evangelos A. Theodorou. "Deep Learning Tubes for Tube MPC". In: *ArXiv* abs/2002.01587 (2020). URL: https://api.semanticscholar.org/CorpusID:211032257.

[26] Alec Farid et al. "Task-relevant failure detection for trajectory predictors in autonomous vehicles". In: *Conference on Robot Learning*. PMLR. 2023, pp. 1959–1969.

[27] Jaime F. Fisac et al. "Hierarchical Game-Theoretic Planning for Autonomous Vehicles". In: *2019 International Conference on Robotics and Automation (ICRA)*. Montreal, QC, Canada: IEEE Press, 2019, pp. 9590–9596. DOI: 10.1109/ICRA.2019.8794007. URL: https://doi.org/10.1109/ICRA.2019.8794007.

[28] Ryan Fonseca. *Autonomous vehicles: Promise and perils on the road to driverless cars*. https://www.latimes.com/california/newsletter/2023-03-01/essential-california-autonomous-vehicles-promise-essential-california. Accessed: 2024-08-19. Mar. 2023.

[29] Isaac Gibbs and Emmanuel Candes. "Adaptive Conformal Inference Under Distribution Shift". In: *Advances in Neural Information Processing Systems*. Ed. by M. Ranzato et al. Vol. 34. Curran Associates, Inc., 2021, pp. 1660–1672.

[30] Isaac Gibbs and Emmanuel J. Candès. "Conformal Inference for Online Prediction with Arbitrary Distribution Shifts". In: *Journal of Machine Learning Research* 25.162 (2024), pp. 1–36. URL: http://jmlr.org/papers/v25/22-1218.html.

[31] Oscar de Groot et al. "Scenario-Based Trajectory Optimization in Uncertain Dynamic Environments". In: *IEEE Robotics and Automation Letters* 6.3 (2021), pp. 5389–5396. DOI: 10.1109/LRA.2021.3074866.

[32] Dirk Helbing and Péter Molnár. “Social force model for pedestrian dynamics”. In: *Phys. Rev. E* 51 (5 May 1995), pp. 4282–4286. DOI: 10.1103/PhysRevE.51.4282. URL: https://link.aps.org/doi/10.1103/PhysRevE.51.4282.

[33] Kai-Chieh Hsu et al. “Interpretable Trajectory Prediction for Autonomous Vehicles via Counterfactual Responsibility”. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2023, pp. 5918–5925. DOI: 10.1109/IROS55552.2023.10341712.

[34] Peter J. Huber. “Robust Estimation of a Location Parameter”. In: *The Annals of Mathematical Statistics* 35.1 (1964). [Online; accessed 2024-07-05], pp. 73–101. URL: https://doi.org/10.1214/aoms/1177703732.

[35] Peter J. Huber. “Robust Estimation of a Location Parameter”. In: *The Annals of Mathematical Statistics* 35.1 (1964), pp. 73–101. DOI: 10.1214/aoms/1177703732. URL: https://doi.org/10.1214/aoms/1177703732.

[36] Eyke Hüllermeier and Willem Waegeman. “Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods”. In: *Machine Learning* 110 (Mar. 2021). DOI: 10.1007/s10994-021-05946-3.

[37] IEEE. *Standard for Assumptions in Safety-Related Models for Automated Driving Systems*. 2022. URL: https://standards.ieee.org/ieee/2846/10831/ (visited on 06/18/2024).

[38] Kristofer D. Kusano et al. *Collision Avoidance Testing of the Waymo Automated Driving System*. 2022. arXiv: 2212.08148 [cs.RO]. URL: https://arxiv.org/abs/2212.08148.

[39] S. M. LaValle. “Better unicycle models”. In: *Planning Algorithms*. Cambridge University Press, 2006, pp. 743–743.

[40] Rikard Laxhammar. “Anomaly detection in trajectory data for surveillance applications”. PhD thesis. Örebro universitet, 2011. URL: https://urn.kb.se/resolve?urn=urn:nbn:se:oru:diva-17235.

[41] Rikard Laxhammar. “Conformal anomaly detection”. In: *Skövde, Sweden: University of Skövde* 2 (2014).

[42] Jordan Lekeufack et al. “Conformal Decision Theory: Safe Autonomous Decisions Without Distributions”. In: *arXiv* (2024).

[43] Lars Lindemann et al. “Safe Planning in Dynamic Environments Using Conformal Prediction”. In: *IEEE Robotics and Automation Letters* 8.8 (2023), pp. 5116–5123. DOI: 10.1109/LRA.2023.3292071.

[44] Brandon Luders et al. “Probabilistically Safe Avoidance of Dynamic Obstacles with Uncertain Motion Patterns”. In: (July 2011).

[45] R. Luo, S. Zhao, J. Kuck, et al. “Sample-efficient safety assurances using conformal prediction”. In: *Algorithmic Foundations of Robotics XV: Proceedings of the Fifteenth Workshop on the Algorithmic Foundations of Robotics*. Springer, 2022, pp. 149–169.

[46] Lyft. *Self-driving Safety Report*. 2020. URL: http://autonomous.lyft.com/wp-content/uploads/2020/06/Safety_Report_2020.pdf (visited on 06/18/2024).

[47] Christoforos Mavrogiannis et al. “Core Challenges of Social Robot Navigation: A Survey”. In: *J. Hum.-Robot Interact.* 12.3 (Apr. 2023). DOI: 10.1145/3583741. URL: https://doi.org/10.1145/3583741.

[48] Alnur Ali Maxime Cauchois Suyash Gupta and John C. Duchi. “Robust Validation: Confident Predictions Even When Distributions Shift”. In: *Journal of the American Statistical Association* 0.0 (2024), pp. 1–66. DOI: 10.1080/01621459.2023.2298037. eprint: https://doi.org/10.1080/01621459.2023.2298037. URL: https://doi.org/10.1080/01621459.2023.2298037.

[49] D.Q. Mayne et al. “Constrained model predictive control: Stability and optimality”. In: *Automatica* 36.6 (2000), pp. 789–814. ISSN: 0005-1098. DOI: https://doi.org/10.1016/S0005-1098(99)00214-9. URL: https://www.sciencedirect.com/science/article/pii/S0005109899002149.

[50] Rhiannon Michelmore, Marta Kwiatkowska, and Yarin Gal. “Evaluating uncertainty quantification in end-to-end autonomous driving control”. In: *arXiv preprint arXiv:1811.06817* (2018).

[51] I.M. Mitchell, A.M. Bayen, and C.J. Tomlin. “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games”. In: *IEEE Transactions on Automatic Control* 50.7 (2005), pp. 947–957. DOI: 10.1109/TAC.2005.851439.

[52] General Motors. *Self-driving Safety Report*. 2018. URL: https://www.gm.com/content/dam/company/docs/us/en/gmcom/gmsafetyreport.pdf (visited on 06/18/2024).

[53] Anish Muthali et al. “Multi-Agent Reachability Calibration with Conformal Prediction”. In: *2023 62nd IEEE Conference on Decision and Control (CDC)*. 2023, pp. 6596–6603. DOI: 10.1109/CDC49753.2023.10383723.

[54] National Highway Traffic Safety Administration. *National Motor Vehicle Crash Causation Survey*. Report to Congress. NHTSA. Springfield: US Department of Transportation, 2008.

[55] D. Nister et al. *An introduction to the safety force field*. https://www.nvidia.com/content/dam/en-zz/Solutions/self-driving-cars/safety-force-field/the-safety-force-field.pdf. 2019.

[56] Marco Pavone. *Building Trust in AI for Autonomous Vehicles*. English. Online video. Presented at GTC Digital Spring. NVIDIA, Mar. 2023. URL: https://www.nvidia.com/en-us/on-demand/session/gtcspring23-s51934/.

[57] S. Pellegrini et al. “You’ll never walk alone: Modeling social behavior for multi-target tracking”. In: *2009 IEEE 12th International Conference on Computer Vision*. 2009, pp. 261–268. DOI: 10.1109/ICCV.2009.5459260.

[58] Dorsa Sadigh et al. “Planning for Autonomous Cars that Leverage Effects on Human Actions”. In: *Robotics: Science and Systems*. 2016. URL: https://api.semanticscholar.org/CorpusID:7087988.

[59] Dorsa Sadigh et al. “Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state”. In: *Autonomous Robots* 42 (Oct. 2018). DOI: 10.1007/s10514-018-9746-1.

[60] Tim Salzmann et al. “Trajectron++: Dynamically-Feasible Trajectory Forecasting with Heterogeneous Data”. In: *Computer Vision – ECCV 2020*. Ed. by Andrea Vedaldi et al. Cham: Springer International Publishing, 2020, pp. 683–700. ISBN: 978-3-030-58523-5.

[61] Wilko Schwarting, Javier Alonso-Mora, and Daniela Rus. “Planning and Decision-Making for Autonomous Vehicles”. In: *Annual Review of Control, Robotics, and Autonomous Systems* 1.Volume 1, 2018 (2018), pp. 187–210. ISSN: 2573-5144. DOI: https://doi.org/10.1146/annurev-control-060117-105157. URL: https://www.annualreviews.org/content/journals/10.1146/annurev-control-060117-105157.

[62] Glenn Shafer and Vladimir Vovk. “A Tutorial on Conformal Prediction”. In: *J. Mach. Learn. Res.* 9 (June 2008), pp. 371–421. ISSN: 1532-4435.

[63] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. *On a Formal Model of Safe and Scalable Self-driving Cars*. 2018. arXiv: 1708.06374 [id='cs.RO'].

[64] Apoorva Sharma, Navid Azizan, and Marco Pavone. “Sketching curvature for efficient out-of-distribution detection for deep neural networks”. In: *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*. Ed. by Cassio de Campos and Marloes H. Maathuis. Vol. 161. Proceedings of Machine Learning Research. PMLR, 27–30 Jul 2021, pp. 1958–1967. URL: https://proceedings.mlr.press/v161/sharma21a.html.

[65] Charis Stamouli, Lars Lindemann, and George Pappas. “Recursively feasible shrinking-horizon MPC in dynamic environments with conformal prediction guarantees”. In: *Proceedings of the 6th Annual Learning for Dynamics amp; Control Conference*. Ed. by Alessandro Abate et al. Vol. 242. Proceedings of Machine Learning Research. PMLR, 15–17 Jul 2024, pp. 1330–1342. URL: https://proceedings.mlr.press/v242/stamouli24a.html.

[66] Benjamin Stoler et al. *SafeShift: Safety-Informed Distribution Shifts for Robust Trajectory Prediction in Autonomous Driving*. 2024. arXiv: 2309.08889 [cs.RO].

[67] Kegan J. Strawn, Nora Ayanian, and Lars Lindemann. “Conformal Predictive Safety Filter for RL Controllers in Dynamic Environments”. In: *IEEE Robotics and Automation Letters* 8.11 (2023), pp. 7833–7840. DOI: 10.1109/LRA.2023.3322644.

[68] Jiankai Sun et al. “Conformal prediction for uncertainty-aware planning with diffusion dynamics model”. In: *Proceedings of the 37th International Conference on Neural Information Processing Systems*. NIPS ’23. New Orleans, LA, USA: Curran Associates Inc., 2024.

[69] Sophia Huiwen Sun and Rose Yu. “Copula Conformal prediction for multi-step time series prediction”. In: *The Twelfth International Conference on Learning Representations*. 2023.

[70] Tesla. *Tesla AI Day 2022*. https://youtu.be/ODSJsviD_SU. Accessed: [Insert date of access here]. 2022.

[71] Tesla. *Tesla Vehicle Safety Report*. 2024. URL: https://www.tesla.com/VehicleSafetyReport (visited on 06/18/2024).

[72] Ryan J Tibshirani et al. “Conformal Prediction Under Covariate Shift”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Wallach et al. Vol. 32. Curran Associates, Inc., 2019.

[73] Ekaterina Tolstaya et al. “Identifying Driver Interactions via Conditional Behavior Prediction”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. 2021, pp. 3473–3479. DOI: 10.1109/ICRA48506.2021.9561967.

[74] Sander Tonkens et al. “Scalable Safe Long-Horizon Planning in Dynamic Environments Leveraging Conformal Prediction and Temporal Correlations”. In: *ICRA*. 2023.

[75] Pete Trautman et al. “Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation”. In: *The International Journal of Robotics Research* 34.3 (2015), pp. 335–356. DOI: 10.1177/0278364914557874. eprint: https://doi.org/10.1177/0278364914557874. URL: https://doi.org/10.1177/0278364914557874.

[76] Peter Trautman and Andreas Krause. “Unfreezing the robot: Navigation in dense, interacting crowds”. In: *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2010, pp. 797–803. DOI: 10.1109/IROS.2010.5654369.

[77] Martin Treiber, Ansgar Hennecke, and Dirk Helbing. “Congested traffic states in empirical observations and microscopic simulations”. In: *Physical Review E* 62.2 (Aug. 2000), pp. 1805–1824. ISSN: 1095-3787. DOI: 10.1103/physreve.62.1805. URL: http://dx.doi.org/10.1103/PhysRevE.62.1805.

[78] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Berlin, Heidelberg: Springer-Verlag, 2005. ISBN: 0387001522.

[79] Vladimir Vovk et al. “Criteria of Efficiency for Conformal Prediction”. In: *Conformal and Probabilistic Prediction with Applications*. Ed. by Alexander Gammerman et al. Cham: Springer International Publishing, 2016, pp. 23–39. ISBN: 978-3-319-33395-3.

[80] Zia Wadud, Don MacKenzie, and Paul Leiby. “Help or hindrance? The travel, energy and carbon impacts of highly automated vehicles”. In: *Transportation Research Part A: Policy and Practice* 86 (2016), pp. 1–18. ISSN: 0965-8564. DOI: https://doi.org/10.1016/j.tra.2015.12.001. URL: https://www.sciencedirect.com/science/article/pii/S0965856415002694.

[81] Qing Wang, Sanjeev R. Kulkarni, and Sergio Verdu. “A Nearest-Neighbor Approach to Estimating Divergence between Continuous Random Vectors”. In: *2006 IEEE International Symposium on Information Theory*. 2006, pp. 242–246. DOI: 10.1109/ISIT.2006.261842.

[82] Qing Wang, Sanjeev R. Kulkarni, and Sergio Verdu. “Divergence Estimation for Multidimensional Densities Via k-Nearest-Neighbor Distances”. In: *IEEE Transactions on Information Theory* 55.5 (2009), pp. 2392–2405. DOI: 10.1109/TIT.2009.2016060.

[83] Wenshuo Wang, Chang Liu, and Ding Zhao. “How Much Data Are Enough? A Statistical Approach With Case Study on Longitudinal Driving Behavior”. In: *IEEE Transactions on Intelligent Vehicles* 2.2 (2017), pp. 85–98. DOI: 10.1109/TIV.2017.2720459.

[84] Waymo. *Sense, solve, and go: The magic of the Waymo driver*. https://youtu.be/hA_-MkU0Nfw.

[85] Waymo. *Waymo Safety Report*. 2021. URL: https://storage.googleapis.com/waymo-uploads/files/documents/safety/2021-12-waymo-safety-report.pdf (visited on 06/18/2024).

[86] Michael T. Wolf and Joel W. Burdick. “Artificial potential functions for highway driving with collision avoidance”. In: *2008 IEEE International Conference on Robotics and Automation*. 2008, pp. 3731–3736. DOI: 10.1109/ROBOT.2008.4543783.

[87] Danfei Xu et al. *BITS: Bi-level Imitation for Traffic Simulation*. 2022. arXiv: 2208.12403 [cs.RO].

[88] Jianpeng Yao et al. *SoNIC: Safe Social Navigation with Adaptive Conformal Inference and Constrained Reinforcement Learning*. 2024. arXiv: 2407.17460 [cs.RO]. URL: https://arxiv.org/abs/2407.17460.

[89] Xinyi Yu et al. *Signal Temporal Logic Control Synthesis among Uncontrollable Dynamic Agents with Conformal Prediction*. Dec. 2023.

[90] Y. Yuan et al. “AgentFormer: Agent-Aware Transformers for Socio-Temporal Multi-Agent Forecasting”. In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. Los Alamitos, CA, USA: IEEE Computer Society, Oct. 2021, pp. 9793–9803. DOI: 10.1109/ICCV48922.2021.00967. URL: https://doi.ieeecomputersociety.org/10.1109/ICCV48922.2021.00967.

[91] Yiqi Zhao et al. “Robust Conformal Prediction for STL Runtime Verification under Distribution Shift”. In: *2024 ACM/IEEE 15th International Conference on Cyber-Physical Systems (ICCPS)*. 2024, pp. 169–179. DOI: 10.1109/ICCPS61052.2024.00022.

[92] Kemin Zhou and John Comstock Doyle. *Essentials of robust control*. Vol. 104. Prentice hall Upper Saddle River, NJ, 1998.

[93] Hai Zhu and Javier Alonso-Mora. “Chance-Constrained Collision Avoidance for MAVs in Dynamic Environments”. In: *IEEE Robotics and Automation Letters* 4.2 (2019), pp. 776–783. DOI: 10.1109/LRA.2019.2893494.

[94] Zoox. *Safety Report Volume 2.0*. 2021. URL: http://autonomous.lyft.com/wp-content/uploads/2020/06/Safety_Report_2020.pdf (visited on 06/18/2024).