

Optimal Subsampling Bootstrap for Massive Data

Yingying Ma, Chenlei Leng & Hansheng Wang

To cite this article: Yingying Ma, Chenlei Leng & Hansheng Wang (2024) Optimal Subsampling Bootstrap for Massive Data, Journal of Business & Economic Statistics, 42:1, 174-186, DOI: [10.1080/07350015.2023.2166514](https://doi.org/10.1080/07350015.2023.2166514)

To link to this article: <https://doi.org/10.1080/07350015.2023.2166514>



© 2023 The Author(s). Published with license by Taylor & Francis Group, LLC.



[View supplementary material](#)



Published online: 14 Feb 2023.



[Submit your article to this journal](#)



Article views: 2261



[View related articles](#)



[View Crossmark data](#)



Citing articles: 1 [View citing articles](#)

Optimal Subsampling Bootstrap for Massive Data

Yingying Ma^a, Chenlei Leng^b, and Hansheng Wang^c

^aSchool of Economics and Management, Beihang University, Beijing, China; ^bDepartment of Statistics, University of Warwick, Coventry, UK; ^cGuanghua School of Management, Peking University, Beijing, China

ABSTRACT

The bootstrap is a widely used procedure for statistical inference because of its simplicity and attractive statistical properties. However, the vanilla version of bootstrap is no longer feasible computationally for many modern massive datasets due to the need to repeatedly resample the entire data. Therefore, several improvements to the bootstrap method have been made in recent years, which assess the quality of estimators by subsampling the full dataset before resampling the subsamples. Naturally, the performance of these modern subsampling methods is influenced by tuning parameters such as the size of subsamples, the number of subsamples, and the number of resamples per subsample. In this article, we develop a novel hyperparameter selection methodology for selecting these tuning parameters. Formulated as an optimization problem to find the optimal value of some measure of accuracy of an estimator subject to computational cost, our framework provides closed-form solutions for the optimal hyperparameter values for subsampled bootstrap, subsampled double bootstrap and bag of little bootstraps, at no or little extra time cost. Using the mean square errors as a proxy of the accuracy measure, we apply our methodology to study, compare and improve the performance of these modern versions of bootstrap developed for massive data through numerical study. The results are promising.

ARTICLE HISTORY

Received August 2021
Accepted December 2022

KEYWORDS

Bag of little bootstraps;
Bootstrap; Computational
cost; Subsampled double
bootstrap; Subsampling

1. Introduction


Real data analysis often faces situations where statistical inference is not tractable. This can happen if the interest is to estimate the variance of an estimator that is not easily estimable or when a robust estimator of the variance is warranted under the suspicion that assumptions are invalid. In this case, the bootstrap method (Efron 1990; Efron and Tibshirani 1994) provides a simple elegant solution for automatic inference, enabling the computation of various inference quantities without the need to know their analytical formula. Under certain conditions, many bootstrap estimators are generally consistent (Van Der Vaart and Wellner 1996) and can be more accurate than those based on asymptotic approximation (Hall 1994).

Traditional Bootstrap (TB) methods were first developed for small datasets for which computation was not an issue. In particular, the vanilla version of the bootstrap estimates parameters of interest by repeatedly resampling N observations with replacement from N data points in the original sample and thus is feasible only if the sample size is relatively modest and the computation is performed on a single computer (Booth and Hall 1994). With the emergence of Big Data, these methods are no longer applicable to modern datasets that are massive in size. First and foremost, it is no longer the case that these datasets can be loaded into the main memory. One approach to mitigate this problem is to use parallel and distributed computing systems to divide a dataset, estimate the parameters of interest

on each computer, and aggregate these estimates on a central machine. This approach usually involves high communication costs between different computer nodes and, thus, is sometimes not desirable (Li, Lin, and Li 2013), although progress has been made to alleviate this (Jordan, Lee, and Yang 2019; Volgushev, Chao, and Cheng 2019; Chen, Liu, and Zhang 2019; Chen and Peng 2021; Fan, Guo, and Wang 2021). The main focus of this article is to develop methods for optimal bootstrap inference on a single computer which will free us from the concern on communication cost.

Recognizing the limitation of TB for big datasets, subsampling methods that only requires repeated computation of the estimator for subsamples with size much smaller than the original dataset, have been developed (Politis, Romano, and Wolf 1999). A leading example is the so called n -out-of- N , bootstrap (Bickel, Götze, and van Zwet 1997), in which each subsample draws n observations out of N points often with $n \ll N$. We will refer to this bootstrap scheme as subsampling bootstrap and abbreviate it as SB for short hereafter. Although SB reduces the size of each subsample from N in TB to n , Bickel and Sakov (2008) showed that its performance is rather sensitive to the choice of subsample size. Moreover, SB must perform a rescaling of their output which requires knowledge and explicit use of the convergence rate of the estimator, making it less automatic to deploy than TB. The above limitations of SB prompted Kleiner et al. (2014) to introduce the method of Bag of Little

CONTACT Chenlei Leng  C.Leng@warwick.ac.uk  Department of Statistics, University of Warwick, Coventry, UK.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/UBES.

© 2023 The Authors. Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Bootstraps (BLB). Similar to SB, BLB starts by subsampling the whole data with the size of each subsampled subset much smaller than N and then follows up by resampling the subsets via simple random sampling with replacement. Crucially, the resampling in the second step of BLB is performed in such a way that the size of each resample is N , the same as the size of the entire data. The computational saving of BLB over that of TB roots in the fact that the maximum number of distinct elements in each resample is bounded by the size of each subsample in the first step. For many estimators such as M -estimators obtained via empirical risk minimization, this means that we just need to optimize a weighted loss function with a smaller number of distinct items in the empirical risk than that of TB. Compared to SB, BLB requires no analytical re-scaling because of the size of the resamples in step two and thus makes a fully automatic method for statistical inference (Kleiner et al. 2014).

Intuitively, the performance of BLB relies on three hyperparameters

- n (the size of the bootstrap subsamples or subsets in step one),
- R (the number of the subsamples or replicates in step one),
- B (the total number of resamples per subset in step two).

With unlimited computational resources, those hyperparameters should be set as large as possible for reliable inference. With limited computational budget, however, these three hyperparameters need to be selected. Indeed, it is found that the performance of BLB can be sensitive to the choice of these parameters as discussed in Kleiner et al. (2014) and Sengupta, Volgushev, and Shao (2016). Because of this, Sengupta, Volgushev, and Shao (2016) raised an open question:

How do we optimally choose n , B and R in BLB to balance statistical accuracy and running time?

In Kleiner et al. (2014), they suggested an adaptive method for selecting R and B . By introducing a new tolerance parameter, they tracked resamples for each subset until that tolerance level was reached. Without knowing the variability of the precision estimate though, it is unclear how this additional hyperparameter should be specified for a given computational cost. As a partial solution, Sengupta, Volgushev, and Shao (2016) proposed an alternative approach named subsampled double bootstrap (SDB), which simply sets $B = 1$ as in BLB. This certainly solves the problem of choosing between R and B . However, the relationship between the subset size n and the number of replicates R remains elusive. In order to choose n and R in SDB, an intuitive approach is to choose n as large as possible as long as n data points can fit on a computer while adjusting the size of R such that the computational budget is not exceeded. As we show however, the choice of n and R does impact the performance of SDB.

In this article, we develop a novel framework to find the optimal balance between statistical efficiency and computational cost for these modern variants of the bootstrap method. To strike this balance, an intuitive procedure is to maximize a certain measure of this efficiency subject to a running time constraint. Intuitively, any measure of efficiency will depend on the hyperparameters n , R , and B and so will the running time. Denoting these dependencies as $\text{Efficiency}(n, R, B)$ and $\text{Time}(n, R, B)$,

respectively, our general framework seeks to identify (n^*, R^*, B^*) such that

$$(n^*, R^*, B^*) = \arg \text{optim}_{n,R,B} \text{Efficiency}(n, R, B),$$

$$\text{subject to } \text{Time}(n, R, B) \leq C_{\max},$$

where C_{\max} is a given computational cost allocated to bootstrap and $\arg \text{optim}$ reads as the argument that optimizes the subsequent function. For the constraint C_{\max} , it is often more convenient to specify it corresponding to the time needed to implement a bootstrap method with some prespecified (n, R, B) combination. In this case, we can immediately understand that our framework aims to improve the performance of BLB with prespecified n, R , and B , by using an optimal combination of these hyperparameters obtained via optimization.

Often for mathematical convenience, it is easier to work with the inefficiency of an estimator based on various bootstrap implementations, for example by examining some loss function or the variance of an estimator. Intuitively the dependence of the inefficiency on these hyperparameters can be characterized by terms usually inversely proportional to n , R , and B or simple functions of them. This is because conditional on data, bootstrap subsamples and resamples can be viewed as independent and identically distributed (Efron 1990; Efron and Tibshirani 1994). The measure of inefficiency of interest often depends on the empirical distribution of these bootstrap subsamples and resamples. This is in sharp contrast to n -out-of- N bootstrap where explicit knowledge of the convergence rate must be known for it to be applicable.

To illustrate our general framework, we use the Mean Square Error (MSE) of an estimator as a proxy for the accuracy measure and discuss the use of various bootstrap methods for estimating the standard errors of the sample mean. The sample mean estimates the population mean which is a quantity of major interest in many statistical problems (He and Shao 1996). By a careful theoretical analysis, we find that the asymptotic efficiency of various bootstrap estimates are closely related to the hyperparameters. In particular, these relationships for BLB and SDB are analytically simple, allowing us to identify the optimal combinations of the hyperparameters in these procedures in closed-form. Thus, we go beyond providing an affirmative answer to the open question raised in Sengupta, Volgushev, and Shao (2016) by presenting closed-form solutions to the optimal hyperparameter triple (n, R, B) for BLB. Furthermore, our procedure provides the optimal pair of parameters in n and R for SB in Sengupta, Volgushev, and Shao (2016). Our approach can be readily extended to deal with multivariate random variables and general parameters and we discuss how this can be done. Although we only consider the independent data, by a similar approach, this method can be extended to block bootstrap designed for dependent data. We show via extensive simulations that our approach improve the performance of BLB and SB with similar computational budget. Thus, our answer to the open question posed in Sengupta, Volgushev, and Shao (2016) is confirmed empirically. Note that an early approach to find the optimal resampling size in the double bootstrap appeared in Booth and Hall (1994) but their method is computationally infeasible for massive data.

The rest of the article is organized as follows. Section 2 introduces different bootstrap methods and their associated

theoretical properties. Our hyperparameter selection approach is presented in Section 3. We discuss the estimation of more general parameters and statistics in Section 3.2. Extensive numerical study is conducted in Section 4. The article is concluded with a discussion in Section 5. All the theoretical conditions, proofs, and additional numerical results are found in the supplementary materials.

The following notations are used throughout the article. For two positive sequences $\{a_n\}$, $\{b_n\}$, we use $a_n \lesssim b_n$ to mean $a_n \leq Cb_n$ for some constant C independent of n , $a_n \gtrsim b_n$ to mean $b_n \lesssim a_n$, and $a_n \asymp b_n$ to mean $a_n \lesssim b_n$ and $b_n \lesssim a_n$. When stating any results, we always assume that the sample size satisfies $N \rightarrow \infty$. We use $U(a, b)$ to denote a uniform random variable on $[a, b]$. Finally, for a matrix $A = (a_{ij}) \in \mathbb{R}^{m \times n}$, we denote $\|A\|_F = (\sum_{i,j} a_{ij}^2)^{1/2}$ as its Frobenius norm.

2. Bootstrap Methods for Univariate Random Variable

For better illustration of our approach, we start by examining the simplest case where the statistic of interest is the mean of a univariate random variable. The theory to be presented hereafter is further developed for multivariate random variables in Section 2.5 and we discuss general statistics in Section 3.2. To facilitate the theoretical development, we need the following technical conditions.

- (C1) (TB Condition) As $N \rightarrow \infty$, we assume $B \rightarrow \infty$.
- (C2) (BLB Condition) As $N \rightarrow \infty$, we assume $\min\{n, R, B\} \rightarrow \infty$. In addition assume that $n < N$, $N = o(nBR)$ and $n \asymp R \asymp B$.
- (C3) (SDB and SB Condition) As $N \rightarrow \infty$, assume $\min\{n, R\} \rightarrow \infty$. In addition assume that $n < N$ and $R \gtrsim n$.

The above conditions are mild and reasonable. By assuming these relationships between different hyperparameters and the whole sample size N , they essentially require that n, R , and B should be large enough to facilitate an asymptotic analysis of higher order terms.

Let $\mathcal{S} = \{X_1, \dots, X_N\}$ be the sample where $X_1, X_2, \dots, X_N \in \mathbb{R}^1$ are independent and identically distributed random variables with mean μ and variance σ^2 . We assume that the centered j th moment of X_i exists such that $E(X_i - \mu)^j = \sigma_j < \infty$ for $3 \leq j \leq 6$. A simple estimator of the mean parameter μ is the same average denoted as $\bar{X} = N^{-1} \sum_{i=1}^N X_i$. The estimation accuracy of \bar{X} is measured by its standard error (SE) which equals σ/\sqrt{N} analytically. This SE can be consistently estimated as

$$\widehat{\text{SE}}^2 = \hat{\sigma}^2/N, \quad (2.1)$$

where $\hat{\sigma}^2 = \sum_{i=1}^N (X_i - \bar{X})^2/N$. We characterize the mean and the variance of SE^2 in the following theorem.

Theorem 1. For the estimator in (2.1), we have

$$E(\widehat{\text{SE}}^2) = \frac{\sigma^2}{N} \left(1 - \frac{1}{N}\right) \text{ and } \text{var}(\widehat{\text{SE}}^2) = \frac{\sigma_4 - \sigma^4}{N^3} \left\{1 + o(1)\right\}.$$

By this theorem, the MSE of $\widehat{\text{SE}}^2$ is dominated by its variance and can be seen as

$$\text{MSE}(\widehat{\text{SE}}^2) = \frac{\sigma_4 - \sigma^4}{N^3} \left\{1 + o(1)\right\}, \quad (2.2)$$

which scales inverse-proportionally to N^3 in the leading order.

2.1. Traditional Bootstrap

We analyze the MSE of the traditional bootstrap or the vanilla bootstrap estimator of SE in this section. Denote B as the total number of bootstrap resamples that are sampled uniformly with replacement from the original sample. That is, for any $b = 1, \dots, B$, the b th bootstrap sample is obtained as $\mathcal{B}^{(b)} = \{X_1^{(b)}, X_2^{(b)}, \dots, X_N^{(b)}\}$, where $X_i^{(b)}$ is independently generated via simple random sampling with replacement from the whole sample \mathcal{S} . From $\mathcal{B}^{(b)}$, the b th bootstrap sample mean can be calculated as $\bar{X}^{(b)} = N^{-1} \sum_{i=1}^N X_i^{(b)}$. Subsequently, an estimate of SE^2 via the bootstrap method is seen as

$$\widehat{\text{SE}}_{\text{TB}}^2 = B^{-1} \sum_{b=1}^B \left(\bar{X}^{(b)} - \bar{X}\right)^2.$$

Conditional on \mathcal{S} , for any b , $X_i^{(b)}$'s are independent and identically distributed due to the bootstrap scheme. This immediately implies that $P(X_i^{(b)} = X_j) = 1/N$ for any $1 \leq j \leq N$. That is, each element in \mathcal{S} has equal probability of being sampled. Accordingly, we have $E(\bar{X}^{(b)}|\mathcal{S}) = E(X_i^{(b)}|\mathcal{S}) = \bar{X}$ and $\text{var}(\bar{X}^{(b)}|\mathcal{S}) = N^{-1} \text{var}(X_i^{(b)}|\mathcal{S}) = N^{-1} \sigma^2$. We have the following theorem for this traditional bootstrap estimator of SE^2 .

Theorem 2. For the traditional bootstrap estimator, assume Condition (C1) holds, we have

$$E(\widehat{\text{SE}}_{\text{TB}}^2) = \frac{\sigma^2}{N} \left(1 - \frac{1}{N}\right) \text{ and } \text{var}(\widehat{\text{SE}}_{\text{TB}}^2) = \text{var}(\widehat{\text{SE}}^2) \left(\frac{2\sigma^4}{\sigma_4 - \sigma^4} \cdot \frac{N}{B} + 1 \right) \left\{1 + o(1)\right\}.$$

From this theorem, we can immediately obtain the MSE of $\widehat{\text{SE}}_{\text{TB}}^2$ as

$$\begin{aligned} \text{MSE}(\widehat{\text{SE}}_{\text{TB}}^2) &= \text{var}(\widehat{\text{SE}}^2) \left(\frac{2\sigma^4}{\sigma_4 - \sigma^4} \cdot \frac{N}{B} + 1 \right) \left\{1 + o(1)\right\} + \frac{\sigma^4}{N^4} \\ &= \left(\frac{\sigma_4 - \sigma^4}{N^3} + \frac{2\sigma^4}{N^2 B} \right) \left\{1 + o(1)\right\}. \end{aligned} \quad (2.3)$$

This suggests that B , the number of bootstrap resamples, needs to be the same order of N or larger, for $\text{MSE}(\widehat{\text{SE}}_{\text{TB}}^2)$ to achieve the same convergence rate as $\text{MSE}(\widehat{\text{SE}}^2)$ in (2.2). Given a computational budget, the most efficient traditional bootstrap estimator is to keep drawing resamples until this budget runs out.

2.2. Bag of Little Bootstraps

We next provide a brief review of the Bag of Little Bootstraps (BLB) method in Kleiner et al. (2014), which is carried out via a two-step procedure including a subsampling and a resampling step.

Step 1. The subsampling step: We draw R subsamples or little bootstrap samples, each of which, denoted as $\mathcal{S}^{(r)} = \{X_1^{(r)}, X_2^{(r)}, \dots, X_n^{(r)}\}$ for $r = 1, \dots, R$, is drawn via simple random sampling with replacement from the whole sample \mathcal{S} . Note that the size of $\mathcal{S}^{(r)}$ is n , which is usually much smaller than N .

Step 2. The resampling step: B weighted resamples are drawn such that the cardinality of each resample is N . Specifically, denote the b th resample as $\mathcal{B}^{(r,b)} = \{X_1^{(r,b)}, X_2^{(r,b)}, \dots, X_N^{(r,b)}\}$. For each r and b , $X_i^{(r,b)}$ is independently drawn via simple random sampling from $\mathcal{S}^{(r)}$ with replacement. This scheme implies that the number of distinct points in each bootstrap resample $\mathcal{B}^{(r,b)}$ is n at maximum.

From $\mathcal{B}^{(r,b)}$, the target statistic \bar{X} can be estimated by the sample average in the resample as

$$\bar{X}^{(r,b)} = N^{-1} \sum_{i=1}^N X_i^{(r,b)} = N^{-1} \sum_{i=1}^n X_i^{(r)} f_i^{(r,b)},$$

where $f_i^{(r,b)}$ is the number of times that $X_i^{(r)}$ appears in $\mathcal{S}^{(r)}$. That is $f_i^{(r,b)} = \sum_{j=1}^N I(X_j^{(r,b)} = X_i^{(r)})$ where $I(\cdot)$ is the indicator function. Obviously, the random vector $f^{(r,b)} = (f_1^{(r,b)}, f_2^{(r,b)}, \dots, f_n^{(r,b)})^\top \in \mathbb{R}^n$ follows a multinomial distribution with parameter N and p with $p = (1/n, 1/n, \dots, 1/n)^\top \in \mathbb{R}^n$.

As the result of the two-step procedure, each bootstrap resample $\mathcal{B}^{(r,b)}$ has cardinality N , but with at most n distinct elements. Equivalently, each $\mathcal{B}^{(r,b)}$ can be seen as a weighted resample of size n . Thus, the BLB avoids the need for repeated computation on resamples having size comparable with that of the original dataset, since each BLB resample contains at most n distinct elements. In a large class of estimators commonly encountered, including M -estimators, computation can take weighted data representation. It is for these estimators that BLB has huge computational advantages. In comparison to the traditional bootstrap, the storage requirement for the BLB method is substantially less demanding as long as $n \ll N$, and for these estimators, the cost of computing the estimator based on the BLB resamples will be substantially lower than that based on the traditional bootstrap resamples.

Given the bootstrap resamples $\bar{X}^{(r,b)}$, SE^2 can be estimated as

$$\widehat{SE}_{BLB}^2 = \frac{1}{RB} \sum_{r=1}^R \sum_{b=1}^B (\bar{X}^{(r,b)} - \bar{X}^{(r)})^2,$$

where $\bar{X}^{(r)} = n^{-1} \sum_{i=1}^n X_i^{(r)}$.

We now derive the MSE of \widehat{SE}_{BLB}^2 . Note that conditional on \mathcal{S} and $\mathcal{S}^{(r)}$, $\bar{X}^{(r,b)}$'s are independent and identically distributed. Accordingly, we have $E(\bar{X}^{(r,b)} | \mathcal{S}, \mathcal{S}^{(r)}) = \bar{X}^{(r)}$ and $\text{var}(\bar{X}^{(r,b)} | \mathcal{S}, \mathcal{S}^{(r)}) = N^{-1} \tilde{\sigma}_r^2$, where $\tilde{\sigma}_r^2 = n^{-1} \sum_{i=1}^n (X_i^{(r)} - \bar{X}^{(r)})^2$. We have the following theorem.

Theorem 3. For the BLB estimator, assume Condition (C2) holds, we have

$$E(\widehat{SE}_{BLB}^2) = \frac{\sigma^2}{N} \left(1 - \frac{1}{n}\right) \left\{1 + o(1)\right\},$$

$$\text{var}(\widehat{SE}_{BLB}^2) = \text{var}(\widehat{SE}^2) \left(\frac{2\sigma^4}{\sigma_4 - \sigma^4} \cdot \frac{N}{RB} + \frac{N}{nR} + 1 \right) \left\{1 + o(1)\right\}.$$

From the above theorem, the MSE of \widehat{SE}_{BLB}^2 is immediately seen as

$$\text{MSE}(\widehat{SE}_{BLB}^2) = \text{var}(\widehat{SE}^2) \left(\frac{2\sigma^4}{\sigma_4 - \sigma^4} \cdot \frac{N}{RB} + \frac{N}{nR} + 1 \right) \left\{1 + o(1)\right\} + \frac{\sigma^4}{N^2 n^2}$$

$$= \left(\frac{\sigma_4 - \sigma^4}{N^3} + \frac{2\sigma^4}{N^2 RB} + \frac{\sigma_4 - \sigma^4}{N^2 nR} + \frac{\sigma^4}{N^2 n^2} \right) \left\{1 + o(1)\right\}. \quad (2.4)$$

Compared with (2.2), for the MSE of \widehat{SE}_{BLB}^2 to be comparable with that of the analytical one, we require $n \gtrsim \sqrt{N}$, $nR \gtrsim N$ and $RB \gtrsim N$ which hold by taking $n \asymp \sqrt{N}$, $R \asymp \sqrt{N}$ and $B \gtrsim \sqrt{N}$ for example. Alternatively, we can take $n \sim N^\omega$ with $\omega \in [0.5, 1]$, $R \sim N^{1-\omega}$ and $B \sim N^\gamma$ with $\gamma > 0.5$.

2.3. Subsampled Bootstrap

The BLB method is closely related to the so-called “ n -out-of- N ” bootstrap method or subsampled bootstrap studied in Bickel, Götze, and van Zwet (1997). With some abuse of notation, let R be the total number of subsamples. For any $1 \leq r \leq R$, we use $\mathcal{S}^{(r)} = \{X_i^{(r)} : 1 \leq i \leq n\}$ to denote the r th subsample, where $X_i^{(r)}$ is generated independently by simple random sampling with replacement from the whole sample \mathcal{S} . Note that the number of distinct elements in each subsample is at most n . Based on $\mathcal{S}^{(r)}$, the target statistic \bar{X} can be computed as $\bar{X}^{(r)} = n^{-1} \sum_{i=1}^n X_i^{(r)}$. Accordingly, SE^2 can be estimated by

$$\widehat{SE}_{SB}^2 = \left(\frac{n}{N}\right) R^{-1} \sum_{r=1}^R (\bar{X}^{(r)} - \bar{X})^2.$$

Comparing \widehat{SE}_{SB}^2 with \widehat{SE}_{TB}^2 , we find that a re-scaling factor n/N is needed for \widehat{SE}_{SB}^2 which requires the knowledge of the convergence rate of the target estimator. This makes the SB less automatic (Kleiner et al. 2014; Sengupta, Volgushev, and Shao 2016).

We now study the MSE of \widehat{SE}_{SB}^2 . Note that conditional on \mathcal{S} , for any given r , $X_i^{(r)}$'s are independent and identically distributed with $P(X_i^{(r)} = X_j) = 1/N$ for any $1 \leq j \leq N$. Accordingly, we have $E(\bar{X}^{(r)} | \mathcal{S}) = E(X_i^{(r)} | \mathcal{S}) = \bar{X}$ and $\text{var}(\bar{X}^{(r)} | \mathcal{S}) = n^{-1} \text{var}(X_i^{(r)} | \mathcal{S}) = n^{-1} \hat{\sigma}^2$. We have the following theorem.

Theorem 4. For the SB estimator, assume Condition (C3) holds, we have

$$E(\widehat{SE}_{SB}^2) = N^{-1} (1 - n^{-1}) \sigma^2,$$

$$\text{var}(\widehat{SE}_{SB}^2) = \text{var}(\widehat{SE}^2) \left(\frac{2\sigma^4}{\sigma_4 - \sigma^4} \cdot \frac{N}{R} + 1 \right) \left\{1 + o(1)\right\}.$$

Thus, the MSE for $\widehat{\text{SE}}_{\text{SB}}^2$ can be expressed as

$$\begin{aligned} \text{MSE}(\widehat{\text{SE}}_{\text{SB}}^2) &= \text{var}(\widehat{\text{SE}}^2) \left(\frac{2\sigma^4}{\sigma_4 - \sigma^4} \cdot \frac{N}{R} + 1 \right) \{1 + o(1)\} + \frac{\sigma^4}{N^2 n^2} \\ &= \left(\frac{\sigma_4 - \sigma^4}{N^3} + \frac{2\sigma^4}{N^2 R} + \frac{\sigma^4}{N^2 n^2} \right) \{1 + o(1)\}. \end{aligned} \quad (2.5)$$

It is seen that for the MSE of the subsampled bootstrap to be comparable to that of the analytical estimator in (2.2), we require $n \gtrsim \sqrt{N}$ and $R \gtrsim N$.

2.4. Subsampled Double Bootstrap

Motivated by the BLB method and the subsampled bootstrap, Sengupta, Volgushev, and Shao (2016) proposed the so-called subsampled double bootstrap (SDB). The SDB takes a similar two-step approach as in the BLB, with the crucial difference of taking a single resample in the resampling step of BLB. Specifically, in the first step, SDB randomly draw R subsamples of the data, denoted as $\mathcal{S}^{(r)} = \{X_1^{(r)}, X_2^{(r)}, \dots, X_n^{(r)}\}$, where $X_i^{(r)}$'s are independently generated by simple random sampling with replacement from the whole sample \mathcal{S} . In the second stage, we only generate one bootstrap resample from each subset $\mathcal{S}^{(r)}$, which is denoted as $\mathcal{B}^{(r,1)} = \{X_1^{(r,1)}, X_2^{(r,1)}, \dots, X_n^{(r,1)}\}$. Here $X_i^{(r,1)}$'s are independently generated from $\mathcal{S}^{(r)}$ by simple random sampling with replacement. Based on $\mathcal{B}^{(r,1)}$, \bar{X} can be estimated as $\bar{X}^{(r,1)} = \frac{1}{n} \sum_{i=1}^n X_i^{(r,1)}$, and SE^2 can be estimated by $\widehat{\text{SE}}_{\text{SDB}}^2 = R^{-1} \sum_{r=1}^R \left(\bar{X}^{(r,1)} - \bar{X}^{(r)} \right)^2$ with $\bar{X}^{(r)} = n^{-1} \sum_{i=1}^n X_i^{(r)}$.

Conditional on \mathcal{S} , $\bar{X}^{(r,1)}$'s are independent and identically distributed for $1 \leq r \leq R$. Accordingly, we have $E(\bar{X}^{(r,1)} | \mathcal{S}, \mathcal{S}^{(r)}) = \bar{X}^{(r)}$ and $\text{var}(\bar{X}^{(r,1)} | \mathcal{S}, \mathcal{S}^{(r)}) = N^{-1} \tilde{\sigma}_r^2$, where $\tilde{\sigma}_r^2 = n^{-1} \sum_{i=1}^n (X_i^{(r)} - \bar{X}^{(r)})^2$. We have the following theorem.

Theorem 5. For the SDB estimator, assume Condition (C3) holds, we have

$$\begin{aligned} E(\widehat{\text{SE}}_{\text{SDB}}^2) &= \left(\frac{\sigma^2}{N} - \frac{\sigma^2}{Nn} \right) \{1 + o(1)\}, \\ \text{var}(\widehat{\text{SE}}_{\text{SDB}}^2) &= \text{var}(\widehat{\text{SE}}^2) \left(\frac{2\sigma^4}{\sigma_4 - \sigma^4} \cdot \frac{N}{R} + 1 \right) \{1 + o(1)\}. \end{aligned}$$

From the above theorem, we can immediately obtain the MSE of $\widehat{\text{SE}}_{\text{SDB}}^2$, which can be expressed as

$$\begin{aligned} \text{MSE}(\widehat{\text{SE}}_{\text{SDB}}^2) &= \text{var}(\widehat{\text{SE}}^2) \left(\frac{2\sigma^4}{\sigma_4 - \sigma^4} \cdot \frac{N}{R} + 1 \right) \{1 + o(1)\} \\ &\quad + \frac{\sigma^4}{N^2 n^2} \{1 + o(1)\} \\ &= \left(\frac{\sigma_4 - \sigma^4}{N^3} + \frac{2\sigma^4}{N^2 R} + \frac{\sigma^4}{N^2 n^2} \right) \{1 + o(1)\}. \end{aligned} \quad (2.6)$$

From the above expression, for the MSE of $\widehat{\text{SE}}_{\text{SDB}}^2$ to be comparable with that of the analytical estimator in (2.2), we will require that $n \gtrsim \sqrt{N}$ and $R \gtrsim N$.

2.5. Bootstrap for Multivariate Random Variables

In this section, we extend the univariate random variable to the multivariate case. With some abuse of notations, we still use $\mathcal{S} = \{X_1, \dots, X_N\}$ to denote the whole sample set where $X_i = (X_{i,1}, \dots, X_{i,p})^\top \in \mathbb{R}^p$ for each $1 \leq i \leq p$, is independent and identically distributed with mean $\mu = (\mu_1, \dots, \mu_p)^\top$ and variance $\Sigma = (\sigma_{j,k})$. We will assume that $p \geq 2$ is fixed for easy exposition. By definition, we have $E(X_{i,j} - \mu_j)(X_{i,k} - \mu_k) = \sigma_{j,k}$ and $E(X_{i,j} - \mu_j)^2 = \sigma_{j,j}$. We assume that the centered fourth moment of $X_{i,j}$ exists; that is, $E(X_{i,j} - \mu_j)^4 = \sigma_{j,j,2} < \infty$ and $E(X_{i,j} - \mu_j)^2(X_{i,k} - \mu_k)^2 = \sigma_{j,k,2} < \infty$ for each $1 \leq j \neq k \leq p$. The estimation accuracy of $\bar{X} = (\bar{X}_1, \dots, \bar{X}_p)^\top \in \mathbb{R}^p$ can be measured by Σ/N , which can be consistently estimated via its sample analogue.

Using the vector notation for X , an estimate of Σ/N via BLB is denoted as $\widehat{\Sigma}_{\text{BLB}}$ where

$$\widehat{\Sigma}_{\text{BLB}} = (RB)^{-1} \sum_{r=1}^R \sum_{b=1}^B (\bar{X}^{(r,b)} - \bar{X}^{(r)}) (\bar{X}^{(r,b)} - \bar{X}^{(r)})^\top$$

with its (j, k) th ($j \neq k$) element denoted by $\widehat{\text{SE}}_{j,k,\text{BLB}}^2$ and the j th diagonal term denoted as $\widehat{\text{SE}}_{j,\text{BLB}}^2$. We now derive the MSE of $\widehat{\text{SE}}_{j,k,\text{BLB}}^2$ in the following theorem.

Theorem 6. For the BLB estimator, assume Condition (C2) holds, we have for $j \neq k$

$$\begin{aligned} \text{MSE}(\widehat{\text{SE}}_{j,k,\text{BLB}}^2) &= \left(\frac{\sigma_{j,k,2} + \sigma_{j,k}^2}{N^3} + \frac{\sigma_{j,j}\sigma_{k,k} + \sigma_{j,k}^2}{N^2 RB} \right. \\ &\quad \left. + \frac{\sigma_{j,k,2} + \sigma_{j,k}^2}{N^2 nR} + \frac{\sigma_{j,k}^2}{N^2 n^2} \right) \{1 + o(1)\}. \end{aligned} \quad (2.7)$$

On the other hand, by the result of Theorem 3, we immediately have the following results for the diagonal terms

$$\begin{aligned} \text{MSE}(\widehat{\text{SE}}_{j,\text{BLB}}^2) &= \left(\frac{\sigma_{j,j,2} - \sigma_{j,j}^2}{N^3} + \frac{2\sigma_{j,j}^2}{N^2 RB} \right. \\ &\quad \left. + \frac{\sigma_{j,j,2} - \sigma_{j,j}^2}{N^2 nR} + \frac{\sigma_{j,j}^2}{N^2 n^2} \right) \{1 + o(1)\}. \end{aligned} \quad (2.8)$$

Define the MSE of $\widehat{\Sigma}_{\text{BLB}}$ as the sum of the MSEs of estimating each individual element. Combining (2.7) and (2.8), we further obtain the MSE of $\widehat{\Sigma}_{\text{BLB}}$ as

$$\begin{aligned} \text{MSE}(\widehat{\Sigma}_{\text{BLB}}) &= \left(\frac{c_1}{N^2 RB} + \frac{c_2}{N^2 nR} + \frac{c_3}{N^2 n^2} + \frac{c_2}{N^3} \right) \{1 + o(1)\}, \end{aligned} \quad (2.9)$$

where

$$c_1 = 2 \sum_{j=1}^p \sigma_{j,j}^2 + \sum_{j_1 \neq j_2} \sigma_{j_1, j_1} \sigma_{j_2, j_2} + \sum_{j_1 \neq j_2} \sigma_{j_1, j_2}^2, \quad (2.10)$$

$$c_2 = \sum_{j=1}^p (\sigma_{j,j,2} - \sigma_{j,j}^2) + \sum_{j_1 \neq j_2} \sigma_{j_1, j_2, 2} + \sum_{j_1 \neq j_2} \sigma_{j_1, j_2}^2, \quad (2.11)$$

$$c_3 = \sum_{j=1}^p \sigma_{j,j}^2 + \sum_{j_1 \neq j_2} \sigma_{j_1, j_2}^2. \quad (2.12)$$

Note that in order to implement our optimal hyperparameter selection method, c_1 , c_2 , and c_3 need to be consistently estimated. Since they only need to be computed once, we know that the computational cost of obtaining them is negligible in comparison to the bootstrap procedure.

Next, the estimate of Σ/N via SB can be seen as

$$\widehat{\Sigma}_{SB} = n(NR)^{-1} \sum_{r=1}^R (\bar{X}^{(r)} - \bar{X})(\bar{X}^{(r)} - \bar{X})^\top,$$

for which we have the following theorem.

Theorem 7. For the SB estimator, assume Condition (C3) holds, we have

$$\text{MSE}(\widehat{SE}_{j,k,SB}^2) = \left\{ \frac{\sigma_{jk}^2}{N^2 n^2} + \frac{\sigma_{jk}^2 + \sigma_{jj}\sigma_{kk}}{N^2 R} + \frac{\sigma_{j,k,2} + \sigma_{jk}^2}{N^3} \right\} \{1 + o(1)\}, \quad 1 \leq j, k \leq p.$$

Moreover, for its diagonal terms, from Theorem 4, we can obtain

$$\text{MSE}(\widehat{SE}_{j,SB}^2) = \left(\frac{2\sigma_{jj}^2}{N^2 R} + \frac{\sigma_{jj}^2}{N^2 n^2} + \frac{\sigma_{j,j,2} - \sigma_{jj}^2}{N^3} \right) \{1 + o(1)\}, \quad 1 \leq j \leq p.$$

Combining the above results, we further obtain the MSE for all the elements in $\widehat{\Sigma}_{SB}$ as

$$\text{MSE}(\widehat{SE}_{SB}^2) = \left(\frac{c_1}{N^2 R} + \frac{c_3}{N^2 n^2} + \frac{c_2}{N^3} \right) \{1 + o(1)\}. \quad (2.13)$$

Lastly, we derive the estimate of Σ/N via SDB denoted as

$$\widehat{\Sigma}_{SDB} = R^{-1} \sum_{r=1}^R (\bar{X}^{(r,1)} - \bar{X}^{(r)})(\bar{X}^{(r,1)} - \bar{X}^{(r)})^\top,$$

with its MSE characterized by the following theorem.

Theorem 8. For the SDB estimator, assume Condition (C3) holds, we have for $j \neq k$

$$\text{MSE}(\widehat{SE}_{j,k,SDB}^2) = \left(\frac{\sigma_{j,k,2} + \sigma_{jk}^2}{N^3} + 3 \frac{\sigma_{j,k,2} - \sigma_{jj}\sigma_{kk}}{N^2 n R} + \frac{\sigma_{jj}\sigma_{kk} + \sigma_{jk}^2}{N^2 R} + \frac{\sigma_{jk}^2}{N^2 n^2} \right) \{1 + o(1)\}.$$

For the diagonal terms, from Theorem 5, we can obtain

$$\text{MSE}(\widehat{SE}_{j,SDB}^2) = \left(\frac{2\sigma_{jj}^2}{N^2 R} + \frac{\sigma_{jj}^2}{N^2 n^2} + \frac{\sigma_{j,j,2} - \sigma_{jj}^2}{N^3} \right) \{1 + o(1)\}.$$

Combining the above results, we immediately obtain the MSE for all the elements in $\widehat{\Sigma}_{SDB}$ as $\text{MSE}(\widehat{SE}_{SDB}^2) = (c_1/N^2 R + c_3/N^2 n^2 + c_2/N^3 + 3c_4/N^2 n R)\{1 + o(1)\}$, where $c_4 = \sum_{j_1 \neq j_2} (\sigma_{j_1, j_2, 2} - \sigma_{j_1, j_1} \sigma_{j_2, j_2})$. When R is of a larger order of n , we have

$$\text{MSE}(\widehat{SE}_{SDB}^2) = \left(\frac{c_1}{N^2 R} + \frac{c_3}{N^2 n^2} + \frac{c_2}{N^3} \right) \{1 + o(1)\}. \quad (2.14)$$

We note that the MSEs of these bootstrap estimators in (2.9), (2.13), and (2.14) based on multivariate random variables are similar to their univariate counterparts but with different constants in the numerators.

3. Optimal Subsampling Bootstraps and General Parameters

3.1. Optimal Subsampling Bootstraps

All the subsampling bootstrap methods discussed so far require the specification of the hyperparameters including the size of the subsamples n , the number of subsamples R , and for the BLB, the number of the resamples B . For these methods to produce satisfactory statistical performance, intuitively these hyperparameters should be set as large as possible. However, setting them unnecessarily large will incur additional computational cost that we want to avoid in the first place when resorting to subsampling bootstrap instead of the bootstrap. Thus, there is a clear tradeoff between statistical accuracy of the estimator and computational budget. Our approach to achieve the optimal tradeoff is to optimize the statistical accuracy subject to a fixed pre-specified computational constraint.

Having analyzed the MSE of the three subsampling bootstrap methods BLB, SB, and SDB in the previous section, we are ready to derive the hyperparameter values that give optimal statistical accuracy in terms of MSE with a fixed computational cost. Since these bootstrap methods are solely developed for massive data, it makes sense to compare their performance when the full sample resides on a hard drive. That is, to conduct any bootstrap scheme, the data will have to be repeatedly read from the hard disk, while the computation happens in memory. Sampling n data points from the disk costs $O(n)$ in terms of time complexity. Denote the computational time for obtaining an estimator on a sample with size n as $t(n)$. For an estimator taking weighted data representation, we assume that the estimation time is $t(n)$ where n is the number of distinct points in the sample. This is the scenario where BLB and SDB can gain computationally over TB. We summarize the sampling and computational cost in Table 1, with more details provided when a specific method is discussed. Later in this article, we focus on those estimators where $t(n) \asymp n$, that is, those estimators that can be obtained in a computational time linear in the sample size. This setup comprises a large class of estimators, including the sample mean and the least-squares estimator of the coefficients in a linear model. For example, to compute the mean discussed so far, the total time required for BLB is $\alpha_1 \cdot nR + \alpha_2 \cdot nRB$, where α_1 and α_2 are two parameters specific to the computer system. See below for more details.

We start by looking at SB first. Since SB needs to sample n data points from the disk R times, its sampling cost is $O(nR)$. To compute \widehat{SE}_{SB}^2 , the leading term is $\bar{X}^{(b)}$, that is, the sample average of each bootstrap subsample, which is evaluated R times, resulting in a computational cost of $O(nR)$. Therefore, the total

Table 1. Sampling time and estimation time of various subsampling bootstrap methods.

Name	Sampling time	Estimation time
Bootstrap	nB	$(B + 1) \times t(n)$
BLB	nR	$R(B + 1) \times t(n)$
SB	nR	$R \times t(n)$
SDB	nR	$2R \times t(n)$

NOTE: N : the total sample size; n : the sample size of subsamples; R : the number of subsamples; B : the number of resamples.

computational cost for the SB is of the order $O(nR)$. Motivated by (2.5) and (2.13), to identify the optimal subsample size and optimal number of subsamples in terms of minimizing the MSE of \widehat{SE}_{SB}^2 , we need to find the (n, R) pair such that

$$\begin{aligned} (n^*, R^*) &= \arg \min_{n, R} \text{MSE}(\widehat{SE}_{SB}^2) \\ &= \arg \min_{n, R} \left\{ \frac{c'_1}{R} + \frac{c'_2}{n^2} \right\}, \quad \text{s.t.} \quad \alpha_{SB} \cdot nR \leq C_{\max}, \end{aligned}$$

where α_{SB} is a constant quantifying the time constant before nR for the SB implementation, and C_{\max} is the computational budget. Here, $c'_1 = 2$ and $c'_2 = 1$ when estimating the mean parameter for univariate random variables. For multivariate random variables, we have $c'_1 = c_1$ and $c'_2 = c_3$ where c_1 and c_3 are defined in (2.10) and (2.12), respectively. Note that for the multivariate case, c_1 and c_3 can be consistently estimated and the computational cost of obtaining them is negligible in comparison to the bootstrap procedure. In practice, if we have the computational budget to subsample R times, each subsample with size n , that is, the computational budget is nR , then solving the above optimization problem gives SB optimal hyperparameters:

$$\begin{aligned} n^* &= \lfloor \{(2c'_2/c'_1) \cdot (C_{\max}/\alpha_{SB})\}^{1/3} \rfloor, \\ R^* &= \lfloor \{c'_2/2c'_1\}^{1/3} \cdot (C_{\max}/\alpha_{SB})^{2/3} \rfloor, \end{aligned} \quad (3.1)$$

where $\lfloor s \rfloor$ stands for the largest integer that is not larger than s .

For the SDB procedure, the MSEs for the univariate and multivariate case as in (2.6) and (2.14) are similar to those of the SB in that the optimal (n, R) pair is given by SDB optimal hyperparameters:

$$\begin{aligned} n^* &= \lfloor \{(2c'_2/c'_1) \cdot (C_{\max}/\alpha_{SDB})\}^{1/3} \rfloor, \\ R^* &= \lfloor \{c'_2/2c'_1\}^{1/3} \cdot (C_{\max}/\alpha_{SDB})^{2/3} \rfloor. \end{aligned} \quad (3.2)$$

In practice, α_{SB} and α_{SDB} can be dependent on the subsample size n because different computer architectures implement their algorithms differently. However, their dependence on R is linear because resampling is just a repetition of the same operation. Our strategy is to tune α_{SB} or α_{SDB} progressively several times. The time needed to tune them can be costly but we minimize its impact by pilot-running experiments on small n and small R . The details can be found in the simulations in Section 4.2–4.3 and also Appendix B.1, supplementary materials.

We are now ready to discuss how to choose optimal hyperparameters for BLB which involves three parameters n , R and B . The sampling cost of BLB is similar to that of SB which is $O(nR)$. The computational cost can be seen as $O(nRB)$ as the leading computation is to calculate $\bar{X}^{(r,b)}$, each based on a resample of size n , for a total of RB times. Motivated by the expression for the MSE in (2.4) and (2.9), to identify the optimal hyperparameter triple, we minimize the MSE of \widehat{SE}_{BLB}^2 subject to the constraint that the computational cost is capped at C_{\max} . More specifically, we solve (n^*, R^*, B^*) equals to

$$\begin{aligned} \arg \min_{n, R, B} \text{MSE}(\widehat{SE}_{BLB}^2) &= \arg \min_{n, R, B} \left\{ \frac{\tilde{c}_1}{RB} + \frac{\tilde{c}_2}{nR} + \frac{\tilde{c}_3}{n^2} \right\}, \\ \text{s.t.} \quad \alpha_1 \cdot nRB + \alpha_2 \cdot nR &\leq C_{\max}, \end{aligned} \quad (3.3)$$

where $\tilde{c}_1 = 2\sigma^4(\sigma_4 - \sigma^4)^{-1}$, $\tilde{c}_2 = 1$, and $\tilde{c}_3 = \sigma^4(\sigma_4 - \sigma^4)^{-1}$ for the univariate mean estimating case and $\tilde{c}_1 = c_1$, $\tilde{c}_2 = c_2$, and $\tilde{c}_3 = c_3$ for the multivariate case with c_1 , c_2 , and c_3 defined in (2.10), (2.11), and (2.12), respectively. In the above optimization problem, α_1 and α_2 are two constants quantifying system specific constants for the sampling cost and the computational cost, respectively. To figure out their values for a specific problem, we can take different combinations of (n, R, B) and record the running time of BLB under these combinations. Fitting a linear model with the running time as the response variable and nRB and nR as the two covariates, we can obtain the estimates of α_1 and α_2 . Our simulation study below suggests that with very few combinations (e.g., 8), we can obtain fairly accurate estimates of α_1 and α_2 with the associated R^2 as large as 98%. In estimating α_1 and α_2 using pilot runs, we can keep R , the number of resamples, small because any dependence of the computational and sampling time on R will be proportional to it. Therefore, the extra time needed to estimate α_1 and α_2 is negligible in comparison to the actual implementation of any bootstrap subsampling method discussed so far because for any bootstrap, R is often very large.

Solving the optimization problem above for BLB amounts to finding the minimum value of $\tilde{c}_1(RB)^{-1} + \tilde{c}_2(nR)^{-1} + \tilde{c}_3(n)^{-2}$ under the constraint $\alpha_1 nRB + \alpha_2 nR = C_{\max}$. To this end, let (n, R, B) be an arbitrary specification such that $\alpha_1 nRB + \alpha_2 nR = C_{\max}$. By Cauchy's inequality, we have

$$\begin{aligned} &\left(\frac{\tilde{c}_1}{RB} + \frac{\tilde{c}_2}{nR} + \frac{\tilde{c}_3}{n^2} \right) (\alpha_1 nRB + \alpha_2 nR) \\ &= \left(\frac{\tilde{c}_1}{RB} + \frac{\tilde{c}_2}{nR} \right) (\alpha_1 nRB + \alpha_2 nR) + \frac{\tilde{c}_3}{n^2} (\alpha_1 nRB + \alpha_2 nR) \\ &\geq \left(\sqrt{\tilde{c}_1 \alpha_1} \sqrt{n} + \sqrt{\tilde{c}_2 \alpha_2} \right)^2 + C_{\max} \frac{\tilde{c}_3}{n^2} := f(n). \end{aligned} \quad (3.4)$$

The first inequality in (3.4) becomes an equality if $\tilde{c}_1 \alpha_2 n/B = \tilde{c}_2 \alpha_1 B$. When n is fixed, this leads to the BLB optimal hyperparameters as

$$\begin{aligned} R^* &= \lfloor C_{\max}/(\alpha_1 n B^* + \alpha_2 n) \rfloor, \\ B^* &= \lfloor (\tilde{c}_1 \alpha_2 / \tilde{c}_2 \alpha_1)^{1/2} \sqrt{n} \rfloor. \end{aligned} \quad (3.5)$$

For n , our simulation results suggest that $n = \lfloor N^{0.7} \rfloor$ is an optimal choice in most cases consistent with the findings in Kleiner et al. (2014) and Sengupta, Volgushev, and Shao (2016). Accordingly, we can directly use $n = \lfloor N^{0.7} \rfloor$ in practice.

3.2. General Parameters and Statistics

We have so far focused on estimating the mean. In this section, we study more general parameters. Without loss of generality, we first assume that the random variables $X_i \in \mathbb{R}^1$, $i = 1, \dots, N$, have mean μ and variance σ^2 . Our interest is to estimate $\theta = g(\mu)$, where $g(\cdot)$ is a known, possibly complicated, but sufficiently smooth function with twice differentiable. Extension to random vectors is straightforward. One simple estimator of θ can be constructed as $\hat{\theta} = g(\bar{X})$ since \bar{X} estimates μ . Asymptotically, we have $\hat{\theta} - \theta = \dot{g}(\mu)(\bar{X} - \mu)\{1 + o_p(1)\}$, where $\dot{g}(\cdot)$ is the first order derivative of $g(\cdot)$. We then have $\sqrt{N}(\hat{\theta} - \theta) \rightarrow_d N(0, \dot{g}^2(\mu)\sigma^2)$. This means that the asymptotic SE^2 of $\hat{\theta}$ is given by $SE^{*2} = \dot{g}^2(\mu)\sigma^2/N$. Accordingly, a natural estimator for SE^{*2}

is given by $\widehat{SE}^{*2} = \dot{g}^2(\bar{X})\widehat{\sigma}^2/N$, which can be easily computed if $\dot{g}^2(\mu)$ is analytically simple.

Next, we discuss how to use the three subsampling bootstraps discussed so far to estimate SE^{*2} . We start with the traditional bootstrap method for which a natural estimator of SE^{*2} is given by

$$\widehat{SE}_{TB}^{*2} = B^{-1} \sum_{b=1}^B \left\{ g(\bar{X}^{(b)}) - g(\bar{X}) \right\}^2. \quad (3.6)$$

It is seen that

$$\begin{aligned} \widehat{SE}_{TB}^{*2} &= B^{-1} \sum_{b=1}^B \left\{ \dot{g}(\bar{X})(\bar{X}^{(b)} - \bar{X}) + 2^{-1}\ddot{g}^2(\bar{X}^*)(\bar{X}^{(b)} - \bar{X})^2 \right\}^2 \\ &= B^{-1} \sum_{b=1}^B \left\{ \dot{g}^2(\mu)(\bar{X}^{(b)} - \bar{X})^2 + \dot{g}(\mu)\ddot{g}^2(\mu)(\bar{X}^{(b)} - \bar{X})^3 \right. \\ &\quad \left. + 4^{-1}\ddot{g}^4(\mu)(\bar{X}^{(b)} - \bar{X})^4 \right\} \left\{ 1 + o_p(1) \right\}, \end{aligned}$$

where $\ddot{g}(\cdot)$ is the second order derivative of $g(\cdot)$ and \bar{X}^* is between \bar{X} and $\bar{X}^{(b)}$. Denote $\widehat{SE}_{TB}^{*2} = (\dot{g}^2(\mu)A_1^* + \dot{g}(\mu)\ddot{g}^2(\mu)A_2^* + 4^{-1}\ddot{g}^4(\mu)A_3^*)\{1 + o_p(1)\}$. According to the results in Section 1.10 of the supplementary material, we have $MSE(A_2^*) = O((N^3B)^{-1})$. If $MSE(A_3^*)$ is ignorable as compared with $MSE(A_1^*)$, it can be verified that $MSE(\widehat{SE}_{TB}^{*2})$ should be determined by $MSE(A_1^*)$ and the reminder terms can be ignored accordingly. By the result of Theorem 2, we have $MSE(A_1^*) = MSE(\widehat{SE}_{TB}^{*2}) = O(N^{-3})$. Hence, we have

$$\begin{aligned} \widehat{SE}_{TB}^{*2} &= \dot{g}^2(\mu)A_1^* \left\{ 1 + o_p(1) \right\} \\ &= \dot{g}^2(\mu) \left\{ B^{-1} \sum_{b=1}^B (\bar{X}^{(b)} - \bar{X})^2 \right\} \left\{ 1 + o_p(1) \right\} \\ &= \dot{g}^2(\mu)\widehat{SE}_{TB}^2 \left\{ 1 + o_p(1) \right\}. \end{aligned}$$

This verifies the fact that the asymptotic behavior of \widehat{SE}_{TB}^{*2} is determined by that of \widehat{SE}_{TB}^2 defined in Section 2.1 and higher order terms can be ignored. Hence, c'_1 , c'_2 , and c'_3 are the same constants defined for the sample mean and can be consistently estimated. Similar analysis can be performed for the BLB, SB, and SDB methods, demonstrating that higher order terms can be ignored. The associated estimators for SE^{*2} are given by

$$\begin{aligned} \widehat{SE}_{BLB}^{*2} &= (RB)^{-1} \sum_{r=1}^R \sum_{b=1}^B \left\{ g(\bar{X}^{(r,b)}) - g(\bar{X}^{(r)}) \right\}^2 \\ &= \dot{g}^2(\mu)\widehat{SE}_{BLB}^2 \left\{ 1 + o_p(1) \right\}, \end{aligned} \quad (3.7)$$

$$\begin{aligned} \widehat{SE}_{SB}^{*2} &= \left(\frac{n}{N}\right)R^{-1} \sum_{r=1}^R \left\{ g(\bar{X}^{(b)}) - g(\bar{X}) \right\}^2 \\ &= \dot{g}^2(\mu)\widehat{SE}_{SB}^2 \left\{ 1 + o_p(1) \right\}, \end{aligned} \quad (3.8)$$

$$\begin{aligned} \widehat{SE}_{SDB}^{*2} &= R^{-1} \sum_{r=1}^R \left\{ g(\bar{X}^{(r,1)}) - g(\bar{X}^{(r)}) \right\}^2 \\ &= \dot{g}^2(\mu)\widehat{SE}_{SDB}^2 \left\{ 1 + o_p(1) \right\}, \end{aligned} \quad (3.9)$$

respectively.

When $X_i \in \mathbb{R}^p$ is a random vector and θ is one-dimensional, we have $\widehat{\theta} - \theta = g(\bar{X}) - g(\mu) = \dot{g}(\mu)^\top (\bar{X} - \mu)\{1 + o_p(1)\}$, leading to $\sqrt{N}(\widehat{\theta} - \theta) \rightarrow_d N(0, \dot{g}(\mu)^\top \Sigma \dot{g}(\mu))$. This means that the asymptotic SE^2 of $\widehat{\theta}$ is given by $SE^{*2} = \dot{g}(\mu)^\top \Sigma \dot{g}(\mu)/N$. Similar to the above calculation, we have $\widehat{SE}^{*2} = N^{-1}\dot{g}(\mu)^\top \widehat{\Sigma} \dot{g}(\mu)\{1 + o_p(1)\}$ and $\widehat{SE}^{*2} - SE^{*2} = N^{-1}\dot{g}(\mu)^\top (\widehat{\Sigma} - \Sigma)\dot{g}(\mu)\{1 + o_p(1)\}$. This suggests that the asymptotic behavior of \widehat{SE}^{*2} is determined by the estimate of Σ .

When $\theta \in \mathbb{R}^d$ for $d > 1$, we use Σ^* to denote the covariance matrix of $\widehat{\theta} = g(\bar{X})$. By a careful calculation, we can find that $\Sigma^* = \dot{g}(\mu)^\top \Sigma \dot{g}(\mu)/N$ still holds. We then discuss how to use the three subsampling bootstraps to estimate Σ^* . We first study the SB method for which a natural estimator is given by

$$\begin{aligned} \widehat{\Sigma}_{SB}^* &= n(NR)^{-1} \sum_{r=1}^R \left\{ g(\bar{X}^{(r)}) - g(\bar{X}) \right\} \left\{ g(\bar{X}^{(r)}) - g(\bar{X}) \right\}^\top \\ &= \left\{ n(NR)^{-1} \sum_{r=1}^R \dot{g}(\mu)^\top (\bar{X}^{(r)} - \bar{X}) (\bar{X}^{(r)} - \bar{X})^\top g(\mu) \right\} \\ &\quad \left\{ 1 + o_p(1) \right\} \\ &= \dot{g}(\mu)^\top \widehat{\Sigma}_{SB} \dot{g}(\mu) \left\{ 1 + o_p(1) \right\}. \end{aligned}$$

This suggests that the asymptotic behavior of $\widehat{\Sigma}_{SB}^*$ is largely determined by that of $\widehat{\Sigma}_{SB}$. Similar analysis can be conducted for the BLB and SDB methods. The associated estimators for Σ^* are given by

$$\begin{aligned} \widehat{\Sigma}_{BLB}^* &= (RB)^{-1} \sum_{r=1}^R \sum_{b=1}^B \left\{ g(\bar{X}^{(r,b)}) - g(\bar{X}^{(r)}) \right\} \left\{ g(\bar{X}^{(r,b)}) - g(\bar{X}^{(r)}) \right\}^\top \\ &= \dot{g}(\mu)^\top \widehat{\Sigma}_{BLB} \dot{g}(\mu) \left\{ 1 + o_p(1) \right\}, \\ \widehat{\Sigma}_{SDB}^* &= R^{-1} \sum_{r=1}^R \left\{ g(\bar{X}^{(r,1)}) - g(\bar{X}^{(r)}) \right\} \left\{ g(\bar{X}^{(r,1)}) - g(\bar{X}^{(r)}) \right\}^\top \\ &= \dot{g}(\mu)^\top \widehat{\Sigma}_{SDB} \dot{g}(\mu) \left\{ 1 + o_p(1) \right\}, \end{aligned}$$

respectively. Again, the asymptotic behaviors of $\widehat{\Sigma}_{BLB}^*$, $\widehat{\Sigma}_{SB}^*$, and $\widehat{\Sigma}_{SDB}^*$ are largely determined by those of $\widehat{\Sigma}_{BLB}$, $\widehat{\Sigma}_{SB}$, and $\widehat{\Sigma}_{SDB}$. Accordingly, the above mentioned optimization approaches in Section 3.1 are still applicable for the general parameters and statistics, as long as the target parameter can be expressed as $\theta = g(\mu)$ where $g(\cdot)$ is a sufficiently smooth function.

Note that the computational advantages of SDB and BLB relative to the traditional bootstrap lies in the fact that the estimator of interest can take weighted data representation as its argument. A large class of commonly used estimators, including M -estimators, can be represented by weighted data. We now discuss a few examples that can be placed in the general parameter framework.

Example 1 (Estimators based on moments). Denote $\mu_j = E(X - \mu)^j$ as the j th moment of a univariate random variable X and its moment estimator as $\widehat{\mu}_j = N^{-1} \sum_{i=1}^N (X_i - \bar{X})^j$. Note that $\widehat{\mu}_j$ is a consistent estimator for μ_j which can take weighted data representation. Thus, the optimal hyperparameter selection method presented in Section 3.1 is applicable for minimizing the MSE of $\widehat{\theta} = g(\widehat{\mu})$ when one uses the subsampling bootstrap for inference.

As another example, we will later apply our method to the problem of estimating the correlation coefficient when missing data is present, a problem considered in Shao and Wang (2002). Denote the dataset as $X = (X_1, \dots, X_N) \in \mathbb{R}^N$ and $Y = (Y_1, \dots, Y_N) \in \mathbb{R}^N$ and the missing indicator as w_i , where $w_i = 1$ if the i th observation (X_i, Y_i) is missed and $w_i = 0$ if not. Shao and Wang (2002) is interested in estimating the correlation coefficient

$$\rho = \frac{\mu_{XY} - \mu_X \mu_Y}{\sqrt{\mu_{X,2} - \mu_X^2} \sqrt{\mu_{Y,2} - \mu_Y^2}},$$

where $\mu_X = E(w_i X_i)$, $\mu_{X,2} = E(w_i X_i^2)$, $\mu_Y = E(w_i Y_i)$, $\mu_{Y,2} = E(w_i Y_i^2)$, and $\mu_{XY} = E(w_i X_i Y_i)$. A simple sample estimator is given by

$$\hat{\rho} = \frac{\sum_{i=1}^N w_i X_i Y_i - \widehat{X} \widehat{Y} / \widehat{N}}{\sqrt{\sum_{i=1}^N (w_i X_i^2 - \widehat{X}^2)} \sqrt{\sum_{i=1}^N (w_i Y_i^2 - \widehat{Y}^2)}}, \quad (3.10)$$

which satisfy $\hat{\rho} = \rho + o(1)$. Write $\mu = (\mu_X, \mu_{X,2}, \mu_Y, \mu_{Y,2}, \mu_{XY})^\top$. Then we can write $\rho = g(\mu)$ and $\hat{\rho} = g(\widehat{\mu})$. The optimization technique for multivariate random variables in Section 3.1 can be directly applied to optimize the estimation accuracy of SE^2 for $\hat{\rho}$.

Example 2 (Linear regression). Consider the problem of estimating the regression coefficient in linear regression $Y_i = X_i^\top \beta + \varepsilon_i$, $i = 1, \dots, N$ where $X_i \in \mathbb{R}^d$ and $\beta \in \mathbb{R}^d$ is the regression coefficient. We use $d = 2$ for illustration purposes. Note the ordinary least-squares estimator of β is given by $\widehat{\beta} = (X^\top X)^{-1} X^\top Y$. Write $\mu = (\mu_1, \mu_2, \mu_3, \mu_4, \mu_5)^\top$ where $\mu_1 = E(X_{i,1}^2)$, $\mu_2 = E(X_{i,2}^2)$, $\mu_3 = E(X_{i,1} X_{i,2})$, $\mu_4 = E(X_{i,1} Y_i)$ and $\mu_5 = E(X_{i,2} Y_i)$. It is easy to see $\widehat{\beta} = g(\widehat{\mu})$ for some function g . Since each element of $\widehat{\mu}$ is a mean estimator, we can apply the results in Section 3.1 to choose the optimal choice for various hyperparameters.

Example 3 (Logistic regression). Consider the problem of estimating the regression coefficient in logistic regression $P(Y_i = 1 | X_i) = p_i(\beta) = \exp(X_i^\top \beta) / (1 + \exp(X_i^\top \beta))$, $i = 1, \dots, N$ where $X_i \in \mathbb{R}^d$ is the covariate and $\beta \in \mathbb{R}^d$ is the regression coefficient. We still use $d = 2$ for illustration. The one-step estimator of β is given by $\widehat{\beta} = \widetilde{\beta} + \{\sum_{i=1}^N \omega_i(\widetilde{\beta}) X_i X_i^\top\}^{-1} \sum_{i=1}^N \{Y_i - p_i(\widetilde{\beta})\} X_i$ with $\omega_i(\beta) = p_i(\beta) \{1 - p_i(\beta)\}$, where $\widetilde{\beta}$ is a pilot, consistent estimator of β using a subsample with sample size n . Write $\mu = (\mu_1, \mu_2, \mu_3, \mu_4, \mu_5)^\top$ where $\mu_1 = E(\omega_i(\beta) X_{i,1}^2)$, $\mu_2 = E(\omega_i(\beta) X_{i,2}^2)$, $\mu_3 = E(\omega_i(\beta) X_{i,1} X_{i,2})$, $\mu_4 = E[X_{i,1} \{Y_i - p_i(\beta)\}]$ and $\mu_5 = E[X_{i,2} \{Y_i - p_i(\beta)\}]$. Then $\widehat{\beta} = g(\widehat{\mu})$ for some function g when β is replaced by $\widetilde{\beta}$.

Example 4 (Two stage least squares). Consider the problem of estimating the regression coefficient in linear regression $Y_i = X_i \beta + u_i$, $i = 1, \dots, N$, where the random error u_i is correlated with $X_i \in \mathbb{R}^1$ such that $\text{cov}(X_i, u_i) \neq 0$. To estimate β , an idea is to identify an instrumental variable, denoted as $Z_i \in \mathbb{R}^1$, and employ the so-called two stage least squares estimator given by $\widehat{\beta}_{IV} = (\sum_{i=1}^N Z_i Y_i) / (\sum_{i=1}^N Z_i X_i)$. Write $\mu = (\mu_1, \mu_2)^\top$ where $\mu_1 = E(Z_i Y_i)$ and $\mu_2 = E(Z_i X_i)$. Clearly, we have $\widehat{\beta} = g(\widehat{\mu})$ for some function g .

4. Numerical Study

We conduct extensive simulation study to assess the performance of the hyperparameter specification method. We will do this by first verifying the MSE results in Section 4.1 for estimating the mean. In Section 4.2, we apply our hyperparameter selection method to linear regression and further extend it to logistic regression in Section 4.3. Due to page limit, we leave the estimation of the correlation coefficient defined for missing data, hypotheses test and comparison with plain subsampling to the supplementary Material. All programs are written in python 3.7 and run on a cloud computing platform called Matpool (<https://www.matpool.com>). The simulation is run on a computer equipped with NVIDIA Tesla P100 data center GPU, 16 GB of graphics memory, 64 GB ram and 500 GB SSD capacity.

4.1. Verify the Formula for the Mean Squared Errors

This section aims to verify the $MSE(\widehat{SE}^2)$ formula derived in Section 2 as expressed in (2.2), (2.3), (2.4), (2.5), and (2.6), when the mean is estimated. Toward this, we repeat each experiment M times and denote $\widehat{SE}^{2(m)}$ as the estimate of SE^2 on the m th simulation replicate. The true mean squared error (MSE) can then be estimated as

$$\widehat{MSE} = M^{-1} \sum_{m=1}^M \left(\widehat{SE}^{2(m)} - SE^{*2} \right)^2 \quad (4.1)$$

where $SE^{*2} = M^{-1} \sum_{m=1}^M (\overline{X}^{(m)} - \overline{X})^2$ and $\overline{X}^{(m)}$ is the estimator of \overline{X} based on the m th simulation replicate. With some abuse of notation, denote MSE^* as the estimator of $MSE(\widehat{SE}^2)$ in (2.2), (2.3), (2.4), (2.5), and (2.6), where σ and σ_4 are estimated as $\widehat{\sigma}^2 = \sum_{i=1}^N (X_i - \overline{X})^2 / N$ and $\widehat{\sigma}_4 = \sum_{i=1}^N (X_i - \overline{X})^4 / N$, respectively. We then compare the difference between \widehat{MSE} and MSE^* by evaluating the ratio \widehat{MSE} / MSE^* .

We generate $X_i \in \mathbb{R}^1$ for $i = 1, \dots, N$ either from a standard normal distribution (Normal) or a centered standard exponential distribution (Exponential). We take $N = 10^5$, $n \in \{[N^{0.4}], [N^{0.5}], [N^{0.6}]\}$, $B \in \{25, 50\}$, and $R \in \{25, 50\}$. For each setup, a total of $M = 1000$ datasets are generated. The simulation results are summarized in Table 2. From this table, we observe that all of the ratios are close to one, suggesting that our analytical formula are fairly accurate. Similarly, we can verify the formula for $MSE(\widehat{SE}^2)$ derived in (2.9), (2.13), and (2.14) with $X_i \in \mathbb{R}^p$. We find that the results are similar to those presented in Table 2 and omit them to save space.

4.2. Hyperparameter Selection for Linear Regression

As an application, we consider a two-dimensional linear regression model $Y_i = X_i^\top \beta + \varepsilon_i$ for $i = 1, \dots, N$, where $\beta_0 = (0.1, 0.1)^\top$ is the regression coefficient to be estimated, X_i follows a bivariate normal distribution $N(0, \Sigma)$ with $\Sigma = (1, 0; 0, 1)$ and ε_i are the independent noise following the standard normal distribution. We denote the ordinary least-squares estimator as $\widehat{\beta}$ whose theoretical covariance is $\text{cov}(\widehat{\beta}) = \Sigma / N$. This covariance can be consistently estimated by the BLB as

Table 2. Simulation results for checking the analytical formula of MSE for various bootstrap methods.

n	Parameter		Normal					Exponential				
	B	R	AF	TB	BLB	SB	SDB	AF	TB	BLB	SB	SDB
[N ^{0.4}]	25	25	1.007	0.982	0.979	1.004	1.002	1.055	1.003	0.969	0.944	1.116
	25	50	1.007	0.982	0.914	1.020	1.010	1.055	1.003	1.001	0.976	1.061
	50	25	1.007	0.927	0.975	1.004	1.002	1.055	1.006	0.929	0.944	1.116
	50	50	1.007	0.927	0.987	1.020	1.010	1.055	1.006	1.023	0.976	1.061
[N ^{0.5}]	25	25	1.007	0.982	1.042	1.075	1.000	1.055	1.003	0.991	0.964	1.069
	25	50	1.007	0.982	0.982	1.075	1.014	1.055	1.003	1.028	1.013	1.032
	50	25	1.007	0.927	0.937	1.014	1.000	1.055	1.006	0.962	0.964	1.069
	50	50	1.007	0.927	0.950	1.014	1.014	1.055	1.006	0.946	1.013	1.032
[N ^{0.6}]	25	25	1.007	0.982	0.946	0.877	0.932	1.055	1.003	1.044	0.967	0.951
	25	50	1.007	0.982	1.003	0.877	1.030	1.055	1.003	0.975	1.121	0.977
	50	25	1.007	0.927	1.059	0.991	0.932	1.055	1.006	0.965	0.967	0.951
	50	50	1.007	0.927	1.082	0.991	1.030	1.055	1.006	1.049	1.121	0.977

Table 3. Simulation results for linear regression.

Original (R, B)	MSE ratios					CPU time ratios				
	κ ₁	κ ₂	κ ₃	κ ₄	κ ₅	κ ₁	κ ₂	κ ₃	κ ₄	κ ₅
(14,4089)	0.705	0.276	0.780	0.580	0.208	0.995	1.075	1.063	0.949	0.946
(15,4097)	0.763	0.289	0.745	0.502	0.194	0.986	1.095	1.074	0.931	0.929
(18,4291)	0.682	0.257	0.718	0.551	0.206	1.002	1.035	1.013	0.969	0.975
(24,3847)	0.694	0.259	0.779	0.593	0.220	0.996	1.076	1.063	0.935	0.931
(27,4430)	0.781	0.319	0.816	0.569	0.214	1.003	1.020	1.000	0.977	0.981
(29,4205)	0.774	0.294	0.770	0.588	0.216	0.997	1.072	1.054	0.938	0.940
(363,65)	0.865	0.276	0.821	0.600	0.223	0.984	1.071	1.044	0.930	0.947
(408,57)	0.799	0.262	0.689	0.571	0.213	0.980	1.057	1.029	0.938	0.956
(414,71)	0.810	0.283	0.809	0.579	0.220	0.978	1.018	0.993	0.962	0.979
(471,66)	0.822	0.307	0.839	0.521	0.200	0.986	0.999	0.976	0.994	1.009
(513,67)	0.804	0.300	0.786	0.545	0.189	0.965	0.998	0.969	0.960	1.074
(584,68)	0.666	0.293	0.981	0.488	0.166	0.953	0.960	0.945	0.984	1.097

NOTE: Here, (R, B) is the original specified setting for BLB with fixed n = [N^{0.7}]. We denote the BLB method using the specified (R, B) combination as BLB, that using the optimal (R, B) combination via our approach as BLB*. Denote κ₁ = BLB*/BLB, κ₂ = SDB*/SDB, κ₃ = SB*/SB, κ₄ = BLB*/SDB*, and κ₅ = BLB*/SB*.

$$\widehat{SE}_{BLB}^{*2} = \frac{1}{RB} \sum_{r=1}^R \sum_{b=1}^B (\widehat{\beta}^{(r,b)} - \widehat{\beta}^{(r)}) (\widehat{\beta}^{(r,b)} - \widehat{\beta}^{(r)})^T,$$

where $\widehat{\beta}^{(r,b)}$ is the ordinary least-squares estimator obtained from the bth resample of the rth subsample and $\widehat{\beta}^{(r)}$ is the ordinary least-squares estimate based on the rth resample. Since we are estimating matrices, we define $MSE(\widehat{SE}_{BLB}^{*2}) = \|\widehat{SE}_{BLB}^{*2} - cov(\widehat{\beta})\|_F$ as the mean square error of the BLB estimate of $cov(\widehat{\beta})$. We can define the MSE of the SB and SDB methods likewise. In the experiments below, we fix n = [N^{0.7}] for the prespecified setting and repeat the simulation under each setting M = 50 times. To compare the performance of BLB, SDB, and SB, we fix the computational budget at the same value of C_{max}. We emphasize that through all of our simulations, for optimal BLB, the actual time budget is always C_{max} - C₀, where C₀ is the cost for estimating coefficients α₁ and α₂.

For this example, we evaluate the relative performance of BLB, SB, and SDB under the same computational budget. This is not always done in the literature due to the difficulty of controlling different bootstrap methods under the same time cost. Here, the original SDB and SB corresponds to the cases with n = [N^{0.7}] and R is further calculated based on the time consumption of BLB. In order to do this, we will need to estimate time coefficients α_{SB}, α_{SDB}, and α₁ and α₂ in the time constraints in the optimization problems in Section 3, using pilot runs.

In our pilot runs, we randomly generate 12 combinations of (R, B) values from [U(1, 10)] × [U(1, 80)] and record the computational time for each combination. Afterwards, we fit a linear regression model with the running time as the response variable and the corresponding (nRB, nR) as the covariates for BLB. This gives ordinary least-squares estimates of α₁ and α₂. For SB and SDB, we use the iterative strategy described in Appendix B to estimate α.

We can now compare BLB, SB, and SDB with some initial specification of their hyperparameters. Toward this, we randomly generate 6 different combinations of (R, B) values from [U(15, 30)] × [U(2500, 5000)] and another six pairs of (R, B) from [U(300, 800)] × [U(50, 80)], as seen in Table 3. Their optimal specifications are obtained via (3.1), (3.2), and (3.5) under the multivariate case with c'₁ = c₁ and c'₂ = c₂ for BLB and c'₁ = c₁ and c'₂ = c₃ for SDB and SB, where c₁, c₂, and c₃ are defined in (2.10), (2.11), and (2.12), respectively. From this table, we can see clearly that the running times of the optimal BLB, SB and SDB are all similar to those of the original setting with prespecified hyperparameters. On the other hand, all three methods with optimally chosen hyperparameters outperform their original settings. Moreover, SB always performs the worst consistent with the results in Kleiner et al. (2014) and Sengupta, Volgushev, and Shao (2016). It is also interesting that the optimal BLB outperforms optimal SDB under the same time cost and the average improve margin is 44.3%.

4.3. Hyperparameter Selection for Logistic Regression

We now consider a logistic regression model by following the notation in Example 3 in Section 3.2. The parameter to be estimated is set as $\beta_0 = (0.5, 0.5)^\top$ and we generate X_i as in the linear regression model in Section 4.2. The covariance of the one-step estimator $\hat{\beta}$ can be consistently estimated by the BLB as

$$\widehat{\text{SE}}_{\text{BLB}}^{*2} = \frac{1}{RB} \sum_{r=1}^R \sum_{b=1}^B (\hat{\beta}^{(r,b)} - \hat{\beta}^{(r)}) (\hat{\beta}^{(r,b)} - \hat{\beta}^{(r)})^\top,$$

where $\hat{\beta}^{(r,b)}$ is the one-step estimator obtained from the b th resample of the r th subsample and $\hat{\beta}^{(r)}$ is the one-step estimate based on the r th resample. The pilot estimator $\tilde{\beta}$ is estimated using $n = \lfloor N^{0.7} \rfloor$ samples only once for each resample. The accuracy of the covariance estimator is evaluated by computing $\text{MSE}(\widehat{\text{SE}}_{\text{BLB}}^{*2}) = \|\widehat{\text{SE}}_{\text{BLB}}^{*2} - \widehat{\text{cov}}(\hat{\beta})\|_F$, where for logistic regression, $\widehat{\text{cov}}(\hat{\beta}) = E\{\sum_{i=1}^N \omega_i(\hat{\beta}) X_i X_i^\top\}^{-1}$ and $\widehat{\text{cov}}(\tilde{\beta}) = M_0^{-1} \sum_{i=1}^{M_0} \{\sum_{i=1}^N \omega_i(\tilde{\beta}) X_i^{(m)} X_i^{(m)\top}\}^{-1}$, where $M_0 = 2000$ and $X_i^{(m)}$ represents that X_i is generated in the m th replication. We can calculate the MSE of the SB and SDB methods likewise.

We then evaluate the relative performance of BLB and SDB under the same computational budget. Here, the original SDB corresponds to the cases with $n = \lfloor N^{0.7} \rfloor$ and R is further calculated based on the time consumption of BLB. We train the optimal hyperparameter the same way with Section 4.2. Toward this, we randomly generate 6 different combinations of (R, B) values from $\lfloor U(15, 30) \rfloor \times \lfloor U(2500, 5000) \rfloor$ and another six pairs of (R, B) from $\lfloor U(150, 300) \rfloor \times \lfloor U(1, 80) \rfloor$, as seen in Table 4. Their optimal specifications are obtained via (3.1), (3.2), and (3.5) under the multivariate case. From this table, we can see clearly that the running times of the optimal BLB and SDB are all similar to those of the original setting with prespecified hyperparameters. On the other hand, both methods with optimally chosen hyperparameters outperform their original settings. Moreover, optimal BLB always performs better than optimal (and also original) SDB and the average improve margin is 44.9%. It is reasonable that optimizing BLB over both parameters is expected to yield better result than just fixing one of the parameters.

4.4. Real Data Analysis

To demonstrate the application of our method in data analysis, we consider a US Airline Dataset publicly available at <http://stat-computing.org>. This dataset contains detailed flight information. For this analysis, we take the data in the year of 2008 with a sample size $N = 492809$, and study how the flight Distance and ArrDelay (arrival delay) affect the ActualElapsedTime (actual elapsed time) via a linear model. We preprocess the two covariates by taking the signed-log-transformation $\log|x| \cdot \text{sign}(x)$ to eliminate the influence of outliers. We evaluate the performance using a procedure similar to Section 4.2, with the critical difference being that the variance of the ordinary least-squares estimator has to be estimated. Toward this, we make use of traditional bootstrap by estimating it as $M_0^{-1} \sum_{m=1}^{M_0} (\hat{\beta}^{(m)} - \hat{\beta})(\hat{\beta}^{(m)} - \hat{\beta})^\top$, where $\hat{\beta}^{(m)}$ is the estimate obtained on the m th sample done by sampling with replacement. The number of bootstrap resamples is taken as $M_0 = 10,000$.

We set $n = \lfloor N^{0.7} \rfloor$, and generate 5 different (R, B) combinations from $\lfloor U(15, 30) \rfloor \times \lfloor U(2500, 5000) \rfloor$ and another 5 pairs of (R, B) from $\lfloor U(250, 500) \rfloor \times \lfloor U(1, 80) \rfloor$, as seen in Table 5. We repeat the experiment 50 times under each setting. Again, we use the MSE defined in Section 4.2 as the performance measure. From Table 5, we can see clearly that the running times of the optimal BLB and SDB are all similar to those under the original settings with prespecified hyperparameters. On the other hand, BLB and SDB with hyperparameters tuned by our approach outperform their original settings. Moreover, the optimal BLB outperforms optimal SDB (and also SB, results not shown) under the same time cost.

5. Conclusion

In this article we propose a hyperparameter selection approach that can be applied to subsampling bootstrap methods. A novelty of the approach is to formulate the problem of finding optimal hyperparameters as an optimization one, for which closed-form solutions can be readily obtained using simple, intuitive, almost back-of-envelope calculations. Our extensive simulation study confirms that our hyperparameter selection approach improves efficiency of an estimator without increasing the computational

Table 4. Simulation results for logistic regression.

Original (R, B)	MSE ratios			CPU time ratios		
	BLB*/BLB	SDB*/SDB	BLB*/SDB*	BLB*/BLB	SDB*/SDB	BLB*/SDB*
(20,3862)	0.868	0.255	0.601	0.997	1.058	0.988
(21,3872)	0.845	0.311	0.457	1.003	1.019	1.040
(22,4114)	0.765	0.286	0.491	0.991	1.013	0.912
(26,3559)	0.752	0.265	0.506	0.999	1.051	0.995
(28,4287)	0.762	0.333	0.475	0.999	1.032	0.954
(29,4006)	0.822	0.315	0.501	0.997	0.968	1.025
(201, 57)	0.873	0.254	0.606	0.992	1.027	0.985
(208, 44)	0.857	0.309	0.543	0.976	1.024	0.975
(222, 18)	0.642	0.258	0.663	0.969	1.011	0.974
(252, 23)	0.682	0.269	0.597	0.985	1.027	0.979
(259, 34)	0.731	0.285	0.656	0.980	1.026	0.976
(297, 56)	0.784	0.317	0.518	0.997	1.057	0.990

NOTE: The explanation of the column headings can be found in the caption of Table 3.

Table 5. Detailed results for the Airline data.

Original (R, B)	MSE ratios			CPU time ratios		
	BLB*/BLB	SDB*/SDB	BLB*/SDB*	BLB*/BLB	SDB*/SDB	BLB*/SDB*
(20,3559)	0.806	0.697	0.509	0.996	0.994	0.967
(21,4287)	0.840	0.723	0.582	0.998	0.993	0.992
(24,3872)	0.881	0.689	0.576	1.005	0.998	1.003
(28,4006)	0.883	0.686	0.630	1.001	1.000	0.988
(29,3862)	0.890	0.736	0.613	1.018	0.953	0.907
(398,74)	0.926	0.611	0.479	0.977	0.979	0.966
(413,45)	0.911	0.670	0.498	0.941	0.945	0.927
(413,15)	0.785	0.635	0.440	0.920	0.933	0.921
(437,17)	0.781	0.696	0.488	0.922	0.923	0.924
(491,25)	0.848	0.681	0.470	0.928	0.922	0.936

NOTE: The explanation of the column headings can be found in the caption of Table 3.

burden. We emphasize, as we have discussed in the Introduction, that bootstrap methods are preferred when analytical formula are unavailable or not easily computable. In view of this, the examples examined in this article should be taken as proof of concept.

This article focuses on those estimators whose computational cost is linear in the number of distinct observations in a dataset. It is straightforward to extend our methodology to estimators whose computational cost is of the order $O(n^\gamma)$ for some $\gamma \geq 1$ based on a dataset of size n . This setup includes all the computationally feasible estimators in statistics that can be computed in polynomial time. To see how the extension can be done, for the SB and SDB, we just need to solve: $\arg \min_{n,R} \{c'_1/R + c'_2/n^2\}$ under the constraint $\alpha n^\gamma R \leq C_{\max}$, where the constraint is on $n^\gamma R$ instead of nR as in the estimators discussed in Section 3. Here, α is either α_{sb} for SB or α_{sdb} for SDB, and c'_1 and c'_2 are estimator specific parameters. This formulation yields the same optimal hyperparameters as in (3.1). Likewise, for the BLB, we replace nRB with (3.3) by $n^\gamma RB$. Derivations similar to (3.3) provide the optimal (B, R) as $R^* = \lfloor C_{\max}/(\alpha_1 n^\gamma B^* + \alpha_2 n) \rfloor$, $B^* = \lfloor (\tilde{c}_1 \alpha_2 / \tilde{c}_2 \alpha_1)^{1/2} n^{1-\gamma/2} \rfloor$, where \tilde{c}_1 and \tilde{c}_2 are estimator specific parameters, and α_1 and α_2 are coefficients determining the computational cost as in (3.3).

We now discuss an extension of our approach to quantile regression. Specifically, we examine linear quantile regression for illustration by considering the model: $Y_i = X_i \beta_\tau + \varepsilon_i$, where β_τ is a d -dimensional coefficient vector, $\varepsilon_i \in \mathbb{R}^1$ is the random noise such that $P(\varepsilon_i \leq 0 | X_i) = \tau$ and $\tau \in (0, 1)$. That is, $X_i \beta_\tau$ is the τ th quantile of Y_i given X_i . Using the quantile regression approaches in Koenker (2005), Wang, Li, and Jiang (2007), and Chen, Liu, and Zhang (2019), we know that $\hat{\beta}_\tau$, the quantile regression estimator, admits the following asymptotic expression:

$$\hat{\beta}_\tau = \beta_\tau + \frac{1}{f(0)} \left(\sum_{i=1}^N X_i X_i^\top \right)^{-1} \left[\sum_{i=1}^N X_i \{ \tau - I(\varepsilon_i \leq 0) \} \right] + o_p(N^{-1/2}). \tag{5.1}$$

Intuitively, if the smaller-order term in (5.1) can be ignored, the quantile regression estimator can be expressed as $g(\mu_1, \mu_2)$, a smooth function of two moment estimators $\mu_1 = \sum_{i=1}^N X_i X_i^\top$ and $\mu_2 = \sum_{i=1}^N X_i \{ \tau - I(\varepsilon_i \leq 0) \}$. Thereafter, our theory in Section 3.1 can be potentially applied, though justifying it is challenging due to the non-smoothness of the loss function in quantile regression. We leave it for a future research topic.

Finally, we find the optimal estimator in terms of its MSE in this article, but it need not be the only choice. Our framework can be extended to other efficiency measures by quantifying the relationship between the choice of efficiency measure and the hyperparameters and then formulating an optimization problem in a manner similar to that used in this article.

Supplementary Materials

Appendix A contains detailed proofs for Theorem 1–8 and technical details for general parameters and statistics. Appendix B contains additional simulation results. Appendix C extends the optimal hyperparameter selection for bootstrap under distributed systems.

Acknowledgments

We are grateful to the Coeditor, Christian Hansen, AE and two referees for their insightful comments and suggestions.

Disclosure Statement

The authors report there are no competing interests to declare.

Funding

Yingying Ma’s research is partially supported by National Natural Science Foundation of China (No.12171020, 11801022). Chenlei Leng’s research is partially supported by EPSRC (EP/X009505/1). Hansheng Wang’s research is partially supported by National Natural Science Foundation of China (12271012, 11831008) and also partially supported by the Open Research Fund of Key Laboratory of Advanced Theory and Application in Statistics and Data Science (KLATASDS-MOE-ECNU-KLATASDS2101).

References

Bickel, P., Götze, F., and van Zwet, W. (1997), “Resampling Fewer than n Observations: Gains, Losses, and Remedies for Losses,” *Statistica Sinica*, 7, 1–31. [174,177]
 Bickel, P. J., and Sakov, A. (2008), “On the Choice of m Out of n Bootstrap and Confidence Bounds for Extrema,” *Statistica Sinica*, 18, 967–985. [174]
 Booth, J. G., and Hall, P. (1994), “Monte Carlo Approximation and the Iterated Bootstrap,” *Biometrika*, 81, 331–340. [174,175]
 Chen, S., and Peng, L. (2021), “Distributed Statistical Inference for Massive Data,” *The Annals of Statistics*, 49, 2851 – 2869. [174]
 Chen, X., Liu, W., and Zhang, Y. (2019), “Quantile Regression under Memory Constraint,” *The Annals of Statistics*, 47, 3244 – 3273. [174,185]

- Efron, B. (1990), "More Efficient Bootstrap Computations," *Journal of the American Statistical Association*, 85, 79–89. [174,175]
- Efron, B., and Tibshirani, R. J. (1994), *An Introduction to the Bootstrap*, Boca Raton, FL: CRC Press. [174,175]
- Fan, J., Guo, Y., and Wang, K. (2021), "Communication-Efficient Accurate Statistical Estimation," *Journal of the American Statistical Association*, DOI: 10.1080/01621459.2021.1969238. [174]
- Hall, P. (1994), "Methodology and Theory for the Bootstrap," *Handbook of Econometrics*, 4, 2341–2381. [174]
- He, X., and Shao, Q.-M. (1996), "A General Bahadur Representation of m -estimators and its Application to Linear Regression with Nonstochastic Designs," *Annals of Statistics*, 24, 2608–2630. [175]
- Jordan, M. I., Lee, J. D., and Yang, Y. (2019), "Communication-Efficient Distributed Statistical Inference," *Journal of the American Statistical Association*, 114, 668–681. [174]
- Kleiner, A., Talwalkar, A., Sarkar, P., and Jordan, M. I. (2014), "A Scalable Bootstrap for Massive Data," *Journal of the Royal Statistical Society, Series B*, 76, 795–816. [174,175,177,180,183]
- Koenker, R. (2005), *Quantile Regression*, Cambridge: Cambridge University Press. [185]
- Li, R., Lin, D. K., and Li, B. (2013), "Statistical Inference in Massive Data Sets," *Applied Stochastic Models in Business and Industry*, 29, 399–409. [174]
- Politis, D., Romano, J., and Wolf, M. (1999), *Subsampling*, New York: Springer. [174]
- Sengupta, S., Volgushev, S., and Shao, X. (2016), "A Subsampled Double Bootstrap for Massive Data," *Journal of the American Statistical Association*, 111, 1222–1232. [175,177,178,180,183]
- Shao, J., and Wang, H. (2002), "Sample Correlation Coefficients based on Survey Data under Regression Imputation," *Journal of the American Statistical Association*, 97, 544–552. [182]
- Van Der Vaart, A. W., and Wellner, J. A. (1996), *Weak Convergence and Empirical Processes*, pp. 16–28, New York: Springer. [174]
- Volgushev, S., Chao, S.-K., and Cheng, G. (2019), "Distributed Inference for Quantile Regression Processes," *The Annals of Statistics*, 47, 1634–1662. [174]
- Wang, H., Li, G., and Jiang, G. (2007), "Robust Regression Shrinkage and Consistent Variable Selection through the lad-lasso," *Journal of Business & Economic Statistics*, 25, 347–355. [185]