Statistica Sinica **36** (2026), 1-22
doi:https://doi.org/10.5705/ss.202022.0131

BIAS, CONSISTENCY, AND ALTERNATIVE PERSPECTIVES OF THE INFINITESIMAL JACKKNIFE

Wei Peng[1], Lucas Mentch*[1] and Len Stefanski[2]

[1] *University of Pittsburgh and* [2] *North Carolina State University*

Though introduced nearly 50 years ago, the infinitesimal jackknife (IJ) remains a popular modern tool for quantifying predictive uncertainty in complex estimation settings. In particular, when supervised learning ensembles are constructed via bootstrap samples, recent work demonstrated that the IJ estimate of variance is particularly convenient and useful. However, despite the algebraic simplicity of its final form, its derivation is rather complex. As a result, studies clarifying the intuition behind the estimator or rigorously investigating its properties have been severely lacking. This work aims to take a step forward on both fronts. We demonstrate that surprisingly, the exact form of the IJ estimator can be obtained via a straightforward linear regression of the individual bootstrap estimates on their respective weights or via the classical jackknife. The latter realization allows us to formally investigate the bias of the IJ variance estimator and better characterize the settings in which its use is appropriate. Finally, we extend these results to the case of U-statistics where base models are constructed via subsampling rather than bootstrapping and provide a consistent estimate of the resulting variance.

Key words and phrases: Bias, bootstrap, consistency, infinitesimal jackknife, random forest, U-statistic.

1. Introduction

Given a sample $X_1, \ldots, X_n \sim \mathrm{P}$, a parameter of interest θ, and an estimator $\hat{\theta} = s(X_1, \ldots, X_n)$, it is often of interest to estimate $\mathrm{Var}(\hat{\theta})$. Given data $\boldsymbol{x} = (x_1, \ldots, x_n)$ and an estimate $\hat{\theta} = s(\boldsymbol{x})$, to provide a bootstrap estimate of variance, we draw B (re)samples of size n with replacement to form bootstrap samples $\boldsymbol{x}_1^*, \ldots, \boldsymbol{x}_B^*$ from which we calculate bootstrap estimates $\hat{\theta}_1, \ldots, \hat{\theta}_B$. The nonparametric bootstrap variance estimate of $\hat{\theta}$ is then taken as the empirical variance of $\hat{\theta}_1, \ldots, \hat{\theta}_B$ (Efron, 1979, 2014). Within this context, given the necessity of calculating $\hat{\theta}_1, \ldots, \hat{\theta}_B$, it is natural to consider the estimator $\tilde{\theta}_B = (1/B)\sum_{b=1}^{B} \hat{\theta}_b$ as a "bootstrap smoothed" or "bagged" alternative of $\hat{\theta}$ (Efron and Tibshirani, 1994; Breiman, 1996).

The standard bootstrap approach to assess the variability of $\tilde{\theta}_B$ is computationally burdensome, requiring bootstrap replicates of not only the original data, but of the bootstrap samples as well. This double bootstrap (Beran, 1988)

*Corresponding author. E-mail: lkm31@pitt.edu