

Multi-Agent Reachability Calibration with Conformal Prediction

Anish Muthali^{1,*}, Haotian Shen^{1,*}, Sampada Deglurkar¹, Michael H. Lim¹,
 Rebecca Roelofs², Aleksandra Faust², Claire Tomlin¹

Abstract—We investigate methods to provide safety assurances for autonomous agents that incorporate learning-based predictions of other, uncontrolled agents’ behavior into their own trajectory planning. Given a learning-based forecasting model that predicts agents’ trajectories, we introduce a method for providing probabilistic assurances on the model’s prediction error with calibrated confidence intervals. Through quantile regression, conformal prediction, and reachability analysis, our method generates probabilistically safe and dynamically feasible prediction sets. We showcase their utility in certifying the safety of planning algorithms, both in simulations using actual autonomous driving data and in an experiment with Boeing vehicles.

I. INTRODUCTION

In safety-critical situations in which an autonomous agent interacts with humans, it is often necessary to predict human behavior in order for the agent to appropriately react. For example, in self-driving tasks, the autonomous car is responsible for assuring the safety of itself and the other vehicles it encounters. State-of-the-art systems try to achieve this through behavior prediction and motion forecasting models, oftentimes black-box neural networks that do not provide interpretation of their inner workings [2]–[4]. However, these models lack rigorous safety assurances, especially in the presence of data distribution shifts. While lower dimensional models can provide safety assurances, these methods rely on parametric assumptions on, for example, a human’s rationality [5]. These assumptions can be inadequate for complex, multi-modal, and noisy decision-making scenarios. In this work, we provide safety assurances for state-of-the-art black-box trajectory forecasting methods by quantifying these models’ uncertainty.

However, uncertainty quantification methods alone are insufficient for providing safety assurances. This is because they do not provide guarantees on model behavior and may be unreliable or uninterpretable [6]–[8]. Additionally, using prediction model uncertainty for assuredly safe decision-making in downstream planning and control is a challenging problem.

*Denotes equal contribution.

¹Hybrid Systems Lab, EECS, UC Berkeley, USA.

²Google Research, USA.

Correspondence to anishmuthali@berkeley.edu.

Appendices are available in [1].

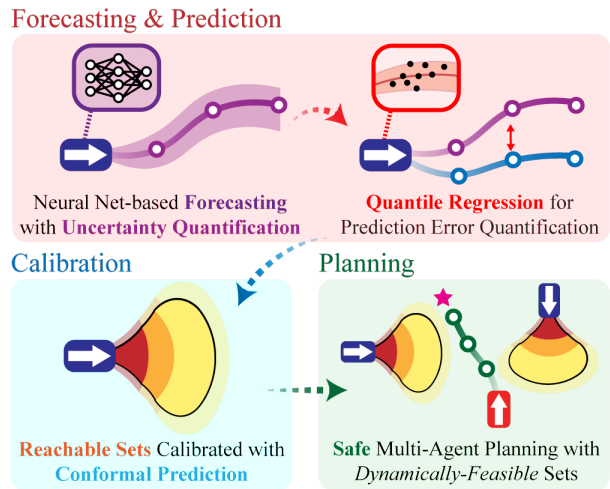


Fig. 1. Our method generates dynamically-feasible, probabilistic confidence sets that are derived from a conformal-calibrated quantile regression model.

In our work, we first utilize conformal prediction to *calibrate* measures of uncertainty [9]. Conformal prediction is a statistical tool that uses a heuristic notion of risk to non-parametrically estimate quantiles of risk given a sequence of past observations [10]. Existing methods that leverage conformal prediction in the context of trajectory forecasting do not explicitly use interpretable metrics of prediction uncertainty [11]–[15]. We propose a method that provides rigorous confidence intervals on model error, a form of probabilistic assurance, given any interpretable heuristics on a trajectory forecasting model’s prediction uncertainty. Our approach also allows us to examine the efficacy of various uncertainty quantification heuristics when attempting to predict model error. In addition, we extend our analysis to multi-agent environments to closely reflect real-world assurance cases. We analyze the problem of a single autonomous agent, commonly described as an “ego agent”, interacting with other, uncontrolled agents.

By producing estimates of model error, we are able to couple statistical assurances with dynamical assurances to allow for safe downstream navigation. In particular, we turn to Hamilton-Jacobi (HJ) reachability analysis [16], which provides guarantees on dynamical systems by means of reachable sets and associated controllers. In HJ reachability, a Hamilton-Jacobi partial differential

equation is solved to obtain an optimal value function and controller. The sub-zero level sets of this value function are the reachable sets (possible states of an agent at a given time) and tubes (possible states of an agent *up to* and including a given time).

Our method can be outlined as follows: given a trajectory forecasting model with an associated uncertainty heuristic, we design a quantile regression model that correlates uncertainty with prediction error, creating an approximate confidence interval on the model’s prediction. We then apply conformal prediction to calibrate the confidence intervals and provide guarantees on miscoverage rate. We map the calibrated intervals in control action space to sets in state space through reachability analysis, and we demonstrate the utility of these confidence sets in planning tasks. The contributions of this paper include:

- 1) A novel way to interpret trajectory forecasting models’ prediction uncertainty and obtain approximate confidence intervals (Section III-B);
- 2) A technique to calibrate the aforementioned intervals using conformal prediction (Section III-C);
- 3) Dynamically-feasible, probabilistic reachable sets using calibrated intervals (Section IV-A);
- 4) A planning framework that leverages assurances developed in the previous steps (Section IV-C).

This paper is organized as follows: Section II discusses related works in conformal prediction and assurances in trajectory forecasting models, Section III and Section IV describe the contributions outlined above, and Section V showcases the safety and performance of our methods compared to baseline methods.

II. RELATED WORKS

A. Conformal Prediction

Conformal prediction [9], [10] is a class of uncertainty quantification methods for constructing prediction sets that satisfy a significance level (false negative rate) requirement. Traditionally, conformal prediction creates empirical histograms of measures of risk, called non-conformity scores, and uses these to estimate prediction intervals. Classical techniques include split conformal prediction, which creates empirical histograms from hold-out sets, and full conformal prediction, which creates empirical histograms using all available data [9]. Inductive conformal prediction, a variant of split conformal prediction, uses a non-conformity score that measures distance between train and test data [17]. These methods require that the data are identically distributed and exchangeable (any permutation of data points are identically distributed). Methods such as [18] relax the requirement for identically distributed data, and [19], [20] relax the exchangeability requirement. Adaptive Conformal Inference [21] and Rolling Risk Control (RollingRC) [22] have been proposed to relax

all assumptions by further calibrating the significance level to match a desired error rate. We adapt RollingRC to provide safety assurances in any multi-agent scenario.

B. Probabilistic Reachability Frameworks

In this work, we introduce a method to generate probabilistic reachable sets to account for agents’ dynamics. Previous work in this space typically involves randomly generating inputs and observing corresponding outputs of a dynamics model, with some associated guarantees in the sampling process [23]. Other methods, much like ours, leverage neural network uncertainty [24]. Specifically, the method of Nakamura and Bansal [24] uses Gaussian mixture models (GMMs) output by a trajectory forecasting model to generate parametric confidence intervals on control actions, which are then used as control bounds in reachability calculations. In our work, we attempt to relax assumptions that the control actions follow any parametric distribution by applying non-parametric inference techniques.

C. Safety Assurances in Trajectory Prediction

Various methods for incorporating uncertainty quantification have been examined for the purposes of providing safety assurances in trajectory prediction problems. Some methods provide probabilistic assurances by inferring parameters of a distribution on an agent’s control actions [5], [25]. Methods such as [11] and [12] estimate confidence intervals with conformal prediction, implicitly leveraging prediction uncertainty through the non-conformity measure. Specifically, the method of Luo *et al.* [11] uses split conformal prediction to create a warning system, alerting drivers of “dangerous” situations. These warnings can be transformed into confidence sets, as shown in [12], which additionally eliminates exchangeability assumptions and incorporates trajectory optimization, much like our approach. Other methods emphasize the design process of the trajectory prediction neural networks, for instance by opting to use ReLU networks [14] or by opting to incorporate conformal prediction in the neural network’s loss function [13]. In contrast to our approach, none of these methods consider the dynamic feasibility of confidence sets, and some methods that investigate conformal prediction, such as [14] and [11], assume exchangeability. Our method relaxes these assumptions while providing interpretability in the uncertainty quantification process and dynamic feasibility in confidence sets.

III. ASSURANCES FROM UNCERTAINTY

Our approach can be summarized in four primary steps: trajectory forecasting with uncertainty quantification (Section III-A), leveraging uncertainty to obtain approximate prediction intervals (Section III-B), calibrating approximate prediction intervals (Section III-C),

and obtaining dynamically feasible prediction sets (Section IV-A). We summarize our algorithm in Section IV-B, and we apply our approach to ego agent planning tasks in Section IV-C.

Running Example: *To motivate and illustrate our method, we introduce a simple running example with two vehicles at an intersection, one of them designated as the “ego” vehicle. In Figure 2, the autonomous ego vehicle, shown in red, aims to safely navigate to the pink-colored star while avoiding the blue vehicle.*

A. Trajectory Forecasting Model

We start by assuming access to a known dynamics model for each agent and a trajectory forecasting model capable of predicting an agent’s control input. This trajectory forecasting model may maintain the capability to predict control actions for multiple agents at once, while considering interactions between agents. We denote the model as $f_T(\cdot) : \mathcal{X} \rightarrow \mathcal{U}$, where \mathcal{X} is some arbitrary input space and \mathcal{U} is a space over control actions. The network predicts $\mathbf{u}_{t:t+h} \in \mathcal{U}$, which is a collection of control action vectors indexed by timesteps t through $t+h$, for each of N total agents. Here, h is a fixed prediction horizon. We also assume the existence of an uncertainty measure on the network’s outputs, denoted as $\sigma_T(\cdot) : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}^d$, with d as the dimension of uncertainty representation. For example, a variance prediction or a variance estimate from inference-time dropout [6] is a valid uncertainty measure. Additionally, some neural network architectures, such as Trajectron++ [2], provide alternative uncertainty measures. This network architecture predicts a GMM over possible control actions, leading to understandings of prediction uncertainty such as the variance of the GMM’s modes.

Running Example: *Given some sequence of the other vehicle’s position history, Trajectron++, our trajectory forecasting model of choice, predicts a GMM in action space, and then integrates the actions to obtain states. We assume that vehicles follow the extended Dubins’ car dynamics model. The state of this system is $\mathbf{x} = [x \ y \ v \ \theta]^\top$, and the dynamics are given by $\dot{\mathbf{x}} = [v \cos(\theta) \ v \sin(\theta) \ u_1 \ u_2]^\top$. The variance of the highest-probability Gaussian component, among other features of the GMM, are incorporated into the design of the uncertainty measure.*

B. Estimating Model Error from Uncertainty

To obtain confidence intervals on a black-box model’s outputs, we estimate the neural network’s confidence in an online manner, correlating its prediction uncertainty with prediction error. Quantile regression models enable us to map heuristic notions of uncertainty to an approximate confidence interval [26] along each action dimension. We choose a linear model since its

simple parametrization allows for fast online updates and interpretability in how it perceives uncertainty. We demonstrate an example of interpretability in Section V-D.1. Intuitively, our quantile regression models are approximately “calibrating” the network’s uncertainty to obtain an estimate of its error.

As we observe new datapoints online, we collect $\mathbf{u}_{t-h:t}$, the last h ground truth control actions prior to timestep t . In practice, we estimate control actions by observing the state history of an agent, and then numerically computing derivatives to estimate actions from an assumed dynamics model. We contrast $\mathbf{u}_{t-h:t}$ with the network’s previous prediction h timesteps ago, i.e., $\hat{\mathbf{u}}_{t-h:t}$, and define $\mathbf{e}_{t-h:t} := \mathbf{u}_{t-h:t} - \hat{\mathbf{u}}_{t-h:t}$ as the prediction error.

Now, suppose we require a $1 - \alpha$ approximate confidence interval on the ground truth control action. We can construct two quantile regression models $\hat{q}_{\frac{\alpha}{2}} : \mathbb{R}^d \rightarrow \mathcal{U}$ and $\hat{q}_{1-\frac{\alpha}{2}} : \mathbb{R}^d \rightarrow \mathcal{U}$ where \hat{q}_ε estimates the ε -quantile on the network’s prediction error from timesteps t to $t+h$ for each agent, denoted $\hat{\mathbf{e}}_\varepsilon$. For notational convenience, let us denote $\mathbb{P}_t(A)$ as the probability of event A conditioned on information until time t . We obtain an approximate $1 - \alpha$ confidence interval as follows:

$$\mathbb{P}_t\left(\hat{\mathbf{e}}_{\frac{\alpha}{2}} \leq \mathbf{e}_{t:t+h} \leq \hat{\mathbf{e}}_{1-\frac{\alpha}{2}}\right) \quad (1)$$

$$= \mathbb{P}_t\left(\hat{\mathbf{e}}_{\frac{\alpha}{2}} \leq \mathbf{u}_{t:t+h} - \hat{\mathbf{u}}_{t:t+h} \leq \hat{\mathbf{e}}_{1-\frac{\alpha}{2}}\right) \quad (2)$$

$$= \mathbb{P}_t\left(\hat{\mathbf{u}}_{t:t+h} + \hat{\mathbf{e}}_{\frac{\alpha}{2}} \leq \mathbf{u}_{t:t+h} \leq \hat{\mathbf{u}}_{t:t+h} + \hat{\mathbf{e}}_{1-\frac{\alpha}{2}}\right) \quad (3)$$

$$\approx 1 - \alpha. \quad (4)$$

Thus, our approximate $1 - \alpha$ confidence interval on $\mathbf{u}_{t:t+h}$ is $\hat{\mathcal{I}}_{t:t+h} = [\hat{\mathbf{u}}_{t:t+h} + \hat{\mathbf{e}}_{\frac{\alpha}{2}}, \hat{\mathbf{u}}_{t:t+h} + \hat{\mathbf{e}}_{1-\frac{\alpha}{2}}]$.

Traditionally, quantile regression models are trained using computationally expensive linear programs [26], so instead, we opt for a faster, online gradient descent approach. We define our loss function for the quantile regression model \hat{q}_ε to be $\mathcal{L}(y, \hat{y}) = (y - \hat{y})\varepsilon \mathbf{1}\{y \geq \hat{y}\} + (\hat{y} - y)(1 - \varepsilon) \mathbf{1}\{y < \hat{y}\}$ (the “pinball loss”) [27]. We set y to be the true model error, \mathbf{e} , and $\hat{y} = \beta^\top \boldsymbol{\sigma}$, where β represents the weights of the regression model, and $\boldsymbol{\sigma}$ is the uncertainty measure from the trajectory forecasting model. We update the weights according to $\beta \leftarrow \beta - \zeta \nabla_{\beta} \mathcal{L}(\mathbf{e}_{t-h:t}, \beta)$, with learning rate ζ .

C. Calibrating Approximate Confidence Intervals

Given that the confidence intervals we obtained in the previous section are merely approximate, we aim to calibrate these intervals. To this end, we apply the RollingRC algorithm [22], which perfectly adapts to the online requirements of our method. We are motivated to use the RollingRC algorithm compared to other conformal prediction methods due to a desire




Fig. 2. Visualization of the running example. The autonomous ego vehicle is shown in red, and the human driver is shown in blue. The ego vehicle aims to navigate to the pink star while avoiding a collision with the human-driven vehicle. Confidence sets for the next three prediction steps are shown. In fig. 2b, the redder regions represent confidence sets for earlier prediction timesteps, and the translucent regions represent conformal prediction’s calibration effect.

to remove the data exchangeability assumption, since we allow for sequentially-dependent data and potential distribution shifts. In addition, we would like to train and calibrate the quantile regression models in a sample-efficient manner. RollingRC guarantees that the error rate deviates from α as $\mathcal{O}(1/T)$, where T is the total number of datapoints provided to the algorithm.

Following the notation from the RollingRC algorithm, we define $\theta_t \in \mathbb{R}$ as our conformal parameter, and $\varphi(\cdot) : \mathbb{R} \rightarrow \mathcal{U}$ as the algorithm’s “stretching function”. Now, we claim that

$$\mathbb{P}_t \left(\widehat{e}_{\frac{\alpha}{2}} - \varphi(\theta) \leq \mathbf{e}_{t:t+h} \leq \widehat{e}_{1-\frac{\alpha}{2}} + \varphi(\theta) \right) \geq 1 - \alpha - \mathcal{O}(1/t). \quad (5)$$

Following similar steps as before, we obtain our newly calibrated confidence interval on $\mathbf{u}_{t:t+h}$ as $\mathcal{I}_{t:t+h} = [\widehat{\mathbf{u}}_{t:t+h} + \widehat{e}_{\frac{\alpha}{2}} - \varphi(\theta), \widehat{\mathbf{u}}_{t:t+h} + \widehat{e}_{1-\frac{\alpha}{2}} + \varphi(\theta)]$.

IV. PROBABILISTIC REACHABILITY AND PLANNING

A. Probabilistic Reachability among Multiple Agents

In the previous sections, we have designed a method to provide confidence intervals on agents’ control actions. However, for some downstream tasks, such as safe planning, confidence sets in spatial dimensions are more desirable. Hence, we use HJ reachability to obtain spatial sets, in the form of forward reachable tubes, on each agent’s location given its dynamics and the probabilistic bound on control [28]. This procedure asserts that an agent’s location will be contained in the produced reachable tube with probability $1 - \alpha$.

Suppose we wish to upper bound the probability that the ego vehicle collides with any agent. Let $\mathbf{x}_t^{(i)}$ be the location of non-ego agent i at timestep t and $\mathcal{S}[t]^{(i)}$ be the corresponding agent’s forward reachable tube, as computed by our algorithm. We define miscoverage rate as the proportion of instances in which the ground truth position of any agent i at time t is outside $\mathcal{S}[t]^{(i)}$.

We aim to obtain an upper bound on miscoverage rate, such that the ego agent can navigate in regions outside of $\mathcal{S}[t]^{(i)}$ for all $i \in \{1, \dots, N\}$ and guarantee that the probability of collision is at most γ , a pre-specified parameter. Consequently, we set the confidence interval significance level α according to our desired total miscoverage rate γ and number of agents N .

THEOREM IV.1 (SIGNIFICANCE LEVEL CORRECTION) Suppose that we wish to have a total miscoverage rate of γ , where total miscoverage rate is an upper bound on the probability that *any* human agent is miscovered:

$$\mathbb{P}_t \left(\bigcup_{i=1}^N \{ \mathbf{x}_t^{(i)} \notin \mathcal{S}[t]^{(i)} \} \right) \leq \gamma. \quad (6)$$

We claim that the following α achieves an (asymptotic) total miscoverage rate of γ for N human agents:

$$\alpha = 1 - (1 - \gamma)^{\frac{1}{N}}. \quad (7)$$

The proof of Theorem IV.1 is available in Appendix A, which uses the fact that the N agents act independently conditioned on past information [1]. Since α must be a fixed quantity in our algorithm, we must also fix N . Hence, we fix our algorithm to only consider the N agents closest to the ego vehicle.

Running Example: Suppose we want a 95% probability safety assurance. Since there is only one other vehicle, we get $\alpha = 0.05$ from Theorem IV.1. Given the previous predictions of the blue agent’s trajectory, we generate uncalibrated, time-indexed intervals on ranges of possible control actions, denoted $\widehat{\mathcal{I}}_t, \widehat{\mathcal{I}}_{t+\Delta t}, \widehat{\mathcal{I}}_{t+2\Delta t}$. We calibrate these using conformal prediction to obtain $\mathcal{I}_t, \mathcal{I}_{t+\Delta t}, \mathcal{I}_{t+2\Delta t}$. As we explain in the next subsection, HJ reachability allows us to take any sequence of intervals on control actions and generate a time-indexed set of states. In Figure 2b, we distinguish the effects of quantile regression and RollingRC’s calibration.

Algorithm 1 Conformal Reachability Calibration.

```
1: procedure GENERATESSETS( $\theta, \hat{\mathbf{u}}_{t:t+h}, \sigma$ )
2:    $\hat{\mathbf{e}}_{\frac{\alpha}{2}} \leftarrow \hat{q}_{\frac{\alpha}{2}}(\sigma) \triangleright$  Obtain lower  $\frac{\alpha}{2}$  quantile
3:    $\hat{\mathbf{e}}_{1-\frac{\alpha}{2}} \leftarrow \hat{q}_{1-\frac{\alpha}{2}}(\sigma) \triangleright$  Obtain upper  $\frac{\alpha}{2}$  quantile
4:    $\mathcal{I}_{t:t+h} \leftarrow \left[ \hat{\mathbf{u}}_{t:t+h} + \hat{\mathbf{e}}_{\frac{\alpha}{2}} - \varphi(\theta), \right.$ 
      $\left. \hat{\mathbf{u}}_{t:t+h} + \hat{\mathbf{e}}_{1-\frac{\alpha}{2}} + \varphi(\theta) \right]$ 
5:    $\mathcal{S} \leftarrow \emptyset$ 
6:   for  $t' \in \{t, t + \Delta t, \dots, t + h\}$  do
7:      $\mathcal{S}[t'] \leftarrow$  HJREACHABILITY( $\mathcal{I}_{t'}$ )
8:   return  $\mathcal{S}, \mathcal{I}_{t:t+h}$ 
9: procedure UPDATE( $\theta, \mathbf{u}_{t-h:t}, \mathcal{I}_{t-h:t}$ )
10:  for  $t' \in \{t, t + \Delta t, \dots, t + h\}$  do
11:     $\theta\{t'\} \leftarrow \theta\{t'\} + \xi(\mathbf{1}\{\mathbf{u}_{t'-h} \notin \mathcal{I}_{t'-h}\} - \alpha)$ 
12:  GRADIENTDESCENT( $\hat{q}_{\frac{\alpha}{2}}, \mathbf{u}_{t-h:t}$ )
13:  GRADIENTDESCENT( $\hat{q}_{1-\frac{\alpha}{2}}, \mathbf{u}_{t-h:t}$ )
14:  return  $\hat{q}_{\frac{\alpha}{2}}, \hat{q}_{1-\frac{\alpha}{2}}, \theta$ 
15: procedure MAIN( $\gamma, N$ )
16:   $\alpha \leftarrow 1 - (1 - \gamma)^{\frac{1}{N}}$ 
17:   $\theta\{t, t + \Delta t, \dots, t + h\} \leftarrow 0$ 
18:   $\hat{q}_{\frac{\alpha}{2}}, \hat{q}_{1-\frac{\alpha}{2}} \leftarrow$  INITIALIZERANDOMWEIGHTS()
19:   $\mathbb{I} \leftarrow \{\}$ 
20:   $t \leftarrow 0$ 
21:  while true do
22:     $\hat{\mathbf{u}}_{t:t+h}, \sigma \leftarrow f_T(\cdot), \sigma_T(\cdot) \triangleright$  Get trajectory predictions
    and uncertainty from model
23:     $\mathcal{S}, \mathcal{I}_{t:t+h} \leftarrow$  GENERATESSETS( $\theta, \hat{\mathbf{u}}_{t:t+h}, \sigma$ )
24:     $\mathbb{I} \leftarrow \mathbb{I} \cup \mathcal{I}_{t:t+h}$ 
25:    if  $t \geq h$  then
26:       $\mathbf{u}_{t-h:t} \leftarrow$  OBSERVEHISTORY()
27:       $\mathcal{I}_{t-h:t} \leftarrow \mathbb{I}[t-h:t]$ 
28:       $\hat{q}_{\frac{\alpha}{2}}, \hat{q}_{1-\frac{\alpha}{2}}, \theta \leftarrow$  UPDATE( $\theta, \mathbf{u}_{t-h:t}, \mathcal{I}_{t-h:t}$ )
29:       $t \leftarrow t + \Delta t$ 
```

B. Full Algorithm

In Algorithm 1, we demonstrate the final algorithm to generate probabilistic reachable sets. The HJREACHABILITY function generates reachable sets given a probabilistic range of control actions, $\mathcal{I}_{t:t+\Delta t}$. Since the range can differ over time (e.g., $\mathcal{I}_t \neq \mathcal{I}_{t+\Delta t}$ necessarily), we iteratively compute time-indexed forward reachable tubes by computing the forward reachable tube over $[t, t + \Delta t]$ and using the reachable set at $t + \Delta t$ as the initial condition to compute the reachable tube over $[t + \Delta t, t + 2\Delta t]$. We also utilize a GRADIENTDESCENT function that updates the weights of the quantile regression models as described in Section III-B. In Algorithm 1, ξ is the “learning rate” associated with the RollingRC algorithm.

C. Safe Planning Framework

Given the time-indexed sets $\mathcal{S}[t] \subseteq \mathcal{S}[t + \Delta t] \subseteq \dots \subseteq \mathcal{S}[t + h]$, we desire that the autonomous agent’s location at time t' is outside $\mathcal{S}[t + k\Delta t]$, where $t + (k - 1)\Delta t \leq t' \leq t + k\Delta t$. We can plan by treating each agent’s time-indexed forward reachable tube as a dynamic obstacle that grows with time. The obstacle-aware planning requirement motivates the application of

a forward reach-avoid tube for the ego agent [16], [29]. We use this to derive an optimal control trajectory by selecting the Hamiltonian-maximizing control trajectory to a desired final state within the forward reach-avoid tube. In practice, this trajectory can involve bang-bang control, so one can track it using a tracker with a provable tracking error bound, such as a constrained iterative linear quadratic regulator. The planner’s output is visualized in Appendix D.

Running Example: From the previous section, we obtained $\mathcal{S}[t], \mathcal{S}[t + \Delta t], \mathcal{S}[t + 2\Delta t]$ as a probabilistic occupancy region on the location of the other vehicle. Now, we can use the time-varying avoidance regions to plan a safe path to the goal in Figure 2c. Notice that the planner allows the ego agent to traverse in the yellow-colored region: it is aware that the ego vehicle would not violate the safety assurance as it can leave the yellow region by the time the other agent would enter it.

V. RESULTS

We compare the empirical safety and efficiency of our contribution to two baselines, Online Update of Safety Assurances Using Confidence-Based Predictions by Nakamura and Bansal [24] and Sample-Efficient Safety Assurances using Conformal Prediction by Luo *et al.* [11]. For both baselines, we perform the significance level correction described in Section IV-A.

For all benchmarking purposes, we use Trajectron++ trained on the relevant datasets. We follow the same architecture and hyperparameters as [2] by using 4 seconds (8 steps) of history to predict 3 seconds (6 steps) into the future. This is consistent with the other baselines’ approaches. Set sizes are shown in square meters. We use a pre-specified total miscoverage rate of $\gamma = 0.05$, and we generate predictions for the closest $N = 3$ agents, which strikes a balance between the speed of our HJ reachability calculations and the practical safety of the system.

A. nuScenes Dataset Results

We compare the coverage rate and efficiency of our method against the two baselines on nuScenes self-driving data [30]. We calculate average coverage rate and average set sizes individually for each forward prediction step $t, t + \Delta t, \dots, t + h$ on 100 randomly sampled scenes. For each scene, we use the first 13 seconds to calibrate each method and make predictions on the last 5.5 seconds. Table I shows step coverage and set sizes at all prediction steps. Note that an ideal algorithm maintains a coverage rate over $1 - \gamma$ while providing the smallest prediction sets.

B. Waymo Open Motion Dataset Results

To demonstrate the planning safety and efficiency of each method, we also perform experiments on the

TABLE I
COVERAGE RATES AND SET SIZES FOR $1 - \gamma = 0.95$.

Coverage Rates for Prediction Step						
Methods	1st (.5s)	2nd (1s)	3rd (1.5s)	4th (2s)	5th (2.5s)	6th (3s)
nuScenes Dataset						
Nakamura and Bansal	0.926 \pm 0.012	0.854 \pm 0.017	0.816 \pm 0.023	0.842 \pm 0.023	0.868 \pm 0.022	0.902 \pm 0.020
Luo <i>et al.</i>	1.000 \pm 0.000	1.000 \pm 0.000	1.000 \pm 0.000	0.998 \pm 0.002	0.989 \pm 0.006	0.968 \pm 0.012
Our Method	0.964 \pm 0.008	0.962 \pm 0.011	0.968 \pm 0.010	0.975 \pm 0.008	0.981 \pm 0.007	0.985 \pm 0.007
Waymo Dataset						
Nakamura and Bansal	0.981 \pm 0.007	0.954 \pm 0.012	0.938 \pm 0.014	0.937 \pm 0.015	0.952 \pm 0.014	0.955 \pm 0.015
Luo <i>et al.</i>	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	0.998 \pm 0.002	0.985 \pm 0.009	0.955 \pm 0.015
Our Method	0.997 \pm 0.002	0.986 \pm 0.006	0.980 \pm 0.008	0.967 \pm 0.011	0.965 \pm 0.011	0.965 \pm 0.013
Set Sizes for Prediction Step						
nuScenes Dataset						
Nakamura and Bansal	57 \pm 2	320 \pm 12	886 \pm 35	2060 \pm 85	3259 \pm 128	4285 \pm 160
Luo <i>et al.</i>	425 \pm 11	523 \pm 16	683 \pm 34	1097 \pm 109	1426 \pm 140	1814 \pm 186
Our Method	39 \pm 3	157 \pm 13	462 \pm 39	1078 \pm 88	2150 \pm 170	3713 \pm 268
Waymo Dataset						
Nakamura and Bansal	64 \pm 7	311 \pm 32	951 \pm 96	2117 \pm 195	3814 \pm 328	5892 \pm 475
Luo <i>et al.</i>	448 \pm 8	736 \pm 43	1110 \pm 94	1568 \pm 169	2126 \pm 237	2687 \pm 301
Our Method	61 \pm 6	246 \pm 27	655 \pm 71	1361 \pm 143	2422 \pm 240	3885 \pm 365

Waymo Open Motion Dataset [31], coupled with the Nocturne simulator [32]. This allows us to apply control actions to the ego vehicle while all other agents replay their respective sequences of control actions from the dataset. We use the same planning method discussed in Section IV-C for all three methods, since neither of the baselines have associated planners. For each scene, we calibrate using the first 7 seconds and use model-predictive control to plan for the last 3 seconds, where the goal is the final position of the ego vehicle in the ground truth data. We measure three quantities: (1) progress to goal, defined as the ratio of the distance from the final state of the ego vehicle to the goal compared to the distance from the start to the goal, subtracted from 1; (2) collision rate; (3) conservatism of each method compared to the ego vehicle’s ground truth trajectory, defined as the ratio of minimum distance between the ego vehicle to other agents at all times as a result of the planner, compared to that of the ground truth. The formulas and computations of these metrics are described in detail in Appendix B [1]. In Appendix C, we additionally demonstrate the impact of the aforementioned theoretical guarantees by providing safety and efficiency metrics in the absence of conformal prediction [1]. We performed the benchmarks on 200 randomly sampled scenes. Table I depicts coverage rates and set sizes for all prediction timesteps. Table II depicts average collision rate, average progress to goal, and conservatism.

TABLE II
WAYMO PLANNING BENCHMARKS

Method	Progress to Goal	Collision Rate	Conservatism
Nakamura and Bansal	0.494 \pm 0.029	0.0	1.504 \pm 0.068
Luo <i>et al.</i>	0.305 \pm 0.028	0.005	1.626 \pm 0.072
Our Method	0.544 \pm 0.028	0.0	1.507 \pm 0.068

C. Discussion of Results

For the nuScenes dataset, we notice that our method achieves more efficient set sizes for initial prediction steps, while Luo *et al.* achieves more efficient set sizes for later prediction timesteps. Nevertheless, neither of these two methods violates the miscoverage requirement of $\gamma = 0.05$. The method of Nakamura and Bansal violates the miscoverage rate, however, supporting the introduction of uncertainty calibration into the algorithm. Hence, calibrating neural network uncertainty is important, not only to provide the desired coverage rate but also to generate efficient prediction sets.

For the Waymo dataset, we notice a very similar phenomenon with set sizes and coverage rates. In the planning benchmarks, our method has the best progress to goal, likely due to the initial-timestep sets being smaller. This is also reflected in the conservatism scores, with reachability-based methods performing the best. The method of Luo *et al.* also encountered one collision scenario in which the produced set was very large and forced the planner to take a sharp avoid action. Thus, we note the importance of initial-timestep sets being small to allow the reachability-based methods to perform better in the planning benchmarks. This allows the ego vehicle to make some progress, whereas a large initial-




Fig. 3. Case Study of Uncertainty Metrics. We demonstrate a simple example in which the choice of uncertainty measure affects the size of sets, with coverage rate held constant.



Fig. 4. Our algorithm is applied to assure safety in potential runway incursion scenarios. Once the ground vehicle is determined to have crossed a designated safety threshold, the aircraft is cleared to land. timestep set would inhibit any progress regardless of the relative size of later timesteps' sets.

D. Case Studies

1) *Understanding Uncertainty Measures:* In this case study, we demonstrate the usefulness of our interpretable quantile regression model when understanding the efficacy of uncertainty metrics. Consider the scene in Figure 3a. We choose the uncertainty measure based on properties of the GMM, including the distance between peaks and the (co)variance of the highest-weighted mode. The learned regression model indicates a positive correlation between prediction error and variance of the most-likely GMM mode. In Figure 3b and Figure 3c, we can visually discern the benefit of including these features.

Overall, this case study shows the importance of understanding the usefulness of different components of the uncertainty measure. A more useful uncertainty metric can provide more efficient sets, since a more accurate quantile regression model would require less calibration (less “stretching” from conformal prediction). Conformal prediction cannot derive confidence intervals conditional on some input, so quantile regression’s accuracy is crucial for providing efficient sets.

2) *Safety in Aerospace Applications:* In this case study, we apply our algorithm to satisfy a real-world

safety assurance requirement by demonstrating our algorithm on Boeing vehicles. We consider the case of an aircraft attempting to land on a runway while accounting for potential runway incursions from ground vehicles. We use our algorithm to provide assurances on the motion of a ground vehicle on the runway. Given a fixed landing plan for the plane, we adapt the sets from our algorithm to design a warning system similar to Luo *et al.*’s original algorithm. If a prediction set intersects the runway, a warning is issued. A visualization of this application is shown in Figure 4. The ground vehicle’s state history is shown in red, and its uncalibrated prediction set is shown in purple. The “stretching” effect from conformal prediction is shown in orange.

VI. DISCUSSION AND FUTURE WORK

In this paper, we introduced a non-parametric approach to using interpretable uncertainty measures from black-box models for generating calibrated prediction intervals. We demonstrated an efficient reachability-based approach to generating prediction sets, and we showed the goal-oriented efficiency and safety of our algorithm in planning tasks for an ego agent through simulations and real-world experiments.

For future investigations, we would be interested in seeing the effects of longer planning horizons. Although many state of the art models cannot provide reliable predictions for agent behavior 1 minute into the future, for example, we would like to see the efficiency of our method compared to the other methods discussed. We would also like to generalize our method to arbitrary measures of risk, instead of only coverage rate. For example, one might want greater confidence in the behavior of nearby or fast-moving agents, than for agents that are far away or stationary. Thus, in certain practical scenarios, a heuristic measure of risk may be more appropriate than miscoverage rate. Along these lines, we would also like to explore adapting the confidence interval significance α for different agents depending

on their properties with respect to maintaining safety. Finally, we are interested in decreasing conservatism across the pipeline to help ensure that the planning framework can always find a feasible solution to the problem.

VII. ACKNOWLEDGEMENTS

This material is based upon work supported by the DARPA Assured Autonomy Program, the SRC CONIX program, Google-BAIR Commons, the NASA ULI program on Safe Aviation Autonomy, and the National Science Foundation Graduate Research Fellowship Program under Grant Nos. DGE 1752814 and DGE 2146752. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any aforementioned organizations. We also thank James Paunicka, Blake Edwards, Dragos Margineantu, Douglas Stuart, and the entire team at Boeing for all their help and contributions in demonstrating our method on Boeing’s runway incursion dataset. Finally, we thank Marius Wiggert for helping with the implementations in the HJ reachability toolbox.

REFERENCES

- [1] A. Muthali, H. Shen, S. Deglurkar, *et al.*, “Multi-agent reachability calibration with conformal prediction,” *arXiv preprint arXiv:2304.00432*, 2023.
- [2] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, “Trajectory++: Dynamically-feasible trajectory forecasting with heterogeneous data,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, Springer, 2020, pp. 683–700.
- [3] B. Varadarajan, A. Hefny, A. Srivastava, *et al.*, “Multipath++: Efficient information fusion and trajectory aggregation for behavior prediction,” in *International Conference on Robotics and Automation*, IEEE, 2022, pp. 7814–7821.
- [4] J. Gu, C. Sun, and H. Zhao, “Densetnt: End-to-end trajectory prediction from dense goal sets,” in *International Conference on Computer Vision*, IEEE, 2021, pp. 15 303–15 312.
- [5] D. Fridovich-Keil, A. Bajcsy, J. F. Fisac, *et al.*, “Confidence-aware motion prediction for real-time collision avoidance,” *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 250–265, 2020.
- [6] H. D. Kabir, A. Khosravi, M. A. Hosen, and S. Nahavandi, “Neural network-based uncertainty quantification: A survey of methodologies and applications,” *IEEE access*, vol. 6, pp. 36 218–36 234, 2018.
- [7] J. Yao, W. Pan, S. Ghosh, and F. Doshi-Velez, “Quality of uncertainty quantification for bayesian neural network inference,” *arXiv preprint 1906.09686*, 2019.
- [8] B. Charpentier, R. Senanayake, M. Kochenderfer, and S. Gunemann, “Disentangling epistemic and aleatoric uncertainty in reinforcement learning,” *arXiv preprint 2206.01558*, 2022.
- [9] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic learning in a random world*. Springer, 2005, vol. 29.
- [10] A. N. Angelopoulos and S. Bates, “A gentle introduction to conformal prediction and distribution-free uncertainty quantification,” *arXiv preprint 2107.07511*, 2021.
- [11] 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.
- [12] A. Dixit, L. Lindemann, S. Wei, M. Cleaveland, G. J. Pappas, and J. W. Burdick, “Adaptive conformal prediction for motion planning among dynamic agents,” *arXiv preprint 2212.00278*, 2022.
- [13] R. Tumu, L. Lindemann, T. Nghiem, and R. Mangharam, “Physics constrained motion prediction with uncertainty quantification,” *arXiv preprint 2302.01060*, 2023.
- [14] Y. Chen, U. Rosolia, C. Fan, A. Ames, and R. Murray, “Reactive motion planning with probabilistic safety guarantees,” in *Conference on Robot Learning*, PMLR, 2021, pp. 1958–1970.
- [15] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, “Safe planning in dynamic environments using conformal prediction,” *IEEE Robotics and Automation Letters*, 2023.
- [16] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, “Hamilton-jacobi reachability: A brief overview and recent advances,” in *2017 IEEE 56th Annual Conference on Decision and Control*, IEEE, 2017, pp. 2242–2253.
- [17] D. Boursinos and X. Koutsoukos, “Assurance monitoring of learning-enabled cyber-physical systems using inductive conformal prediction based on distance learning,” *AI EDAM*, vol. 35, no. 2, pp. 251–264, 2021.
- [18] R. J. Tibshirani, R. Foygel Barber, E. Candes, and A. Ramdas, “Conformal prediction under covariate shift,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [19] C. Xu and Y. Xie, “Conformal prediction interval for dynamic time-series,” in *International Conference on Machine Learning*, PMLR, 2021, pp. 11 559–11 569.
- [20] R. F. Barber, E. J. Candes, A. Ramdas, and R. J. Tibshirani, “Conformal prediction beyond exchangeability,” *arXiv preprint 2202.13415*, 2022.
- [21] I. Gibbs and E. Candes, “Adaptive conformal inference under distribution shift,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 1660–1672, 2021.
- [22] S. Feldman, L. Ringel, S. Bates, and Y. Romano, “Achieving risk control in online learning settings,” *arXiv preprint 2205.09095*, 2023.
- [23] A. Devonport, F. Yang, L. El Ghaoui, and M. Arcak, “Data-driven reachability analysis with christoffel functions,” in *2021 60th IEEE Conference on Decision and Control*, IEEE, 2021, pp. 5067–5072.
- [24] K. Nakamura and S. Bansal, “Online update of safety assurances using confidence-based predictions,” *arXiv preprint 2210.01199*, 2022.
- [25] A. Bajcsy, S. L. Herbert, D. Fridovich-Keil, *et al.*, “A scalable framework for real-time multi-robot, multi-human collision avoidance,” in *2019 International Conference on Robotics and Automation*, IEEE, 2019, pp. 936–943.
- [26] R. Koenker, *Quantile regression*. Cambridge university press, 2005, vol. 38.
- [27] I. Steinwart and A. Christmann, “Estimating conditional quantiles with the help of the pinball loss,” *arXiv preprint 1102.2101*, 2011.
- [28] E. Schmerling, *Hj_reachability*, https://github.com/StanfordASL/hj_reachability, 2021.
- [29] J. F. Fisac, M. Chen, C. J. Tomlin, and S. S. Sastry, “Reach-avoid problems with time-varying dynamics, targets and constraints,” in *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*, 2015, pp. 11–20.
- [30] H. Caesar, V. Bankiti, A. H. Lang, *et al.*, “nuScenes: A multimodal dataset for autonomous driving,” in *Computer Vision and Pattern Recognition*, IEEE, 2020.
- [31] S. Ettinger, S. Cheng, B. Caine, *et al.*, “Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset,” in *International Conference on Computer Vision*, IEEE, 2021, pp. 9710–9719.
- [32] E. Vinitzky, N. Lichtlé, X. Yang, B. Amos, and J. Foerster, “Nocturne: A scalable driving benchmark for bringing multi-agent learning one step closer to the real world,” *arXiv preprint 2206.09889*, 2022.