

Conformal Decision Theory: Safe Autonomous Decisions from Imperfect Predictions

Jordan Lekeufack^{1,*} Anastasios N. Angelopoulos^{2,*} Andrea Bajcsy^{3,*} Michael I. Jordan^{1,2,**} Jitendra Malik^{2,**}




Fig. 1: Robot planner using a conformal controller on the Stanford Drone Dataset [1]. The future trajectories of humans are predicted online by a machine learning algorithm (not visualized). The robot planner finds an optimal spline through the scene and is penalized for being close to humans. This penalty is proportional to a conformal control variable, λ_t , which is adjusted online by the conformal controller so the average distance from a human is no less than two meters. The orange, red, and blue curves are the robot trajectory with different planners: the conformal controller, an aggressive planner with $\lambda = 0$ (i.e., no reward for avoiding humans), and a conservative planner with a large negative value of λ (i.e., a large reward for avoiding humans). The darkness of the lines indicates the passage of time. Illustrative pedestrian trajectories are plotted as arrows; only the yellow pedestrians affect the spline planner. Details in Section V-A and videos on [project website](#)[†].

Abstract— We introduce *Conformal Decision Theory*, a framework for producing safe autonomous decisions despite imperfect machine learning predictions. Examples of such decisions are ubiquitous, from robot planning algorithms that rely on pedestrian predictions, to calibrating autonomous manufacturing to exhibit high throughput and low error, to the choice of trusting a nominal policy versus switching to a safe backup policy at run-time. The decisions produced by our algorithms are safe in the sense that they come with provable statistical guarantees of having low risk without any assumptions on the world model whatsoever; the observations need not be I.I.D. and can even be adversarial. The theory extends results from conformal prediction to calibrate decisions directly, without requiring the construction of prediction sets. Experiments demonstrate the utility of our approach in robot motion planning around humans and robot manufacturing.

I. INTRODUCTION

Autonomous systems increasingly rely on complex learned models to supply predictions that are the basis for decision-making. Self-driving cars rely on deep neural networks [2]–[5] to plan paths around nearby pedestrians, robotic

manipulators leverage learned grasp models [6] to plan high-throughput pick-and-place maneuvers in factories, and AI-enabled trading agents optimize the financial future of investors [7]. There is a conceptual gap between prediction and decision-making, and it remains a challenge to ensure that systems make *good decisions* despite *imperfect predictions*.

One increasingly popular strategy is to quantify the uncertainty in the predictions independently of their downstream effect on the decision via conformal prediction (CP) [8]–[13]. This approach has become popular because, when used to provide simultaneous prediction sets on all outcomes, conformal prediction provides statistical guarantees of safe autonomous behavior without any assumption on the underlying distribution or model. This application of CP has shown impact in robot navigation [14]–[17], early warning systems [18], out-of-distribution detection [19], [20], probabilistic pose estimation [21], and for large language models [22]. However, the requirement of simultaneous coverage is challenging to satisfy and for many decision systems is excessive. What if we could provide statistical guarantees, as in CP, *directly* on our decisions, bypassing the need to construct prediction sets?

This work presents Conformal Decision Theory, a theoretical and algorithmic framework that unifies predictive uncertainty and safe decision-making. Our key idea is

instead of calibrating prediction sets for coverage, we directly calibrate decisions for low risk.

¹Department of Statistics, UC Berkeley. ²Department of Electrical Engineering and Computer Science, UC Berkeley. Emails: {jordan.lekeufack, angelopoulos, michael.jordan, malik}@berkeley.edu ³Robotics Institute, Carnegie Mellon University. Email: {abajcsy}@cmu.edu * Equal contribution. ** Equal senior authorship.

[†] <https://conformal-decision.github.io>

Our main algorithmic innovation is a class of algorithms called *conformal controllers*. A conformal controller starts with a conformal control variable, λ_t , which determines the decision-maker’s conservatism or aggressiveness. Then, it dynamically adjusts λ_t to balance risk and performance in such a way that guarantees a low risk. The main practical benefit of this approach is its *emergent ability to ignore irrelevant uncertainty*, only accounting for that which *affects decisions*. This can be much less conservative than the prediction-set strategy. For example, in Figure 1, the planner only considers the humans that pose a collision risk.

The contributions of this paper are threefold:

- We introduce Conformal Decision Theory, the idea of directly calibrating decisions with conformal controllers. This extends the line of work in online adversarial conformal prediction [10], [13], [23], [24] to the decision-making setting.
- We prove finite-time risk bounds for conformal controllers. Even when applied to prediction sets, these results are stronger than any previously known results for online adversarial conformal prediction.
- We show the utility of the framework in three simulations where Conformal Decision Theory is applied to robot navigation: the Stanford Drone Dataset [1], a stock trading simulation, and a robot manufacturing example.

The main potential impact of this work is to broaden the scope of conformal prediction. Our methods are more appropriate for disciplines that focus on decision-making, such as control theory, reinforcement learning, and logistics. In these disciplines, algorithms are ultimately evaluated by the decisions, not the predictions, that they make. Furthermore, there are many settings where it does not make sense to construct prediction sets, and our technique can provide a distribution-free outlook for such problems (see, e.g., Section V-B).

II. RELATED WORK

Decision-Making Under Predictive Uncertainty. Within the machine learning and statistics community, uncertainty quantification of prediction models has been studied widely, from conformal prediction to Bayesian neural networks to ensembles [25]–[29]. Instead of focusing on prediction calibration alone, the controls and optimization community have coupled prediction uncertainty with safe (i.e., risk bounded) decision-making via chance-constrained optimal control [30], [31] and scenario optimization [32], [33]. The former typically constructs prediction sets that are used as constraints while the latter safeguards against samples drawn from the prediction model. Instead of directly calibrating the output of upstream prediction modules or solving decision-making problems under probabilistic constraints, this work presents a theoretical and algorithmic approach to tuning the robot’s decision risk directly as a function of historical decision-making performance.

Online Learning & Nonstochastic Control. The method herein is reminiscent of online learning, and specifically

the online gradient descent (OGD) update of [34]. The connection is most apparent when examining the forthcoming Equation (5) with $\ell_t = \mathbf{1}\{y_t \in \mathcal{C}(x_t)\}$; this recovers the ACI algorithm of [10], which is OGD on the quantile loss [29]. However, the update in (5) is substantially more general because it incorporates arbitrary decision rules, and reframing it as OGD on an analytic loss function is generally impossible. Furthermore, the guarantees in [10] are not to our knowledge recoverable by existing regret analyses from online convex optimization and nonstochastic control, e.g., [35]–[37]. However, the guarantees do share a retrospective flavor, in that, like regret analyses, they provide guarantees on average over the observed history.

III. CONFORMAL DECISION THEORY

Conformal Decision Theory (CDT) is an approach for calibrating an agent’s decisions to achieve statistical guarantees for the realized average loss of those decisions. Consider a decision-making agent whose input space is \mathcal{X} and action space is \mathcal{U} . In our running example of robot navigation, $x_t \in \mathcal{X}$ captures the current state of the robot, the current scene information (e.g., environment geometry), and the agent information (e.g., pedestrian predictions) while $u_t \in \mathcal{U}$ is the action that the ego vehicle plans at the current time t . At time t , the agent has access to a family of *decision functions*

$$\mathcal{D}_t := \{D_t^\lambda : \mathcal{X} \rightarrow \mathcal{U}, \lambda \in \mathbb{R}\}, \quad (1)$$

parameterized by λ , which we call a *conformal control variable*. One should think of λ as indexing the decisions from least to most conservative. In Figure 1, \mathcal{D}_t is the set of dynamically feasible splines at time t , λ is the coefficient of the reward term for avoiding humans, and D_t^λ is the spline maximizing the total reward given λ .

Assessing the quality of an agent’s decision depends on a space of *targets* \mathcal{Y} . Importantly, the realizations of these targets are *unknown* at the time of the decision; the agent only observes them at deployment time, after decisions are made, and in an online fashion. For example, the robot in Figure 1 does not know the true future state of nearby pedestrians; at any current time t , it only knows the (potentially erroneous) pedestrian predictions. In this example, \mathcal{Y} is the space of pedestrian states (e.g., 2D positions) and $y_t \in \mathcal{Y}$ is the *true* state that the pedestrian moves to at time t .

Mathematically, the quality of the decision-making is quantified by a *loss function* $\mathcal{L} : \mathcal{U} \times \mathcal{Y} \rightarrow [0, 1]$.¹ Often, the loss is more likely to be large when aggressive decisions are taken—i.e., when λ is large. Aggressive decisions can be unsafe, but taking λ too small yields conservative and under-performing decisions.

We seek an algorithm for adapting λ_t (and thus the corresponding decision D_t^λ) at each time step such that the average loss is controlled in hindsight for *any* realization of an input-target sequence $\{(x_t, y_t)\}_{t=1}^T$. This is commonly

¹The framework works for any bounded loss, but we assume the loss to be in $[0, 1]$ for simplicity.

known as the *adversarial sequence model* [10], [38]. Here, our goal is to set $\lambda_{1:T}$ to achieve a *long-term risk bound*:

$$\text{find } \lambda_{1:T} \text{ s.t. } \hat{R}_T(\mathcal{D}_{1:T}, \lambda_{1:T}) \leq \varepsilon + \frac{C \cdot h(T)}{T}, \quad (2)$$

where ε is a pre-defined risk level in $[0, 1]$, C is a (small) constant, $h(T)$ is any sublinear function; i.e., one where $h(T)/T \rightarrow 0$ as $T \rightarrow \infty$, and

$$\hat{R}_T(\mathcal{D}_{1:T}, \lambda_{1:T}) := \frac{1}{T} \sum_{t=1}^T \mathcal{L}(D_t^{\lambda_t}(x_t), y_t) \quad \text{and} \quad \hat{R}_0 = 0. \quad (3)$$

We will omit $\mathcal{D}_{1:T}$ in the notation of risk when the sequence of family of decision functions is clear from the context

IV. THEORY & CONFORMAL CONTROLLER ALGORITHM

In this section, we prove the core theoretical results behind Conformal Decision Theory. Specifically, we show that any sequence of families of decision functions $\mathcal{D}_{1:T}$ that are *eventually safe* can be calibrated online to achieve bounded long-term risk. We then introduce an example of conformal controller which solves Equation (2) under the assumption of eventual safety.

Definition 1 (Eventually Safe). *In the setting above, we say that $\mathcal{D}_{1:T}$ is eventually safe if $\exists \varepsilon^{\text{safe}} \in [0, 1], \lambda^{\text{safe}} \in \mathbb{R}$ and a time horizon $K > 0$ such that uniformly over all sequences $\lambda_{1:K}$ and $\{(x_1, y_1), \dots, (x_k, y_k)\} \in \mathcal{X} \times \mathcal{Y}$,*

$$\begin{aligned} & \{\forall k \in [K], \lambda_k \leq \lambda^{\text{safe}}\} \\ & \implies \frac{1}{K} \sum_{k=1}^K \mathcal{L}(D_k^{\lambda_k}(x_k), y_k) \leq \varepsilon^{\text{safe}}. \end{aligned} \quad (4)$$

Intuitively, this condition says that there exists a safe value λ^{safe} such that if the conformal control variable lands below that value, it will incur a low risk $\varepsilon^{\text{safe}}$ after no more than K time steps. For example, even the most conservative robot planner may not be able to change its trajectory fast enough in a *single timestep*, but it could possibly do so in K time steps. For general decision-making, the existence of a safe decision function is not guaranteed, and requires domain-specific knowledge (e.g., when the loss function captures the distance between agents [39]–[41]). But when the decision is a *prediction set*, a safe decision function is trivial because you can always output the entire space. Note that the *eventually safe* is a strictly weaker assumption than that used for the proofs in other works, such as [10], [24], [42], which require $K = 1$. Moreover, conformal controllers are simple yet efficient algorithms that solve the Conformal Decision Theory problem stated in Equation (2). An example is below.

Theorem 1 (Conformal Controller). *Consider the following update rule for $\lambda_{1:T}$:*

$$\lambda_{t+1} = \lambda_t + \eta(\varepsilon - \ell_t), \quad \forall t \in [T] \quad (5)$$

where $\eta > 0$ and $\ell_t := \mathcal{L}(D_t^{\lambda_t}(x_t), y_t)$.

If $\lambda_1 \geq \lambda^{\text{safe}} - \eta$ and $\mathcal{D}_{1:T}$ satisfies Definition 1 for a given $K \geq 1$ and $\varepsilon^{\text{safe}} \leq \varepsilon$, then for any realization of the data, the empirical risk is bounded:

$$\hat{R}_t(\lambda_{1:t}) \leq \varepsilon + \frac{(\lambda_1 - \lambda^{\text{safe}})/\eta + K}{t}, \quad (6)$$

for all $t \in [K, \dots, T]$.

The update in (5) resembles ACI [10] and is a hybrid between the RollingRC update [13], and the P-controller update [24]. The difference is that the update is applied to λ and not the conformal quantile or quantile level.

Proof of Theorem 1. By the definition of the update rule,

$$\lambda_{t+1} = \lambda_1 + \eta \sum_{s=1}^t (\varepsilon - \ell_s). \quad (7)$$

By isolating $\sum_{s=1}^t \ell_s$ on one side and moving all other terms to the right-hand side, we obtain:

$$\hat{R}_t(\lambda_{1:t}) = \frac{1}{t} \sum_{s=1}^t \ell_s = \varepsilon + \frac{\lambda_1 - \lambda_{t+1}}{\eta t}.$$

To conclude, we just need to lower bound λ_t by a constant w.r.t t which is done in the following Lemma 1.1. \square

Lemma 1.1. *For the sequence in Equation 5, with $\lambda_1 \geq \lambda^{\text{safe}} - \eta$ we have that the parameter λ_t is bounded below by $\lambda_t \geq \lambda^{\text{safe}} - K\eta$, for all $t \in [T + 1]$.*

Proof. First note that the maximal change in the parameter is $\sup_{s \in [T]} |\lambda_{s+1} - \lambda_s| < \eta$, because $\ell_s \in [0, 1]$ and $\varepsilon \in [0, 1]$. We will then proceed by contradiction: Assume that $\inf_{s \in [T+1]} \lambda_s < \lambda^{\text{safe}} - K\eta$. Denote $t = \arg \min_{s \in [T+1]} \{\lambda_s : \lambda_s < \lambda^{\text{safe}} - K\eta\}$. That is, t is the first instant when the parameter goes below that lower bound. Then, by definition of t , $\forall s < t, \lambda_s < \lambda^{\text{safe}} - K\eta \leq \lambda_s$.

Because the max difference between successive steps is η , we can prove recursively that $\forall k \in \{0, \dots, K\}, \lambda_{t-k} < \lambda^{\text{safe}} - (K-k)\eta$. Note that, from those inequalities, we deduce that $t > K$ since $\lambda_1 \geq \lambda^{\text{safe}} - \eta$. By recursively applying the update rule $\lambda_t = \lambda_{t-K} + K\eta(\varepsilon - \frac{1}{K} \sum_{k=1}^K \ell_{t-k})$, we have:

$$\begin{aligned} & (\forall k \in \{0, \dots, K-1\}, \lambda_{t-k} < \lambda^{\text{safe}}) \\ & \implies \frac{1}{K} \sum_{k=1}^K \ell_{t-k} \leq \varepsilon^{\text{safe}} \quad (\text{Definition 1}) \\ & \implies \lambda_t = \lambda_{t-K} + K\eta \left(\varepsilon - \frac{1}{K} \sum_{k=1}^K \ell_{t-k} \right) \\ & \implies \lambda_t \geq \lambda_{t-K} + K\eta(\varepsilon - \varepsilon^{\text{safe}}) \\ & \implies \lambda_t \geq \lambda_{t-K}. \end{aligned}$$

Since t is the first ever timestep to go below $\lambda^{\text{safe}} - K\eta$, this is a contradiction. \square

Remark 1. *The assumption $\lambda_1 \geq \lambda^{\text{safe}} - \eta$ is not necessary to prove that $\hat{R}_t(\lambda_{1:t}) \leq \varepsilon + O(1/t)$. Intuitively, two scenarios can occur:*

- 1) *If $\forall k \in [K], \lambda_k \leq \lambda^{\text{safe}}$, then the empirical risk over the first K steps will be upper bounded by $\varepsilon^{\text{safe}}$. In this*

case, we need only to upper bound the risk between $K+1$ and T , which can be achieved using the previous theorem or this remark.

- 2) If there exists a $k \in [K]$ such that $\lambda_k > \lambda^{\text{safe}}$, we can apply the previous theorem to upper bound the risk between k and $T+1$. The cumulative loss between 1 and k is upper bounded by k , which is $o(1)$.

Conformal Decision Theory in Batch

Conformal decision theory can also be applied in the so-called batch setting, wherein a separate calibration dataset is available for learning a safe decision. Here, a dataset or simulator allows for offline experimentation to quantify the risk of different decisions, e.g., offline RL. This requires a different statistical setup. Consider the case of $n+1$ exchangeable decision functions $D_1(\lambda), \dots, D_{n+1}(\lambda)$ and an associated loss function \mathcal{L} taking a decision and returning a value in $[0, 1]$. The first n decision functions will be used for calibration of a parameter $\hat{\lambda}$ that will be used in the final decision. These exchangeable decision functions may be produced, for example, by applying a single decision function to a sequence of exchangeable data points. For the sake of simplicity, we assume that the decisions have monotone loss, i.e., that for all i ,

$$\lambda_1 \leq \lambda_2 \implies \mathcal{L}(D_i(\lambda_1)) \leq \mathcal{L}(D_i(\lambda_2)). \quad (8)$$

Following [43], the conformal control variable can be chosen as

$$\hat{\lambda} = \sup \left\{ \lambda : \frac{1}{n} \sum_{i=1}^n \mathcal{L}(D_i(\lambda)) \leq \epsilon - \frac{1-\epsilon}{n} \right\}. \quad (9)$$

This will give a risk guarantee as a corollary of Theorem 1 of [43].

Corollary 2. *With the choice of $\hat{\lambda}$ above,*

$$\mathbb{E}[\mathcal{L}(D_{n+1}(\hat{\lambda}))] \leq \epsilon. \quad (10)$$

Though the validity of the algorithm follows from the theory of conformal risk control, it is substantially different in practice and deserves further study. Specifically, unlike the previous methods, in order to calculate $\hat{\lambda}$, one must iterate through a sequence of counterfactual decisions (possible values of λ) and evaluate what the loss would have been. This restricts the applications of the batch algorithm and also presents an opportunity for future work to make it more efficient and expand its scope.

V. EXPERIMENTS

We demonstrate Conformal Decision Theory in two autonomous decision-making domains, which exhibit two of the many different ways in which a conformal controller can be instantiated. First, we consider a robot-navigation-around-humans example in the Stanford Drone Dataset [1], where CDT tunes the robot’s reward function in an online manner to be safe but efficient. Then, we model a manufacturing setting where CDT directly calibrates the speed of the conveyor belt under a robot to achieve high-throughput and successful robot grasps.

A. Robot Navigation in Stanford Drone Dataset

Robot navigation around people must balance safety (i.e., not colliding with humans) and efficiency (i.e., the robot makes progress towards a goal). To ensure that the risk of collision is low while still making progress to the goal, the robot will calibrate its cost function at run-time using a conformal controller (CC).

Decision Function & Parameterization. The robot plans via model predictive control, where at each timestep it fits a minimum-cost spline subject to its dynamic constraints, which are modeled as a nonlinear Dubins car [44]. Let $g := [g_x, g_y] \in \mathbb{R}^2$ be the robot’s goal location. Let t be the current time, $H < T$ be the planning horizon, and $u_{t:t+H} \in \mathbb{R}^{H \times 3}$ be a spline consisting of the robot’s planar position and orientation. The robot also gets as input the current set of short-horizon predictions of each human’s state, $x_{t:t+H} \in \mathcal{P}_t$, generated by an autoregressive predictive model [45]. Note that this set \mathcal{P}_t can include predictions for *multiple* humans in the scene (as shown in Figure 1). The robot’s planning objective is

$$J(u_{t:t+H}; \mathcal{P}_t, \lambda) := \sum_{\tau=t}^{t+H} \underbrace{\|u_{\tau}^{\text{pos}} - g\|}_{\text{Goal distance}} + \lambda \cdot \underbrace{\left(- \inf_{x_{\tau} \in \mathcal{P}_t} \|u_{\tau}^{\text{pos}} - x_{\tau}\|_2^2 \right)}_{\text{Human avoidance}}, \quad (11)$$

where the notation $u_{\tau}^{\text{pos}} \in \mathbb{R}^2$ indicates the xy -positional entries of the robot’s state at time τ . Note that the conformal control variable λ scales the cost of staying far away from predicted human states: if $\lambda = 0$ the robot only cares about reaching the goal; if $\lambda > 0$ then the robot is increasingly penalized for intersecting with predicted human trajectories. The decision function outputs the minimum-cost trajectory for the robot

$$D_t^{\lambda} := \arg \min_{u_{t:t+H} \in \mathcal{U}} J(u_{t:t+H}; \mathcal{P}_t, \lambda), \quad (12)$$

where \mathcal{U} is the set of feasible splines (ones that are dynamically feasible for the robot and also do not intersect with environment obstacles). At the next timestep, the robot re-predicts the human trajectory (i.e., generates \mathcal{P}_{t+1}) and re-plans the decision D_{t+1}^{λ} .

Loss Function. Let $\mathcal{Y} \subset \mathbb{R}^2$ and the targets $y_t^1, \dots, y_t^M \in \mathcal{Y}$ be the actual xy positions of each of the M humans that the robot observes at time t . The loss function is defined as the negative distance to the nearest human,

$$\mathcal{L} := - \inf_{i \in [M]} \|y_t^i - u_t^{\text{pos}}\|_2, \quad (13)$$

where u_t^{pos} is the robot’s current position. To make this value bounded, we clip the loss to the size of the video. Note that because we use a negative loss, we also changed λ so that the larger λ , the more conservative the decision.

Metrics. We measure a boolean *safe* variable indicating if the robot did not ever collide with a human. We also measure a boolean *success* variable if the robot reached the goal location by the end of the interaction episode (i.e., length of video in the dataset). We also measure the time

TABLE I: **Stanford Drone Dataset: Quantitative Results.** Results on the `nexus_4` scenario from SDD [1]. The robot’s goal is to cross the nexus while avoiding pedestrians. Safety was violated if the robot collided with a human. At all learning rates η , the conformal controller is more efficient at navigation than ACI in terms of time. It remains safe so long as the learning rate is set high enough so that the robot planner can quickly adapt to nearby humans; when the learning rate is set too low (near zero), proximity to humans is effectively not penalized, leading to collisions.

		Metrics								
Method	η	success	time (s)	safe	min dist (m)	avg dist (m)	5% dist (m)	10% dist (m)	25% dist (m)	50% dist (m)
Aggressive	n/a	✓	8.567	✗	0.1595	4.058	1.253	1.546	2.495	4.021
ACI ($\alpha = 0.01$)	0	✓	27.17	✗	0.07612	5.201	1.842	2.415	3.9	5.614
	0.01	✓	26.67	✗	0.8026	4.575	2.261	3.014	3.507	4.574
	0.1	✓	24.73	✗	0.7906	4.771	2.284	2.825	3.561	4.78
Conformal Controller ($\varepsilon = 2m$)	50	✓	20.03	✗	0.6122	3.299	0.8688	1.426	2.022	2.978
	100	✓	17.4	✓	1.142	3.794	1.678	1.811	2.378	3.262
	500	✓	17.33	✓	1.116	3.989	1.69	1.812	2.452	3.795
	1000	✓	16.17	✓	1.265	3.599	1.698	1.81	2.282	3.303
Conservative	n/a	✗	∞	✓	2.268	6.291	3.801	3.982	4.982	5.993

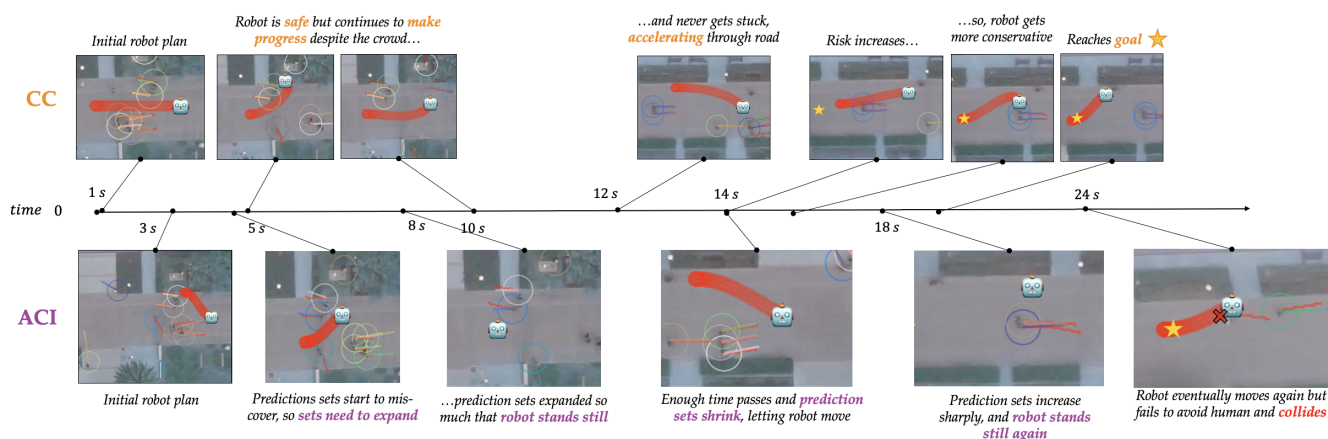


Fig. 3: **Stanford Drone Dataset: Qualitative Results.** Visualization of interaction over time (left to right). (Top) With our conformal controller (CC), the robot always makes progress towards its goal while remaining safe, even when blocked by crowds of people. (Bottom) The ACI baseline calibrates the prediction sets. As soon as a mis-prediction happens, ACI expands the prediction sets to obtain coverage, but this frequently blocks the robot from moving anywhere (see $t = 10s$), even though the mis-predictions occurred for a pedestrian who was far away and not interfering with the robot’s plan.

to reach the goal location and the minimum, mean, and $\{5\%, 10\%, 25\%, 50\%\}$ quantiles of the distance to the nearest human.

Experimental Setup. All methods are evaluated on interactions from the `nexus_4` video in the Stanford Drone Dataset (SDD) [1]. The risk threshold is $\varepsilon = 2m$ (i.e., radius around human). The robot always starts from the same initial condition and moves to the same goal. This scenario has a high density of pedestrians, making the risk-performance tradeoff for the robot nontrivial. Our approach (CC) adapts the reward weight λ_t on the human collision cost based on Equation 5 so that the decision risk is calibrated. Our baseline robot planners: **conservative** which always uses the safe decision function $D_t^{\lambda=1}$, **aggressive** which uses $D_t^{\lambda=0}$, and **ACI** [16] which first uses adaptive conformal prediction to calibrate prediction sets and then plans to avoid these sets.

Results. Quantitative results shown in Table I and qualitative results in Figure 1. Because the **conformal controller** calibrates the robot’s decisions directly, it is substantially (\sim




Fig. 4: **Stanford Drone Dataset.** (Top) Trajectories of λ_t (calibrated by CC) and α_t (from ACI to calibrate sets). When $\alpha_t \leq 0$, ACI returns infinite set and the robot stops. (Bottom) Distance to the nearest human over time. λ_t is large when the robot is close to human, while α_t is unrelated. The λ_t trajectory is shorter because it reaches the goal faster.

29%) faster at reaching the goal than the ACI algorithm (see visualization over time in Figure 3). While the **aggressive** baseline reaches the goal fastest, it consistently violates the safety threshold. On the other hand, the **conservative**




Fig. 5: **Manufacturing Assembly Line Robot: Quantitative Results.** (Left) Illustrative example: Robot must adjust the speed so that it grasps the most items while minimizing grasp failure. (Right) Empirical risk, \hat{R}_T , and average utility (i.e., successful grasps), \hat{V}_T on 1000 runs. Our method is denoted by (CC). Dashed red line is target risk $\varepsilon = 0.05$.

baseline never completes the task, getting stuck far away from the crowds of pedestrians. The **conformal controller** ensures safety so long as the learning rate is fast enough for the robot planner to quickly adapt to changes in nearby human behavior (see Figure 4). Note that **ACI** can result in collisions for two reasons: 1) the prediction sets do not adapt fast enough for the spline planner to react and swerve out of the way of the pedestrian, 2) if the prediction sets become so large that there is no feasible spline and the robot must stand in place, the pedestrians sometimes run into the robot. This issue was independently observed in [16].

B. Manufacturing Assembly Line Robot

Consider a factory assembly line where a robot has to grab items from a conveyor belt (left, Figure 5). As the speed increases, the throughput of items increases but so does the ratio of robot grasp failures. The agent must calibrate the speed so that the ratio of failures over time stays below ε .

Decision Function & Parameterization. The agent directly modifies the speed, thus the action is defined as $u_t := \lambda_t$. Here we take $\lambda_t \in [0, 1]$.

Risk Function. For a given conveyor belt speed λ , the robot will attempt to grab $n(\lambda)$ items, among which $d(\lambda)$ are failed grasps. The loss received by the robot will be $\mathcal{L}(\lambda) := d(\lambda)/n(\lambda)$.

Metrics. We measure average utility (i.e., # of successful grasps), $\hat{V}_T := \frac{1}{T} \sum_{t=1}^T V(\lambda_t)$, and empirical risk, $\hat{R}_T(\lambda_{1:T})$.

Experimental Setup. We assume that the number of items $n(\lambda)$ the robot attempts to grab is drawn as $\text{Pois}(C \cdot \sqrt{\lambda})$. The number of failed grasps conditioned on the total number of items is $d(\lambda)|n \sim \text{Bin}(n, C' \cdot \lambda)$. Importantly, the distributions of n, d , and the parameters C, C' are all *unknown to the agent*. Our conformal controller method (**CC**) adjusts the speed λ_t based on the update rule from Equation 5. In addition to the risk function, we also track a utility function which is the number of successful grasps $V(\lambda) := n(\lambda) - d(\lambda)$.

We compare our method with two baselines: A bandit algorithm running the upper confidence bound algorithm (**UCB**) [46] to maximize the utility V and another algorithm running the lower confidence bound algorithm (**LCB**) to minimize the loss \mathcal{L} . We also add two methods with oracle

access to the otherwise unknown parameters: **Oracle-Value** selects the best speed to maximize grasp success $\lambda_V^* := \arg \max_{\lambda} \mathbb{E}[V(\lambda)]$ and **Oracle-Loss** selects the best speed $\lambda_{\mathcal{L}}^*$ such that $\mathbb{E}[\mathcal{L}(\lambda_{\mathcal{L}}^*)] := \varepsilon$. The values selected for the parameters are in Figure 5. We run all methods for a horizon $T = 2000$, set $C = 10$, $C' = 0.2$, and the target risk is $\varepsilon = 0.05$ (i.e., $\leq 5\%$ failed grasp).

Results. We run the simulation $N = 1000$ times, and calculate the average empirical risk and the average number of successful grasps. In Figure 5, we find that our method performs as well as the **Oracle-Loss**, ensuring that the empirical risk of grasps never exceeds $\varepsilon = 0.05$, while still ensuring high throughput of successfully grasped items. **UCB** and **LCB** both violate the empirical risk threshold: **UCB** incurs this risk but achieves a higher number of successful grasps, while **LCB** is slow to learn its target, resulting in a higher risk over the time horizon.

VI. DISCUSSION & CONCLUSION

In this paper, we introduce *Conformal Decision Theory*, a framework for producing safe decisions with machine-learning predictions. The method works both online and in batch; the online algorithm is computationally easy while the batch algorithm requires evaluating a large amount of *counterfactual* decisions—speeding up the latter and extending it beyond exchangeability (e.g. using [47]) is a good direction for future work. Finally, despite λ being 1-dimensional, our procedure can index an arbitrary set of decisions. Consider a set of decisions \mathcal{D} , a utility predictor $\hat{u}(d; x)$ where $d \in \mathcal{D}$, and a loss predictor $\hat{\mathcal{L}}(d; x)$, we can maximize utility subject to the constraint that our predicted loss is controlled:

$$D_t = \underset{d \in \mathcal{D}}{\operatorname{argmax}} \quad \hat{u}(d; x_t) \\ \text{s.t.} \quad \hat{\mathcal{L}}(d; x_t) \leq \lambda_t.$$

This will work as long as we revert to a safe decision if $\lambda_t \leq \lambda^{\text{safe}}$; where the sequence λ_t is defined in Equation 5. However, no guarantees on utility are provided. This topic would be a great avenue for future work, bringing conformal prediction closer to the classical statistical decision theory of Lehmann [48], von Neuman and Morgenstern [49], and others.

REFERENCES

- [1] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese, "Learning social etiquette: Human trajectory understanding in crowded scenes," in *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*. Springer, 2016, pp. 549–565.
- [2] 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.
- [3] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena, "Structural-RNN: Deep learning on spatio-temporal graphs," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 5308–5317.
- [4] A. Vemula, K. Muelling, and J. Oh, "Social attention: Modeling attention in human crowds," in *2018 IEEE international Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4601–4607.
- [5] 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.
- [6] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," *arXiv preprint arXiv:1703.09312*, 2017.
- [7] T. Yin, C. Liu, F. Ding, Z. Feng, B. Yuan, and N. Zhang, "Graph-based stock correlation and prediction for high-frequency trading systems," *Pattern Recognition*, vol. 122, p. 108209, 2022.
- [8] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic learning in a random world*. Springer, 2005, vol. 29.
- [9] V. Vovk and C. Bendtsen, "Conformal predictive decision making," in *Conformal and Probabilistic Prediction and Applications*. PMLR, 2018, pp. 52–62.
- [10] I. Gibbs and E. Candes, "Adaptive conformal inference under distribution shift," in *Advances in Neural Information Processing Systems*, vol. 34. Curran Associates, Inc., 2021, pp. 1660–1672.
- [11] M. Zaffran, O. Féron, Y. Goude, J. Josse, and A. Dieuleveut, "Adaptive conformal predictions for time series," in *International Conference on Machine Learning*. PMLR, 2022, pp. 25 834–25 866.
- [12] A. N. Angelopoulos and S. Bates, "A gentle introduction to conformal prediction and distribution-free uncertainty quantification," *arXiv preprint arXiv:2107.07511*, 2021.
- [13] S. Feldman, L. Ringel, S. Bates, and Y. Romano, "Achieving risk control in online learning settings," 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] A. Dixit, L. Lindemann, S. X. Wei, M. Cleaveland, G. J. Pappas, and J. W. Burdick, "Adaptive conformal prediction for motion planning among dynamic agents," in *Learning for Dynamics and Control Conference*. PMLR, 2023, pp. 300–314.
- [17] 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.
- [18] R. Luo, S. Zhao, J. Kuck, B. Ivanovic, S. Savarese, E. Schmerling, and M. Pavone, "Sample-efficient safety assurances using conformal prediction," in *International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2022, pp. 149–169.
- [19] F. Cai and X. Koutsoukos, "Real-time out-of-distribution detection in learning-enabled cyber-physical systems," in *2020 ACM/IEEE 11th International Conference on Cyber-Physical Systems (ICCP)*. IEEE, 2020, pp. 174–183.
- [20] R. Sinha, E. Schmerling, and M. Pavone, "Closing the loop on runtime monitors with fallback-safe mpc," *Conference on Decision and Control*, 2023.
- [21] H. Yang and M. Pavone, "Object pose estimation with statistical guarantees: Conformal keypoint detection and geometric uncertainty propagation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 8947–8958.
- [22] A. Z. Ren, A. Dixit, A. Bodrova, S. Singh, S. Tu, N. Brown, P. Xu, L. Takayama, F. Xia, J. Varley *et al.*, "Robots that ask for help: Uncertainty alignment for large language model planners," *Conference on Robot Learning*, 2023.
- [23] O. Bastani, V. Gupta, C. Jung, G. Noarov, R. Ramalingam, and A. Roth, "Practical adversarial multivalid conformal prediction," *Advances in Neural Information Processing Systems*, vol. 35, pp. 29 362–29 373, 2022.
- [24] A. N. Angelopoulos, E. J. Candès, and R. J. Tibshirani, "Conformal pid control for time series prediction," *arXiv preprint arXiv:2307.16895*, 2023.
- [25] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning*. PMLR, 2016, pp. 1050–1059.
- [26] J. Lampinen and A. Vehtari, "Bayesian approach for neural networks—review and case studies," *Neural networks*, vol. 14, no. 3, pp. 257–274, 2001.
- [27] E. Goan and C. Fookes, "Bayesian neural networks: An introduction and survey," *Case Studies in Applied Bayesian Data Science: CIRM Jean-Morlet Chair, Fall 2018*, pp. 45–87, 2020.
- [28] 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.
- [29] R. Koehler and G. Bassett Jr, "Regression quantiles," *Econometrica: Journal of the Econometric Society*, pp. 33–50, 1978.
- [30] N. E. Du Toit and J. W. Burdick, "Robot motion planning in dynamic, uncertain environments," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 101–115, 2011.
- [31] M. Bujarbaruah, X. Zhang, M. Tanaskovic, and F. Borrelli, "Adaptive stochastic mpc under time-varying uncertainty," *IEEE Transactions on Automatic Control*, vol. 66, no. 6, pp. 2840–2845, 2020.
- [32] M. C. Campi, S. Garatti, and M. Prandini, "The scenario approach for systems and control design," *Annual Reviews in Control*, vol. 33, no. 2, pp. 149–157, 2009.
- [33] O. de Groot, L. Ferranti, D. Gavrilu, and J. Alonso-Mora, "Scenario-based motion planning with bounded probability of collision," *arXiv preprint arXiv:2307.01070*, 2023.
- [34] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *Proceedings of the 20th international conference on machine learning (icml-03)*, 2003, pp. 928–936.
- [35] E. Hazan *et al.*, "Introduction to online convex optimization," *Foundations and Trends® in Optimization*, vol. 2, no. 3–4, pp. 157–325, 2016.
- [36] S. Bubeck, "Introduction to online optimization," *Lecture notes*, vol. 2, pp. 1–86, 2011.
- [37] E. Hazan and K. Singh, "Introduction to online nonstochastic control," *arXiv preprint arXiv:2211.09619*, 2022.
- [38] A. P. Dawid, "The well-calibrated Bayesian," *Journal of the American Statistical Association*, vol. 77, no. 379, pp. 605–610, 1982.
- [39] 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 (CDC)*, 2017, pp. 2242–2253.
- [40] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger, "Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Systems Magazine*, vol. 43, no. 5, pp. 137–177, 2023.
- [41] K.-C. Hsu, H. Hu, and J. F. Fisac, "The safety filter: A unified view of safety-critical control in autonomous systems," 2023.
- [42] A. Bhatnagar, H. Wang, C. Xiong, and Y. Bai, "Improved online conformal prediction via strongly adaptive online learning," *arXiv preprint arXiv:2302.07869*, 2023.
- [43] A. N. Angelopoulos, S. Bates, A. Fisch, L. Lei, and T. Schuster, "Conformal risk control," *arXiv preprint arXiv:2208.02814*, 2022.
- [44] R. Walmbe, N. Agarwal, S. Kale, and V. Joshi, "Optimal trajectory generation for car-type mobile robot using spline interpolation," *IFAC-PapersOnLine*, vol. 49, no. 1, pp. 601–606, 2016.
- [45] R. J. Hyndman and G. Athanasopoulos, *Forecasting: principles and practice*. OTexts, 2018.
- [46] T. Lattimore and C. Szepesvári, *Bandit Algorithms*. Cambridge University Press, 2020.
- [47] A. Farinhas, C. Zerva, D. Ulmer, and A. F. Martins, "Non-exchangeable conformal risk control," *arXiv preprint arXiv:2310.01262*, 2023.

- [48] E. L. Lehmann, "Some principles of the theory of testing hypotheses," in *Selected works of EL Lehmann*. Springer, 2011, pp. 139–164.
- [49] J. Von Neumann and O. Morgenstern, *Theory of games and economic behavior (60th Anniversary Commemorative Edition)*. Princeton university press, 2007.