




When Control Changes the Data: Safety under Interaction-Driven Distribution Shift

Lars Lindemann
Automatic Control Laboratory

Our vision for autonomous systems

Control systems that operate autonomously in complex and open-ended environments.

Autonomous driving



Robotics




Wildfire prevention




Disaster recovery



Space exploration



Warehouses




Agriculture



Smart Cities



Autonomous systems are a reality!



Boston Dynamics




Tesla



Kumar Lab (UPenn)


But are they safe?



Boston Dynamics



Tesla



Kumar Lab (UPenn)

Autonomous systems in the wild



December 3, 2025, 3:15 PM

Waymo driverless taxi takes passenger into apparent police standoff



Los Angeles Times

Sept. 25, 2025

A delivery robot collided with a disabled man on L.A. street. The aftermath is getting ugly



sciencealert

Viral Footage of Robot Headbutting Woman Raises Safety Questions


TECH 28 February 2025 By CARL STRATHEARN, THE CONVERSATION



AI's Achilles Heel: Can We *Prove* Plans Before They Execute?

DEV Arvind SundaraRajan
Posted on Nov 12, 2025


#ai #programming #research #logic



Medium


Apr 14, 2020

Distribution-Shift — the hidden reason self-driving cars aren't safe yet.



New UK study reveals how pedestrians react to self-driving cars

October 30, 2025



Safe Control in Dynamic Environments with Conformal Prediction

The safe control problem in dynamic environments

Compute **control inputs** so that the system avoids **dynamic agents** with a probability of at least $1 - \delta$.

Challenges:

- Uncertain and interactive agents
- Imperfect learning-enabled predictors

**uncertainty quantification
in control is challenging**



Existing work (references omitted):

- “Vanilla” motion planner (e.g., graph/sampling-based)
- Reactive control (e.g., navigation functions, DWA, CBF)
- Predictive interactive and non-interactive methods (e.g., deep RL, GP-based control)

ignores dynamic agents

**no predictions,
simple systems**

**no safety guarantees, or
simplifying assumptions**




How can we quantify uncertainty of **learning-enabled predictors** and design **safe control laws**?

Uncertainty quantification of learning-enabled systems

How do we reason about safety of high-dimensional systems?

How do we reason about safety of stochastic system?



Existing work (references omitted):

- SMT/MILP techniques or hybrid systems reachability ← scalability issues
- SDP/LP relaxations or compositional approaches ← overly conservative

Can we use efficient statistical tools to reason about safety of learning-enabled systems?

Conformal Prediction

Conformal prediction in a nutshell

Uncertainty quantification under minimal assumptions:

- Distribution-free
- No assumptions on the predictor needed

$$(u, z) \sim \mathcal{D}$$

$$\hat{z} = \mu(u)$$

no assumptions needed

no assumptions needed


Assumption: Availability of i.i.d. calibration data $\{(u^{(i)}, z^{(i)})\}_{i=1}^k$

Goal¹:

For a failure probability of $\delta \in (0, 1)$, we want to obtain a prediction region C s.t.

$$\text{Prob}(z \in C(u)) \geq 1 - \delta$$

Example²:



$\left\{ \begin{array}{l} \text{fox} \quad \text{squirrel} \\ 0.99 \end{array} \right\}$



$\left\{ \begin{array}{l} \text{fox} \quad \text{gray} \quad \text{bucket} \quad \text{rain} \\ \text{squirrel}, \text{fox}, \text{barrel} \\ 0.82 \quad 0.03 \quad 0.02 \quad 0.02 \end{array} \right\}$



$\left\{ \begin{array}{l} \text{marmot}, \text{squirrel}, \text{mink}, \text{weasel}, \text{beaver}, \text{polecat} \\ 0.30 \quad 0.22 \quad 0.18 \quad 0.16 \quad 0.03 \quad 0.01 \end{array} \right\}$


¹Vovk and Shafer, “A Tutorial on Conformal Prediction”, *JMLR*, 2008.

²Angelopoulos, “A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification”, *subm.*, 2021.

Conformal prediction in a nutshell

test data calibration data nonconformity score (e.g., prediction error)

Quantile lemma: Let $R^{(0)}, R^{(1)}, \dots, R^{(k)}$ be $k + 1$ i.i.d. random variables. It holds that

$$\text{Prob}(R^{(0)} \leq \text{Quantile}_{1-\delta}(R^{(1)}, \dots, R^{(k)}, \infty)) \geq 1 - \delta$$


The $p := \lceil (k + 1)(1 - \delta) \rceil$ -th smallest nonconformity score $\rightarrow$ computationally efficient

- Implicit data requirement: $p \leq k$

How much data do we need?

$$(k + 1)(1 - \delta) \leq k$$

$$\Leftrightarrow \frac{1 - \delta}{\delta} \leq k$$


more details



Formal Verification and Control With Conformal Prediction

PRACTICAL SAFETY GUARANTEES FOR AUTONOMOUS SYSTEMS

LARS LINDEMANN¹, YIQI ZHAO², XINYI YU²,
GEORGE J. PAPPAS², and JYOTIRMOY V. DESHMUKH²



Confidence meets control — verifiably safe systems with conformal prediction.

Engineering nonconformity scores




- **Example (regression):** the nonconformity score $R^{(i)} := |z^{(i)} - \mu(u^{(i)})|$ results in

$$\text{test data} \rightarrow \text{Prob}(|z - \mu(u)| \leq \text{Quantile}_{1-\delta}(R^{(1)}, \dots, R^{(k)}, \infty)) \geq 1 - \delta \leftarrow \text{calibration data}$$

$$\Leftrightarrow \text{Prob}(z \in [\mu(u) - \text{Quantile}_{1-\delta}, \mu(u) + \text{Quantile}_{1-\delta}]) \geq 1 - \delta$$

valid prediction region

- **Simple, general, and efficient** with the right nonconformity score¹:



tight prediction regions

multi modal distributions

¹Tumu, Cleaveland, Mangharam, Pappas, and Lindemann, "Multi-Modal Conformal Prediction Regions by Optimizing Convex Shape Templates", *L4DC*, 2024.

Marginal and calibration conditional guarantees

- Conformal prediction provides **marginal coverage** guarantees:

$$\text{Prob}(R^{(0)} \leq \text{Quantile}_{1-\delta}(R^{(1)}, \dots, R^{(k)}, \infty)) \geq 1 - \delta$$

captures randomness in the draw over test and calibration data $R^{(0)}, R^{(1)}, \dots, R^{(k)}$

- This is in contrast to **calibration conditional coverage** guarantees:

conditional probability

$$\text{Prob}(R^{(0)} \leq \text{Quantile}_{1-\delta}(R^{(1)}, \dots, R^{(k)}, \infty) | R^{(1)}, \dots, R^{(k)}) \geq 1 - \delta$$


captures randomness in the draw over test data $R^{(0)}$ → impossible to obtain

Let $R^{(0)}, R^{(1)}, \dots, R^{(k)}$ be $k + 1$ i.i.d. **continuous** random variables. It holds that¹

$$\text{Prob}(R^{(0)} \leq \text{Quantile}_{1-\delta}(R^{(1)}, \dots, R^{(k)}, \infty) | R^{(1)}, \dots, R^{(k)}) \sim \text{Beta}(1 - \delta, 1/k)$$

conditional probability

mean variance




¹Vovk, “Conditional Validity of Inductive Conformal Predictors”, *ACML*, 2012.

Safe Control in Dynamic Environments



no interaction

Safe control using conformal prediction




- **System Dynamics:** well understood and can be modeled accurately
- **Dynamic Agents:** in general hard to model
- **Data-rich:** large datasets describing behavior of dynamic agents
- **Prediction:** learning-enabled algorithms

We use conformal prediction to design **safe and computationally efficient control** algorithms.

Overview

Offline trajectory dataset



Goal: Design a probabilistically safe controller.

Conformal Prediction



Dynamics

$$x_{t+1} = f(x_t, u_t)$$

Online observations

learned, e.g., RNN

Trajectory predictor



Prediction region

$$\text{Prob}(y \in \bullet) \geq 1 - \delta$$



safety constraint

Model Predictive Control
 $\text{Prob}(c(x, y) \geq 0) \geq 1 - \delta$

How good are these estimates?


Key features:

- Computationally lightweight
- Probabilistic safety guarantees
- Distribution shift
- Temporal logic constraints

Uncertainty representation in dynamic environments

- Let $(y_0, y_1, \dots) \sim \mathcal{D}$ be a random trajectory

static for now, i.e, no interaction $\mathcal{D}(x)$




- Consider a trajectory predictor that maps $(y_0, \dots, y_t)$ to $(\hat{y}_{t+1}, \dots, \hat{y}_T)$

observations


mission horizon
predictions of $(y_{t+1}, \dots, y_T)$

- For instance, social LSTM



Model-agnostic: generative or physics-based models could be used

- Task:** quantify uncertainty at the state prediction level





Uncertainty representation in dynamic environments

- Define the nonconformity score $R^{(i)} := \max_{\tau \in \{t+1, \dots, T\}} \alpha_{\tau} \|y_{\tau}^{(i)} - \hat{y}_{\tau}^{(i)}\|$
prediction error
normalization constant (see reference¹)


and apply conformal prediction $C = \text{Quantile}_{1-\delta}(R^{(1)}, \dots, R^{(k)}, \infty)$

Multi-step prediction regions: $\text{Prob}(\|y_{\tau} - \hat{y}_{\tau}\| \leq C/\alpha_{\tau}, \forall \tau \in \{t+1, \dots, T\}) \geq 1 - \delta$

Example:



1.25 second ahead prediction error



1.8 second ahead prediction error

¹Cleaveland, Lee, Pappas, and Lindemann, “Conformal Prediction for Time Series using Linear Complementarity Programming”, *AAAI*, 2024.

Uncertainty-aware predictive control

Model predictive control¹: $\min_{(u_t, \dots, u_{T-1})} J(x, u)$

s.t. $x_{\tau+1} = f(x_\tau, u_\tau), \quad \tau \in \{t, \dots, T-1\}$

enforced via $c(x_\tau, \hat{y}_\tau) \geq LC_\tau$

robust safety constraint


$\inf_{y_\tau \in \mathcal{B}_\tau} c(x_\tau, y_\tau) \geq 0, \quad \tau \in \{t+1, \dots, t+H\}$

$\mathcal{B}_\tau := \{y \text{ s.t. } \|y - \hat{y}_\tau\| \leq C_\tau\}$

$u_\tau \in \mathcal{U}, x_{\tau+1} \in \mathcal{X}, \quad \tau \in \{t, \dots, T-1\}$

If the MPC is recursively feasible^{2,3}, we find a safe controller with a probability of at least $1 - \delta$.

Toy example:



Animated:

¹Lindemann, Cleaveland, Shim, and Pappas, “Safe Planning in Dynamic Environments using Conformal Prediction”, *RA-L*, 2023.

²Stamouli, Lindemann, and Pappas, “Recursively Feasible Shrinking-Horizon MPC in Dynamic Environments with Conformal Prediction Guarantees”, *L4DC*, 2024.

³Yu, Zhao, Yin, and Lindemann, “STL Control Synthesis among Uncontrollable Agents with Conformal Prediction”, *Automatica*, 2026.

Case Study: CARLA

Results:


- 99.985 % correct one-step ahead predictions (at least 99.83 % expected)
- 99/100 test runs without constraint violation (at most 5 expected)



Comparison with Gaussian Process:

- Largely undercovers

one instance where “control changes the data”!



Safe Control in Interactive Environments

Interaction-induced distribution shifts

What if interaction between the control system and the dynamic agents changes the distribution?

- **Example:** self-driving vehicles among pedestrians

no safety guarantees



- **Chicken-and-egg problem:**

Dynamics
 $x_{t+1} = f(x_t, u_t)$




$$x = (x_1, x_2, \dots)$$

a function of x

$$u = \pi(\text{Pred.Region}(x))$$

Distribution
 $\mathcal{D}(x)$


distribution shift




Can we iteratively update our controller to break circularity?

Iteratively breaking circularity

Collect data under $\pi(C_0)$




Iteratively breaking circularity




Iteratively breaking circularity

Collect data under $\pi(C_0)$




Iteratively breaking circularity

Collect data under $\pi(C_1)$



Iteratively breaking circularity

Collect data under $\pi(C_1)$



Safe iterative control design

- **Iterative controller updates:** $\pi(C_j) := \underset{(u_0, \dots, u_{T-1})}{\operatorname{argmin}} J(x, u)$

episode j

$$\text{s.t. } x_{\tau+1} = f(x_\tau, u_\tau),$$

iterative safety constraint

$$\inf_{y_\tau \in \mathcal{B}_{j,\tau}} c(x_\tau, y_\tau) \geq 0$$

Controller-agnostic: any controller $\pi(C_j)$ could be used!

$$\mathcal{B}_{j,\tau} := \{y \text{ s.t. } \|y - \hat{y}_\tau\| \leq C_j\}$$

- **Design goals:** update the radius C_j such that:

trajectories under controller $\pi(C_j)$

- (1) $\operatorname{Prob} \{ c(x_j, y_j) \geq 0 \} \geq 1 - \delta$ for all j , ← episodic safety
- (2) $J_{j+1} \leq J_j - \epsilon_j$ for all j , ← improved performance
- (3) $C_{j+1} \approx C_j$ as $j \rightarrow \infty$ ← convergence

Can we use **robust conformal prediction** to update the radius C_j safely?

Robust conformal prediction for non i.i.d. data

Let $R^{(0)}, R^{(1)}, \dots, R^{(k)}$ be $k + 1$ i.i.d. random variables. It holds that


$$\text{Prob}(R^{(0)} \leq \text{Quantile}_{1-\delta}(R^{(1)}, \dots, R^{(k)}, \infty)) \geq 1 - \delta$$

What if the data is dependent or not identically distributed?

Conformal prediction in the non i.i.d. case:

- Adversarially robust conformal prediction
- Distributionally robust conformal prediction
- Adaptive conformal prediction
- Weighted conformal prediction
- ...


summary in



Formal Verification and Control With Conformal Prediction

PRACTICAL SAFETY GUARANTEES FOR AUTONOMOUS SYSTEMS

LARS LINDEMANN¹, YIQI ZHAO², XINYI YU²,
GEORGE J. PAPPAS², and JYOTIRMOY V. DESHMUKH²



Confidence meets control — verifiably safe systems with conformal prediction.

Today!






Safe control with robust conformal prediction

- **Adversarial conformal prediction¹:** Let $\sup_{\Delta} R(y + \Delta) \leq R(y) + M$ for all trajectories y .

Δ

It holds that $\text{Prob}(R(y^{(0)} + \Delta) \leq q + M, \forall \Delta) \geq 1 - \delta$ with $q := \text{Quantile}_{1-\delta}(R(y^{(1)}), \dots, R(y^{(k)}))$

- **Our approach²:** iterative robustification




¹Gendler, Weng, Daniel, Romano, “Adversarially Robust Conformal Prediction”, ICLR, 2022.

²Mirzaeedodangeh, Shekhtman, Matni, **Lindemann**, “Safe Planning in Interactive Environments via Iterative Policy Updates and Adversarially Robust Conformal Prediction”, *L4DC*, 2026. 20


Algorithm for controller updates

- Collect new calibration data $y_j^{(1)}, \dots, y_j^{(k)} \sim \mathcal{D}(x_j)$




Algorithm for controller updates

- Collect new calibration data $y_j^{(1)}, \dots, y_j^{(k)} \sim \mathcal{D}(x_j)$



Algorithm for controller updates

- Collect new calibration data $y_j^{(1)}, \dots, y_j^{(k)} \sim \mathcal{D}(x_j)$
 - Apply conformal prediction $q_j = \text{Quantile}_{1-\delta}(R_j^{(1)}, \dots, R_j^{(k)}, \infty)$
- $R_j^{(i)} := \max_{\tau \in \{1, \dots, T\}} \|y_{j,\tau}^{(i)} - \hat{y}_\tau^{(i)}\|$
nonconformity score




Algorithm for controller updates

- Collect new calibration data $y_j^{(1)}, \dots, y_j^{(k)} \sim \mathcal{D}(x_j)$
- Apply conformal prediction $q_j = \text{Quantile}_{1-\delta}(R_j^{(1)}, \dots, R_j^{(k)}, \infty)$

$$R_j^{(i)} := \max_{\tau \in \{1, \dots, T\}} \|y_{j,\tau}^{(i)} - \hat{y}_\tau^{(i)}\|$$


nonconformity score



Algorithm for controller updates

- Collect new calibration data $y_j^{(1)}, \dots, y_j^{(k)} \sim \mathcal{D}(x_j)$
 - Apply conformal prediction $q_j = \text{Quantile}_{1-\delta}(R_j^{(1)}, \dots, R_j^{(k)}, \infty)$
 - Observe that $R_{j+1}^{(i)} - R_j^{(i)} \leq L_y \|\pi(C_{j+1}) - \pi(C_j)\|_\infty =: M_{j+1}$
- $R_j^{(i)} := \max_{\tau \in \{1, \dots, T\}} \|y_{j,\tau}^{(i)} - \hat{y}_\tau^{(i)}\|$
nonconformity score
- Lipschitz constant of dynamics


Adversarially conformal prediction guarantees that $\text{Prob}(R_{j+1} \leq q_j + M_{j+1}) \geq 1 - \delta$



Algorithm for controller updates

- Collect new calibration data $y_j^{(1)}, \dots, y_j^{(k)} \sim \mathcal{D}(x_j)$
 - Apply conformal prediction $q_j = \text{Quantile}_{1-\delta}(R_j^{(1)}, \dots, R_j^{(k)}, \infty)$
 - Observe that $R_{j+1}^{(i)} - R_j^{(i)} \leq L_y \|\pi(C_{j+1}) - \pi(C_j)\|_\infty =: M_{j+1}$
- $R_j^{(i)} := \max_{\tau \in \{1, \dots, T\}} \|y_{j,\tau}^{(i)} - \hat{y}_\tau^{(i)}\|$
nonconformity score
- Lipschitz constant of dynamics

Adversarially conformal prediction guarantees that $\text{Prob}(R_{j+1} \leq q_j + M_{j+1}) \geq 1 - \delta$



Algorithm for controller updates


- Collect new calibration data $y_j^{(1)}, \dots, y_j^{(k)} \sim \mathcal{D}(x_j)$
 - Apply conformal prediction $q_j = \text{Quantile}_{1-\delta}(R_j^{(1)}, \dots, R_j^{(k)}, \infty)$
 - Observe that $R_{j+1}^{(i)} - R_j^{(i)} \leq L_y \|\pi(C_{j+1}) - \pi(C_j)\|_\infty =: M_{j+1}$
- $R_j^{(i)} := \max_{\tau \in \{1, \dots, T\}} \|y_{j,\tau}^{(i)} - \hat{y}_\tau^{(i)}\|$
nonconformity score
- Lipschitz constant of dynamics

Adversarially conformal prediction guarantees that $\text{Prob}(R_{j+1} \leq q_j + M_{j+1}) \geq 1 - \delta$



- This motivates the rule $C_{j+1} \geq q_j + M_{j+1}$

The policy $\pi(C_{j+1})$ guarantees episodic safety.

1. Solve via numerical solver (e.g., bisection)
2. Solve via $C_{j+1} \geq q_j + L_\pi L_y |C_{j+1} - C_j|$





Academic example: car-pedestrian collision avoidance



$$\text{Prob}(c(x_j, y_j) \geq 0) \geq 1 - \delta$$


$$\text{Prob}\left(\max_{\tau} \|\hat{y}_{j,\tau} - y_j\| \leq C_j\right) \geq 1 - \delta$$







convergence conditions
presented after simulations

Quadswarm drone simulator

- We use Quadswarm¹ with five uncontrollable dynamic RL agents
- Ego drone deploys $\pi(C_j)$ filtered through a tracking CBF







• Metrics:



$$\text{Prob}\left(\max_{\tau} \|\hat{y}_{j,\tau} - y_j\| \leq C_j\right) \geq 1 - \delta$$

$$\text{Prob}(c(x_j, y_j) \geq 0) \geq 1 - \delta$$

• Agent 2:



¹Huang et al., “QuadSwarm: A Modular Multi-Quadrotor Simulator for Deep Reinforcement Learning with Direct Thrust Control”, *ICRA (Workshop)*, 2023.

Radius update analysis

- Recall the radius update condition $C_{j+1} \geq q_j + L_y \|\pi(C_{j+1}) - \pi(C_j)\|_\infty$
- Consider the explicit update rule $C_{j+1} \geq C_j + L_\pi L_y |C_{j+1} - C_j|$

“closed-loop” gain $\kappa := L_\pi L_y$

small gain condition

If $\kappa < 1$, then we have
$$C_{j+1} = \begin{cases} \frac{q_j + \kappa C_j}{1 + \kappa} & \text{if } q_j \leq C_j \quad (\text{Shrinkage Branch}) \\ \frac{q_j - \kappa C_j}{1 - \kappa} & \text{if } q_j > C_j \quad (\text{Expansion Branch}) \end{cases}$$

$C_{j+1} \leq C_j$

$C_{j+1} > C_j$

Can we use these insights to analyze **convergence of our algorithm**?

Convergence analysis

- Consider the true quantile $Q(C) := \inf\{t \mid \text{Prob}(R(y) \leq t) \geq 1 - \delta\}$

$$y \sim \mathcal{D}(\pi(C))$$

under some additional
continuity assumptions

Assume that $Q(C)$ admits the fixed point $Q(C^*) = C^*$. Then

$$|C_j - C^*| \leq \left(\frac{2\kappa}{1 - \kappa}\right)^{j+1} |C_0 - C^*| + \frac{1}{1 - \kappa} \sum_{m=0}^j \left(\frac{2\kappa}{1 - \kappa}\right)^{j-m} |Q(C_j) - q_j|$$

can be bounded with high probability (see paper¹)

- We can derive a performance bound $|J_j - J(C^*)|$ if $J(C)$ is monotone/continuous

¹Mirzaeododangeh, Shekhtman, Matni, **Lindemann**, “Safe Planning in Interactive Environments via Iterative Policy Updates and Adversarially Robust Conformal Prediction”, *L4DC*, 2026. 25

Selected References

Main Collaborators

Survey article:

- **L. Lindemann**, Y. Zhao, X. Yu, G. J. Pappas, and J. V. Deshmukh, “Formal Verification and Control with Conformal Prediction”, CSM, 2025

Main references:

- **L. Lindemann**, M. Cleaveland, G. Shim, and G. J. Pappas, “Safe Planning in Dynamic Environments using Conformal Prediction”, RA-L, 2023.
- M. Cleaveland, I. Lee, G. J. Pappas, and **L. Lindemann**, “Conformal Prediction Regions for Time Series using Linear Complementarity Programming”, AAAI, 2024.
- O. Mirzaeedodangeh, E. Shekhtman, N. Matni, **L. Lindemann**, “Safe Planning in Interactive Environments via Iterative Policy Updates and Adversarially Robust Conformal Prediction”, L4DC, 2026.
- O. Mirzaeedodangeh, E. Shekhtman, N. Matni, **L. Lindemann**, “Robust Conformal CBF and CLF Controllers via Iterative Policy Updates”, CDC (subm.), 2026.



Omid
Mirzaeedodangeh



Matthew
Cleaveland




Eliot
Shekhtman



George Pappas



Jyotirmoy
Deshmukh



Nikolai Matni