Conformal Prediction for Testing and Evaluation of Intelligent/ML Systems

AI4SE & SE4AI Workshop

September 27, 2023

Agenda

- Problem Formulation
- Introduction to Conformal Prediction
- Why Conformal?
- Extensions

Problem Formulation

- Predict Amount of Fuel Needed
- Observations (x_i, y_i) for $i = 1, \ldots, 100$
- Target – RMSE < 6gal

Vehicles (x1000)	Gas Consumed Gal (x1000)
2.953	17.363
2.901	15.517
1.125	8.158
5.984	27.345
1.856	11.646

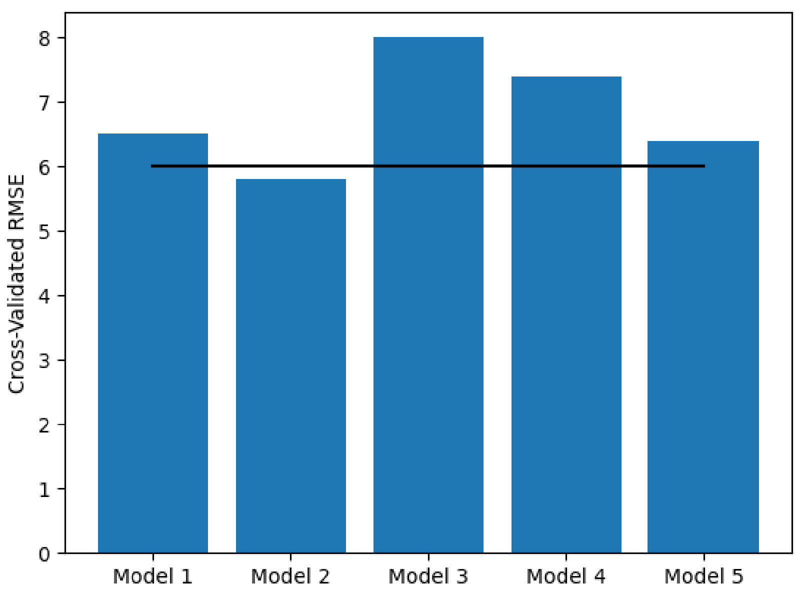

Problem Formulation

- Validation - The process of determining the degree to which a model is an accurate representation of the real world from the perspective of the intended uses of the model*

*Ben H Thacker, Scott W Doebling, Francois M Hemez, Mark C Anderson, Jason E Pepin, and Edward A Rodriguez. 2004. Concepts of model verification and validation. (2004).

Georgia Tech
CREATING THE NEXT

Problem Formulation

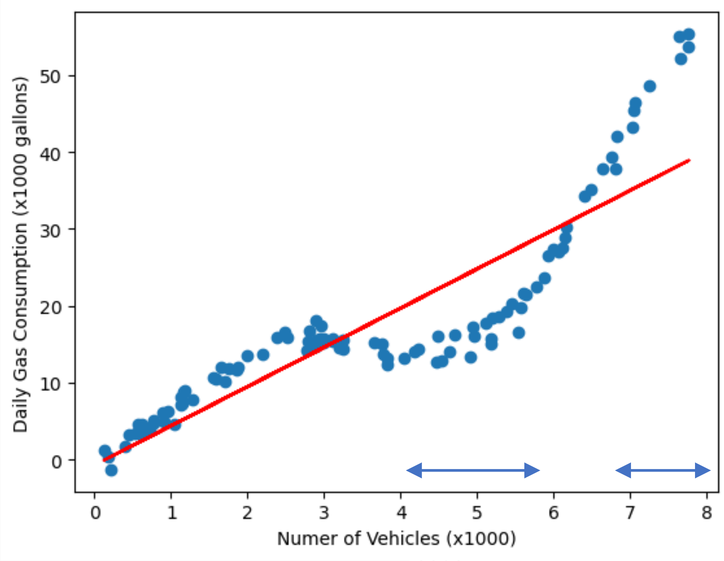

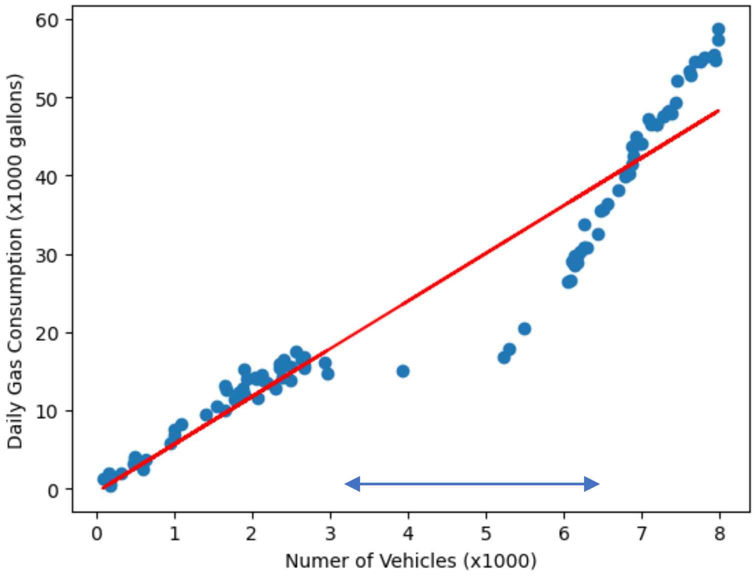

- What do you do when you move to higher dimensions?

Problem Formulation

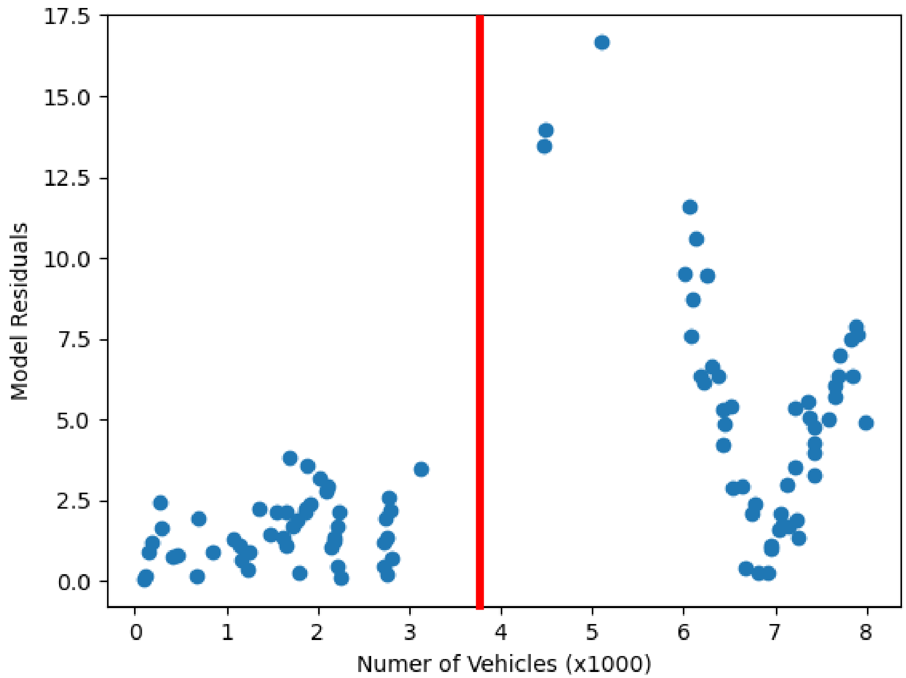

Range 1

Range 2

Range 3

Range 4

Problem Formulation

	Model 1	**Model 2**
	- Golden Retriever	- Golden Retriever
	- Golden Retriever	- Golden Retriever
	- Golden Retriever	- Golden Retriever

Georgia Tech
CREATING THE NEXT

Problem Formulation

	Model 1	**Model 2**
	- Golden Retriever: **75%**	- Golden Retriever: **95%**
	- Golden Retriever: **70%**	- Golden Retriever: **90%**
	- Golden Retriever: **80%**	- Golden Retriever: **30%**

Problem Formulation

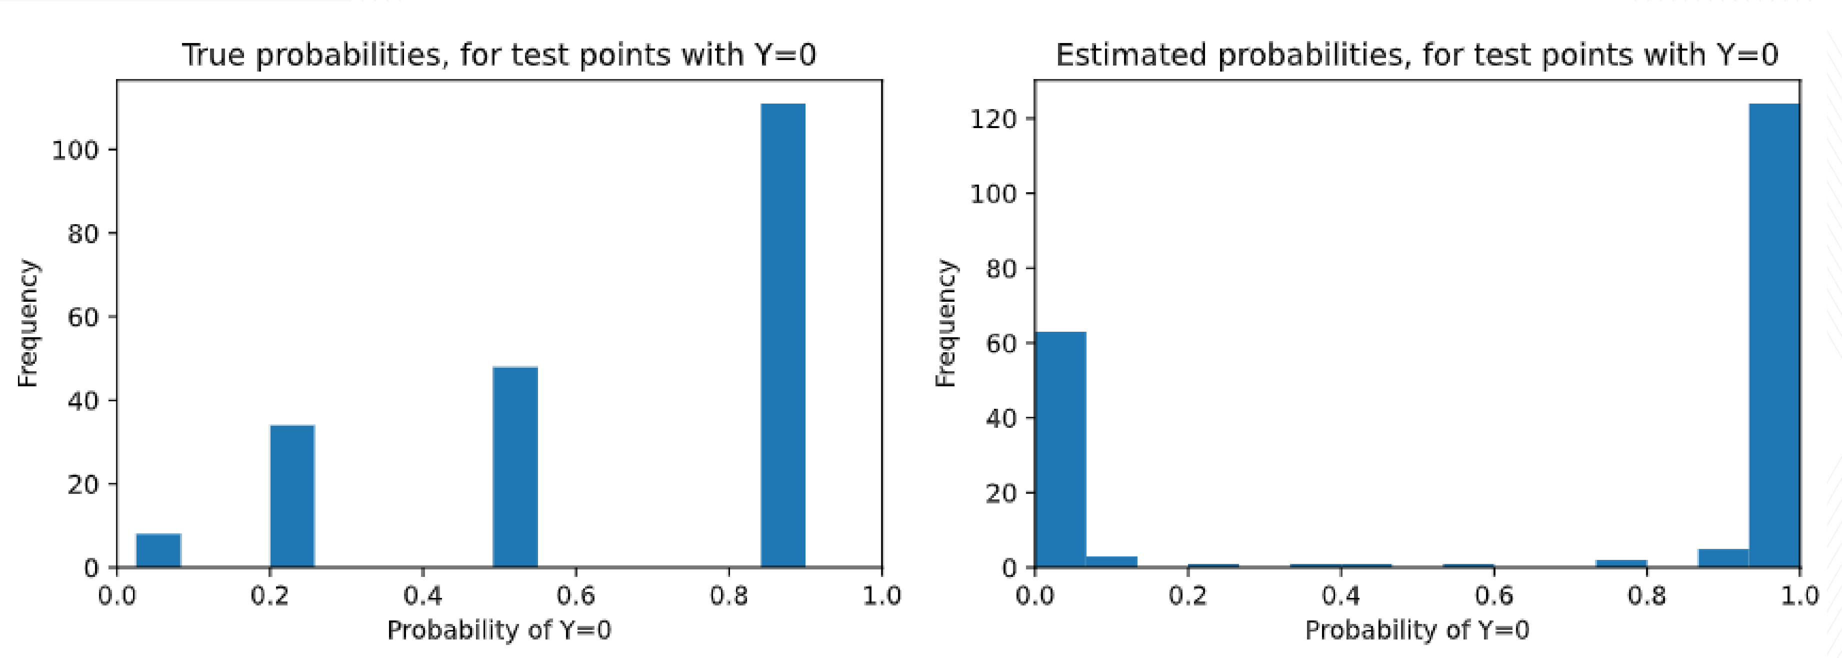

- Calibration* - the probability associated with a predicted label should reflect its true likelihood

*Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In International conference on machine learning. PMLR, 1321–1330.

Introduction to Conformal Prediction

- Y_i – true label
- $\hat{C}_\alpha(X_i)$ – prediction interval or set
- α – confidence value

$$\mathbb{P}(Y_i \in \hat{C}_\alpha(X_i)) \geq 1 - \alpha$$

- Two method branches:
  - Full and $\boxed{\text{Split}}$

V. Vovk, A. Gammerman, and G. Shafer, Algorithmic Learning in a Random World. Springer, 2005.

Introduction to Conformal Prediction

- y_i – true label
- $\hat{q}$ – (1 - α) quantile for observed scores

$$\hat{C}_\alpha(x_i) = \{y | s(x_i, y) \leq \hat{q}\}$$

$$s(x_i, y_i) = (\hat{f}(x_i) - y_i)^2$$

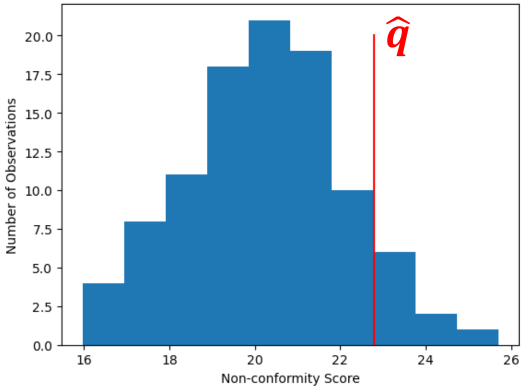

https://www.uv.es/~bernardo/Exchangeability.pdf

Introduction to Conformal Prediction

- Marginal Coverage Guarantee

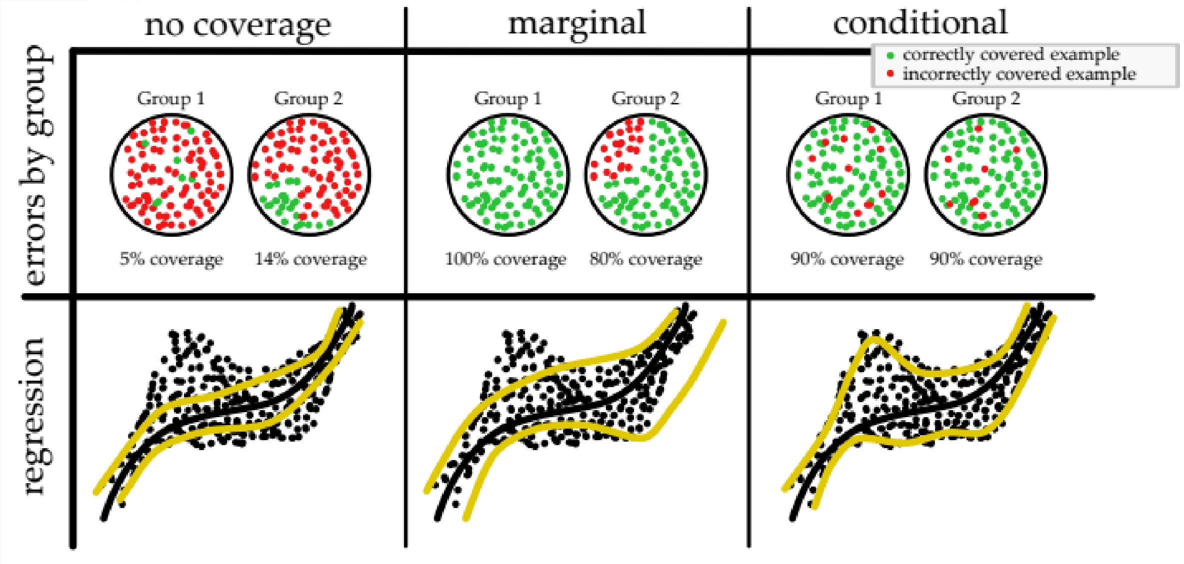

Anastasios N. Angelopoulos and Stephen Bates. 2022. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. arXiv:2107.07511

Introduction to Conformal Prediction

- Additional data

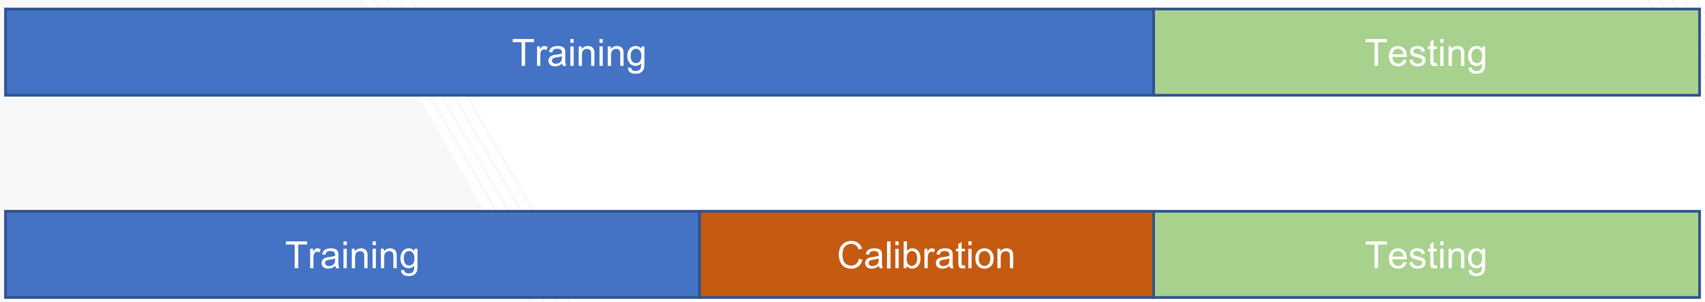

- Dependent on underlying algorithm

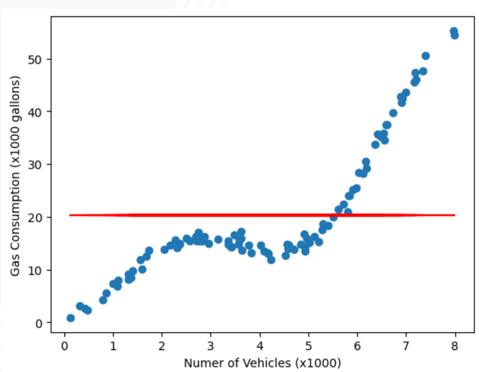

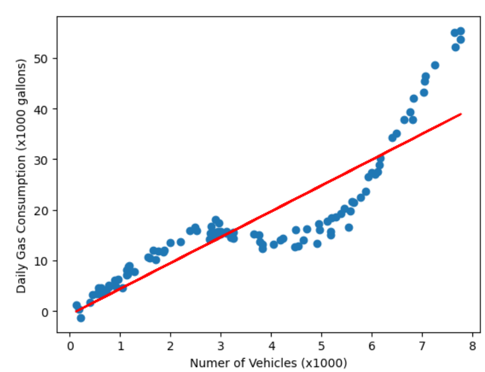

Why (Split) conformal?

- Calibrated with guarantee

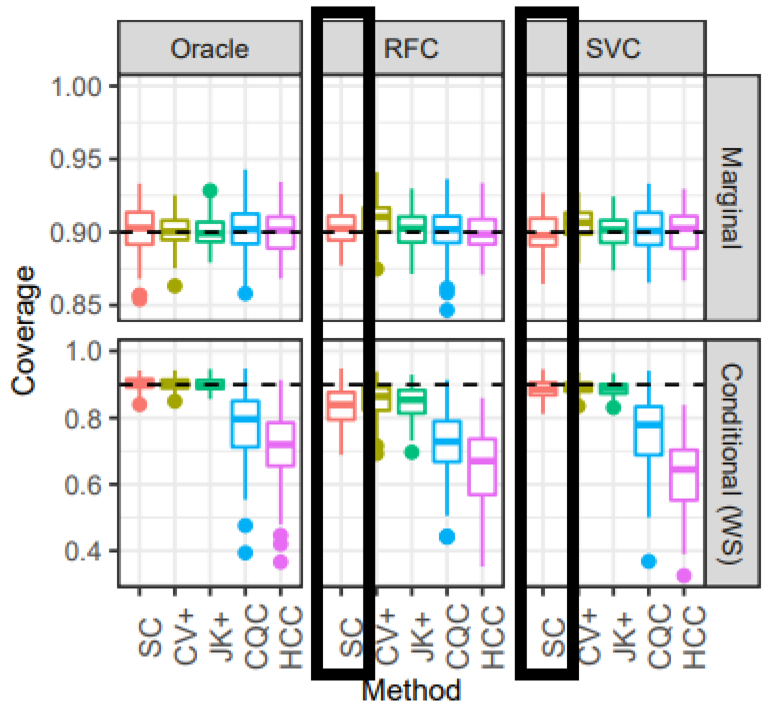

Yaniv Romano, Matteo Sesia, and Emmanuel Candes. 2020. Classification with valid and adaptive coverage. Advances in Neural Information Processing Systems 33 (2020), 3581–3591.

Why (Split) conformal?

- No additional model training
- Wrapper for any model

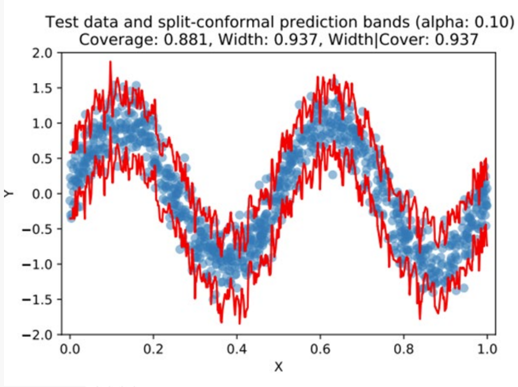

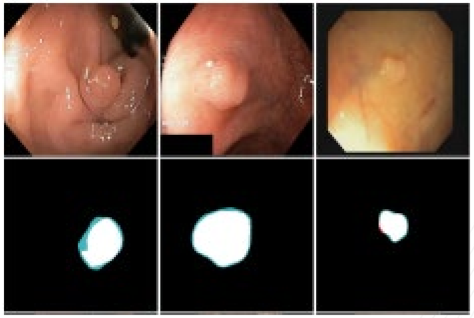

1,2,4 - Anastasios N. Angelopoulos and Stephen Bates. 2022. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification. arXiv:2107.07511

Why (Split) conformal?

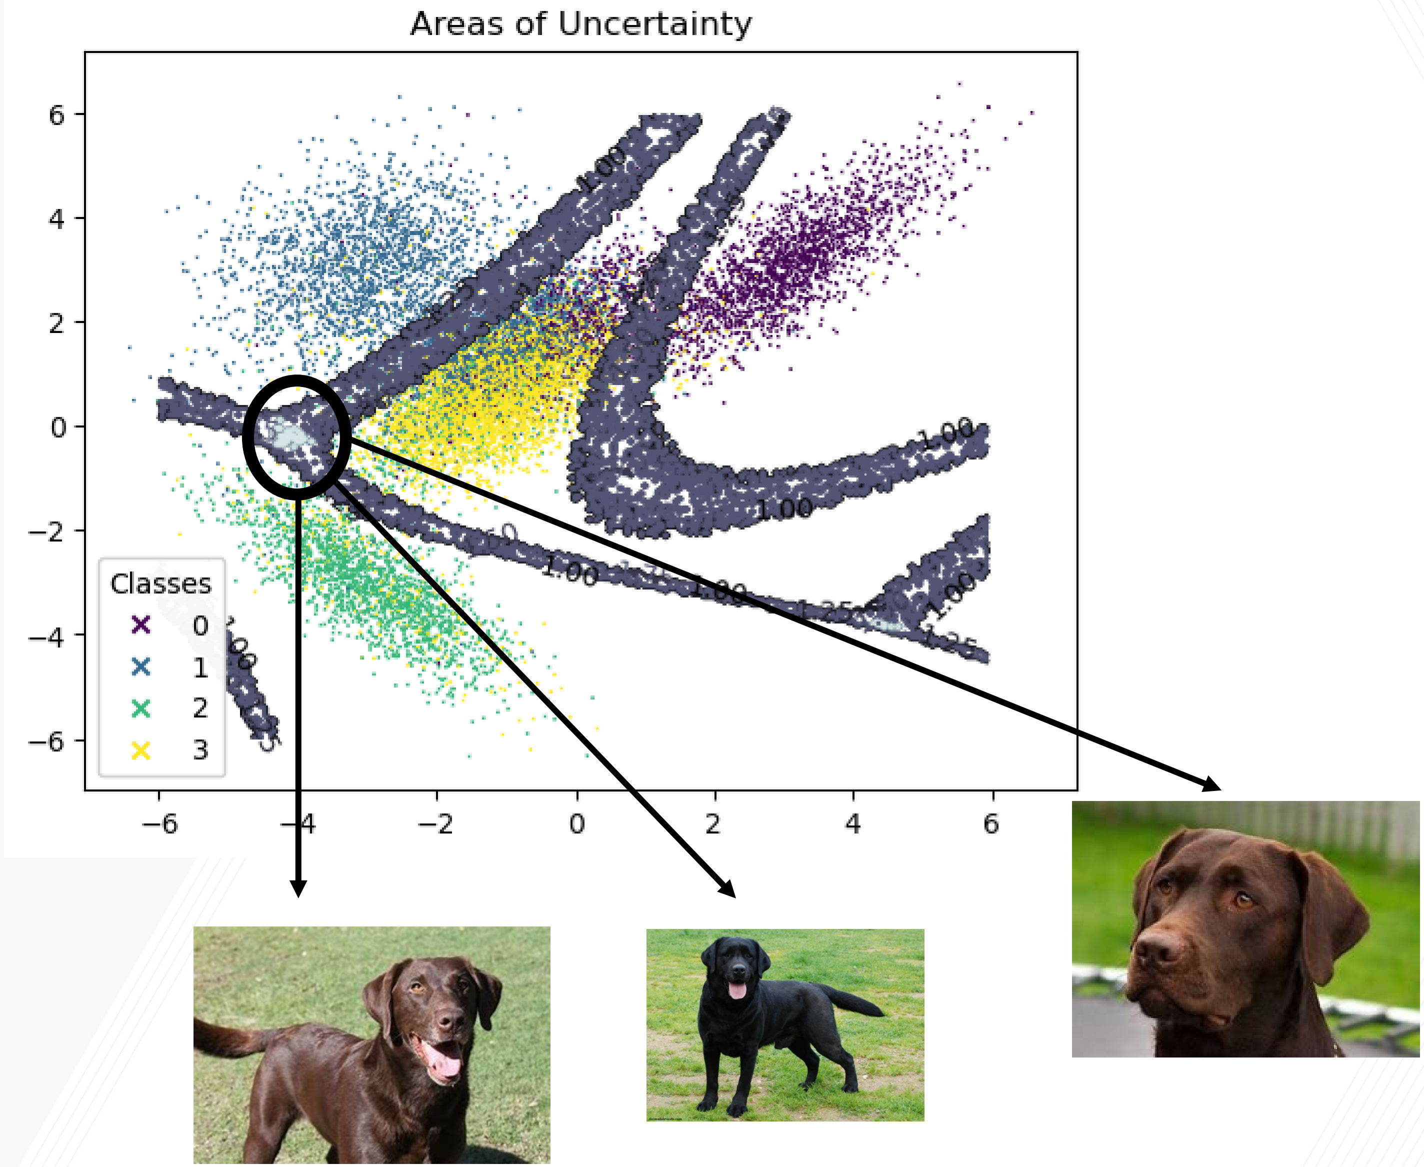

Why (Split) conformal?

- 90% Confidence Interval through Conformal Prediction

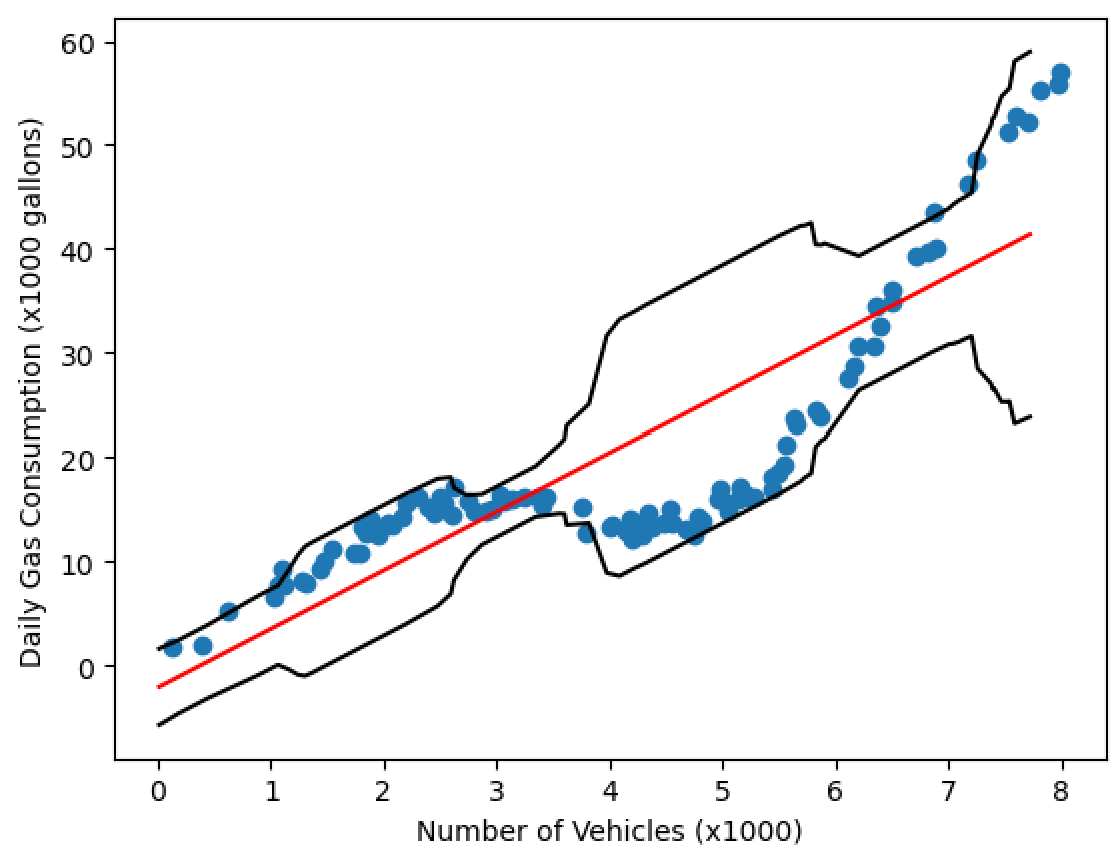

Georgia Tech
CREATING THE NEXT

Why (Split) conformal?

- {Golden Retriever}

- {Golden Retriever, Tibetan Mastiff, Irish Setter, Great Pyrenees}

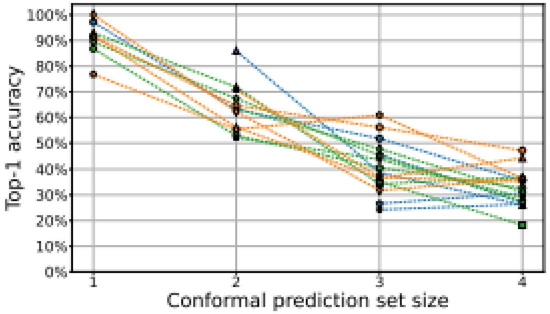

Kumar, Bhawesh, et al. "Conformal Prediction with Large Language Models for Multi-Choice Question Answering." arXiv preprint arXiv:2305.18404 (2023).

Extensions

- Active Learning

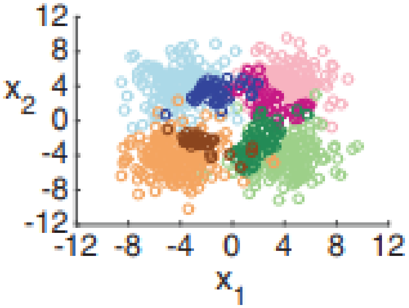

Sergio Matiz and Kenneth E. Barner. 2020. Conformal prediction based active learning by linear regression optimization. Neurocomputing 388 (2020), 157–169.

- Set size regularization

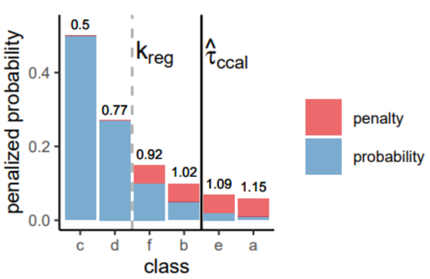

Angelopoulos, A., Bates, S., Malik, J., & Jordan, M. I. (2020). Uncertainty sets for image classifiers using conformal prediction. *arXiv preprint arXiv:2009.14193*.

Extensions

- Covariate Shift

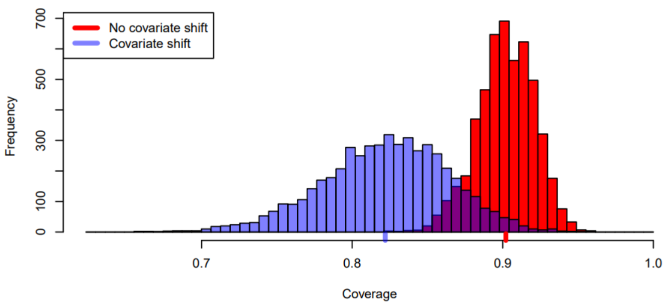

Tibshirani, Ryan J., et al. "Conformal prediction under covariate shift." Advances in neural information processing systems 32 (2019).

- Time Series

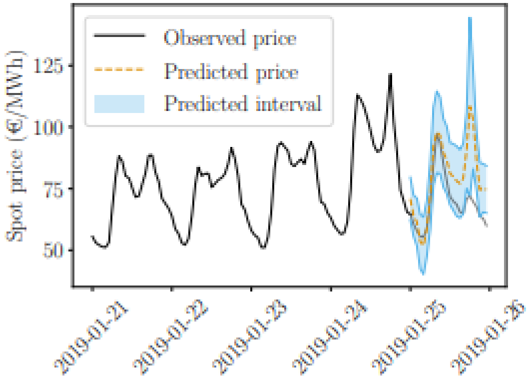

Margaux Zaffran, Olivier Feron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. 2022. Adaptive Conformal Predictions for Time Series. In Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 162). PMLR, 25834–25866. https://proceedings.mlr.press/v162/zaffran22a.html

Georgia Tech
CREATING THE NEXT

Summary

- Empirical risk has gaps for model validation
- Model generated notions of uncertainty are uncalibrated
- Conformal prediction provides a guaranteed form of uncertainty quantification
- Guarantee applies to marginal coverage
- Requires exchangeable (IID) data
- Wraps around previously trained model with no additional training

Georgia Tech
CREATING THE NEXT

Acknowledgement

- This work was partially supported by NSF grant DMS-2015405