

Abstract

Recently, new methods for model assessment, based on subsampling and posterior approximations, have been proposed for scaling leave-one-out cross-validation (LOO) to large datasets. Although these methods work well for estimating predictive performance for individual models, they are less powerful in model comparison. We propose an efficient method for estimating differences in predictive performance by combining fast approximate LOO surrogates with exact LOO subsampling using the difference estimator and supply proofs with regards to scaling characteristics. The resulting approach can be orders of magnitude more efficient than previous approaches, as well as being better suited to model comparison.