Scalable Safe Long-Horizon Planning in Dynamic Environments Leveraging Conformal Prediction and Temporal Correlations

Sander Tonkens*, Sophia Sun*, Rose Yu, and Sylvia Herbert

***Abstract*— Safe and effective planning in cluttered and diverse scenes that include the presence of other agents, requires a robot that is (1) equipped with a state-of-the-art autonomy stack and (2) offers strict guarantees on its failure rate. We expand upon prior work in planning with probabilistic safety guarantees using conformal prediction, a distribution-free uncertainty quantification technique. Specifically, to have high performance and provide guarantees for long horizon planning, we use copulas to model the temporal correlations of an agent's uncertainty along a trajectory. Additionally, we highlight the versatility of conformal prediction and its potential for safe planning and control. Our proposed algorithm improves task performance (average distance to target) by 40% compared to current methods, while retaining strict safety guarantees.**

I. Introduction

Deploying autonomus systems in dense, unstructured, and uncertain environments requires systematic reasoning of uncertainty [1]. Recent advances in robotics & machine learning have enabled the deployment of autonomous systems in a variety of domains, from automating warehouses to space exploration. However, these domains are intentionally very structured and devoid of humans. In contrast, many promising future application domains for autonomous systems, ranging from last-mile ground and aerial deliveries to autonomous driving in cities, involve environments with much more complexity and uncertainty.

In this work, we explore planning algorithms that are robust to uncertainty arising from both the environment dynamics and from upstream models in the autonomy stack. Specifically, we utilize Conformal Prediction [2], [3], a distribution-free uncertainty quantification (UQ) method that (1) makes no assumptions about the prediction model or the underlying data distribution, and (2) provides finite-sample statistical guarantees. Existing UQ methods, in contrast, do not provide statistical guarantees, or produce error bounds that are too conservative to be useful, which prohibits their use in safety-critical applications. By incorporating provably valid UQ methods for trajectory prediction, we develop probabilistically safe, yet not overly conservative, long-horizon motion plans in cluttered multi-agent environments.

A. Related Work

Historically, the control under uncertainty literature assumes that uncertainties are either (1) bounded, or (2) modeled by common distributions, e.g. Gaussian. The former case leads to robust control techniques that assume worst-case uncertainty [4], whereas the latter is typically paired with chance-constraint optimal control [5]. Both of these assumptions are typically invalid and/or overly conservative in dynamic and cluttered environments. Existing works on probabilistically safe planning commonly use Gaussian processes [6] or consider uncertainty in a model's parameters using a Bayesian framework [5]–[8]. In realistic settings, the parameters of the distributions are approximately fit to the data and the model which jointly quantify the uncertainty (over e.g., the environment, the system's dynamics, or the system's state). Hence, its safety guarantees hinge on strong assumptions on the underlying data distribution.

In contrast, the properties of conformal prediction (CP) makes it a particularly useful technique for safety assurance. These properties notably include (1) being agnostic to the underlying (prediction) model, (2) being agnostic to the underlying data distribution, and (3) providing finite sample guarantees. Additionally, it can calibrate heuristic uncertainty metrics provided by the prediction model for better calibrated UQ of models. The robotics community has started exploring conformal prediction, to bound tracking errors [9], for vision-based robot control [10], and for robust optimization [11].

In the context of safe planning and control, *Muthali et. al.* [12] use CP-calibrated quantile regression to provide probabilistically safe control actions leveraging reachability analysis [13]. Closer to our work, *Lindemann et. al.* [14] use CP to generate prediction sets of agents moving in an environment, and then use these sets for planning with model predictive control (MPC). However, planning over long time horizons with probabilistic safety requires providing a coverage guarantee over *all* future timesteps. Distribution-free methods, including [14], often resort to union bounding [15] which leads to very conservative prediction regions. Recently, tighter bounds have been obtained leveraging Copulas [16] and linear complementarity programming [17].

B. Contributions

We propose a planning framework with valid probabilistic safety assurances. We leverage tools from conformal prediction to obtain valid prediction regions that quantify the uncertainty of trajectory predictions over multiple timesteps. Notably, we account for the temporal correlations by considering the timesteps jointly using copulas. We provide extensive experiments that demonstrate that accounting for temporal correlations leads to much tighter, yet valid, uncertainty estimates, enabling less-conservative long-horizon planning.

II. Methods

A. Problem Setup

Let $\mathcal{D} = \{z^i = (x^i, y^i)\}_{i=1}^n$ be a dataset with input $x_i \in \mathcal{X}$ and output $y^i \in \mathcal{Y}$ such that each data point $z^i \in \mathcal{Z} := \mathcal{X} \times \mathcal{Y}$ is drawn i.i.d. from an unkonwn distribution $\mathcal{Z}$.

The authors are with the University of California, San Diego. {sander, shs066, roseyu, sherbert}@ucsd.edu. * contributed equally.

Specifically, for time series we consider $z^i = (x^i_{1:m}, y^i_{1:k})$, where $x^i_{1:m} \in \mathbb{R}^{m \times d_x}$ is m input time steps of dimension d_x, and $y^i_{1:k} \in \mathbb{R}^{t \times d_y}$ is k prediction time steps of dimension d_y. Commonly in time series forecasting x and y describe the same variable and given the previous m states we predict the k future states. We will use superscript x^i to index samples of $\mathcal{D}$, and subscript x_t to index time steps.

We consider an ego-agent with discrete-time dynamics of the form $s_{t+1} = f(s_t, a_t)$ with $s_t \in \mathbb{R}^{n_s}$ the state and $a_t \in \mathbb{R}^{n_a}$ the action (or control). We define the observation variables $o_t = g(s_t) \in \mathbb{R}^{d_y}$ to be in the observation space (which describes the same quantity as the model output space) and g a mapping from the state space to observation space. For example, in autonomous navigation, the subset of the state describing the 2D position corresponds to the observation variables (with g the position mask). We are interested in planning under constraints and hence define a time-varying obstacle region $\mathcal{O}^{\text{obs}}_t$. Then, given a desired confidence level $1-\epsilon$, a sequence of observations $o \in \mathbb{R}^{N \times d_y}$ is $1-\epsilon$ probabilistically safe if

$$\mathbb{P}\left(\bigwedge_{t=1}^{N} o_t \notin \mathcal{O}^{\text{obs}}_t\right) \geq 1 - \epsilon, \tag{1}$$

with N the problem duration. Note that this probability has to hold jointly over all time steps to be valid.

B. Conformal Prediction

We will briefly present the algorithm and theoretical results for conformal prediction, and refer readers to [18] for a thorough introduction. The goal of conformal prediction is to produce a *valid* confidence region (Def. 1) for any underlying prediction model.

Definition 1 (Validity)**.** *Given a new data point (x, y) and a desired confidence $1-\epsilon \in (0,1)$, the confidence region $\Gamma^{1-\epsilon}(x)$ is a subset of $\mathcal{Y}$ containing probable outputs $\tilde{y} \in \mathcal{Y}$ given x. The region $\Gamma^{1-\epsilon}$ is valid if*

$$\mathbb{P}(y \in \Gamma^{1-\epsilon}(x)) \geq 1 - \epsilon \tag{2}$$

We begin the algorithm by splitting the dataset into a proper training set $\mathcal{D}_{train}$, a calibration set $\mathcal{D}_{cal}$, and a test set $\mathcal{D}_{test}$. A prediction model $h : \mathcal{X} \to \tilde{\mathcal{Y}}$ is trained on $\mathcal{D}_{train}$. Note that the prediction space $\tilde{\mathcal{Y}}$ can be the same as the output space $\mathcal{Y}$, or can contain auxiliary information, such as a heuristic measure of uncertainty.

Next, we pick a *nonconformity score* $R : \mathcal{Z}^{|\mathcal{D}_{train}|} \times \mathcal{Z} \to \mathbb{R}$ that quantifies how well a data sample from calibration *conforms* to the training dataset. Typically, we leverage the prediction model $h(x)$ and choose a metric of disagreement between the prediction and the true label as the non-conformity score, such as the Euclidean distance.

$$R(\mathcal{D}_{train}, (x, y)) \overset{\text{e.g.}}{=} d(y, h(x)) \overset{\text{e.g.}}{=} \|y - h(x)\|_2$$

For consiseness, we write $R(\mathcal{D}_{train}, (x^i, y^i))$ as $R(z^i)$ in rest of the paper. Let $\mathcal{R}_{cal}$ denote the set of nonconformity scores of all data in $\mathcal{D}_{cal}$.

Given a new test data sample $z' = (x', y')$ and a target confidence level $1 - \epsilon \in (0, 1)$, conformal prediction constructs the confidence regions as follows:

$$\Gamma^{1-\epsilon}(x') := \{y : d(h(x'), y) \leq Q(1-\epsilon, \mathcal{R}_{cal} \cup \{\infty\})\} \tag{3}$$

with $Q(p, \mathcal{R})$ is the quantile function that finds the p-quantile in a set of scalars $\mathcal{R}$. We call a data point *covered* when the confidence region contains the true label: $y^i \in \Gamma^{1-\epsilon}(x^i)$.

The confidence region obtained above with conformal prediction is provably valid if the data sample z is exchangeable with $\mathcal{D}_{cal}$ [2]:

Definition 2 (Exchangeability)**.** *In a dataset $\{z^1, z^2, \ldots, z^n\}$ of size n, any of its $n!$ permutations are equally probable.*

Importantly, this statistical guarantee holds regardless of the underlying prediction model and choice of nonconformity score.

The algorithm introduced above is known as *inductive* or *split* conformal prediction [2], [19]; it is commonly used for uncertainty quantification for machine learning models [18], because of its flexibility and computational efficiency.

C. Copula Conformal Prediction for multistep time series

To leverage CP for multi-step safe planning, a coverage guarantee over *all* future time steps is desired. Formally, for a prediction horizon of length k, we want to produce k confidence regions $\Gamma^{1-\epsilon_t}_t(x)$ for $t = 1, \ldots, k$ such that

$$\mathbb{P}[\, y_t \in \Gamma^{1-\epsilon_t}_t(x) \; \forall t \in \{1, \ldots, k\}] \geq 1 - \epsilon. \tag{4}$$

Note that $\epsilon_t \neq \epsilon$. However, we can naively apply Boole's inequality and obtain $\epsilon_t = \epsilon/k \; \forall t$. This is referred to as union bounding [15], in which finding the $1-\epsilon/k$ confidence region for each time step ensures the joint coverage rate over k future predictions to be greater than $1 - \epsilon$. The resulting union bound is a worst-case bound that assumes the residuals of each time step of a trajectory to be independent, which is rarely the case. Consequently, this approach results in very large confidence regions, especially for long prediction horizons (large k) or for multivariate data (e.g. 2D coordinates), which poses challenges for downstream planning.

We propose mitigating this challenge by incorporating copulas [16]. A copula (Def. 3) is a function that combines marginal probabilities of multiple random variables into a joint cumulative distribution function (CDF).

Definition 3 (Copula)**.** *Consider a random vector $(X_1, \ldots X_k)$. We denote the marginal CDF for each variable X_t, $t \in \{1, \ldots, k\}$ as*

$$F_t(x) = \mathbb{P}[X_t \leq x]$$

The copula of $(X_1, \ldots X_k)$, written as $C : [0,1]^k \to [0,1]$, is defined as the joint CDF of $(F_1(X_1), \ldots, F_k(X_k))$:

$$C(u_1, \ldots, u_k) = \mathbb{P}\left[F_1(X_1) \leq u_1, \ldots, F_t(X_k) \leq u_k\right]$$

The copula function C captures the correlation between distributions of the variables $X_1, \ldots, X_k$. Sklar's theorem [16, Theorem 1] proves the existence of a copula function for any arbitrary collection of variables.

We will outline our algorithm for estimating the copula and generating valid confidence regions for the entire trajectory. We divide $\mathcal{D}_{cal}$ further into equal-sized $\mathcal{D}_{cal,1}$ and $\mathcal{D}_{cal,2}$. For each time step $t = 1, \dots, k$, we first estimate an empirical CDF on the nonconformity scores:

$$\hat{F}_t(r) = \mathbb{P}_{z \in \mathcal{D}_{cal,1}}[R(z) \leq r] = \frac{\sum_{z_t^i \in \mathcal{D}_{cal,1}} \mathbb{1}_{R(z_t^i) \leq r}}{|\mathcal{D}_{cal,1}|} \quad (5)$$

With the empirical CDFs, we can calculate the marginal probability vectors $\mathbf{u}^i = (\hat{F}_1(R(z_1^i)), \dots, \hat{F}_t(R(z_k^i))) \in [0,1]^k$ for all $z^i \in \mathcal{D}_{cal,2}$. To model the joint probability, we estimate an empirical copula [20] as follows:

$$C_{\text{empirical}}(u_1, \dots, u_t) = \frac{1}{|\mathcal{D}_{cal,2}|} \sum_{z_t^i \in \mathcal{D}_{cal,2}} \prod_{t=1}^{k} \mathbb{1}_{\mathbf{u}_t^i < u_t}. \quad (6)$$

Lastly, we can use a search algorithm to find

$$\underset{r_1,\dots,r_k}{\text{argmin}} \quad C_{\text{empirical}}(\hat{F}_1(r_1), \dots, \hat{F}_1(r_k)) \geq 1 - \epsilon. \quad (7)$$

The thresholds $r_1, \dots, r_k$ are then used to generate confidence sets for planning. We refer readers to [16] for details on the copula CP algorithm and proof of its validity. By estimating a copula, we leverage the correlation between the time steps to produce sharper confidence regions that still guarantees coverage over the full prediction horizon.

D. Planning

To highlight the performance of our approach, we consider long-horizon open-loop planning tasks. We assume we have access to a prediction model that at minimum provides a point-estimate of future agent trajectories in a scene over the full horizon. In practice, it is common to leverage past trajectories (which we assume we have observed) to predict future trajectories. A popular approach that provides safety assurances is trajectory optimization, which casts the task as an optimal control problem (OCP). Critically, for autonomy applications, we require algorithms that guarantee satisfaction of all constraints over an entire trajectory with high probability. To provide probabilistic safety assurances, a popular approach uses chance constraints, and requires solving a chance-constrained OCP [21], which are typically much harder to solve than deterministic OCPs. However, using conformal prediction enables transforming a chance-constrained OCP into a deterministic OCP by planning a robust path with respect to the CP-generated confidence sets. Specifically, based on the model prediction (which includes the predicted trajectory and optionally a heuristic uncertainty metric) we compute the Γ_t^ϵ confidence set for each agent over all time steps t and include a constraint to avoid intersection with the confidence sets. In particular, we solve the following optimization problem:

$$\begin{aligned} \underset{s_{1:N+1}, a_{1:N}}{\text{min.}} \quad & J(s_{1:N+1}, a_{1:N}), \\ \text{s.t.} \quad & s_{t+1} = f(s_t, a_t) \qquad t \in \{1, \dots, N\}, \\ & s_1 = s_{\text{init}}, \quad s_{N+1} \in \mathcal{S}_{\text{final}}, \\ & a_t \in \mathcal{A}, \ s_t \in \mathcal{S}, \ g(s_t) \notin \mathcal{O}_t^{\text{obs}} \ \ t \in \{1, \dots, N\}, \end{aligned} \quad (8)$$

with $J(s_{1:N+1}, a_{1:N})$ the objective of the ego agent, $\mathcal{A}$ the action constraints, and $\mathcal{S}$ the state constraints. s_{init} denotes the starting state and $\mathcal{S}_{\text{final}}$ denotes the desired final state region. Our complete pipeline is outline in Algorithm 1.

Algorithm 1 Safe Planning with Copula CP UQ

Input: Dataset $\mathcal{D} = \{z^i\}_{i=1}^n$, target significant level $1-\epsilon$, Test sample $x_{1:m} \sim \mathcal{D}_{test}$, initial state s_{init}, $\mathcal{S}_{\text{final}}$

```
// Copula Conformal Prediction
Split dataset D into D_train, D_cal,1, and D_cal,2.
Train prediction model h on D_train.
F̂_1, ..., F̂_k ← Eq. (5)
C_empirical(·) ← Eq. (6)
r_1, ..., r_k ← Eq. (7)
// Planning with Confidence Sets
ỹ_{1:k} ← h(x_{1:m})
for t = {1, ... N} do
    O_t^obs ← {y : d(y, ỹ_t) ≤ r_t}
end for
a_1, ..., a_N ← Eq. (8)
Apply a_1, ... a_N to ego-agent
```

III. Experiments

For our experiments we use the TrajNet++ dataset (Update 4.0) [22], which is a compiled set of pedestrian trajectories captured in both indoor and outdoor locations such as in universities, hotels, retail stores, and train stations. The samples are 21 time steps long with a data frequency of 2.5Hz. Our algorithm predicts 2D spatial positions for each pedestrian over the last $k = 12$ time steps given the first $m = 9$. The dataset contains $240,896$ samples, which we split $70/10/20$ into train, calibration, and test sets.

In this work, we do not model interactions between the ego-agent and other agents, and hence implicitly assume that the ego-agent's motion will not affect the pedestrian's movement. This is a standard assumption considered in prior work [12], [14]. We leave joint modeling of human-robot motion to future work.

We consider the following baseline comparisons: (1) Certainty equivalence (CE), in which we only leverage the mean point prediction and do not account for uncertainty, (2) the probabilistic output of the SocialLSTM [23] model, which provides a multi-variate Gaussian (for which we take the $1-\epsilon_h$ confidence sets), and (3) union bounding CP as in [14].

A. Results

To quantitatively compare different UQ methods, we compute the empirical coverage over the entire prediction horizon (see Eq. 4) and the size of the confidence area predicted summed over a trajectory. These metrics correspond to validity and sharpness respectively. The results are shown in Table I. By leveraging the temporal correlations over a trajectory we see that our method (CopulaCP) provides a sharper UQ estimate (50% lower area) than union bounding CP while maintaining validity. Additionally, we observe that the uncertainty obtained from the multi-variate Gaussian estimate of the SocialLSTM model is overly confident, producing small regions that do not provide adequate coverage.

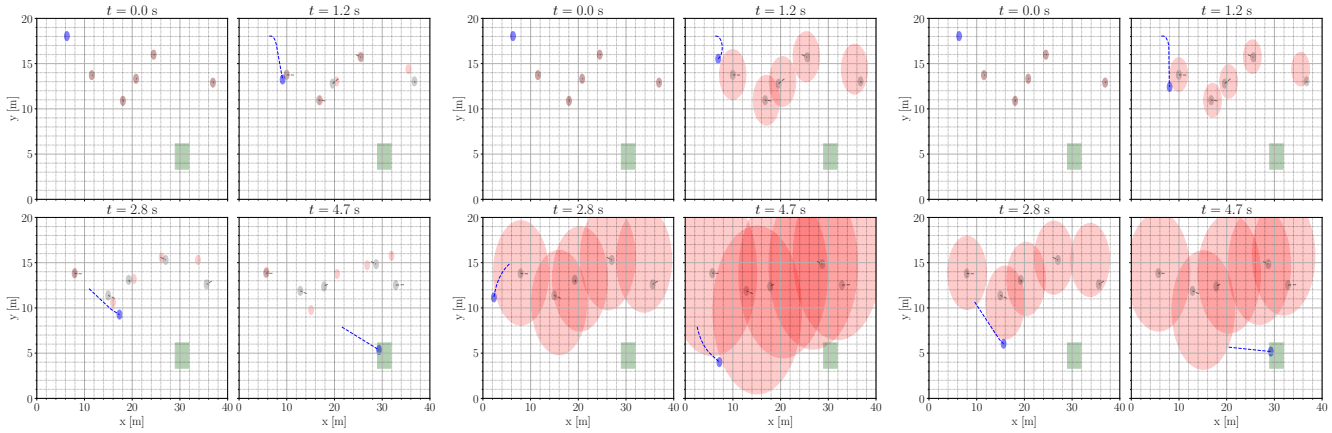

(a) Certainty equivalent planning. **(b)** Planning with union bounding CP UEs. **(c)** Planning with Copula CP UEs.

Fig. 1: A scene with uncertainty bounds for the other agents in red, the actual position of other agents and their trajectory in dark gray, the ego agent's position and its trajectory in blue and the goal region in green. We observe the CE baseline crashing, the union bounding CP unable to reach the goal due to conservative uncertainty estimates (UEs), and the copula CP method reaching the goal successfully.

TABLE I: Comparison of Uncertainty Quantification methods. Methods that are *invalid* (coverage below 90%) are greyed out. Compared to baseline methods, CopulaCP achieves high level of calibration (coverage is close to 90%) while producing sharper confidence regions.

Model	SocialLSTM	CP	CopulaCP
Coverage [%] (90%)	2.86	98.4	92.6
Sum of Area [m^2] (↓)	21.5	403.8	204.2

TABLE II: Collision rate and mean distance to the goal over 1000 test trajectories with 5 agents per scene. We empirically observe that CE and the SocialLSTM planners do not provide probabilistic safety, while CP and CopulaCP do. Additionally, our proposed approach improves planning performance by 40%.

Model	CE	SocialLSTM	CP	CopulaCP
Collision rate [%]	37.7	30.0	0.8	1.1
Distance to goal at $t = N$ [m]	0.6	0.6	9.5	5.8

Next, we utilize the uncertainty estimates to quantify the obstacle sets $\mathcal{O}_{\text{obs}}$. We compare CopulaCP to the aforementioned baselines when they are used for safe planning with time-varying obstacles. We consider the collision rate (between the ego agent and the actual future agent trajectories $y \sim \mathcal{D}_{test}$) and the distance to the goal region $\mathcal{S}_{\text{final}}$ as evaluation metrics, see Table II. As expected, the CE approach and the SocialLSTM baselines do not provide probabilistic safety assurances (higher than 10% collision rate). We visualize the open-loop planning solution for all 4 methods for a descriptive test scene and the uncertainty bounds in Figure 1. We observe that the CE baseline collides, that the uncertainty estimates are too large using union bounding CP and that only CopulaCP reaches the goal safely.

B. Discussion

The usefulness of the prediction sets obtained with conformal prediction is strongly reliant on leveraging problem structure and model characteristics. In this work, we tackle the former by modeling the temporal correlations over the length of a trajectory prediction. Beyond this, obtaining sharp uncertainty estimates is driven by the choice of nonconformity score, as it aims to encode all information about the problem and the data [18].

One popular technique uses quantile regression, as in e.g., [12], but this requires pre-specifying the quantile to regress to and incorporating quantile loss in the loss function for the prediction module, which makes it less amenable to using large pre-trained models. Instead, a popular choice is to design the non-conformity score to incorporate an uncertainty metric from the model directly (or through an ensemble of models). For example, if the model outputs a heatmap over the possible output values, one can choose the nonconformity score to be the peak probability of the model output multiplied by the estimation error [24]. Alternatively, a popular choice is a measure of entropy, such as the (generalized) variance over an ensemble of models or the predicted variance of a Bayesian model.

As previously discussed, the majority of state-of-the art trajectory prediction models predict a distribution of potential future trajectories [25], e.g., a multi-variate Gaussian or mixture of Gaussians (GMM). The distributions usually overfit dramatically due to the nature of model training (see [26], [27], and as demonstrated in Table I), hence are unfit to be used as UQ directly. However, they can be used as heuristics to improve distribution free calibration methods such as CP. We foresee that, e.g., decoupling longitudinal and lateral movement and uncertainty estimates can greatly improve the UQ predictions to provide tighter, yet still valid bounds. In turn, this reduces the conservativeness of the downstream planning model that leverages the CPs.

C. Future work

We plan on leveraging better trajectory prediction models that provide richer notions of uncertainty, such as Trajectron++ [25]. Additionally, we plan to model the spatial correlations and interaction between agents to provide sharper UQ estimates. To scale to more complex tasks, we plan on using receding horizon planning approaches as an uncertainty reduction tool. To ensure successive iterations of the planner and solvable and safe, i.e. persistent feasibility, we will rely on a library of fallback strategies for the ego agent. Additionally, quantifying the effect of future ego-actions on other agents and the environment is a key enabler for scaling to more complex tasks.

REFERENCES

[1] N. Roy, I. Posner, T. Barfoot, P. Beaudoin, Y. Bengio, J. Bohg, O. Brock, I. Depatie, D. Fox, D. Koditschek *et al.*, "From machine learning to robotics: challenges and opportunities for embodied intelligence," *arXiv preprint arXiv:2110.15245*, 2021.

[2] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic learning in a random world*. Springer Science & Business Media, 2005.

[3] J. Lei, M. G'Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman, "Distribution-free predictive inference for regression," *Journal of the American Statistical Association*, vol. 113, no. 523, pp. 1094–1111, 2018.

[4] K. Zhou, J. C. Doyle, and K. Glover, *Robust and optimal control*. Prentice Hall, 1996.

[5] T. Lew, A. Sharma, J. Harrison, A. Bylard, and M. Pavone, "Safe active dynamics learning and control: A sequential exploration-exploitation framework," *IEEE Transactions on Robotics*, 2022, in Press.

[6] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-based model predictive control: Toward safe learning in control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, no. 1, pp. 269–296, 2020.

[7] J. F. Fisac, A. 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," in *Robotics: Science and Systems*, 2018.

[8] V. Dhiman, M. J. Khojasteh, M. Franceschetti, and N. Atanasov, "Control barriers in bayesian learning of system dynamics," *IEEE Transactions on Automatic Control*, 2021.

[9] D. Sun, S. Jha, and C. Fan, "Learning certified control using contraction metric," *arXiv preprint arXiv:2011.12569*, 2020.

[10] A. Farid, D. Snyder, A. Z. Ren, and A. Majumdar, "Failure prediction with statistical guarantees for vision-based robot control," *arXiv preprint arXiv:2202.05894*, 2022.

[11] C. Johnstone and B. Cox, "Conformal uncertainty sets for robust optimization," in *Conformal and Probabilistic Prediction and Applications*. PMLR, 2021, pp. 72–90.

[12] A. Muthali, H. Shen, S. Deglurkar, M. H. Lim, R. Roelofs, A. Faust, and C. Tomlin, "Multi-agent reachability calibration with conformal prediction," *arXiv preprint arXiv:2304.00432*, 2023.

[13] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-Jacobi reachability: A brief overview and recent advances," in *Proc. IEEE Conf. on Decision and Control*, 2017.

[14] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, "Safe planning in dynamic environments using conformal prediction," *arXiv preprint arXiv:2210.10254*, 2022.

[15] K. Stankevičiūtė, A. Alaa, and M. van der Schaar, "Conformal time-series forecasting," in *Advances in Neural Information Processing Systems*, 2021.

[16] S. Sun and R. Yu, "Copula conformal prediction for multi-step time series forecasting," *arXiv preprint arXiv:2212.03281*, 2022.

[17] M. Cleaveland, I. Lee, G. J. Pappas, and L. Lindemann. (2023) Conformal prediction regions for time series using linear complementarity programming. Available at https://arxiv.org/abs/2304.01075.

[18] A. N. Angelopoulos and S. Bates, "A gentle introduction to conformal prediction and distribution-free uncertainty quantification," *arXiv preprint arXiv:2107.07511*, 2021.

[19] J. Lei and L. Wasserman, "Distribution free prediction bands," *arXiv preprint arXiv:1203.5422*, 2012.

[20] L. Ruschendorf, "Asymptotic distributions of multivariate rank order statistics," *The Annals of Statistics*, pp. 912–923, 1976.

[21] L. Blackmore, M. Ono, and B. C. Williams, "Chance-constrained optimal path planning with obstacles," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1080–1094, 2011.

[22] P. Kothari, S. Kreiss, and A. Alahi, "Human trajectory forecasting in crowds: A deep learning perspective," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 7386–7400, 2021.

[23] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social lstm: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 961–971.

[24] H. Yang and M. Pavone, "Object pose estimation with statistical guarantees: Conformal keypoint detection and geometric uncertainty propagation," *arXiv preprint arXiv:2303.12246*, 2023.

[25] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *European Conf. on Computer Vision*, 2020.

[26] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *Advances in neural information processing systems*, vol. 30, 2017.

[27] A. Alaa and M. van der Schaar, "Frequentist uncertainty in recurrent neural networks via blockwise influence functions," in *ICML*, 2020.