

Adaptive Conformal Prediction for Motion Planning among Dynamic Agents

Anushri Dixit^{*1}

ADIXIT@CALTECH.EDU

Lars Lindemann^{*2}

LARSL@SEAS.UPENN.EDU

Skylar X. Wei¹

SWEI@CALTECH.EDU

Matthew Cleaveland²

MCLEAV@SEAS.UPENN.EDU

George J. Pappas²

PAPPASG@SEAS.UPENN.EDU

Joel W. Burdick¹

JWB@ROBOTICS.CALTECH.EDU

¹California Institute of Technology, Pasadena, CA

²University of Pennsylvania, Philadelphia, PA

^{*} Indicates equal contribution

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

This paper proposes an algorithm for motion planning among dynamic agents using adaptive conformal prediction. We consider a deterministic control system and use trajectory predictors to predict the dynamic agents' future motion, which is assumed to follow an unknown distribution. We then leverage ideas from adaptive conformal prediction to dynamically quantify prediction uncertainty from an online data stream. Particularly, we provide an online algorithm that uses delayed agent observations to obtain uncertainty sets for multistep-ahead predictions with probabilistic coverage. These uncertainty sets are used within a model predictive controller to safely navigate among dynamic agents. While most existing data-driven prediction approaches quantify prediction uncertainty heuristically, we quantify the true prediction uncertainty in a distribution-free, adaptive manner that even allows to capture changes in prediction quality and the agents' motion. We empirically evaluate our algorithm on a case study where a drone avoids a flying frisbee.

Keywords: MPC, dynamic environments, uncertainty quantification, and conformal prediction.


1. Introduction

Motion planning of autonomous systems in dynamic environments requires the system to reason about uncertainty in its environment, e.g., a self-driving car needs to reason about uncertainty in the motion of other vehicles, and a mobile robot navigating a crowded space needs to assess uncertainty of nearby pedestrians. These applications are safety critical, as the agents' intentions are unknown, and systems must be able to plan reactive behaviors in response to an increase in uncertainty.

Existing works include predictive and reactive approaches, e.g., multi-agent navigation via the dynamic window approach [Fox et al. \(1997\)](#); [Mitsch et al. \(2013\)](#) or navigation functions [Dimarogonas et al. \(2006\)](#); [Tanner et al. \(2003\)](#). Reactive approaches typically consider simplified dynamics and do not optimize performance. Predictive approaches incorporate predictions of the agents' future motion and can optimize performance. Interactive approaches take inter-agent interaction into account [Kretzschmar et al. \(2016\)](#); [Everett et al. \(2021\)](#), while non-interactive approaches ignore potential interactions [Trautman and Krause \(2010\)](#); [Du Toit and Burdick \(2011\)](#).

While many prior works assume perfect knowledge of the environment, an important challenge is to account for uncertainty in prediction. Existing works address the problem by making simplifying assumptions, such as linear system dynamics and bounded or Gaussian uncertainty distributions [Aoude et al. \(2013\)](#); [Thomas et al. \(2021\)](#); [Renganathan et al. \(2020\)](#). However, addressing the problem in its full generality for nonlinear dynamics and arbitrary distributions is an open problem.

In this paper, we use trajectory predictors to predict the agents' future motion, and quantify prediction uncertainty in an adaptive and online manner from past agent observations of a single trajectory. Particularly, we use tools from the adaptive conformal prediction (ACP) literature [Gibbs and Candès \(2021\)](#); [Gibbs and Candès \(2022\)](#); [Zaffran et al. \(2022\)](#) to construct prediction regions that quantify multistep-ahead prediction uncertainty. Based on this quantification, we formulate an uncertainty-informed motion planner. Our contributions are as follows:



- We propose an algorithm that adaptively quantifies uncertainty of trajectory predictors using ACP. Our algorithm is distribution-free and applies to a broad class of trajectory predictors, providing average probabilistic coverage.
- We propose a model predictive controller (MPC) that leverages uncertainty quantifications to plan probabilistically safe paths around dynamic obstacles. Importantly, our adaptive algorithm enables us to capture and react to changes in prediction quality and the agents' motion.
- We provide empirical evaluations of a drone avoiding a flying frisbee.

1.1. Related Work

Planning in dynamic environments has found broad interest, and non-interactive sampling-based motion planner were presented in [Phillips and Likhachev \(2011\)](#); [Renganathan et al. \(2022\)](#); [Aoude et al. \(2013\)](#); [Majd et al. \(2021\)](#); [Kalluraya et al. \(2022\)](#), while [Du Toit and Burdick \(2011\)](#); [Wei et al. \(2022\)](#); [Wang et al. \(2022\)](#); [Thomas et al. \(2021\)](#) propose non-interactive receding horizon planning algorithms. However, accounting for uncertainty in the agent motion is challenging.

Intent-driven models for planning among human agents have estimated agent uncertainty using Bayesian inference [Fisac et al. \(2018\)](#); [Nakamura and Bansal \(2022\)](#); [Fridovich-Keil et al. \(2020\)](#); [Bansal et al. \(2020\)](#). Model predictive control was also used in a stochastic setting to account for uncertainty under the assumption of bounded or Gaussian uncertainty [Fan et al. \(2021\)](#); [Nair et al. \(2022\)](#); [Yoon et al. \(2021\)](#). Data-driven trajectory predictors can provide mean and variance information of the predictions, which can be approximated as a Gaussian distribution [Busch et al. \(2022\)](#) and used within stochastic planning frameworks [Choi et al. \(2017\)](#); [Omainiska et al. \(2021\)](#); [Fulgenzi et al. \(2008\)](#). These approaches quantify prediction uncertainty in a heuristic manner for real systems as the authors make certain assumptions on prediction algorithms and agent models and its distribution, e.g., being Gaussian. Distributionally robust approaches such as [Wei et al. \(2022\)](#) are distribution free and can ensure safety at the cost of conservatism.

Data-driven trajectory predictors, such as RNNs or LSTMs, provide no information about prediction uncertainty which can lead to unsafe decisions. For this reason, prediction monitors were recently presented in [Farid et al. \(2022\)](#); [Luo et al. \(2021\)](#) to monitor prediction quality. Especially [Luo et al. \(2021\)](#) used conformal prediction to obtain guarantees on the predictor's false negative

rate. Conformal prediction was further used to obtain estimates on constraint satisfaction via neural network predictors [Dietterich and Hostetler \(2022\)](#); [Bortolussi et al. \(2019\)](#); [Qin et al. \(2022\)](#); [Lindemann et al. \(2022b\)](#). Conceptually closest to our work are [Chen et al. \(2020\)](#); [Lindemann et al. \(2022a\)](#) where prediction uncertainty quantifications are obtained using conformal prediction, and then utilized to design model predictive controllers. While the algorithm in [Chen et al. \(2020\)](#) can not provide end-to-end safety guarantees, [Lindemann et al. \(2022a\)](#) can provide probabilistic safety guarantees for the planner. However, changes in the distribution that describes the agents' motion can not be accounted for, e.g., when the agents' motion changes depending on the motion of the control system. Another distinct difference is that offline trajectory data is needed, while we obtain uncertainty quantifications in an adaptive manner from past agent observations of a single trajectory.

2. Problem Formulation and Preliminaries

The dynamics of our autonomous system are governed by the discrete-time dynamical system,

$$x_{t+1} = f(x_t, u_t), \quad x_0 := \zeta \quad (1)$$

where $x_t \in \mathcal{X} \subseteq \mathbb{R}^n$ and $u_t \in \mathcal{U} \subseteq \mathbb{R}^m$ denote the state and the control input at time $t \in \mathbb{N} \cup \{0\}$, respectively. The sets $\mathcal{U}$ and $\mathcal{X}$ denote the set of permissible control inputs and the workspace of the system, respectively. The measurable function $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ describes the system dynamics and $\zeta \in \mathbb{R}^n$ is the initial condition of the system. For brevity, let $x := (x_0, x_1, \dots)$ denote the trajectory of (1) under a given control sequence $u := (u_0, u_1, \dots)$.

The system operates in an environment with N dynamic agents whose trajectories are a priori unknown. Let $\mathcal{D}$ be an unknown distribution over agent trajectories, i.e., let $Y := (Y_0, Y_1, \dots) \sim \mathcal{D}$ describe a random trajectory where the joint agent state $Y_t := (Y_{t,1}, \dots, Y_{t,N})$ at times $t \in \mathbb{N} \cup \{0\}$ is drawn from $\mathbb{R}^{Nn}$, i.e., $Y_{t,j}$ is the state of agent j at time t . For instance, Y_t can denote the uncertain two-dimensional positions of N pedestrians at time t . Modeling dynamic agents by a distribution $\mathcal{D}$ provides great flexibility, and $\mathcal{D}$ can generally describe the motion of Markov decision processes. We make no other assumptions on the distribution $\mathcal{D}$, and in our proposed algorithm we will predict states $(Y_{t+1}, \dots, Y_{t+H})$ for a prediction horizon of H from $(Y_0, \dots, Y_t)$ and quantify prediction uncertainty using ideas from ACP.

Problem 1 *Given the system in (1), the unknown random trajectories $Y \sim \mathcal{D}$, and a failure probability $\delta \in (0, 1)$, design the control inputs u_t such that the Lipschitz continuous constraint function $c : \mathbb{R}^n \times \mathbb{R}^{nN} \rightarrow \mathbb{R}$ is satisfied¹ with a probability of at least $1 - \delta$ on average, i.e., that*

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \text{Prob}(c(x_t, Y_t) \geq 0) \geq 1 - \delta \quad (2)$$

We note that our previous work [Lindemann et al. \(2022a\)](#) considers a similar problem formulation, but where we aim for pointwise satisfaction of the constraint function c , i.e., $\text{Prob}(c(x_t, Y_t) \geq 0) \geq 1 - \delta$ for all $t \geq 0$. In [Lindemann et al. \(2022a\)](#), however, it is assumed that an offline calibration dataset drawn from $\mathcal{D}$ is available, which we will not assume here. This makes the problem

1. For an obstacle avoidance constraint, like $c(x, y) := \|x - y\| - 0.5 \geq 0$, the Lipschitz constant is 1. We implicitly assume that the constraint function is initially satisfied, i.e., that $c(x_0, y_0) \geq 0$.

more challenging so that we aim for the average probabilistic guarantee in equation (2). Additionally, we allow the distribution $\mathcal{D}$ to depend on x , e.g., a pedestrian’s behavior may change if a car comes too close, which is not possible in Lindemann et al. (2022a). This work is a step towards the implementation of a general framework that can adapt to such changes in the agent distribution.

To address Problem 1, we use trajectory predictors to predict the motion of the agents $(Y_0, Y_1, \dots)$ to enforce the constraint (2) within a MPC framework. In Lindemann et al. (2022a), we assumed the availability of calibration data from $\mathcal{D}$ to build prediction regions that quantify uncertainty of trajectory predictors. In this setting, we can collect data online to adapt our uncertainty sets based on past performance of our predictor using ACP without any assumptions on the distribution of the uncertainty and exchangeability of the validation and training dataset.

Remark 1 *We could parameterize the distribution $\mathcal{D}$ by the trajectory x to model potential interactions between system and agents. This way, we can react to cases where the trajectory predictor (introduced next) is trained without information of x , i.e., without taking interactions into account.*

Trajectory Predictors: Given observations $(Y_0, \dots, Y_t)$ at time t , we want to predict future states $(Y_{t+1}, \dots, Y_{t+H})$ for a prediction horizon of H . Assume that PREDICT is a function that maps observations $(Y_0, \dots, Y_t)$ to predictions $(\hat{Y}_t^1, \dots, \hat{Y}_t^H)$ of $(Y_{t+1}, \dots, Y_{t+H})$. Note that t in $\hat{Y}_t^\tau$ denotes the time at which the prediction is made, while τ indicates how many steps we predict ahead. In principle, PREDICT can be a classical auto-regressive model or a neural network based method.

While our proposed problem solution is compatible with any trajectory predictor PREDICT, we focus in the case studies on real-time updating strategies like sliding linear predictors with extended Kalman filter. Extracting a dynamics model from data is challenging, especially when the available data is limited, noisy, and partial. Takens (1981) showed that the method of delays can be used to reconstruct qualitative features of the full-state, phase space from delayed partial observations. By building on our previous work using time delay embedding in dynamic obstacle avoidance (Wei et al. (2022)), we employ a linear predictor based on spatio-temporal factorization of the delayed partial observations as the pairing trajectory predictor (Dixit et al., 2022, Appendix).

Adaptive Conformal Prediction (ACP): Conformal prediction is used to obtain prediction regions for predictive models, e.g., neural networks, without making assumptions on the underlying distribution or the predictive model Vovk et al. (2005); Shafer and Vovk (2008); Angelopoulos and Bates (2021). Let $R_1, \dots, R_{t+1}$ be $t+1$ independent and identically distributed (i.i.d.) random variables. The goal in conformal prediction is to obtain a prediction region of R_{t+1} based on $R_1, \dots, R_t$. Formally, given a failure probability $\delta \in (0, 1)$, we want to obtain a prediction region $R_{t+1} \leq C$ such that

$$\text{Prob}(R_{t+1} \leq C) \geq 1 - \delta.$$

We refer to R_i also as the nonconformity score. For supervised learning, we can select $R_i := \|Z_i - \mu(X_i)\|$ where μ is the predictor so that a large nonconformity score indicates a poor predictive model. By a quantile argument, see (Tibshirani et al., 2019, Lemma 1), we can obtain C to be the $(1 - \delta)$ th quantile of the empirical distribution of the values $R_1, \dots, R_t$ and ∞ . Calculating the $(1 - \delta)$ th quantile can be done by assuming that $\bar{R}_1, \dots, \bar{R}_t$ correspond to the values of $R_1, \dots, R_t$, but instead sorted in non-decreasing order ($\bar{R}$ refers to the order statistic of R), i.e., for each $\bar{R}_i$ there exists exactly one R_j such that $\bar{R}_i = R_j$ and $\bar{R}_{i+1} \geq \bar{R}_i$. By setting $q := \lceil (t+1)(1 - \delta) \rceil \leq t$, we obtain the $(1 - \delta)$ th quantile as $C := \bar{R}_q$, i.e., the q^{th} smallest nonconformity score.

The underlying assumption in conformal prediction is that $R_1, \dots, R_{t+1}$ are exchangeable (exchangeability includes i.i.d. data). This is an unreasonable assumption for time-series prediction where R_t may denote the nonconformity score at time t . To address this issue, ACP was studied in [Gibbs and Candes \(2021\)](#); [Gibbs and Candès \(2022\)](#); [Zaffran et al. \(2022\)](#); [Bastani et al. \(2022\)](#). The idea is now to obtain a prediction region $R_{t+1} \leq C_{t+1}$ adaptively so that $\text{Prob}(R_{t+1} \leq C_{t+1}) \geq 1 - \delta$ for each time t . In fact, the prediction region is now obtained using $C_{t+1} := \bar{R}_{q_{t+1}}$ where $q_{t+1} := \lceil (t+1)(1 - \delta_{t+1}) \rceil$ depends on the variable δ_{t+1} that is adapted online based on observed data. In this way, C_{t+1} becomes a tuneable parameter by the choice of δ_{t+1} . To adaptively obtain the parameter δ_{t+1} , ideas from online learning are used and we update δ_{t+1} as

$$\delta_{t+1} := \delta_t + \gamma(\delta - e_t) \quad \text{with} \quad e_t := \begin{cases} 0 & \text{if } r_t \leq C_t \\ 1 & \text{otherwise} \end{cases} \quad (3)$$

where we denote by r_t the observed realization of R_t and where γ is a learning rate. The idea is to use δ_{t+1} to adapt to changes in the distribution of $R_1, \dots, R_{t+1}$ over time by using information on how much C_t overcovered ($r_t \ll C_t$) or undercovered ($r_t \gg C_t$) in the past.

Remark 2 *One of the main performance enhancers is the proper choice of γ . In [Gibbs and Candès \(2022\)](#), the authors present fully adaptive conformal prediction (FACP) where a set of learning rates $\{\gamma_i\}_{1 \leq i \leq k}$ is used in parallel from which the best γ is selected adaptively. Based on past performance (using a reweighting scheme that evaluates which γ_i provided the best coverage), the authors maintain a belief $p_t^{(i)}$ at each time step t for each $\{\delta_t^{(i)}\}_{1 \leq i \leq k}$. The new update laws are*

$$\delta_{t+1}^{(i)} := \delta_t^{(i)} + \gamma_i(\delta - e_t^{(i)}) \quad \text{with} \quad e_t^{(i)} := \begin{cases} 0 & \text{if } r_t \leq C_t^{(i)} \\ 1 & \text{otherwise} \end{cases}$$

where $C_t^{(i)} := \bar{R}_{q_t^{(i)}}$ with $q_t^{(i)} := \lceil (t+1)(1 - \delta_t^{(i)}) \rceil$, while $C_t := \bar{R}_{q_t}$ with $q_t := \lceil (t+1)(1 - \sum_{i=1}^k p_t^{(i)} \delta_t^{(i)}) \rceil$.

3. Adaptive Conformal Prediction Regions for Trajectory Predictions

Recall that we can obtain predictions $(\hat{Y}_t^1, \dots, \hat{Y}_t^H)$ at time t of future agent states $(Y_{t+1}, \dots, Y_{t+H})$ from past observations $(Y_0, \dots, Y_t)$ using the PREDICT function. Note, however, that these point predictions contain no information about prediction uncertainty and can hence not be used to reason about the safety constraint (2). To tackle this issue, we aim to construct prediction regions for $(Y_{t+1}, \dots, Y_{t+H})$ using ideas from ACP.

To obtain prediction regions for $(Y_{t+1}, \dots, Y_{t+H})$, we could consider the nonconformity score $\|Y_{t+\tau} - \hat{Y}_t^\tau\|$ at time t that captures the multistep-ahead prediction error for each $\tau \in \{1, \dots, H\}$. A large nonconformity score indicates that the prediction $\hat{Y}_t^\tau$ of $Y_{t+\tau}$ is not accurate, while a small score indicates an accurate prediction. For each τ , we wish to obtain a prediction region using C_t^τ that is again defined by an update variable δ_t^τ . Note, however, that we can not evaluate $\|Y_{t+\tau} - \hat{Y}_t^\tau\|$ at time t as only measurements $(Y_0, \dots, Y_t)$ are known, but not $(Y_{t+1}, \dots, Y_{t+H})$. Consequently, we cannot use the update rule (3) to update δ_t^τ , as the error e_t^τ would depend on checking if $\|Y_{t+\tau} - \hat{Y}_t^\tau\| \leq C_t^\tau$. To address this issue, we define the time lagged nonconformity score

$$R_t^\tau := \|Y_t - \hat{Y}_{t-\tau}^\tau\|$$

that we can evaluate at time t so that we can use the update rule (3). This nonconformity score R_t^τ is time lagged in the sense that, at time t , we evaluate the τ step-ahead prediction error that was made τ time steps ago. We can now update the parameter δ_{t+1}^τ that defines C_{t+1}^τ as

$$\delta_{t+1}^\tau := \delta_t^\tau + \gamma(\delta - e_t^\tau) \quad \text{with} \quad e_t^\tau := \begin{cases} 0 & \text{if } \|Y_t - \hat{Y}_{t-\tau}^\tau\| \leq C_t^\tau \\ 1 & \text{otherwise.} \end{cases} \quad (4)$$

To compute C_{t+1}^τ , note that we can not compute $R_1^\tau, \dots, R_{\tau-1}^\tau$. Therefore, with minor change, we let C_{t+1}^τ be the $\lceil (t - \tau + 1)(1 - \delta_{t+1}^\tau) \rceil^{\text{th}}$ smallest value of $(R_\tau^\tau, \dots, R_t^\tau)$ ².

By obtaining a prediction region for R_{t+1}^τ using ACP, we obtain a prediction region for the τ step-ahead prediction error that was made $\tau - 1$ time steps ago, i.e., for $\|Y_{t+1} - \hat{Y}_{t+1-\tau}^\tau\|$. Under the assumption that R_{t+1}^τ and $R_{t+\tau}^\tau$ are independent and identically distributed, R_{t+1}^τ serves as a prediction region for τ step-ahead prediction error that was made 0 time steps ago (now at time t), i.e., for $R_{t+\tau}^\tau$ which encodes $\|Y_{t+\tau} - \hat{Y}_t^\tau\|$. Naturally, in our setting R_{t+1}^τ and $R_{t+\tau}^\tau$ are not independent and identically distributed. However, by adapting the prediction regions based on the performance of the multi-step predictions using ACP, we can still obtain on-average coverage guarantees, albeit weaker than the point-wise guarantees provided by conformal predictions. We remark that for the theoretical guarantees that we provide in the next section, only the one step-ahead prediction errors are relevant.

Corollary 3 *Let γ be a learning rate, $\delta_0^1 \in (0, 1)$ be an initial value for the recursion (4), and T be the number of times that we compute the recursion (4). Then, for the onestep-ahead prediction errors, it holds that*

$$1 - \delta - p_1 \leq \frac{1}{T} \sum_{t=0}^{T-1} \text{Prob}(\|Y_{t+1} - \hat{Y}_t^1\| \leq C_{t+1}^1) \leq 1 - \delta + p_2 \quad (5)$$

with constants $p_1 := \frac{\delta_0^1 + \gamma}{T\gamma}$, $p_2 := \frac{(1 - \delta_0^1) + \gamma}{T\gamma}$ so that $\lim_{T \rightarrow \infty} p_1 = 0$ and $\lim_{T \rightarrow \infty} p_2 = 0$.

Proof Since the probability of an event is equivalent to the expected value of the indicator function of that event, it follows by the definition of the error e_{t+1}^1 that

$$\text{Prob}(\|Y_{t+1} - \hat{Y}_t^1\| \leq C_{t+1}^1) = \mathbb{E}[1 - e_{t+1}^1] = 1 - \mathbb{E}[e_{t+1}^1]. \quad (6)$$

For a given initialization δ_0^τ and learning rate γ , we know from (Gibbs and Candes, 2021, Proposition 4.1) that the following bound holds (with probability one) for the misclassification errors

$$\frac{-(1 - \delta_0^1) + \gamma}{T\gamma} \leq \frac{1}{T} \sum_{t=0}^{T-1} e_{t+1}^1 - \delta \leq \frac{\delta_0^1 + \gamma}{T\gamma} \implies \left| \frac{1}{T} \sum_{t=0}^{T-1} e_{t+1}^1 - \delta \right| \leq \frac{\max(\delta_0^1, 1 - \delta_0^1) + \gamma}{T\gamma}.$$

2. Instead of keeping track of all data, we will choose a sliding window of the N most recent data. For all prediction regions, we will then consider $(R_{t-N}^\tau, \dots, R_t^\tau)$ and compute C_{t+1}^τ as the $\lceil (N + 1)(1 - \delta_{t+1}^\tau) \rceil^{\text{th}}$ smallest value.

Hence, taking the expectation of the above two-sided inequality, we get that

$$\begin{aligned}
 \frac{-(1 - \delta_0^1) + \gamma}{T\gamma} &\leq \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[e_{t+1}^1] - \delta \leq \frac{\delta_0^1 + \gamma}{T\gamma}, \\
 \stackrel{(a)}{\Leftrightarrow} \frac{-(1 - \delta_0^1) + \gamma}{T\gamma} &\leq \frac{1}{T} \sum_{t=0}^{T-1} (1 - \text{Prob}(\|Y_{t+1} - \hat{Y}_t^1\| \leq C_{t+1}^1)) - \delta \leq \frac{\delta_0^1 + \gamma}{T\gamma}, \\
 \Leftrightarrow 1 - \delta + \frac{(1 - \delta_0^1) + \gamma}{T\gamma} &\geq \frac{1}{T} \sum_{t=0}^{T-1} \text{Prob}(\|Y_{t+1} - \hat{Y}_t^1\| \leq C_{t+1}^1) \geq 1 - \delta - \frac{\delta_0^1 + \gamma}{T\gamma},
 \end{aligned}$$

where we used equation (6) for the equivalence in (a). ■

Remark 4 The above result can be similarly extended to the FACP case with a set of candidate learning rates, γ , (Gibbs and Candès, 2022, Theorem 3.2).

Example 1 To illustrate these multistep-ahead prediction regions, consider a planar double pendulum whose dynamics are governed by chaotic, nonlinear dynamics that are sensitive to the initial condition (Shinbrot et al., 1992). We study the predictions made by a linear predictor that uses noisy observations of the position of the double pendulum (Dixit et al., 2022, Appendix) and use ACP to predict the uncertainty in the predictions. Both the trajectory predictor and the uncertainty quantification using ACP use online data from a single trajectory. ACP provides the multi-step errors in the linear predictions with a coverage level of $\delta = 0.1$, and learning rates $\gamma = (0.0008 \ 0.0015 \ 0.003 \ 0.005 \ 0.009 \ 0.017 \ 0.03 \ 0.05 \ 0.08)$.

Figure 1 compares the 1-step and 6-step ahead error prediction regions to the true multi-step errors for two states, the second mass position, x_2, y_2 . The percentages of one-step errors that are incorrectly predicted, i.e., $e_t^1 = 1$, for the positions of each mass, x_1, x_2, y_1, y_2 are 2.36%, 0.94%, 1.57%, 1.73% respectively. We can see the effects of adaptation as the ACP prediction regions are larger in areas of poor performance of the linear predictor (and consequently higher error in the prediction) and smaller in regions where the linear predictor performs well. We note that the miscoverage levels are lower than the expected miscoverage when $\delta = 0.1$. This is because the learning rates γ used are small, because of which the adaptation is slower. As the learning rate becomes larger, the adaptation is faster and the miscoverage levels will be closer to 10%. In practice, we found that higher learning rates lead to high variations in the prediction sets.

4. Uncertainty-Informed Model Predictive Control

Based on the obtained uncertainty quantification from the previous section, we propose an uncertainty-informed model predictive controller (MPC) that uses predictions $\hat{Y}_t^\tau$ and adaptive predictions C_{t+1}^τ . The underlying optimization problem that is solved at every time step t is:

$$\min_{(u_t, \dots, u_{t+H-1})} \sum_{k=t}^{t+H-1} J(x_{k+1}, u_k) \tag{7a}$$

$$\text{s.t.} \quad x_{k+1} = f(x_k, u_k), \quad k \in \{t, \dots, t+H-1\} \tag{7b}$$

$$c(x_{t+\tau}, \hat{Y}_t^\tau) \geq LC_{t+1}^\tau, \quad \tau \in \{1, \dots, H\} \tag{7c}$$

$$u_k \in \mathcal{U}, x_{k+1} \in \mathcal{X}, \quad k \in \{t, \dots, t+H-1\} \tag{7d}$$




Figure 1: The multi-step prediction errors are shown for two of the six states of a double pendulum (x_2, y_2) . ACP can correctly predict regions of high and low error (90% coverage regions) by adjusting the prediction quantile using update law (3). The orange lines are the true multi-step prediction errors and the blue areas are the error regions predicted by ACP.

where L is the Lipschitz constant of the constraint function c , J is a step-wise cost function, and $u_t, \dots, u_{t+H-1}$ is the control sequence. The optimization problem in (7) is convex if the functions J and f are convex, the function c is convex in its first argument, and the sets $\mathcal{U}$ and $\mathcal{X}$ are convex.

Based on this optimization problem, we propose a receding horizon control strategy in Algorithm 1. In line 1 of Algorithm 1, we initialize the parameter δ_0^t simply to δ . Lines 2-11 present the real-time planning loop by: 1) updating the states x_t and y_t and calculating new predictions $\hat{Y}_t^\tau$ (lines 3-4), 2) computing the adaptive nonconformity scores C_{t+1}^τ (lines 5-9), and 3) solving the optimization problem in (7) of which we apply only u_t (lines 10-11).

Algorithm 1 MPC with ACP Regions

Input: Failure probability δ , prediction horizon H , learning rate γ

Output: Control input $u_t(x_t, Y_0, \dots, Y_t)$ at each time t

- 1: $\delta_0^\tau \leftarrow \delta$ for $\tau \in \{1, \dots, H\}$
 - 2: **for** t from 0 to ∞ **do** # real-time motion planning loop
 - 3: Update x_t and Y_t
 - 4: Obtain predictions $\hat{Y}_t^\tau$ for $\tau \in \{1, \dots, H\}$
 - 5: **for** τ from 1 to H **do** # compute ACP regions
 - 6: $\delta_{t+1}^\tau \leftarrow \delta_t^\tau + \gamma(\delta - e_t^\tau)$
 - 7: $R_t^\tau := \|Y_t - \hat{Y}_{t-\tau}^\tau\|$
 - 8: $q \leftarrow \lceil (t+1)(1 - \delta_{t+1}^\tau) \rceil$
 - 9: Set C_{t+1}^τ as the q th smallest value of $(R_t^\tau, \dots, R_t^\tau)$
 - 10: Calculate controls $u_t, \dots, u_{t+H-1}$ as the solution of (7)
 - 11: Apply u_t to (1)
-

Remark 5 While Algorithm 1 uses a single learning rate, one can similarly extend the above algorithm to be fully adaptive using a candidate set of $\{\gamma_i\}_{1 \leq i \leq k}$ without loss of generality.

Remark 6 *Gibbs and Candes (2021)* assume that when $\delta_{t+1} \leq 0$, $C_{t+1} \rightarrow \infty$. This means that when the algorithm requires robust behavior, the ∞ -prediction region ensures that any prediction at the next time-step should be correctly classified. For a physical system, there are limits on how much the dynamic obstacle can accelerate in one time-step which gives us an upper bound $R_{\max} < \infty$ on the worst-case error. In practice, we enforce $0 \leq \delta_{t+1} \leq 1$ with $C_{t+1} \leq R_{\max}$.

Theorem 7 Let γ be a learning rate, $\delta_0^1 \in (0, 1)$ be an initial value for the recursion (4), and T be the number of times that we compute the recursion (4). If the optimization problem (7) in Algorithm 1 is recursively feasible, then Algorithm 1 will lead to

$$\frac{1}{T} \sum_{t=0}^{T-1} \text{Prob}(c(x_{t+1}, Y_{t+1}) \geq 0) \geq 1 - \delta - p_1 \quad (8)$$

with constant $p_1 := \frac{\delta_0^1 + \gamma}{T\gamma}$ so that $\lim_{T \rightarrow \infty} p_1 = 0$.

Proof By assumption, the optimization problem in (7) is feasible at each time $t \in \{0, 1, \dots\}$. Due to constraint (7c) and Lipschitz continuity of c , it hence holds that

$$0 \leq c(x_{t+1}, \hat{Y}_t^1) - LC_{t+1}^1 \leq c(x_{t+1}, Y_{t+1}) + L\|Y_{t+1} - \hat{Y}_t^1\| - LC_{t+1}^1 \quad (9)$$

at each time $t \in \{0, 1, \dots\}$. Consequently, note that $\|Y_{t+1} - \hat{Y}_t^1\| \leq C_{t+1}^1$ is a sufficient condition for $c(x_{t+1}, Y_{t+1}) \geq 0$. In a next step, we can derive that

$$\begin{aligned} \text{Prob}(c(x_{t+1}, Y_{t+1}) \geq 0) &\stackrel{(a)}{=} \text{Prob}(c(x_{t+1}, Y_{t+1}) \geq 0 \mid \|Y_{t+1} - \hat{Y}_t^1\| \leq C_t^1) \text{Prob}(\|Y_{t+1} - \hat{Y}_t^1\| \leq C_t^1) \\ &\quad + \text{Prob}(c(x_{t+1}, Y_{t+1}) \geq 0 \mid \|Y_{t+1} - \hat{Y}_t^1\| > C_t^1) \text{Prob}(\|Y_{t+1} - \hat{Y}_t^1\| > C_t^1) \\ &\stackrel{(b)}{\geq} \text{Prob}(c(x_{t+1}, Y_{t+1}) \geq 0 \mid \|Y_{t+1} - \hat{Y}_t^1\| \leq C_t^1) \text{Prob}(\|Y_{t+1} - \hat{Y}_t^1\| \leq C_t^1) \\ &\stackrel{(c)}{=} \text{Prob}(\|Y_{t+1} - \hat{Y}_t^1\| \leq C_t^1) \end{aligned}$$

where the equality in (a) follows from the law of total probability, while the inequality in (b) follows from the nonnegativity of probabilities. The equality in (c) follows as $\text{Prob}(c(x_{t+1}, Y_{t+1}) \geq 0 \mid \|Y_{t+1} - \hat{Y}_t^1\| \leq C_t^1) = 1$ since $\|Y_{t+1} - \hat{Y}_t^1\| \leq C_t^1$ implies $c(x_{t+1}, Y_{t+1}) \geq 0$ according to (9). We now use the result from Corollary 3 to complete the proof. $\blacksquare$

5. Case Studies: Multirotor operating in small angle regime dodging a flying frisbee

We compare the performance of the MPC with ACP uncertainty prediction regions with our past work that uses a distributionally robust approach to uncertainty quantification (Wei et al., 2022). We use the same multirotor operating in the presence of a moving obstacle example with a MPC planner. The multirotor is constrained to operate within the state constraints $\theta \in [-0.45, 0.45]$ radians and $\varphi \in [-0.45, 0.45]$ radians. We use the following standard multirotor linear dynamics,

$$\ddot{x} = -g\theta, \quad \ddot{y} = g\varphi, \quad \ddot{z} = u_1 - g, \quad \ddot{\varphi} = \frac{u_2}{I_{xx}}, \quad \ddot{\theta} = \frac{u_3}{I_{yy}}, \quad \ddot{\psi} = \frac{u_4}{I_{zz}}, \quad (10)$$

where the planner control inputs u_1, u_2, u_3, u_4 correspond to the thrust force in the body frame and three moments. The vehicle's moments of inertia are $I_{xx} = 0.0075 \text{kgm}^2$, $I_{yy} = 0.0075 \text{kgm}^2$, $I_{zz} =$

Case	δ	0.025			0.05		
	UQ method	Proposed	Wei et al. (2022)	w/EKF	Proposed	Wei et al. (2022)	w/EKF
Frisbee w/drag	%Feas.	83.8	87.4	97.1	80.9	90.3	97.6
	%Succ.	99.2	100	100	100	100	100
	$\bar{d}_{min}$	2.91	14.2	5.27	2.74	4.97	4.25
	$\sigma(d_{min})$	1.25	2.04	1.28	1.3	1.97	1.11

Table 1: Summary of results from MC simulations of system (10). We used FACP for predicting uncertainty sets with learning rates $\gamma = \{0.0008, 0.0015, 0.003, 0.005, 0.009, 0.017, 0.03, 0.05, 0.08, 0.13\}$ and using the last 30 measurements of the obstacle.

$0.013kgm^2$. The MPC planner has a horizon length of 10 steps and the planner is updated at 20 Hz. It is implemented through a Sequential Convex Programming approach (Morgan et al., 2014).

Numerical simulations of the proposed MPC planner with ACP regions and dynamics (10) are presented as it avoids a Frisbee that is thrown at the drone from various initial positions, velocities, and rotation speed. The Frisbee is modeled following Hummel (2003), and we implement linear predictions of the trajectory arising from its nonlinear dynamics.

We conducted 1000 Monte Carlo simulations per allowed failure probability level δ to compare the numerical feasibility, percentage of success in obstacle avoidance (if the MPC planner is feasible), and the planner’s conservativeness, as measured by the minimum distance between the obstacle and agent centers, i.e., $\bar{d}_{min}$ and $\sigma(d_{min})$ describe the average and standard deviation of this minimum distance across simulations, respectively. We compare three uncertainty quantification techniques in Table 1, (1) The proposed ACP method (Algorithm 1), (2) empirical bootstrap prediction that accounts for the prediction uncertainty using the empirical bootstrap variance (Wei et al., 2022), and (3) the sliding linear predictor with an Extended Kalman Filter (EKF) that provides Gaussian approximations of the obstacle prediction uncertainty (Dixit et al., 2022, Appendix).

Discussion: Table 1 shows that our proposed method can successfully avoid the Frisbee, while using a significantly smaller average divergence distance ($d_{min}, \sigma(d_{min})$) from the Frisbee. I.e., our approach avoids the conservatism of other approaches due to the adaptivity of the uncertainty sets. Our method can usefully adjust the prediction sets when the underlying uncertainty distribution is shifting (due to discrepancy in the linear dynamic predicted and the true nonlinear obstacle motion). We also note that the feasibility of the MPC optimization is worse for our method compared to Wei et al. (2022) and the EKF predictor. This issue arises during sudden changes in the size of the uncertainty sets when the learning rate γ is chosen too large. We will investigate this issue in future work by considering tools to ensure recursive feasibility (Hewing et al., 2020) or by providing backup controllers (Singletary et al., 2022; Tordesillas et al., 2020) when the MPC is infeasible.

6. Conclusion

We presented an algorithm for safe motion planning in an environment with other dynamic agents using ACP. Specifically, we considered a deterministic control system that uses state predictors to estimate the future motion of dynamic agents. We then leveraged ideas from ACP to dynamically quantify prediction uncertainty from an online data stream, and designed an uncertainty-informed model predictive controller to safely navigate among dynamic agents. In contrast to other data-driven prediction models that quantify prediction uncertainty in a heuristic manner, we quantify the true prediction uncertainty in a distribution-free, adaptive manner that even allows to capture changes in prediction quality and the agents’ motion.

Acknowledgments

Lars Lindemann, Matthew Cleaveland, and George J. Pappas were generously supported by NSF award CPS-2038873. The work of Anushri Dixit and Skylar Wei was supported in part by DARPA, through the LINC program.

References

- Anastasios N Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*, 2021.
- Georges S Aoude, Brandon D Luders, Joshua M Joseph, Nicholas Roy, and Jonathan P How. Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns. *Autonomous Robots*, 35(1):51–76, 2013.
- Somil Bansal, Andrea Bajcsy, Ellis Ratner, Anca D Dragan, and Claire J Tomlin. A hamilton-jacobi reachability-based framework for predicting and analyzing human motion for safe planning. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7149–7155. IEEE, 2020.
- Osbert Bastani, Varun Gupta, Christopher Jung, Georgy Noarov, Ramya Ramalingam, and Aaron Roth. Practical adversarial multivalid conformal prediction. *arXiv preprint arXiv:2206.01067*, 2022.
- Luca Bortolussi, Francesca Cairolì, Nicola Paoletti, Scott A Smolka, and Scott D Stoller. Neural predictive monitoring. In *International Conference on Runtime Verification*, pages 129–147. Springer, 2019.
- Finn Lukas Busch, Jake Johnson, Edward L. Zhu, and Francesco Borrelli. A gaussian process model for opponent prediction in autonomous racing, 2022. URL <https://arxiv.org/abs/2204.12533>.
- Yuxiao Chen, Ugo Rosolia, Chuchu Fan, Aaron D Ames, and Richard Murray. Reactive motion planning with probabilistic safety guarantees. *arXiv preprint arXiv:2011.03590*, 2020.
- Sungjoon Choi, Eunwoo Kim, Kyungjae Lee, and Songhwai Oh. Real-time nonparametric reactive navigation of mobile robots in dynamic environments. *Robotics and Autonomous Systems*, 91: 11–24, 2017.
- Thomas G Dietterich and Jesse Hostetler. Conformal prediction intervals for markov decision process trajectories. *arXiv preprint arXiv:2206.04860*, 2022.
- Anushri Dixit, Lars Lindemann, Skylar Wei, Matthew Cleaveland, George J Pappas, and Joel W Burdick. Adaptive conformal prediction for motion planning among dynamic agents. *arXiv preprint arXiv:2212.00278*, 2022.
- Noel E Du Toit and Joel W Burdick. Robot motion planning in dynamic, uncertain environments. *IEEE Transactions on Robotics*, 28(1):101–115, 2011.

- Michael Everett, Yu Fan Chen, and Jonathan P How. Collision avoidance in pedestrian-rich environments with deep reinforcement learning. *IEEE Access*, 9:10357–10377, 2021.
- David D. Fan, Kyohei Otsu, Yuki Kubo, Anushri Dixit, Joel W. Burdick, and Ali-Akbar Agha-Mohammadi. Step: Stochastic traversability evaluation and planning for risk-aware off-road navigation. In *Robot.: Sci. Syst.*, 2021.
- Alec Farid, Sushant Veer, Boris Ivanovic, Karen Leung, and Marco Pavone. Task-relevant failure detection for trajectory predictors in autonomous vehicles. *arXiv preprint arXiv:2207.12380*, 2022.
- Jaime F Fisac, Andrea Bajcsy, Sylvia L Herbert, David Fridovich-Keil, Steven Wang, Claire J Tomlin, and Anca D Dragan. Probabilistically safe robot planning with confidence-based human predictions. In *14th Robotics: Science and Systems, RSS 2018*. MIT Press Journals, 2018.
- Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, 4(1):23–33, 1997.
- David Fridovich-Keil, Andrea Bajcsy, Jaime F Fisac, Sylvia L Herbert, Steven Wang, Anca D Dragan, and Claire J Tomlin. Confidence-aware motion prediction for real-time collision avoidance1. *The International Journal of Robotics Research*, 39(2-3):250–265, 2020.
- Chiara Fulgenzi, Christopher Tay, Anne Spalanzani, and Christian Laugier. Probabilistic navigation in dynamic environment using rapidly-exploring random trees and gaussian processes. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1056–1062. IEEE, 2008.
- Isaac Gibbs and Emmanuel Candes. Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34:1660–1672, 2021.
- Isaac Gibbs and Emmanuel Candès. Conformal inference for online prediction with arbitrary distribution shifts. *arXiv preprint arXiv:2208.08401*, 2022.
- Lukas Hewing, Kim P. Wabersich, and Melanie N. Zeilinger. Recursively feasible stochastic model predictive control using indirect feedback. *Automatica*, 119:109095, 2020. ISSN 0005-1098. doi: <https://doi.org/10.1016/j.automatica.2020.109095>. URL <https://www.sciencedirect.com/science/article/pii/S0005109820302934>.
- Sarah A. Hummel. *Frisbee flight simulation and throw biomechanics*. University of California, Davis, 2003.
- Samarth Kalluraya, George J. Pappas, and Yiannis Kantaros. Multi-robot mission planning in dynamic semantic environments. *arXiv preprint arXiv:2209.06323*, 2022.
- Henrik Kretschmar, Markus Spies, Christoph Sprunk, and Wolfram Burgard. Socially compliant mobile robot navigation via inverse reinforcement learning. *The International Journal of Robotics Research*, 35(11):1289–1307, 2016.
- Lars Lindemann, Matthew Cleaveland, Gihyun Shim, and George J Pappas. Safe planning in dynamic environments using conformal prediction. *arXiv preprint arXiv:2210.10254*, 2022a.

- Lars Lindemann, Xin Qin, Jyotirmoy V Deshmukh, and George J Pappas. Conformal prediction for stl runtime verification. *arXiv preprint arXiv:2211.01539*, 2022b.
- Rachel Luo, Shengjia Zhao, Jonathan Kuck, Boris Ivanovic, Silvio Savarese, Edward Schmerling, and Marco Pavone. Sample-efficient safety assurances using conformal prediction. *arXiv preprint arXiv:2109.14082*, 2021.
- Keyvan Majd, Shakiba Yaghoubi, Tomoya Yamaguchi, Bardh Hoxha, Danil Prokhorov, and Georgios Fainekos. Safe navigation in human occupied environments using sampling and control barrier functions. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5794–5800. IEEE, 2021.
- Stefan Mitsch, Khalil Ghorbal, and André Platzer. On provably safe obstacle avoidance for autonomous robotic ground vehicles. In *Robotics: Science and Systems IX, Technische Universität Berlin, Berlin, Germany, June 24-June 28, 2013*, 2013.
- Daniel Morgan, Soon-Jo Chung, and Fred Y. Hadaegh. Model predictive control of swarms of spacecraft using sequential convex programming. *Journal of Guidance, Control, and Dynamics*, 37:1–16, 04 2014. doi: 10.2514/1.G000218.
- Siddharth H. Nair, Vijay Govindarajan, Theresa Lin, Chris Meissen, H. Eric Tseng, and Francesco Borrelli. Stochastic mpc with multi-modal predictions for traffic intersections. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pages 635–640, 2022. doi: 10.1109/ITSC55140.2022.9921751.
- Kensuke Nakamura and Somil Bansal. Online update of safety assurances using confidence-based predictions. *arXiv preprint arXiv:2210.01199*, 2022.
- Marco Omainka, Junya Yamauchi, Thomas Beckers, Takeshi Hatanaka, Sandra Hirche, and Masayuki Fujita. Gaussian process-based visual pursuit control with unknown target motion learning in three dimensions. *SICE Journal of Control, Measurement, and System Integration*, 14 (1):116–127, 2021.
- Mike Phillips and Maxim Likhachev. Sipp: Safe interval path planning for dynamic environments. In *2011 IEEE International Conference on Robotics and Automation*, pages 5628–5635. IEEE, 2011.
- Xin Qin, Yuan Xian, Aditya Zutshi, Chuchu Fan, and Jyotirmoy V Deshmukh. Statistical verification of cyber-physical systems using surrogate models and conformal inference. In *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPS)*, pages 116–126. IEEE, 2022.
- Venkatraman Renganathan, Iman Shames, and Tyler H Summers. Towards integrated perception and motion planning with distributionally robust risk constraints. *IFAC-PapersOnLine*, 53(2): 15530–15536, 2020.
- Venkatraman Renganathan, Sleiman Safaoui, Aadi Kothari, Benjamin Gravell, Iman Shames, and Tyler Summers. Risk bounded nonlinear robot motion planning with integrated perception & control. *arXiv preprint arXiv:2201.01483*, 2022.

- Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3), 2008.
- Troy Shinbrot, Celso A Grebogi, Jack Wisdom, and James A Yorke. Chaos in a double pendulum. In *American Journal of Physics, American Association of Physics Teachers*. 1992.
- Andrew Singletary, Aiden Swann, Ivan Dario Jimenez Rodriguez, and Aaron D. Ames. Safe drone flight with time-varying backup controllers, 2022. URL <https://arxiv.org/abs/2207.05220>.
- Floris Takens. Detecting strange attractors in turbulence. In *Dynamical systems and turbulence*, pages 366–381. Springer, 1981.
- Herbert G Tanner, Savvas G Loizou, and Kostas J Kyriakopoulos. Nonholonomic navigation and control of cooperating mobile manipulators. *IEEE Trans. Robotics and Automation*, 19(1):53–64, 2003.
- Dimos V. Dimarogonas *et al.* A feedback stabilization and collision avoidance scheme for multiple independent non-point agents. *Automatica*, 42(2):229–243, 2006.
- Antony Thomas, Fulvio Mastrogiovanni, and Marco Baglietto. Probabilistic collision constraint for motion planning in dynamic environments. In *International Conference on Intelligent Autonomous Systems*, pages 141–154. Springer, 2021.
- Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. *Advances in neural information processing systems*, 32, 2019.
- Jesus Tordesillas, Brett T. Lopez, Michael Everett, and Jonathan P. How. Faster: Fast and safe trajectory planner for navigation in unknown environments, 2020. URL <https://arxiv.org/abs/2001.04420>.
- Peter Trautman and Andreas Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 797–803. IEEE, 2010.
- Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*. Springer Science & Business Media, 2005.
- Allan Wang, Christoforos Mavrogiannis, and Aaron Steinfeld. Group-based motion prediction for navigation in crowded environments. In *Conference on Robot Learning*, pages 871–882. PMLR, 2022.
- Skylar X. Wei, Anushri Dixit, Shashank Tomar, and Joel W. Burdick. Moving obstacle avoidance: A data-driven risk-aware approach. *IEEE Control Systems Letters*, 7:289–294, 2022. doi: 10.1109/LCSYS.2022.3181191.
- Youngmin Yoon, Changhee Kim, Jongmin Lee, and Kyongsu Yi. Interaction-aware probabilistic trajectory prediction of cut-in vehicles using gaussian process for proactive control of autonomous vehicles. *IEEE Access*, 9:63440–63455, 2021. doi: 10.1109/ACCESS.2021.3075677.

Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *International Conference on Machine Learning*, pages 25834–25866. PMLR, 2022.