Moment Multicalibration for Uncertainty Estimation*

Christopher Jung[1], Changhwa Lee[2], Mallesh M. Pai[3], Aaron Roth[1], and Rakesh Vohra[4]

[1]University of Pennsylvania Department of Computer and Information Science
[2]University of Pennsylvania Department of Economics
[3]Rice University Department of Economics
[4]University of Pennsylvania Department of Economics and Electrical and Systems Engineering

We show how to achieve the notion of "multicalibration" from [1] not just for means, but also for variances and other higher moments. Informally, this means that we can find regression functions which, given a data point, can make point predictions not just for the expectation of its label, but for higher moments of its label distribution as well—and those predictions match the true distribution quantities when averaged not just over the population as a whole, but also when averaged over an enormous number of finely defined subgroups. It yields a principled way to estimate the uncertainty of predictions on many different subgroups—and to diagnose potential sources of unfairness in the predictive power of features across subgroups. As an application, we show that our moment estimates can be used to derive marginal prediction intervals that are simultaneously valid as averaged over all of the (sufficiently large) subgroups for which moment multicalibration has been obtained.

References

[1] Ú. Hébert-Johnson, M. Kim, O. Reingold, and G. Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *International Conference on Machine Learning*, pages 1939–1948, 2018.

[2] C. Jung, C. Lee, M. M. Pai, A. Roth, and R. Vohra. Moment multicalibration for uncertainty estimation. *arXiv preprint arXiv:2008.08037*, 2020.

*The full version of the paper is available on arXiv [2]