

An Efficient Variance Estimator of AUC and Its Applications to Binary Classification

Wendy Wang^{1*}

¹Mathematics Department, Wellesley College, Wellesley, MA 02481, U.S.A
qwang@wellesley.edu

Abstract

The area under the ROC (receiver operating characteristic) curve, AUC, is one of the most commonly used measures to evaluate the performance of a binary classifier. Due to sampling variation, the model with the largest observed AUC score is not necessarily optimal, so it is crucial to assess the variation of AUC estimate. We extend the proposal by Wang and Lindsay (2014) and devise an unbiased variance estimator of AUC estimate that is of a two-sample U-statistic form. The proposal can be easily generalized to estimate the variance of a K -sample U-statistic ($K \geq 2$). To make our developed variance estimator more applicable, we employ a partition resampling scheme that is computationally efficient. Simulation studies suggest that the developed AUC variance estimator yields much better or comparable performance to jackknife and bootstrap variance estimators, and computational times that are about 10 to 30 times faster than the times of its counterparts. In practice, the proposal can be used in the one-standard-error rule for model selection, or to construct an asymptotic confidence interval of AUC in binary classification. In addition to conducting simulation studies, we illustrate its practical applications using a real data set in medical sciences. This work was done in collaboration with Alexandria Guo ('19, Wellesley College).

*Corresponding author.