Adaptive Bounding Box Uncertainties via Two-Step Conformal Prediction

Alexander Timans

Amsterdam Machine Learning Lab,
University of Amsterdam
In collab. with the Bosch Center for AI

October 1, 2024

Collaborators

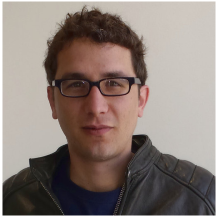

Christoph-Nikolas Straehle,
Bosch Center for AI

Kaspar Sakmann,
Bosch Center for AI

Eric Nalisnick,
Johns Hopkins
University

Why uncertainty for bounding boxes?

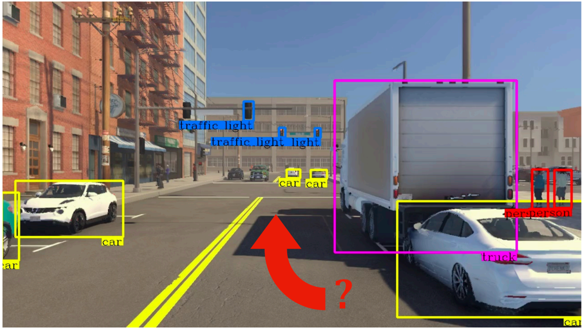

https://developer.nvidia.com/blog/deploying-a-scalable-object-detection-pipeline-the-inferencing-process-part-2/

Why uncertainty for bounding boxes?

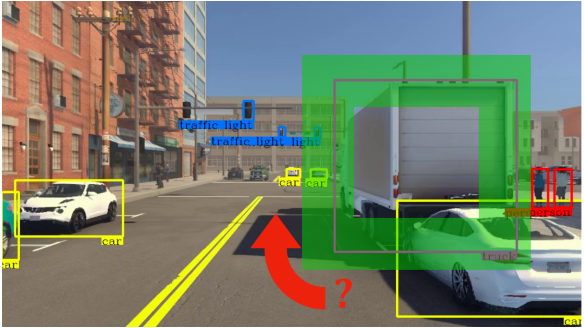

https://developer.nvidia.com/blog/deploying-a-scalable-object-detection-pipeline-the-inferencing-process-part-2/

Why uncertainty for bounding boxes?

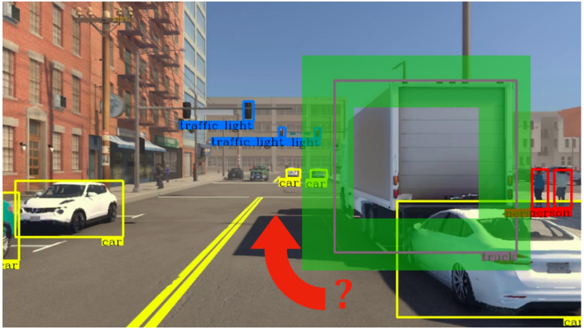

https://developer.nvidia.com/blog/deploying-a-scalable-object-detection-pipeline-the-inferencing-process-part-2/

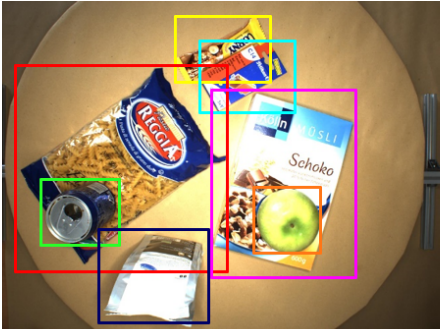

https://www.mvtec.com/de/technologien/deep-learning/deep-learning-methoden/objektdetektion

Why uncertainty for bounding boxes?

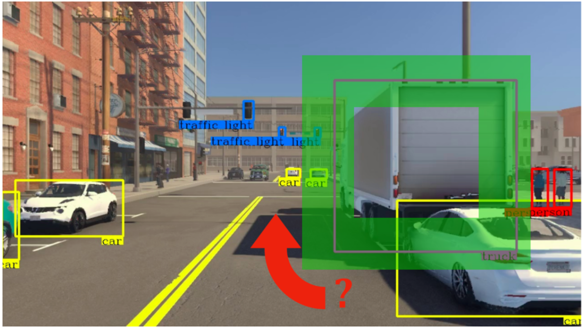

https://developer.nvidia.com/blog/deploying-a-scalable-object-detection-pipeline-the-inferencing-process-part-2/

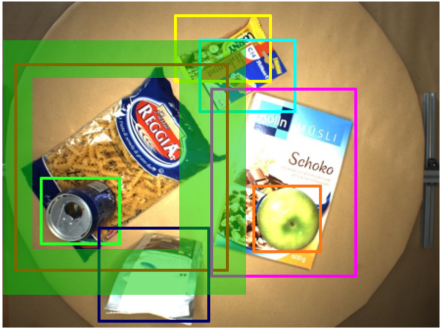

https://research.nvidia.com/publication/2018-09-deep-object-pose-estimation-semantic-robotic-grasping-household-objects

Conformal prediction

Statistical framework for uncertainty quantification via prediction sets or intervals.

Conformal prediction

Statistical framework for uncertainty quantification via prediction sets or intervals.

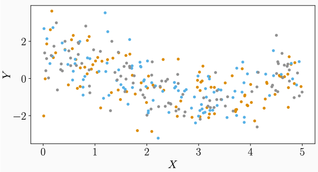

1. Data splitting

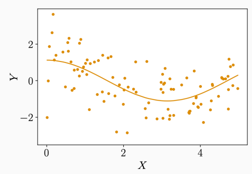

2. Train predictor

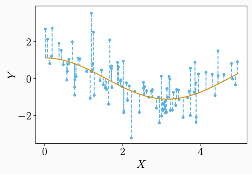

3. Calibrate predictor

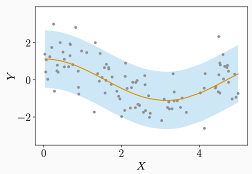

4. Test prediction intervals

Images from Zaffran, Margaux, et al. Adaptive conformal predictions for time series. ICML, 2022.

Conformal prediction

Statistical framework for uncertainty quantification via prediction sets or intervals.

Pretrained model $\hat{f}(x)$

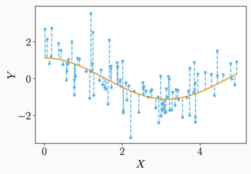

3. Calibrate predictor

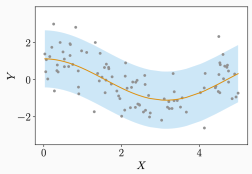

4. Test prediction intervals

Images from Zaffran, Margaux, et al. Adaptive conformal predictions for time series. ICML, 2022.

Conformal prediction

Statistical framework for uncertainty quantification via prediction sets or intervals.

Pretrained model $\hat{f}(x)$

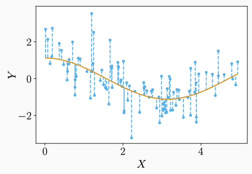

3. Calibrate predictor

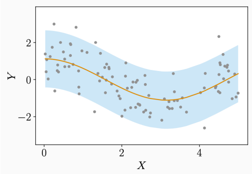

4. Test prediction intervals

Images from Zaffran, Margaux, et al. Adaptive conformal predictions for time series. ICML, 2022.

Coverage guarantee (Vovk, 2005)

“On average, the conformal prediction interval contains the true test point with user-defined probability $(1-\alpha)$*”*

Application to object detection

Data

Image: X_i

Object bounding box: $Y_i = (c^1, c^2, c^3, c^4, \ell)_i$

One image $\longrightarrow$ multiple objects and classes

Application to object detection

Data

Image: X_i

Object bounding box: $Y_i = (c^1, c^2, c^3, c^4, \ell)_i$

One image $\longrightarrow$ multiple objects and classes

Our coverage guarantee for bounding boxes

*"On average **per object class**, the conformal bounding box interval covers the object's true bounding box with user-defined probability* $(1-\alpha)$*"*

Protecting against misclassification

Challenge: The object's class label prediction also exhibits uncertainty, which we want to incorporate when constructing the bounding box intervals.

Protecting against misclassification

Challenge: The object's class label prediction also exhibits uncertainty, which we want to incorporate when constructing the bounding box intervals.

Solution: First quantify predictive uncertainty in class labels with conformal prediction (again), resulting in a **two-step** conformal approach robust to object misclassification.

Protecting against misclassification

Challenge: The object's class label prediction also exhibits uncertainty, which we want to incorporate when constructing the bounding box intervals.

Solution: First quantify predictive uncertainty in class labels with conformal prediction (again), resulting in a **two-step** conformal approach robust to object misclassification.

Pass test image through predictor.

Step 1: Get conformal label set.

Step 2: Use Step 1 to build conformal bounding box interval.

Results: visual

Bounding box intervals are tight and scale adaptively with class instances.

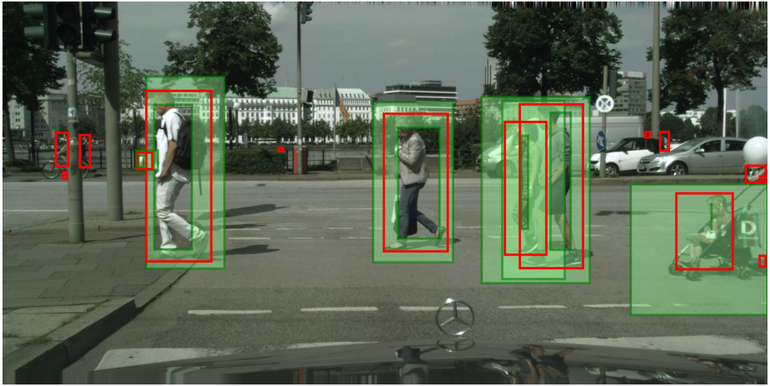

On Cityscapes test image for class 'person'.

Results: visual

More visuals for classes on COCO, Cityscapes and BDD100k.

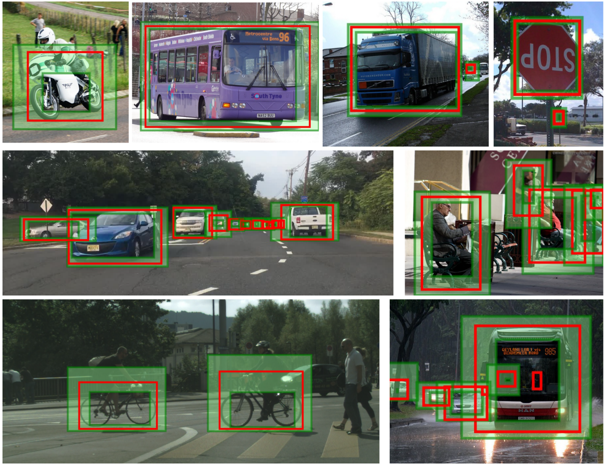

Results: coverage

Check: Is the coverage guarantee of $(1 - \alpha) = 90\%$ satisfied?

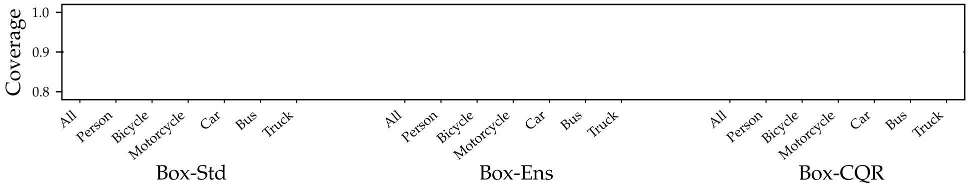

Results: coverage

Yes, guarantees hold per class on all datasets.

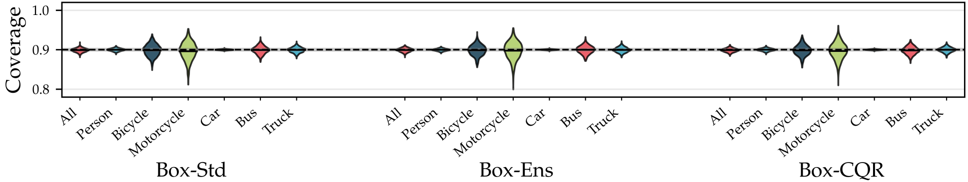

Empirical coverage vs. target coverage $(1-\alpha)=0.9$ for BDD100k.

Results: model comparison

Check: Are the guarantees model-agnostic?
Yes, guarantees hold across different 'black-box' pretrained object detectors.

Uncertainty method	Object detector	Two-sided box intervals *MPIW*	Two-sided box intervals *Cov*	One-sided box intervals *MPIW*	One-sided box intervals *Cov*
DeepEns	5× Faster R-CNN	12.31 ± 0.47	0.21 ± 0.01	74.15 ± 2.01	0.49 ± 0.01
GaussianYOLO	YOLOv3	7.00 ± 0.14	0.08 ± 0.01	87.07 ± 4.25	0.35 ± 0.01
Andéol *et al.* (Best)	Faster R-CNN	N/A		87.62 ± 1.79	0.91 ± 0.01
	YOLOv3	N/A		107.93 ± 4.85	0.92 ± 0.02
	DETR	N/A		82.21 ± 1.64	0.90 ± 0.01
	Sparse R-CNN	N/A		79.35 ± 1.78	0.91 ± 0.01
Box-Std (Ours)	Faster R-CNN	55.47 ± 2.97	0.88 ± 0.02	85.42 ± 1.99	0.88 ± 0.02
	YOLOv3	61.73 ± 3.66	0.88 ± 0.02	103.12 ± 3.95	0.88 ± 0.02
	DETR	45.34 ± 3.33	0.88 ± 0.02	80.57 ± 1.78	0.88 ± 0.01
	Sparse R-CNN	41.92 ± 2.16	0.89 ± 0.01	77.33 ± 1.72	0.89 ± 0.01

Conclusion

If your point estimates are alone and in need of reliable uncertainty intervals, consider **conformal prediction** – model-agnostic, distribution-free and efficient.

COCO-val image #522

Thanks for your attention!

To learn more:

Visit our poster in the next session (Poster session #1, Tue 10:30 – 12:30)

Take a look at our paper on arXiv

Go to arXiv

Catch me for a chat ☺

https://alextimans.github.io/

The conformal recipe

Given:

- a prediction model $\hat{f}$ fitted on some training dataset $\mathcal{D}_{train}$
- an unseen calibration dataset $\{(X_i, Y_i)\}_{i=1}^n$ and an unseen test sample (X_{n+1}, Y_{n+1})
- a scoring function $s : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ and resulting scores $s_i = s(X_i, Y_i)$, $i = 1, \ldots, n$
- a conformal quantile $\hat{q}$ defined as the $\left\lceil \frac{(n+1)(1-\alpha)}{n} \right\rceil$ quantile of $s_1, \ldots, s_n$
- a desired coverage rate $(1 - \alpha) \in [0, 1]$
- a prediction set for X defined as $\hat{C}(X) = \{y \in \mathcal{Y} : s(X, y) \leq \hat{q}\}$

Marginal coverage guarantee

Assuming the samples $\{(X_i, Y_i)\}_{i=1}^{n+1}$ are exchangeable, we have that

$$\mathbb{P}(Y_{n+1} \in \hat{C}(X_{n+1})) \geq 1 - \alpha.$$

Proof (simplified notation)

- Without loss of generality assume sorted scores $s_1 < \cdots < s_n$
- By exchangeability of $\{(X_i, Y_i)\}_{i=1}^{n+1}$ we observe that for any choice of $k \in \{1, \ldots, n\}$ we have $\mathbb{P}(s_{n+1} \leq s_k) = \frac{k}{n+1}$
- Then, it follows:

$$
\begin{aligned}
\mathbb{P}(Y_{n+1} \in \hat{C}(X_{n+1})) &= \mathbb{P}(s_{n+1} \leq \hat{q}) \\
&= \mathbb{P}(s_{n+1} \leq s_{\lceil (n+1)(1-\alpha) \rceil}) \\
&= \left\lceil \frac{(n+1)(1-\alpha)}{n+1} \right\rceil \\
&\geq \frac{(n+1)(1-\alpha)}{n+1} = 1 - \alpha.
\end{aligned}
$$

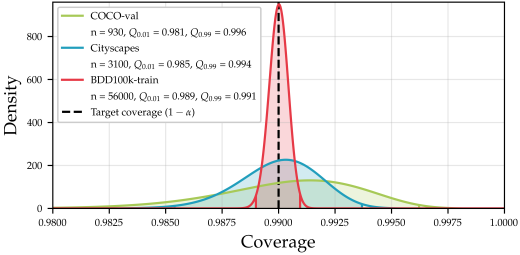

Coverage distribution conformal label sets

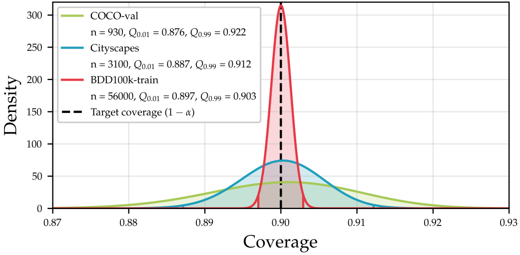

Coverage distribution conformal box intervals