

Calibrated inference: Statistical inference that accounts for both sampling uncertainty and distributional uncertainty

How can we draw trustworthy scientific conclusions? One criterion is that a study can be replicated by independent teams. While replication is critically important, it is arguably insufficient. If a study is biased for some reason and other studies recapitulate the approach then findings might be consistently incorrect. It has been argued that trustworthy scientific conclusions require disparate sources of evidence. However, different methods might have shared biases, making it difficult to judge the trustworthiness of a result. We formalize this issue by introducing a "distributional uncertainty model", which captures biases in the data collection process. Distributional uncertainty is related to other concepts in statistics, ranging from correlated data to selection bias and confounding. We show that a stability analysis on a single data set allows to construct confidence intervals that account for both sampling uncertainty and distributional uncertainty. We introduce an R package that allows to draw data under distributional uncertainty and calibrate inference in (generalized) linear models.

This is joint work with Yujin Jeong.