

CREST - GENES

Cours doctoraux 2023 – 2024

AN INTRODUCTION TO CONFORMAL PREDICTION AND DISTRIBUTION-FREE INFERENCE

Rina Foygel Barber
University of Chicago

SCHEDULE	Monday	18th March 2024	From 13:00 to 16:15	Room 2033
	Thursday	21st March 2024	From 13:00 to 16:15	Room 2033

Aims and objectives

This short course will introduce the framework of distribution-free statistical inference, and will provide an in-depth overview of both theoretical foundations and practical methodologies in this field. We will cover methods including holdout set methods, conformal prediction, cross-validation based methods, calibration procedures, and more, with emphasis on how these methods can be adapted to a range of settings to achieve robust uncertainty quantification without compromising on accuracy. The course will also introduce the theoretical results behind these methods, including the role of exchangeability (and its variants) in establishing the distribution-free validity of these methods, as well as more classical theoretical results establishing how these distribution-free methods relate to the answers we would obtain via parametric models or other classical assumption-based techniques. Our theoretical overview will also cover hardness results that carve out the space of inference questions that are possible or impossible to answer within the distribution-free framework.

Outline

1. Introduction (approx. 0.5hrs -- day 1, part 1)
 - (a) From classical inference to distribution-free inference
 - (b) Introduction to holdout set & cross-validation based methods
2. Conformal prediction methodology (approx. 2.5hrs -- day 1, part 2)
 - (a) Introduction to exchangeability
 - (b) Split conformal prediction
 - (c) Full conformal prediction

- (d) Jackknife, jackknife+, cross-conformal
- (e) Comparing conformal and classical methods
- (f) Testing for outliers / distribution drift

3. Limits of distribution-free inference (approx. 1.5hrs -- day 2, part 1)

- (a) Conditional predictive inference
- (b) Distribution-free regression
- (c) Distribution-free tests of conditional independence
- (d) Algorithmic stability

4. Extensions of conformal prediction (approx. 1.5hrs -- day 2, part 2)

- (a) Weighted conformal prediction: covariate shift, label shift, localized conformal prediction
- (b) Conformal prediction in the streaming setting
- (c) Conformal prediction under distribution drift
- (d) Calibration methods

Pre-requisites

This course will assume familiarity with standard regression and classification methods (e.g., linear regression, logistic regression, nearest neighbor methods, kernel methods), and undergraduate level probability theory (e.g., conditional distributions).

Some related literature

Angelopoulos, A. N., & Bates, S. (2023). Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4), 494-591. <https://arxiv.org/abs/2107.07511>

Barber, R. F., Candes, E. J., Ramdas, A., & Tibshirani, R. J. (2021). Predictive inference with the jackknife+. *Annals of Statistics* 49(1): 486-507. <https://arxiv.org/abs/1905.02928>

Barber, R. F., Candes, E. J., Ramdas, A., & Tibshirani, R. J. (2023+). Conformal prediction beyond exchangeability. To appear in *Annals of Statistics*. <https://arxiv.org/abs/2202.13415>

Chernozhukov, V., Wüthrich, K., & Zhu, Y. (2021). Distributional conformal prediction. *Proceedings of the National Academy of Sciences*, 118(48), e2107794118. <https://arxiv.org/abs/1909.07889>

Gupta, C., Podkopaev, A., & Ramdas, A. (2020). Distribution-free binary classification: prediction sets, confidence intervals and calibration. *Advances in Neural Information Processing Systems*, 33, 3711-3723. <https://arxiv.org/abs/2006.10564>

Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R. J., & Wasserman, L. (2018). Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523), 1094-1111. <https://arxiv.org/abs/1604.04173>

Tibshirani, R. J., Barber, R. F., Candes, E., & Ramdas, A. (2019). Conformal prediction under covariate shift. *Advances in neural information processing systems*, 32. <https://arxiv.org/abs/1904.06019>

Vovk, V. (2012). Conditional validity of inductive conformal predictors. In Asian conference on machine learning (pp. 475-490). PMLR. <https://arxiv.org/abs/1209.2673>

Vovk, V., Gammerman, A., & Shafer, G. (2005). Algorithmic learning in a random world (Vol. 29). New York: Springer. <http://www.alrw.net/>

Vovk, V., Nouretdinov, I., Manokhin, V., & Gammerman, A. (2018). Cross-conformal predictive distributions. In Conformal and Probabilistic Prediction and Applications (pp. 37-51). PMLR.