

The HulC: Hull based Confidence Regions

Arun Kumar Kuchibhotla*, Sivaraman Balakrishnan, and Larry Wasserman

{arunku, siva, larry}@stat.cmu.edu

Department of Statistics & Data Science, Carnegie Mellon University

Abstract

We develop and analyze the HulC, an intuitive and general method for constructing confidence sets using the convex hull of estimates constructed from subsets of the data. Unlike classical methods which are based on estimating the (limiting) distribution of an estimator, the HulC is often simpler to use and effectively bypasses this step. In comparison to the bootstrap, the HulC requires fewer regularity conditions and succeeds in many examples where the bootstrap provably fails. Unlike subsampling, the HulC does not require knowledge of the rate of convergence of the estimators on which it is based. The validity of the HulC requires knowledge of the (asymptotic) median-bias of the estimators. We further analyze a variant of our basic method, called the Adaptive HulC, which is fully data-driven and estimates the median-bias using subsampling. We show that the Adaptive HulC retains the aforementioned strengths of the HulC. In certain cases where the underlying estimators are pathologically asymmetric, the HulC and Adaptive HulC can fail to provide useful confidence sets. We discuss these methods in the context of several challenging inferential problems which arise in parametric, semi-parametric and non-parametric inference. Although our focus is on validity under weak regularity conditions, we also provide some general results on the width of the HulC confidence sets, showing that in many cases the HulC confidence sets have near-optimal width.

*presenter