Robust Yet Efficient Conformal Prediction Sets

Xiaoyu (Joy) Liang

Beihang University

Meeting with: **Prof. Aleksandar Bojchevski**

University of Cologne

2025 Apr 14

Problem

Conformal Prediction: Returns prediction sets with a distribution-free guarantee to cover the true label.

$$\Pr[\text{airplane} \in \mathcal{C}(\ \ )] \geq 1 - \alpha$$

Attacks: Break the guarantee by perturbing the test inputs (evasion) or poisoning the calibration data.

$$\Pr[\text{airplane} \in \mathcal{C}(\ \ + 0.001\ \ )] \not\geq 1 - \alpha$$

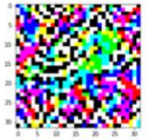

How can we extend robustness certificates to conformal prediction sets?

Robust CP: Recover the same guarantee for the worst-case perturbed input and calibration.

$$\Pr[\text{airplane} \in \bar{\mathcal{C}}(\ \ + 0.001\ \ )] \geq 1 - \alpha$$

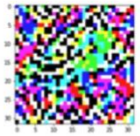

[1] https://icml.cc/media/PosterPDFs/ICML%202024/34224.png?t=1721216753.8029292

Intuition

Conformal Prediction

- Conformity score function s: $\mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, e.g. TPS, APS
- Calibration set $\mathcal{D}_{\text{cal}} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^{n}$
- Coverage probability $1 - \alpha$, user-specified

Calibration. α-quantile $q = \text{Quant}(\alpha; \{s(\boldsymbol{x}_i, y_i)\}_{i=1}^{n})$.

Prediction. Define the set $\mathcal{C}(x_{n+1}) = \{y: s(\boldsymbol{x}_{n+1}, y)\} \geq q\}$.

Guarantee. If test and calibration are exchangeable then we have:

$$\Pr[y_{n+1} \in \mathcal{C}(x_{n+1})] \geq 1 - \alpha.$$

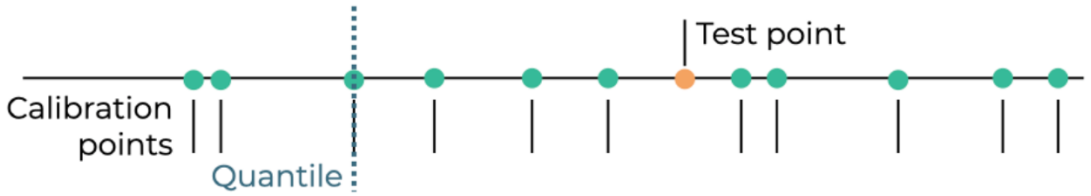

Randomized Smoothing Conformal Prediction

- Bound score: $\bar{s}(\boldsymbol{\xi}(\tilde{x}), y) \geq s(x, y), \tilde{x} \in \mathcal{B}(x)$

$$\bar{s}(x, y) = \Phi^{-1}(\mathbb{E}_{\delta \sim \mathcal{N}(0, \sigma^2 I)}[s(x + \delta, y)])$$

1. Loose bound 2. Continuous

3. Poison attack

Calibration. $q = \text{Quant}(\alpha; \{\bar{s}(\boldsymbol{x}_i, y_i)\}_{i=1}^{n}) - \boldsymbol{M_{\mathcal{B}}}$.

Prediction. $\bar{\mathcal{C}}(x_{n+1}) = \{y: \bar{s}(\boldsymbol{\xi}(\boldsymbol{x_{n+1}}), y)\} \geq q\}$.

Guarantee.

$$\Pr\left[y_{n+1} \in \bar{\mathcal{C}}\left(\boldsymbol{\xi}(\boldsymbol{x_{n+1}})\right)\right] \geq 1 - \alpha.$$

4. Finite sample error

1. CDF Aware smooth scores.
2. Sparse smoothing.
3. Lower bounds on calibration for poison.
4. Finite sample correction

[1] https://icml.cc/media/PosterPDFs/ICML%202024/34224.png?t=1721216753.8029292

Robust CP

How can we build conservative prediction set against evasion and poison?

Robustness to Evasion Attacks

Goal: $\Pr[y_{n+1} \in \mathcal{C}(\tilde{x}_{n+1})] \geq 1 - \alpha$

Method: Upper bound of score $\bar{s}(\xi(\tilde{x}), y)$

Conservative prediction set

$$\bar{\mathcal{C}}(x_{n+1}) = \{y : \bar{s}(\xi(\tilde{x}_{n+1}), y)\} \geq q\}.$$

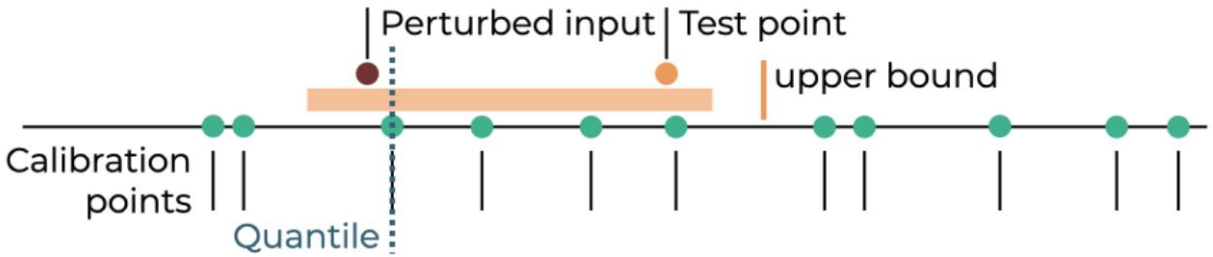

Robustness to Poison Attacks

Method: Lower bound of calibration

Feature $\underline{q_\alpha} = \min_{z_i \in \mathcal{X}} \text{Quant}(\alpha; \{s(z_i, y)\}_{i=1}^n)$

Label $\underline{q_\alpha} = \min_{z_i \in \mathcal{Y}} \text{Quant}(\alpha; \{s(x_i, z_i)\}_{i=1}^n)$

MILP

Conservative prediction set

$$\bar{\mathcal{C}}(x_{n+1}) = \{y : s(x_{n+1}, y)\} \geq \underline{q_\alpha}\}.$$

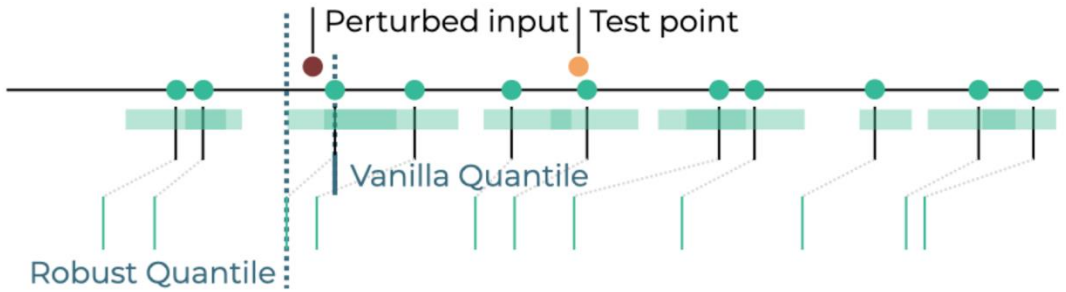

[1] https://icml.cc/media/PosterPDFs/ICML%202024/34224.png?t=1721216753.8029292

Randomized Smoothing Bounds

Smooth scores

Gaussian/sparse smoothing score: $\hat{s}(x, y) \coloneqq \mathbb{E}(s(\xi(x), y))$

Bound score for worst case input: $\max_{\tilde{x} \in \mathcal{B}(x)} \mathbb{E}(s(\xi(\tilde{x}), y))$

Relax by search worst case score function:

$$\max_{\tilde{x} \in \mathcal{B}(x)} \mathbb{E}(s(\xi(\tilde{x}), y)) \leq \max_{\tilde{x} \in \mathcal{B}(x), h \in \mathcal{H}} \mathbb{E}(h(\xi(\tilde{x}), y))$$

Bound h with mean: $\mathbb{E}\big(h(\xi(\tilde{x}), y)\big) = \mathbb{E}\big(s(\xi(x), y)\big)$

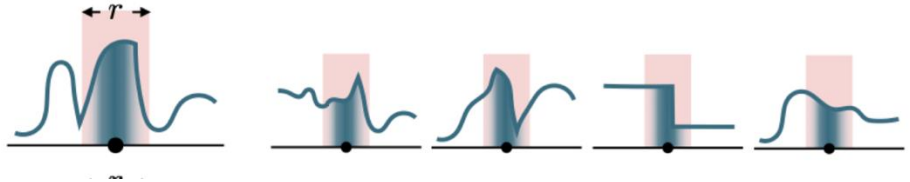

Bound h with **CDF**: $\forall b_i : \Pr[h(\xi(\tilde{x}), y) \leq b_i] = p_i$

$$a = b_1 < b_2 \cdots b_m = b, p_i = \Pr[s(\xi(x), y) \leq b_i]$$

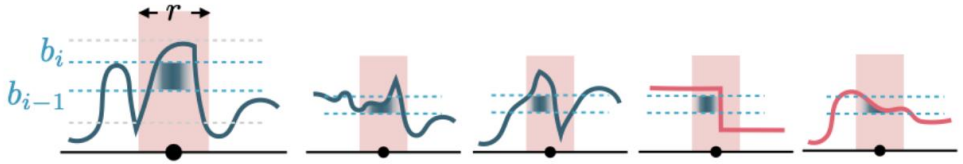

How can we build smaller conservative sets?

Bounds for smooth scores

Gaussian smoothing

$$\bar{s}_{\text{mean}}(\boldsymbol{x}, y) = \Phi(\Phi^{-1}(p) + r)$$

$$\bar{s}_{\text{cdf}}(\boldsymbol{x}, y) = b_m - \sum_{j=2}^{m-1} \Phi(\Phi^{-1}(p_j) - r)(b_{j+1} - b_j)$$

Sparse smoothing

$\bar{s}_{\text{mean}}(\boldsymbol{x}, y)$: Greedily distribute the p mass to each constant likelihood region.

$\bar{s}_{\text{cdf}}(\boldsymbol{x}, y)$: Distribute the p_j masses in each region and each bin $[b_j, b_j + 1]$.

$$\max_{\boldsymbol{H}} b_m - \boldsymbol{H}\tilde{\boldsymbol{t}}\,\boldsymbol{d} \;\text{ s.t. } \boldsymbol{H}\boldsymbol{t} = \boldsymbol{p}, 0 \leq \boldsymbol{H} \leq 1$$

[1] https://icml.cc/media/PosterPDFs/ICML%202024/34224.png?t=1721216753.8029292

Robust Yet Efficient CP

How can we build efficient and correct conservative sets?

CAS: CDF-Aware Sets

Algorithm 1 CDF-Aware Sets (CAS, Evasion)

$q_\alpha = \mathrm{Quant}\left(\alpha; \{\hat{s}(\boldsymbol{x}, y)\}_{(\boldsymbol{x},y)\in\mathcal{D}_{\mathrm{cal}}}\right)$ ▷ Clean quantile
Compute $\bar{s}_{\mathrm{cdf}}(\boldsymbol{x}, y)$, e.g. with Eq. 10 ▷ Upper bound
Return $\bar{C}_\alpha = \{y : \bar{s}_{\mathrm{cdf}}(\boldsymbol{x}, y) \geq q_\alpha\}$ ▷ Conservative set

Calibration-time variant: Compare the smooth test score $\hat{s}(\boldsymbol{x}, y)$ against conservative (lower) quantile.

$$\underline{q_\alpha} = \mathrm{Quant}(\alpha; \{\underline{s}(\boldsymbol{x_i}, y_i)\})$$

Poisoning: Conservative threshold $\underline{q_\alpha} = \min_{\mathrm{z}} \mathrm{Quant}$

Combined: $\bar{C}_\alpha = \{y: \bar{s}_{\mathrm{cdf}}(x, y) \geq \underline{q_\alpha}\}$

Finite Sample Correction

Monte-Carlo sample for mean/CDF: $\boldsymbol{1-\eta}$ confidence

Calibrate $\boldsymbol{\alpha' = \alpha - \eta}$

Corrected CDF score: $\underline{s}_{\mathrm{cdf}+}(\boldsymbol{x_i}, y_i) \leq \underline{s}_{\mathrm{cdf}}$ with $1 - \eta/(2|\mathcal{D}_{\mathrm{cal}}|)$

DKW inequality

Corrected smooth score: $\hat{s}_+(\tilde{\boldsymbol{x}}_i, y_i) \geq \hat{s}$ with $1 - \eta/(2|\mathcal{Y}|)$

Bernstein bound

Conservative threshold: $\underline{q_{\alpha+}} = \mathrm{Quant}(\alpha - \eta; \{\underline{s}_{\mathrm{cdf}+}(\boldsymbol{x_i}, y_i)\})$

Conservative set: $\bar{C}_{\alpha+}(x_{n+1}) = \{y: \hat{s}_+(x_{n+1}, y) \geq \underline{q_{\alpha+}}\}$

Experiments

CAS returns smaller prediction set with the same guarantee. CAS results in larger “for free” radius.

All radii
All dataset
All guarantee

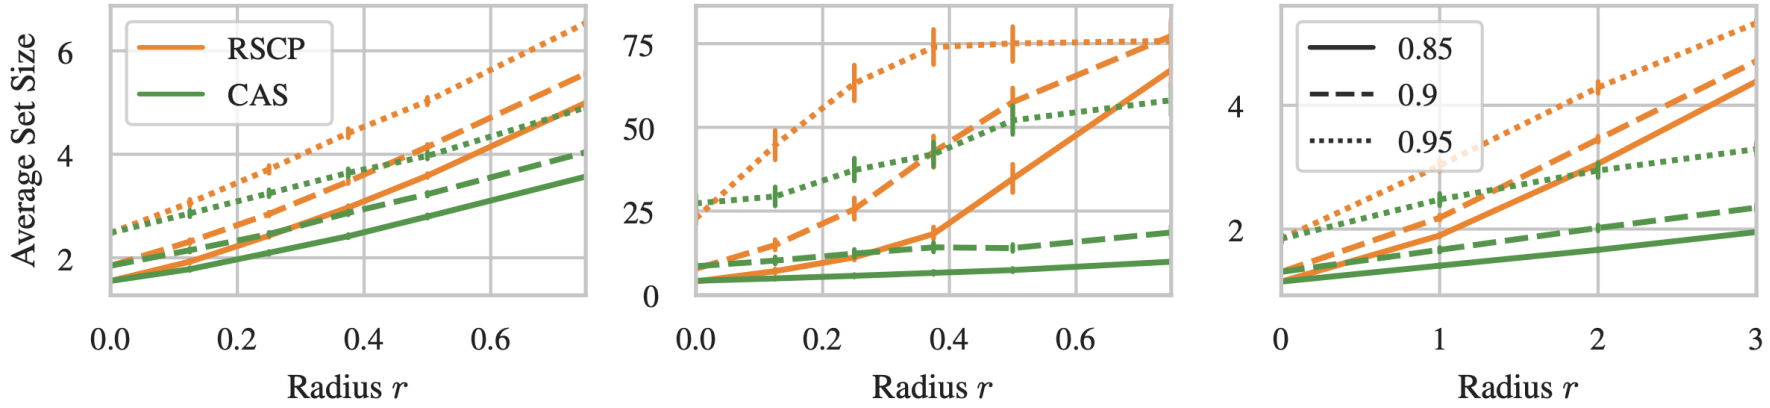

All score
All smoothing

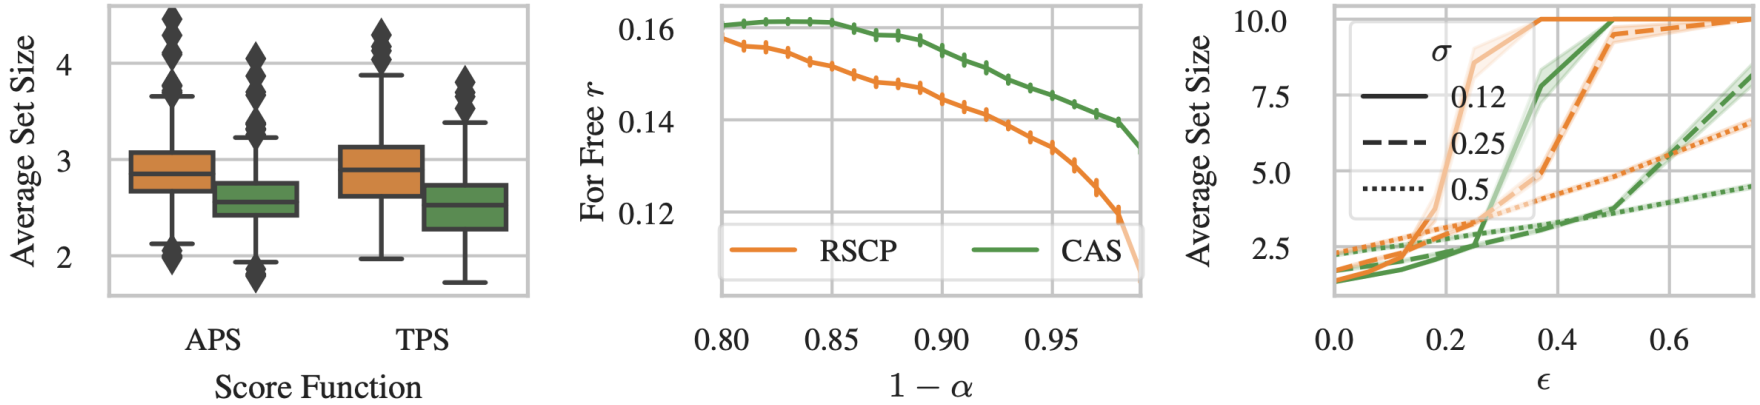

Experiments

CAS certifies a larger lower bound for vanilla CP. CAS shows smaller sets for the cal-time certificate and with correction.

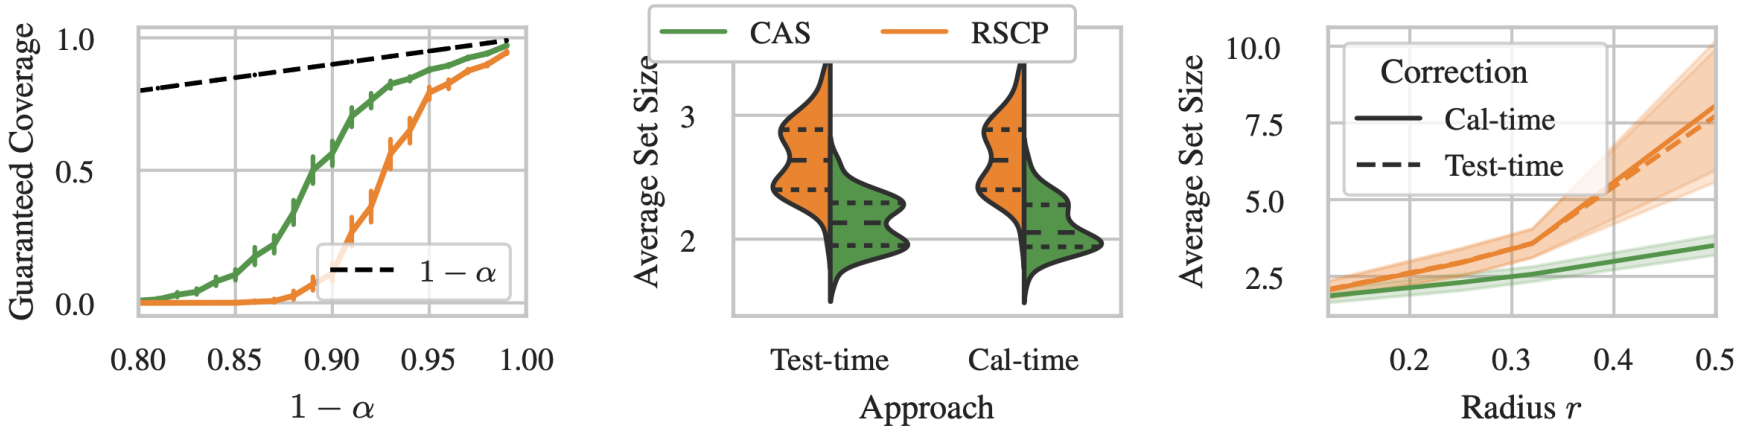

CAS is less conservative since it is closer to the nominal $1-\alpha$, and has smaller sets.

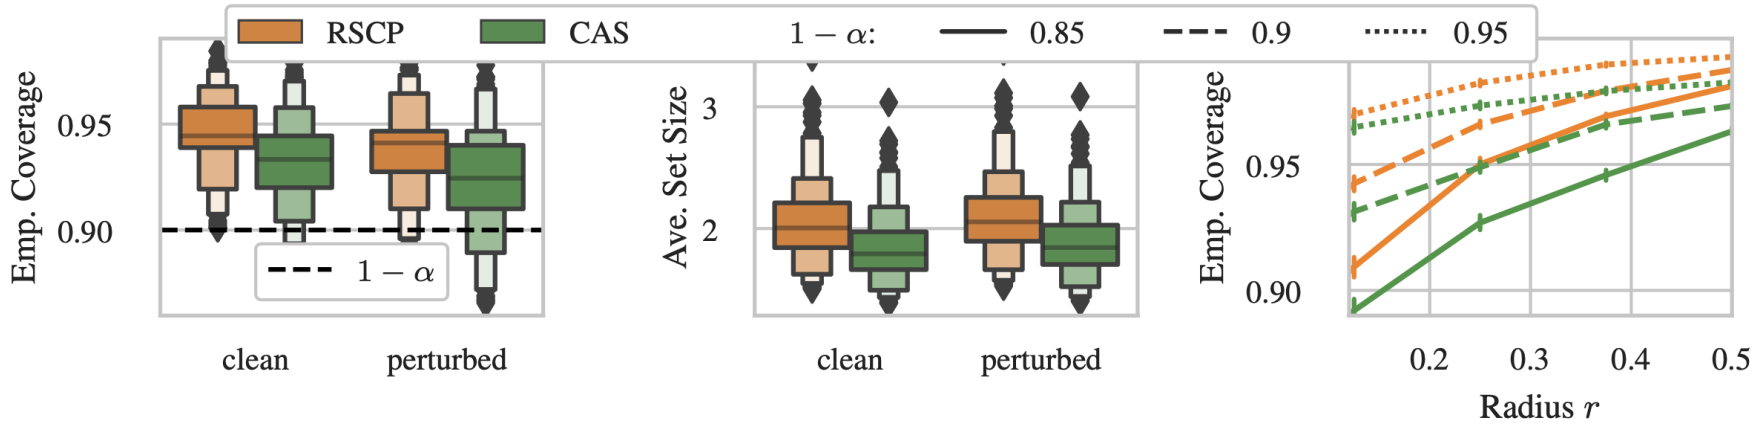

Experiments

CAS makes CP robust to feature poisoning, coming at only a small cost

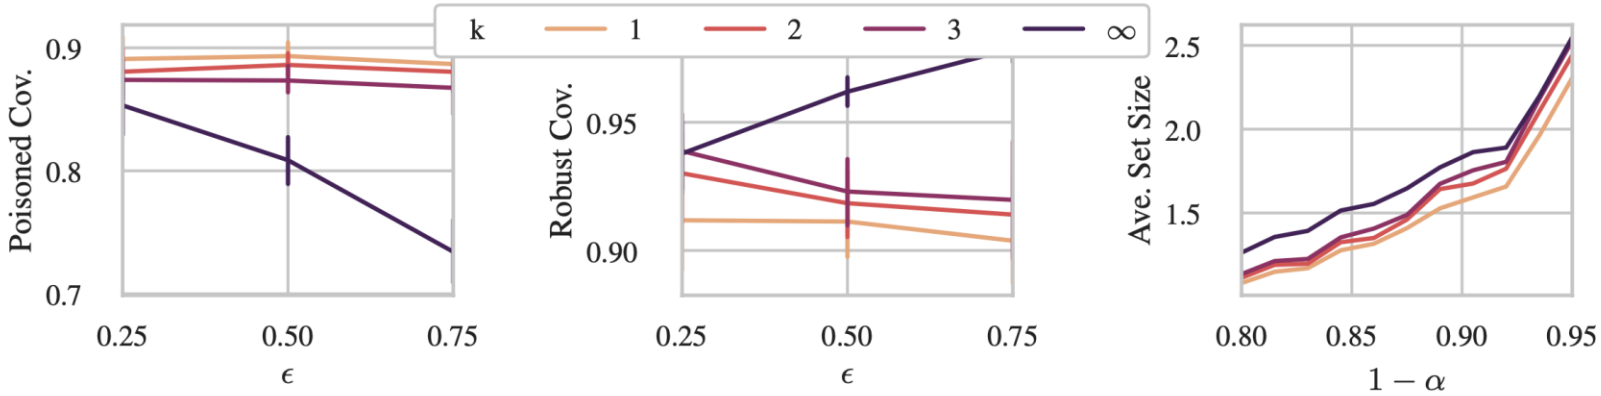

CAS makes CP robust to label poisoning, with a predictably larger set size. CAS is also faster.

Table 1. Label poisoning for `CIFAR-10`.

k	Cov. (Clean)	Vanilla Cov. (Pert)	Robust Cov. (Pert)	Set Size (Clean)
0	0.897	0.897	0.897	1.41
1	0.916	0.872	0.900	1.58
2	0.923	0.859	0.901	1.62

	Time (seconds)		No. Data
	Baseline	CAS (Ours)	
Calib.	0.15	0.79	204
Testing	2.93	0.15	100
Total	3.08	0.94	

Conclusion

Robust Yet Efficient Conformal Prediction Sets

- We provide certified robustness for conformal prediction both for evasion and poisoning attacks.

  Upper bound of score & Lower bound of calibration

- We propose a CDF-aware bound on the conformity scores under adversarial perturbation.

  $Pr[h(\xi(\tilde{\boldsymbol{x}}), y) \leq b_i] = Pr[s(\xi(\boldsymbol{x}), y) \leq b_i]$

- We generalize both results to discrete and binary (sparse) data.

  Sparse smoothing

- We show how we can correct for finite-sample error.

  Tighter DKW & Bernstein bound

Robust Yet Efficient Conformal Prediction Sets

Thank you for your attention. Looking forward to your questions!

2025 Apr 14

Supplement

Upper-bound score function

Proposition 3.1. *Define $\overline{s}(\boldsymbol{x}, y)$ as the upper bound for $\{s(\tilde{\boldsymbol{x}}, y) : \tilde{\boldsymbol{x}} \in \mathcal{B}(\boldsymbol{x})\}$. With q_α as the α-quantile of the true (clean) calibration scores, let $\overline{\mathcal{C}}_\alpha(\boldsymbol{x}) = \{y : \overline{s}(\boldsymbol{x}, y) \geq q_\alpha\}$. For all $\tilde{\boldsymbol{x}}_{n+1} \in \mathcal{B}(\boldsymbol{x}_{n+1})$, if $(\boldsymbol{x}_{n+1}, y_{n+1})$ is exchangeable with $\mathcal{D}_{\text{cal}}$ then we have $\Pr\left[y_{n+1} \in \overline{\mathcal{C}}_\alpha(\tilde{\boldsymbol{x}}_{n+1})\right] \geq 1 - \alpha$.*

Lower-bound quantile

Proposition 3.2. *Let $\underline{q_\alpha}$ to be the solution to the optimization problem in Eq. 3. With the conservative prediction sets*

$$\overline{\mathcal{C}}_\alpha(\boldsymbol{x}_{n+1}) = \left\{y_i : s(\boldsymbol{x}_{n+1}, y_i) \geq \underline{q_\alpha}\right\} \tag{4}$$

for any $(\boldsymbol{x}_{n+1}, y_{n+1})$ exchangeable with (clean) $\mathcal{D}_{\text{cal}}$ we have $\Pr\left[y_{n+1} \in \overline{\mathcal{C}}_\alpha(\boldsymbol{x}_{n+1})\right] \geq 1 - \alpha$.

Score function mean bound

Baseline bound. A straightforward approach only incorporates the expected smoothed score (mean) for the given input $\boldsymbol{x}$. Let $p = \mathbb{E}[s(\xi(\boldsymbol{x}), y)]$ for simplicity. With $\tilde{\boldsymbol{x}} \in \mathcal{B}(\boldsymbol{x})$ the baseline upper-bound for $\hat{s}(\tilde{\boldsymbol{x}}, y) = \mathbb{E}[s(\xi(\tilde{\boldsymbol{x}}), y)]$ is determined by the following problem:

$$\begin{aligned} \overline{s}_{\text{mean}}(\boldsymbol{x}, y) = \max_{h \in \mathcal{H}} \quad & \mathbb{E}[h(\xi(\tilde{\boldsymbol{x}}), y)] \\ \text{s.t.} \quad & \mathbb{E}[h(\xi(\boldsymbol{x}), y)] = p \end{aligned} \tag{7}$$

Score function CDF bound

CDF-based bound. Let $a = b_1 < b_2 \leq \cdots \leq b_{m-1} < b_m = b$ be m real numbers that partition the output space. Let $p_i = \Pr\left[s(\xi(x), y) \leq b_i\right]$. We define the problem:

$$\begin{aligned} \overline{s}_{\text{cdf}}(\tilde{\boldsymbol{x}}, y) = \max_{h \in \mathcal{H}} \quad & \mathbb{E}[h(\xi(\tilde{\boldsymbol{x}}), y)] \\ \text{s.t.} \quad \forall b_i : \quad & \Pr\left[h(\xi(\boldsymbol{x}), y) \leq b_i\right] = p_i \end{aligned} \tag{8}$$

Supplement

CDF Aware Sets (CAS)

Bound Through Smoothing. We search over the space of all possible functions (due to black-box access).

$$\max_{\tilde{\boldsymbol{x}} \in \mathcal{B}(\boldsymbol{x})} \mathbb{E}[s(\tilde{\boldsymbol{x}} + \boldsymbol{\epsilon}, y)] \leq \max_{\tilde{\boldsymbol{x}} \in \mathcal{B}(\boldsymbol{x}), h \in \mathcal{H}} \mathbb{E}[h(\tilde{\boldsymbol{x}} + \boldsymbol{\epsilon}, y)]$$

Mean Constraint. We limit the search space to all functions with the same mean around the reference point. $\mathbb{E}[h(\boldsymbol{x} + \boldsymbol{\epsilon}, y)] = \mathbb{E}[s(\boldsymbol{x} + \boldsymbol{\epsilon}, y)]$

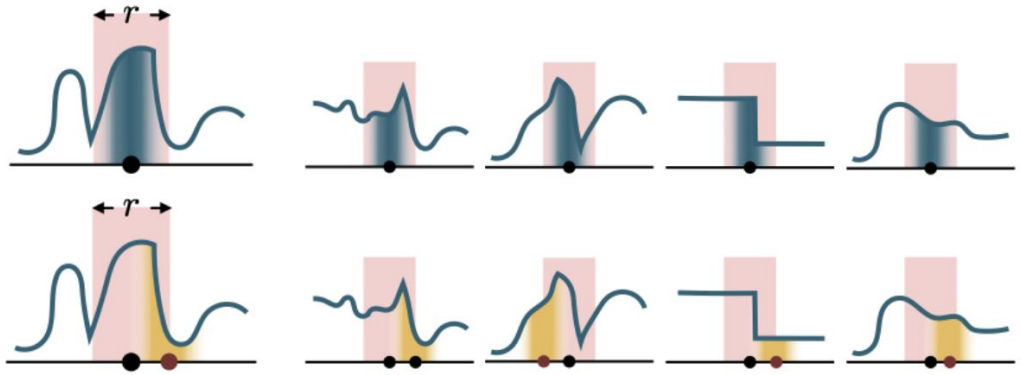

CDF Constraint. We limit the space to all functions with same CDF around the reference point. $\Pr[h(\boldsymbol{x} + \boldsymbol{\epsilon}, y) \leq b_i] = \Pr[s(\boldsymbol{x} + \boldsymbol{\epsilon}, y) \leq b_i]$

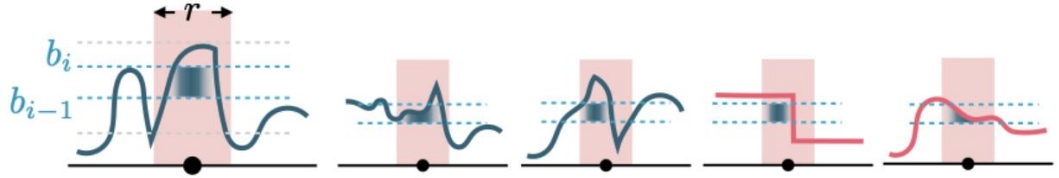

Supplement

Bound sparse smoothing

For the baseline bound we can rewrite Eq. 7 as a linear program by partitioning the input space $\mathcal{X}$ into regions of constant likelihood ratio (Lee et al., 2019). Let $\mathcal{X} = \bigcup_i \mathcal{R}_i$ and $\mathcal{R}_i \bigcap \mathcal{R}_j = \emptyset$ be a partitioning into disjoint regions of constant likelihood ratio such that for ever $\boldsymbol{z} \in \mathcal{R}_i$ it holds $\frac{\Pr[\xi(\boldsymbol{x})=\boldsymbol{z}]}{\Pr[\xi(\tilde{\boldsymbol{x}})=\boldsymbol{z}]} = c_i$ for some constant c_i. Let $t_i = \Pr[\xi(\boldsymbol{x}) \in \mathcal{R}_i]$ and $\tilde{t}_i = \Pr[\xi(\tilde{\boldsymbol{x}}) \in \mathcal{R}_i]$ for for each region $\mathcal{R}_i$. Then Eq. 7 is equivalent to:

$$\max_{\boldsymbol{h}} \boldsymbol{h}^T \tilde{\boldsymbol{t}} \quad \text{s.t.} \quad \boldsymbol{h}^T \boldsymbol{t} = p, \quad 0 \leq \boldsymbol{h} \leq 1 \tag{17}$$

where $\boldsymbol{h} \in [0,1]^I$ is the vector that we are optimizing over corresponding to the score function $h \in \mathcal{H}$, $\boldsymbol{t}$ and $\tilde{\boldsymbol{t}}$ are the vectors with t_i and $\tilde{t}_i$ as elements, and $I = r_a + r_d + 1$ is the number of regions. Note, that by replacing the constraint with $a \leq \boldsymbol{h} \leq b$ we can handle score functions that are bounded in $[a, b]$. The exact solution to this LP can be easily obtained with a simple algorithm. We visit each region in increasing order w.r.t. c_i where

$$c_i = \left[\frac{p_0}{1-p_1}\right]^{i-r_d} \left[\frac{p_1}{1-p_0}\right]^{i-r_a} \tag{18}$$

and assign $h_i = 1$ for all regions $\mathcal{R}_i$ until the budget constraint is met, and $h_i = 0$ for the remaining regions, with the exception of the region in between where h_i is a value between 0 and 1 such that the equality constraint is exactly met. Since the likelihood ratios c_i are monotonic in i, the regions are automatically sorted so the solution to the LP can be obtain in linear $O(I)$ time. See Bojchevski et al. (2020) for more details and the pseudo-code.

For the CDF-based bound we can similarly rewrite Eq. 8 as the following linear program:

$$\max_{\boldsymbol{H}} b_m - \boldsymbol{H}\tilde{\boldsymbol{t}}\boldsymbol{d} \quad \text{s.t.} \quad \boldsymbol{H}\boldsymbol{t} = \boldsymbol{p}, \quad 0 \leq \boldsymbol{H} \leq 1 \tag{19}$$

where $\boldsymbol{H} \in [0,1]^{(m-1)\times I}$ is the matrix that we are optimizing over with H_{ji} being the score that we assign to the j-th bin and the i-th region, $\boldsymbol{d}$ is the vector of bin widths such that $d_j = b_j - b_{j-1}$, and $\boldsymbol{p}$ is a vector where $p_i = \Pr[s(\xi(\boldsymbol{x}), y) \leq b_i]$. Intuitively, for each bin and each region the worst-case score function $h \in \mathcal{H}$ assigns the same score to all $\boldsymbol{z}$ in that region since the likelihood ratio is constant. As before we have a simple algorithm to obtain the exact solution to this LP. Observe that Eq. 19 can be decomposed into $m-1$ separate LPs similar to Eq. 17 which can be solved in parallel using the same algorithm as above. The reason is that there is no interaction between the different bins (different rows of $\boldsymbol{H}$) in neither the constraint nor the objective function. Therefore, the solution can be obtain in $O(m \times I)$ with serial computation and $O(I)$ with parallel computation.

Supplement

Calibration-time

Proposition 5.1. *For $\tilde{\boldsymbol{x}}_{n+1} \in \mathcal{B}(\boldsymbol{x}_{n+1})$ and $(\boldsymbol{x}_{n+1}, y_{n+1})$ exchangeable with $\mathcal{D}_{\text{cal}}$, define*

$$\underline{q}_{\alpha} = \text{Quant}\left(\alpha; \{\underline{s}(\boldsymbol{x}_i, y_i) : (\boldsymbol{x}_i, y_i) \in \mathcal{D}_{\text{cal}}\}\right) \tag{11}$$

For prediction sets $\overline{\mathcal{C}}_{\alpha}(\tilde{\boldsymbol{x}}_{n+1}) = \{y : \hat{s}(\tilde{\boldsymbol{x}}_{n+1}, y) \geq \underline{q}_{\alpha}\}$ we have $\Pr[y_{n+1} \in \overline{\mathcal{C}}_{\alpha}(\tilde{\boldsymbol{x}}_{n+1})] \geq 1 - \alpha$. Moreover, the vanilla CP covers the true label with probability $\geq 1 - \beta$ for

$$\beta = \text{Quant}^{-1}\left(q_{\alpha}; \{\underline{s}(\boldsymbol{x}_i, y_i) : (\boldsymbol{x}_i, y_i) \in \mathcal{D}_{\text{cal}}\}\right) \tag{12}$$

and $\text{Quant}^{-1}\left(t; A\right) = \min\{\tau' : \text{Quant}\left(\tau'; A\right) \geq t\}$.

Finite sample correction - Evasion

Proposition 6.1. *Let $\underline{s}_{\text{cdf}+}(\boldsymbol{x}_i, y_i) \leq \underline{s}_{\text{cdf}}(\boldsymbol{x}_i, y_i)$ hold with $1 - \eta/(2|\mathcal{D}_{\text{cal}}|)$ probability for each $(\boldsymbol{x}_i, y_i) \in \mathcal{D}_{\text{cal}}$, and $\hat{s}_{+}(\tilde{\boldsymbol{x}}_{n+1}, y_{n+1}) \geq \hat{s}(\tilde{\boldsymbol{x}}_{n+1}, y_{n+1})$ hold with $1 - \eta/(2|\mathcal{Y}|)$ probability. Define the conservative $\underline{q}_{\alpha+} = \text{Quant}\left(\alpha - \eta; \left\{\underline{s}_{\text{cdf}+}(\boldsymbol{x}_i, y_i) : (\boldsymbol{x}_i, y_i) \in \mathcal{D}_{\text{cal}}\right\}\right)$ and $\overline{\mathcal{C}}_{\alpha+}(\boldsymbol{x}_{n+1}) = \{y : \hat{s}_{+}(\boldsymbol{x}_{n+1}, y) \geq \underline{q}_{\alpha+}\}$. Then*

$$\Pr[y_{n+1} \in \overline{\mathcal{C}}_{\alpha+}(\tilde{\boldsymbol{x}}_{n+1})] \geq 1 - \alpha \tag{13}$$

Finite sample correction - Poison

Proposition 6.2. *Let $\underline{s}_{\text{cdf}+}(\tilde{\boldsymbol{x}}_i, y_i) \leq \underline{s}_{\text{cdf}}(\tilde{\boldsymbol{x}}_i, y_i)$ hold with $1 - \eta/(2|\mathcal{D}_{\text{cal}}|)$ probability for all $(\tilde{\boldsymbol{x}}_i, y_i) \in \mathcal{D}_{\text{cal}}$. Let $\underline{q}_{\alpha+}$ be the solution to Eq. 24 (Eq. 3 with CDF bounds) for $\alpha = \alpha' - \eta$ with $\underline{s}_i = \underline{s}_{\text{cdf}+}(\tilde{\boldsymbol{x}}_i, y_i) - \epsilon_{\text{hoef}}$. Then for each new test point $\boldsymbol{x}_{n+1}$ exchangeable with $\mathcal{D}_{\text{cal}}$ the prediction set defined as $\overline{\mathcal{C}}(\boldsymbol{x}_{n+1}) = \{y : s_{\text{mc}}(\boldsymbol{x}_{n+1}, y) \geq \underline{q}_{\alpha+}\}$ has $1 - \alpha'$ coverage.*

Supplement

Finite sample correction bound

any adjustable $1-\eta$ probability;

$$|\mathbb{E}[z] - \mathbb{E}_m[z]| \leq \sqrt{\frac{\log(\frac{2}{\eta})}{2m}}$$

This bound only accesses to the empirical mean and not the samples. Therefore, in cases where we want to account for the distance of empirical mean, and the true expectation for an unknown variable, we can use this bound. An example of this case is the test-time correction where the upper bound on the mean of the unseen point is computed while the empirical mean is not computable (since there are no samples).

Let σ_m^2 be the variance of the MC samples, then empirical Bernstein inequality produces variance-dependent confidence intervals as following:

$$|\mathbb{E}[z] - \mathbb{E}_m[z]| \leq \sqrt{2\sigma_m^2 \frac{\ln(\frac{4}{\eta})}{m}} + \frac{7\ln(\frac{4}{\eta})}{3(m-1)}$$

Similar to the mean, the empirical CDF is also bounded between an upper and a lower CDF, via the Dvoretzky–Kiefer–Wolfowitz (DKW) inequality (Dvoretzky et al., 1956). Let $F(b_i) = \Pr[z \leq b_i]$ and $F_m(b_i) = \sum_{j=1}^{m} \mathbf{1}[z_j \leq b_i]$,

$$|F(b_i) - F_m(b_i)| \leq \sqrt{\frac{\log(\frac{2}{\eta})}{2m}}$$

The above inequality holds simultaneously for all b_i.

For Eq. 7 we use the Bernstein inequality as is has shown a better empirical result compare to Hoeffding's inequality. For Eq. 8 we use the DKW inequality to find confidence intervals the empirical CDF.