Leveraging Geometry for Conformal Prediction via Canonicalization

uai2025

Putri van der Linden, Alexander Timans, Erik Bekkers

Amsterdam Machine Learning Lab (AMLab), University of Amsterdam

Motivation

Under geometric distribution shifts, pretrained models often deteriorate, resulting in poor performance and uninformative prediction intervals.

	mAP	C4-mAP
Mask-RCNN	47.81	12.79

Tab: Zero-shot Mask-RCNN segmentation performance (mAP) on regular and C4-rotated COCO data.

We propose using geometric information to preserve exchangeability in conformal prediction under distribution shifts, acquired via *canonicalization* [1,2].

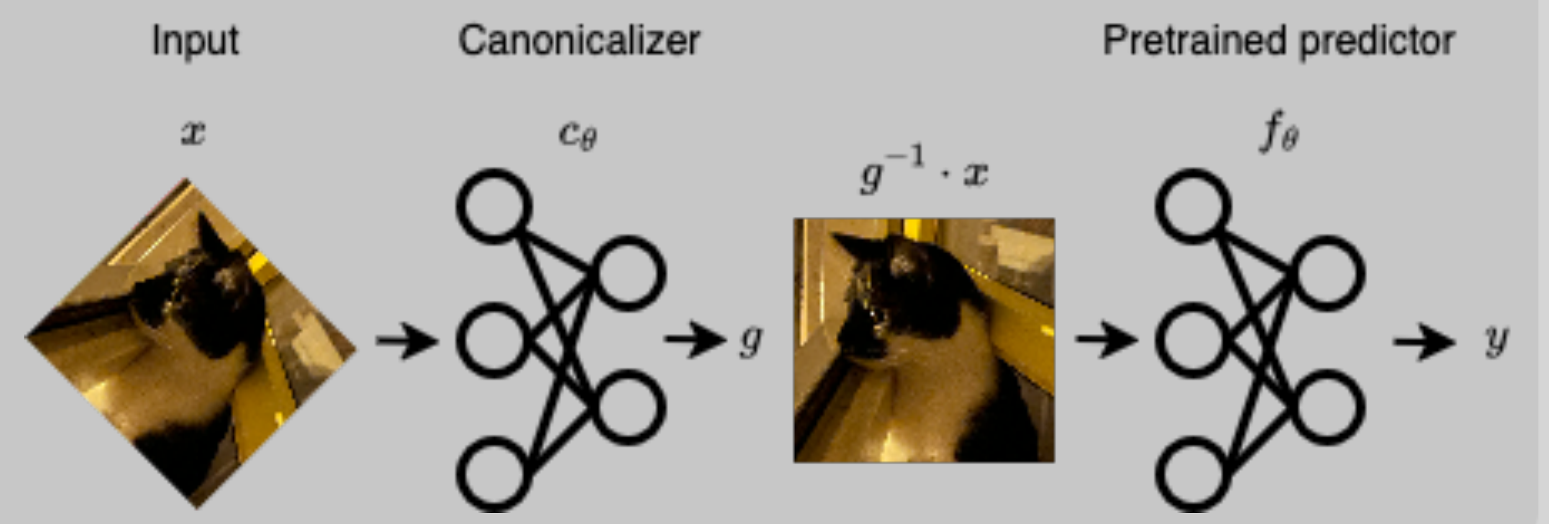

Geometry for Conformal Prediction

- Conformal prediction requires *exchangeability* — invariance to the permutation group $\mathbb{S}_n$ — between calibration and test data to guarantee valid coverage.
- Under *geometric shifts*, calibration and test samples undergo transformations $g_i \in G$, leading to a distribution over G^n-shifted data: $x_i \to g_i \cdot x_i$
- Consistent calibration and test transformations maintain (shifted) exchangeability, but conformal prediction intervals become uninformative due to poor performance of the pretrained predictor (data—model misalignment).
- A *canonicalization network* c_θ restores predictor invariance to G by neutralizing the transformation: $c_\theta(g_i \cdot x_i)^{-1} \cdot g_i \cdot x_i = c_\theta(x_i)^{-1} \cdot x_i$
- This enforces exchangeability over $\mathbb{S}_n \times G^n$, and ensures that the distribution over canonicalized samples remains invariant under G^n.

Experiments

Estimated group elements g can inform and robustify Conformal Prediction in different ways.

Train	Calibration	Test	Robustness
$\delta(e)$	$P_{G\|x}$	$P_{G\|x}$	Canon + CP
$\delta(e)$	$P_{G\|x}$	$w(x) \cdot P_{G\|x}$	Canon + CP
$\delta(e)$	$P_{G\|x}$	$P_{G_{new}\|x}$	Canon + WCP

Robustness to Geometric Data Shifts

Model	No Shift			$C8$ Rotation Shift		
	Acc ↑	Coverage	Set size ↓	Acc ↑	Coverage	Set size ↓
$\hat{f}_\theta$	71.66	95.09±.3	6.21±.1	33.68	95.63±.1	61.40±.5
$\hat{f}_\theta + SO(2)$	60.13	95.02±.5	11.36±.4	58.27	95.00±.4	11.21±.3
$\hat{f}_\theta + C8$	62.53	95.38±.5	10.76±.4	60.82	95.26±.2	10.97±.2
$CP^2\,(G=8)$	65.46	95.34±.4	11.20±.4	63.94	94.82±.3	11.30±.4

Tab: Accuracy and conformal metrics on CIFAR-100 for target coverage of 95% under G-shift. See our paper for additional CIFAR-10 (image) and ModelNet40 (point cloud) results.

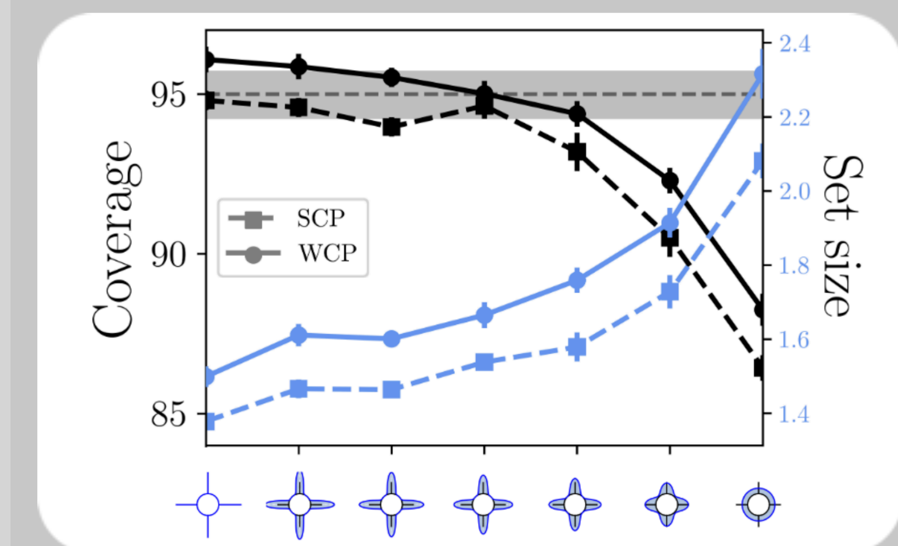

Weighted Conformal Prediction

Can be used to select calibration samples with similar orientation, delaying coverage breakdown under additional geometric shifts.

Fig: Induced gradual shift from C4 to SO(2).

Conditional Group Distributions

Conditional G-shifts can be detected (data diagnostics) and used downstream for better partition-conditional coverage (by proxy).

Fig: True (top) and recovered (bottom) conditional shifts.

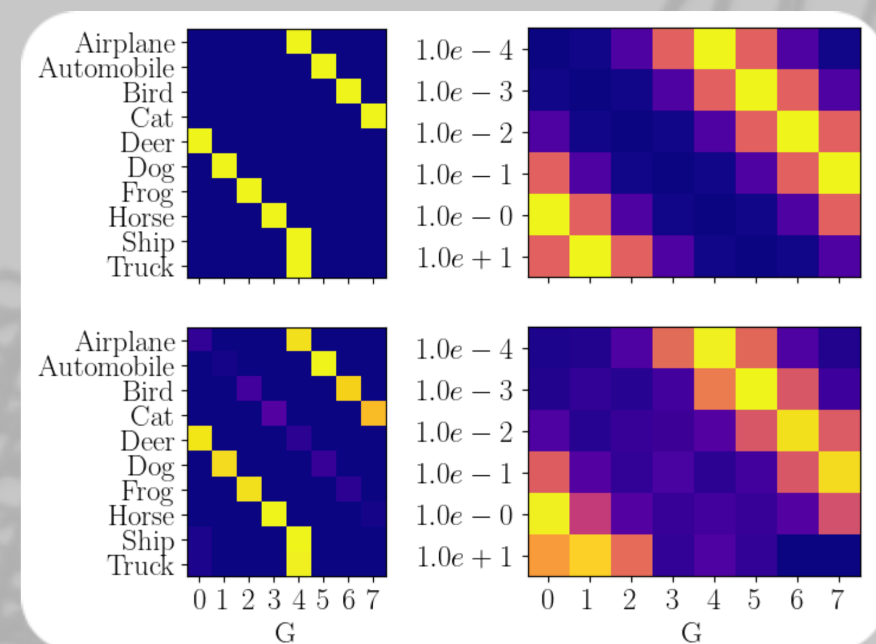

References

[1] Equivariance with Learned Canonicalization Functions. S. O. Kaba, A. K. Mondal, Y. Zhang, Y. Bengio, S. Ravanbakhsh (ICML 2023)
[2] Equivariant Adaption of Large Pretrained Models. A. K. Mondal, S. S. Panigrahi, S. O. Kaba, S. Rajeswar, S. Ravanbakhsh (NeurIPS 2023)
[3] A tutorial on Conformal Prediction. G. Shafer, V. Vovk (JMLR 2008)

AMLAB Amsterdam Machine Learning Lab

University of Amsterdam