Statistica Sinica Preprint No: SS-2021-0257

Title	Subsampling and Jackknifing: A Practically Convenient Solution for Large Data Analysis With Limited Computational Resources
Manuscript ID	SS-2021-0257
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202021.0257
Complete List of Authors	Shuyuan Wu, Xuening Zhu and Hansheng Wang
Corresponding Author	Xuening Zhu
E-mail	xueningzhu@fudan.edu.cn
Notice: Accepted version subject to English editing.	

Subsampling and Jackknifing: A Practically Convenient Solution for Large Data Analysis with Limited Computational Resources

Shuyuan Wu¹, Xuening Zhu², and Hansheng Wang¹

¹*Peking University* and ²*Fudan University*

Abstract: Modern statistical analysis often encounters datasets with large sizes.

For these datasets, conventional estimation methods can hardly be used immediately because practitioners often suffer from limited computational resources.

In most cases, they do not have powerful computational resources (e.g., Hadoop or Spark). How to practically analyze large datasets with limited computational resources then becomes a problem of great importance. To solve this problem, we propose here a novel subsampling-based method with jackknifing. The key idea is to treat the whole sample data as if they were the population. Then, multiple subsamples with greatly reduced sizes are obtained by the method of simple random sampling with replacement. It is remarkable that we do not recommend sampling methods without replacement because this would incur a significant cost for data processing on the hard drive. Such cost does not exist if the data are processed in memory. Because subsampled data have relatively small sizes, they can be comfortably read into computer memory as a whole and then processed easily. Based on subsampled datasets, jackknife-debiased estimators can

be obtained for the target parameter. The resulting estimators are statistically consistent, with an extremely small bias. Finally, the jackknife-debiased estimators from different subsamples are averaged together to form the final estimator. We theoretically show that the final estimator is consistent and asymptotically normal. Its asymptotic statistical efficiency can be as good as that of the whole sample estimator under very mild conditions. The proposed method is simple enough to be easily implemented on most practical computer systems and thus should have very wide applicability.

Key words and phrases: GPU, Jackknife, Large Dataset, Subsampling.

1. INTRODUCTION

Modern statistical analysis often encounters datasets with large sizes. Meanwhile, most researchers possess very limited computational resources. In most cases, they do not have a powerful computation system, such as a distributed computation system like Hadoop or Spark. As a consequence, they must rely on their handy computational resources (e.g., a personal computer) for large data analysis. Thus, how to practically analyze large datasets with limited computational resources becomes a problem of great importance.

To solve this problem, various subsampling methods have been proposed ([Mahoney, 2011](#); [Drineas et al., 2011](#); [Ma et al., 2015](#); [Wang et al.,](#)

2018; Wang, 2019; Yu et al., 2020; Ma et al., 2020). The key idea of most existing methods is to design a novel sampling strategy so that excellent statistical efficiency can be achieved with small sample sizes. For example, Ma et al. (2015) developed a novel method to select the optimal subsample according to leverage scores. Wang et al. (2018) studied a similar problem and proposed an A -optimality criterion. Yu et al. (2020) developed an optimal Poisson subsampling approach. Ma et al. (2020) derived the asymptotic distribution of the sampling estimator based on the linear regression. Despite their usefulness, these pioneer methods suffer from two limitations. First, specific sampling strategies must be carefully designed for different analysis purposes. Second, they are not computationally inexpensive. A significant computation cost is required for practical implementation. In most cases, the sampling cost should be at least $O(N)$, where N represents the whole sample size.

To overcome these challenges, here, we aim to develop a novel method with the following unique features. First, our method is simple enough to be easily implemented on most practical computer systems. We argue that the simplicity is particularly relevant and important because simplicity implies wider applicability. Second, due to jackknifing, our estimators lead to significant bias reduction compared with other methods. As a result, the

same asymptotic efficiency can be achieved with a much reduced subsample size as long as the number of subsamples is large enough. Moreover, our method supports fully automatic and unified inference. For most real applications, valid statistical inferences (e.g. confidence interval) are inevitably needed. However, the analytical formula for the asymptotic distribution of the estimator could be too complicated to be analytically derived. Our proposal is automatic in the sense that the standard errors of various statistics can be automatically computed without referring to the analytical formula of their asymptotic distributions. In addition, our proposal is unified in the sense that it can be readily applied to many different statistics.

Specifically, we develop here a subsampling method with jackknifing. To implement our method, multiple subsamples are obtained by simple random sampling with replacement. For each subsampled dataset, a jackknife-debiased estimator is computed for the parameter of interest. Subsequently, these jackknife-debiased estimators are further averaged. This leads to the final estimator. We show theoretically that the resulting estimator is consistent and asymptotically normal. Its statistical efficiency can be asymptotically as good as the whole sample estimator under very mild conditions. This useful property remains valid even if the subsample size is very small. The desirable property is mainly attributed to jackknifing. As a byproduct,

a jackknife estimator for the standard error of the proposed estimator can be obtained. This enables automatic statistical inference. For practical implementation, a GPU-based algorithm is developed. Empirical experiments suggest that it is extremely computationally efficient. Extensive numerical studies are presented to demonstrate the finite sample performance.

Despite the usefulness, the proposed method also suffers from obvious limitations. The main limitation is that it is computationally less efficient as compared to the one-pass-full-sample-mean estimator computed by the distributed approaches (Suresh et al., 2017). However, our proposal carries its unique value because it could be a practically more convenient alternative under the following two important situations. The first situation is that the whole sample size N is extremely large. In this case, a significant amount of clock-time cost has to be paid for the whole sample computation (e.g., computing the one-pass-full-sample-mean). This is particularly true if no powerful distributed computation system is available. However, for most practical data analysis, the practical demand for estimation precision is limited. On the contrary, the budget for time spending as measured by clock-time cost is extremely valuable. Then, it might be more appealing to sacrifice the statistical efficiency to some extent to trade for less clock-time cost. Accordingly, we do NOT expect our method to be implemented with

a very large subsample size n and a very large number of subsamples K . Instead, they should be implemented with reasonably large n and K , as long as the desired statistical precision can be achieved.

The second situation is that automatic statistical inferences are required as we previously mentioned. In this case, if the one-pass-full-sample-mean is used, then the analytical formula for the asymptotic distribution of the estimator has to be manually derived. It is then preferable to have an automatic and unified solution for statistical inference. This is another case where our method could be a practically more convenient solution.

The rest of the paper is organized as follows. Section 2 develops the proposed estimators and their asymptotic properties. The numerical studies are presented in Section 3, including the GPU-based algorithm, simulation experiments, and real dataset analysis. Finally, the article concludes with a brief discussion in Section 4. All technical details are delegated to the supplement.

2. THE METHODOLOGY

2.1 Model and Notations

Let X_i be an independent random variable observed from the i th subject, where $1 \leq i \leq N$ and N is the whole sample size. Let $\mathbb{S} = \{1, \dots, N\}$

2.1 Model and Notations

be the index set of the whole sample. Let μ be one particular moment about X_i . For simplicity, we can assume $\mu = E(X_i)$ to be a scalar. The theory to be presented hereafter can be easily extended to a more general situation with multivariate moments and M estimators. Let $\theta = g(\mu)$ be the parameter of interest, where $g(\cdot)$ is a known nonlinear function. We assume that $g(\cdot)$ is sufficiently smooth. To estimate θ , one can use a sample moment estimator $\hat{\theta} = g(\hat{\mu})$, where $\hat{\mu} = N^{-1} \sum_{i \in \mathbb{S}} X_i$.

For convenience, we refer to $\hat{\theta}$ as the whole sample (WS) estimator to emphasize the fact that this is an estimator computed based on the whole sample. The merit of the WS estimator $\hat{\theta}$ is that it offers excellent statistical efficiency. However, it could be difficult to compute if the whole sample size N is too large. This is particularly true if researchers are given very limited computational resources. Accordingly, one must consider other estimation methods that are more computationally feasible. In this regard, here, we study one particular type of subsampling method ([Mahoney, 2011](#); [Drineas et al., 2011](#); [Ma et al., 2015](#); [Wang et al., 2018](#); [Wang, 2019](#); [Yu et al., 2020](#); [Ma et al., 2020](#)) as an excellent and practical solution.

Let n be the subsample size, which is typically much smaller than N . Let K be the number of subsamples. Write $\mathcal{S}_k = \{i_1^{(k)}, \dots, i_n^{(k)}\} \subset \mathbb{S}$ as the k th subsample set, where $i_m^{(k)}$ s (for any $1 \leq m \leq n$, $1 \leq k \leq K$)

2.2 Variance and Bias Analysis of the SOS Estimator

are generated independently from $\mathbb{S}$ by the method of simple random sampling with replacement. In other words, conditional on $\mathbb{S}$, $i_m^{(k)}$ s are independently and identically distributed with probability $P(i_m^{(k)} = j) = N^{-1}$ for any $j \in \mathbb{S}$. Accordingly, a moment estimator based on $\mathcal{S}_k$ can be computed as $\hat{\theta}^{(k)} = g(\hat{\mu}^{(k)})$, where $\hat{\mu}^{(k)} = n^{-1} \sum_{i \in \mathcal{S}_k} X_i$. One can then combine these subsample estimators together to form a more accurate one as $\hat{\theta}_{\text{SOS}} = K^{-1} \sum_{k=1}^K \hat{\theta}^{(k)}$. This is then referred to as a subsample one-shot (SOS) estimator. It is similar to the so-called one-shot estimator developed for distributed systems ([Mcdonald et al., 2009](#); [Zinkevich et al., 2011](#); [Zhang et al., 2013](#)). However, the key difference is that the subsamples used by a standard one-shot estimator should not have any overlap with each other. In contrast, the subsamples used by our proposed subsampling methods are allowed to be partially overlapped.

2.2 Variance and Bias Analysis of the SOS Estimator

To motivate our method, we offer an informal analysis of the bias and variance of the SOS estimator $\hat{\theta}_{\text{SOS}}$. The formal theoretical results are provided in Section 2.4. Specifically, by Taylor's expansion, we can approx-

2.2 Variance and Bias Analysis of the SOS Estimator

imate $\hat{\theta}^{(k)}$ as

$$\hat{\theta}^{(k)} \approx \theta + \dot{g}(\mu) (\hat{\mu}^{(k)} - \mu) + \frac{1}{2} \ddot{g}(\mu) (\hat{\mu}^{(k)} - \mu)^2,$$

where $\dot{g}(\mu)$ and $\ddot{g}(\mu)$ are the first- and second-order derivatives of $g(\mu)$ with respect to μ , respectively. Accordingly, we have

$$\hat{\theta}_{sos} = \frac{1}{K} \sum_{k=1}^K \hat{\theta}^{(k)} \approx \theta + \frac{\dot{g}(\mu)}{K} \sum_{k=1}^K (\hat{\mu}^{(k)} - \mu) + \frac{\ddot{g}(\mu)}{2K} \sum_{k=1}^K (\hat{\mu}^{(k)} - \mu)^2. \quad (2.1)$$

By equation (2.1) we know that $\text{var}(\hat{\theta}_{sos})$ can be approximated by the variance of $\dot{g}(\mu) K^{-1} \sum_{k=1}^K (\hat{\mu}^{(k)} - \mu)$. Let $\hat{\sigma}^2 = N^{-1} \sum_{i \in \mathbb{S}} (X_i - \hat{\mu})^2$. With a slight abuse of notation, we use $\mathbb{S}$ to represent the information contained in the whole sample, that is, the σ -field generated by $\{X_1, \dots, X_N\}$. Recall that, conditional on $\mathbb{S}$, X_i s are independent and identically distributed for any $i \in \mathcal{S}_k$ and $1 \leq k \leq K$. We then have $E(\hat{\mu}^{(k)} | \mathbb{S}) = n^{-1} E(\sum_{i \in \mathcal{S}_k} X_i | \mathbb{S}) = \hat{\mu}$ and $\text{var}(\hat{\mu}^{(k)} | \mathbb{S}) = n^{-2} \text{var}(\sum_{i \in \mathcal{S}_k} X_i | \mathbb{S}) = n^{-1} \hat{\sigma}^2$. Assume that the second moment of X_1 is finite with $\sigma^2 = \text{var}(X_1)$. We then have

$$\begin{aligned} E \left\{ \text{var}(\hat{\theta}_{sos} | \mathbb{S}) \right\} &\approx \frac{\dot{g}(\mu)^2}{K} E \left\{ \text{var}(\hat{\mu}^{(k)} | \mathbb{S}) \right\} \approx \frac{\dot{g}(\mu)^2}{nK} \sigma^2, \\ \text{var} \left\{ E(\hat{\theta}_{sos} | \mathbb{S}) \right\} &\approx \dot{g}(\mu)^2 \text{var}(\hat{\mu} - \mu) = \frac{\dot{g}(\mu)^2}{N} \sigma^2. \end{aligned} \quad (2.2)$$

2.2 Variance and Bias Analysis of the SOS Estimator

By equation (2.2), we find that $\text{var}(\hat{\theta}_{\text{SOS}})$ can be approximated by $\tau_1\{1/(nK)+1/N\}$, with $\tau_1 = \dot{g}(\mu)^2\sigma^2$. Under the condition $nK \gg N$, we then determine that the variance of the subsample estimator can be further approximated by τ_1/N , which is the asymptotic variance of the WS estimator $\hat{\theta}$.

Next, we study the bias of $\hat{\theta}_{\text{SOS}}$. We define the bias of T_n as $\text{Bias}(T_n) = E(T_n) - \theta$ for any estimator T_n of θ . Then, by equation (2.1), we have

$$\text{Bias}(\hat{\theta}_{\text{SOS}}) = E\left(\frac{1}{K} \sum_{k=1}^K \hat{\theta}^{(k)}\right) - \theta = E(\hat{\theta}^{(k)}) - \theta \approx \frac{\ddot{g}(\mu)}{2n}\sigma^2 + \frac{\ddot{g}(\mu)}{2N}\sigma^2. \quad (2.3)$$

The leading term of $\text{Bias}(\hat{\theta}_{\text{SOS}})$ is given by τ_2/n , with $\tau_2 = \ddot{g}(\mu)\sigma^2/2$. Unfortunately, it does not improve as K increases. This indicates that the bias of $\hat{\theta}_{\text{SOS}}$ is of an order $O(n^{-1})$. This is a smaller order term as compared with $\hat{\theta} - \theta = O_p(1/\sqrt{N})$, as long as $n \gg \sqrt{N}$. This condition seems to be quite reasonable for a distributed system (Huang and Huo, 2015; Jordan et al., 2019). In that case, K is the number of distributed computers. As a consequence, K is typically much smaller than n , where n is the subsample size allocated to each distributed computer. However, this condition could be problematic for a subsampling method. In this case, K is the total number of subsamples and could be very large. In contrast, for computation

2.3 The Jackknife Estimators

convenience, the subsample size n could be much smaller than $\sqrt{N}$. This makes the bias introduced in equation (2.3) possibly non-negligible. To fix this problem, we are motivated to search for an improved estimator for θ so that its bias can be greatly reduced. In this regard, jackknife is a well known method to reduce the bias of estimators (Quenouille, 1949; Efron and Stein, 1981; Cameron and Trivedi, 2005). However, the performance of Jackknife in subsampling scenario is not clear. This leads to the novel jackknife estimators presented in the next subsection.

2.3 The Jackknife Estimators

The objective of this subsection is two-fold. The first goal is to develop a jackknife debiased subsample (JDS) estimator for θ . The second goal is to propose a jackknife standard error (JSE) estimator for the JDS estimator.

First, we develop the JDS estimator to reduce the estimation bias. To this end, we define a jackknife estimator $\hat{\theta}_{-j}^{(k)}$ for the k th subsample as follows:

$$\hat{\theta}_{-j}^{(k)} = g\left(\hat{\mu}_{-j}^{(k)}\right), \text{ where } \hat{\mu}_{-j}^{(k)} = \frac{1}{n-1} \sum_{i \in \mathcal{S}_k, i \neq j} X_i.$$

By similar analysis to that for equation (2.1), we know that $\text{Bias}(\hat{\theta}_{-j}^{(k)})$ approximately equals $\tau_2/(n-1)$. Then, $n^{-1} \sum_{j \in \mathcal{S}_k} \text{Bias}(\hat{\theta}_{-j}^{(k)}) \approx \tau_2/(n-1)$, and $E(n^{-1} \sum_{j \in \mathcal{S}_k} \hat{\theta}_{-j}^{(k)} - \hat{\theta}^{(k)}) \approx \tau_2/\{n(n-1)\}$. This inspires an estimator

2.3 The Jackknife Estimators

for the bias, which is given by $\widehat{\text{Bias}}^{(k)} = (n-1)n^{-1} \sum_{j \in S_k} \hat{\theta}_{-j}^{(k)} - (n-1)\hat{\theta}^{(k)}$. Accordingly, we can propose a bias-corrected estimator for the k th subsample as $\hat{\theta}_{JDS}^{(k)} = \hat{\theta}^{(k)} - \widehat{\text{Bias}}^{(k)}$. Thereafter, $\hat{\theta}_{JDS}^{(k)}$ s can be further averaged across different k . As a consequence, we obtain the final JDS estimator $\hat{\theta}_{JDS} = K^{-1} \sum_{k=1}^K \hat{\theta}_{JDS}^{(k)}$. Subsequently, we rigorously verify that $\text{Bias}(\hat{\theta}_{JDS})$ is much smaller than that of $\hat{\theta}_{sos}$. Specifically, $\text{Bias}(\hat{\theta}_{JDS}) = O(1/n^2) + O(1/N)$ and $\text{Bias}(\hat{\theta}_{sos}) \approx \tau_2/n$; see equation (2.3). Furthermore, we can theoretically prove that the asymptotic variance of $\hat{\theta}_{JDS}$ remains the same as that of the WS estimator. As a consequence, assuming that K is large enough, excellent statistical efficiency can be achieved by $\hat{\theta}_{JDS}$ with a very small subsample size n .

Other than bias correction, the jackknife method can also serve as an excellent estimator for the standard error, that is, the standard deviation of the JDS estimator $\hat{\theta}_{JDS}$. The basic idea is given as follows. Recall that by equation (2.2), we know that

$$\text{var}(\hat{\theta}_{sos}) \approx \dot{g}(\mu)^2 \sigma^2 \left(\frac{1}{nK} + \frac{1}{N} \right). \quad (2.4)$$

Because N, n and K are all known to the user, the key objective here is to estimate the unknown parameter $\tau_1 = \dot{g}(\mu)^2 \sigma^2$. Moreover, by the definition

2.4 Theoretical Properties

of the jackknife estimator and Taylor's expansion, we have

$$\widehat{\theta}_{-j}^{(k)} - \widehat{\theta}^{(k)} \approx \dot{g}(\mu) \left(\widehat{\mu}_{-j}^{(k)} - \widehat{\mu}^{(k)} \right) = \frac{\dot{g}(\mu)}{n-1} (\widehat{\mu}^{(k)} - X_j)$$

for any $j \in \mathcal{S}_k$ and $1 \leq k \leq K$. We know immediately that

$$E \left\{ \left(\widehat{\theta}_{-j}^{(k)} - \widehat{\theta}^{(k)} \right)^2 \right\} = E \left[E \left\{ \left(\widehat{\theta}_{-j}^{(k)} - \widehat{\theta}^{(k)} \right)^2 \mid \mathbb{S} \right\} \right] \approx \frac{\dot{g}(\mu)^2 \sigma^2}{n(n-1)} = \frac{\tau_1}{n(n-1)},$$

which is closely related to the unknown parameter τ_1 in equation (2.4).

Note that the sample mean of $\left(\widehat{\theta}_{-j}^{(k)} - \widehat{\theta}^{(k)} \right)^2$ across different j and k is a reasonable approximation of $E \left\{ \left(\widehat{\theta}_{-j}^{(k)} - \widehat{\theta}^{(k)} \right)^2 \right\}$. This inspires the following JSE estimator $\widehat{\text{SE}}$:

$$\widehat{\text{SE}}^2 = \left(\frac{1}{K} + \frac{n}{N} \right) \frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{S}_k} \left(\widehat{\theta}_{-j}^{(k)} - \widehat{\theta}^{(k)} \right)^2.$$

We will theoretically prove that $\widehat{\text{SE}}^2$ is a consistent estimator of $\text{var}(\widehat{\theta}_{JDS})$.

In addition, $\text{var}(\widehat{\theta}_{JDS}) / \text{var}(\widehat{\theta}_{SOS}) = 1 + o(1)$. Consequently, $\widehat{\text{SE}}^2$ is also a consistent estimator of $\text{var}(\widehat{\theta}_{SOS})$.

2.4 Theoretical Properties

In this subsection, we study the theoretical properties of the three esti-

2.4 Theoretical Properties

mators (i.e., the SOS, JDS and JSE estimators). To this end, the following standard technical conditions are needed.

(C1) (SUB-GAUSSIAN DISTRIBUTION) Assume X_i follow a sub-Gaussian distribution, i.e., there exists positive constants C, ν such that $P(|X_i| > t) \leq C \exp\{-\nu t^2\}$ for every $t > 0$.

(C2) (SMOOTHNESS CONDITION) Define $g^{(k)}(\cdot)$ as the k th order derivative function of $g(\cdot)$ and assume $g^{(k)}(\cdot)$ is a continuous function for $k \leq 8$.

(C3) (SUBSAMPLING CONDITION) As $N \rightarrow \infty$, the subsample size $n \rightarrow \infty$.

In addition, assume that $n < N$, $N = o(n^4)$ and $\log K = o(\sqrt{n})$.

Condition (C1) is a classical and flexible assumption on covariates (Jordan et al., 2019; Zhu et al., 2021). Condition (C2) requires the g -function to be sufficiently smooth so that a Taylor's expansion can be obtained around μ . We require slightly stronger condition since we will derive the asymptotic bias in more explicit forms. The condition can be relaxed to requiring $g(\cdot)$ to be fourth continuously differentiable function to guarantee the asymptotic normality (Wu, 1986; Lehmann and Casella, 2006). Lastly, Condition (C3) states the relationships between n , N and K . It requires that the subsample size should be large enough to facilitate the asymptotic analysis of higher

2.4 Theoretical Properties

order terms. In addition, we require $\log K = o(\sqrt{n})$ to guarantee a uniform convergence for all subsamples, which is easy to satisfy in practice.

We next consider how to understand the asymptotic behavior of various subsample estimators without a finite moment constraints. Inspired by the asymptotic theory of Shao (2003), we adopt here a Taylor's expansion approach. Specifically, take $\hat{\theta}_{sos}$ as an example, by the Taylor's expansion, we have $\hat{\theta}_{sos} = K^{-1} \sum_{k=1}^K \hat{\theta}^{(k)} = K^{-1} \sum_{k=1}^K g(\hat{\mu}^{(k)}) = \theta + \hat{\Delta}_{sos}^{(1)} + \hat{\Delta}_{sos}^{(2)} + \mathcal{O}$, where $\hat{\Delta}_{sos}^{(1)} = \dot{g}(\mu) K^{-1} \sum_{k=1}^K (\hat{\mu}^{(k)} - \mu)$, $\hat{\Delta}_{sos}^{(2)} = K^{-1} \sum_{k=1}^K \{\ddot{g}(\mu) (\hat{\mu}^{(k)} - \mu)^2/2 + g^{(3)}(\mu) (\hat{\mu}^{(k)} - \mu)^3/6\}$, and $\mathcal{O}$ stands for some higher order terms. As we have discussed informally in Section 2.2, it suggests that the asymptotic behavior of $\hat{\theta}_{sos} - \theta$ could be fully determined by $\hat{\Delta}_{sos}^{(1)}$ and $\hat{\Delta}_{sos}^{(2)}$. Here $\hat{\Delta}_{sos}^{(1)}$ is unbiased and mainly contributes the variance, while $\hat{\Delta}_{sos}^{(2)}$ has ignorable variance and mainly controls the bias. We can accordingly understand the asymptotic performance of $\hat{\theta}_{sos}$'s bias by $E(\hat{\Delta}_{sos}^{(2)})$ and $\hat{\theta}_{sos}$'s variance by $\text{var}(\hat{\Delta}_{sos}^{(1)})$. Specifically, we have the following theorem.

Theorem 1. Assume conditions (C1)-(C3) hold, then we have $\hat{\theta}_{sos} - \theta = \hat{\Delta}_{sos}^{(1)} + \hat{\Delta}_{sos}^{(2)} + \mathcal{O}$ with $E(\hat{\Delta}_{sos}^{(1)}) = 0$, $\text{var}(\hat{\Delta}_{sos}^{(2)}) = o\{1/(nK) + 1/N\}$, $\mathcal{O} = o_p(1/n + \sqrt{1/(nK) + 1/N})$, and

2.4 Theoretical Properties

$$E(\widehat{\Delta}_{SOS}^{(2)}) = \tau_2 \left(\frac{1}{n} + \frac{1}{N} \right) + o\left(\frac{1}{n}\right) \quad (2.5)$$

$$\text{var}(\widehat{\Delta}_{SOS}^{(1)}) = \tau_1 \left(\frac{1}{nK} + \frac{1}{N} \right) + o\left(\frac{1}{nK} + \frac{1}{N}\right). \quad (2.6)$$

By Theorem 1, first we find that the higher order terms $\mathcal{O}$ could be ignorable compared with $\widehat{\Delta}_{SOS}^{(1)}$ and $\widehat{\Delta}_{SOS}^{(2)}$. In addition, the asymptotic bias behavior of $\widehat{\theta}_{SOS}$ is decided by $\widehat{\Delta}_{SOS}^{(2)}$, while the asymptotic variance behavior of $\widehat{\theta}_{SOS}$ is determined by $\widehat{\Delta}_{SOS}^{(1)}$. Then by equation (2.5), we know that the bias of $\widehat{\Delta}_{SOS}^{(2)}$ is affected by both N and n . The $1/n$ and $1/N$ terms represent the asymptotic bias due to the subsampling and overall sampling errors, respectively. The leading term of variance for $\widehat{\Delta}_{SOS}^{(1)}$ also includes two quantities. They are the $1/(nK)$ and $1/N$ terms. The first term is due to the subsampling error, and the second term is due to the overall sampling error. Recall that the asymptotic variance of the WS estimator $\widehat{\theta}$ approximately equals τ_1/N . Then, for the SOS estimator to achieve the same asymptotic efficiency as $\widehat{\theta}$, we must have $nK/N \rightarrow \infty$. Unfortunately, the subsampling error term of $\text{Bias}(\widehat{\Delta}_{SOS}^{(2)})$ is $O(1/n)$, which does not reduce at all as $K \rightarrow \infty$. Consequently, we need to have $n \gg \sqrt{N}$ so that the asymptotic bias is of $o(1/\sqrt{N})$. Otherwise, the SOS estimator can never be asymptotically as

2.4 Theoretical Properties

efficient as the WS estimator $\hat{\theta}$. Similar with $\hat{\theta}_{SOS}$, we could express $\hat{\theta}_{JDS}$ by the Taylor's expansion as $\hat{\theta}_{JDS} = \hat{\Delta}_{JDS}^{(1)} + \hat{\Delta}_{JDS}^{(2)} + \mathcal{O}$, the detailed expression is given in Appendix B. Define $\tau_3 = g^{(3)}(\mu)\mu_3/6$, $\tau_4 = g^{(4)}(\mu)\sigma^4/8$, and $\mu_3 = E(X_i - \mu)^3$. We next analyze the properties of the JDS estimator in the following theorem.

Theorem 2. *Assume conditions (C1)-(C3) hold, then we have $\hat{\theta}_{JDS} - \theta = \hat{\Delta}_{JDS}^{(1)} + \hat{\Delta}_{JDS}^{(2)} + \mathcal{O}$ with $E(\hat{\Delta}_{JDS}^{(1)}) = 0$, $\text{var}(\hat{\Delta}_{JDS}^{(2)}) = o\{1/(nK) + 1/N\}$, $\mathcal{O} = o_p(1/n^2 + 1/N + \sqrt{1/(nK) + 1/N})$ and*

$$E(\hat{\Delta}_{JDS}^{(2)}) = \frac{\tau_2}{N} + \frac{\tau_3 + \tau_4}{n^2} + o\left(\frac{1}{N} + \frac{1}{n^2}\right) \quad (2.7)$$

$$\text{var}(\hat{\Delta}_{JDS}^{(1)}) = \tau_1 \left(\frac{1}{nK} + \frac{1}{N}\right) + o\left(\frac{1}{nK} + \frac{1}{N}\right). \quad (2.8)$$

Comparing (2.5) and (2.7), we find that for the JDS estimator, the bias term due to the subsampling error is substantially reduced. It is only of the order $O(1/n^2)$. In contrast, that of the SOS estimator is much larger and is of the order $O(1/n)$. Comparing (2.6) and (2.8), we conclude that the leading terms for the variance of both estimators are identical. They can be consistently estimated by the proposed JSE estimator $\widehat{\text{SE}}$. Its asymptotic property is given as follows.

Theorem 3. *Define $\tau^2 = \tau_1 \{1/(nK) + 1/N\}$, and further assume condi-*

2.4 Theoretical Properties

tions (C1)–(C3) hold. The JSE estimator is then ratio consistent for τ , that is, $\widehat{\text{SE}}^2/\tau^2 \rightarrow_p 1$, where “ $\rightarrow_p$ ” stands for “convergence in probability”.

Lastly, for valid asymptotic inference, we need to study the asymptotic distributions of the JDS estimator $\hat{\theta}_{JDS}$ and the SOS estimator $\hat{\theta}_{SOS}$. Consequently, we develop the following theorem to establish the asymptotic normality for both $\hat{\theta}_{JDS}$ and $\hat{\theta}_{SOS}$.

Theorem 4. *Assume conditions (C1)–(C3) hold. The JDS estimator $\hat{\theta}_{JDS}$ is then asymptotically normal with $(\hat{\theta}_{JDS} - \theta)/\tau \rightarrow_d N(0, 1)$, where “ $\rightarrow_d$ ” represents “convergence in distribution”. If one can impose the stronger condition that $n/N^{1/2} \rightarrow \infty$, then the SOS estimator $\hat{\theta}_{SOS}$ is also asymptotically normal with $(\hat{\theta}_{SOS} - \theta)/\tau \rightarrow_d N(0, 1)$.*

From Theorem 4, we know that both the SOS and JDS estimators are asymptotically normal. However, the technical conditions required by both estimators are different. The JDS estimator requires $n/N^{1/4} \rightarrow \infty$. This is a condition that can be very easily satisfied. However, for the SOS estimator, a much stronger condition (i.e., $n/N^{1/2} \rightarrow \infty$) is required (Huang and Huo, 2015; Jordan et al., 2019; Wang et al., 2020).

3. NUMERICAL ANALYSIS

3.1 Why Sampling with Replacement

We aim to develop a GPU-based algorithm for the proposed method with data being placed on the hard drive. Thus, it is important to understand the sampling mechanism on the hard drive. In particular, we want to carefully elaborate computational efficiency between different sampling mechanisms (i.e., simple random sampling with replacement and simple random sampling without replacement on the hard drive) in the following steps.

- (1) First, we assume that there are a total of N data points (representing a massive dataset) placed on the hard drive. They are displayed in the top left of Figure 6. It contains a total of 2 columns. The first column is the sample ID ($ID = 1, 2, 3, 4, 5, \dots, N$) and the second column is the interested variable $Y = (Y_1, \dots, Y_N)$, which represents the interested information.
- (2) Second, to conduct random sampling, we can randomly generate an integer between 1 and N . This determines which data line should be sampled. Without loss of generality, assume that the sampled unit is i^* . Then we read Y_{i^*} into memory. (We should note that this

3.1 Why Sampling with Replacement

sampling procedure is a simplified version. In practice, we cannot access a data line by its sample ID on the hard drive. Instead, we refer to it according to its physical address on the hard drive. This is also not a very straightforward operation and fairly sophisticated).

We then update the index set $\mathcal{S}_0$ from $\mathcal{S}_0 = \{\emptyset\}$ to $\mathcal{S}_0 = \{i^*\}$.

- (3) Third, we should explain how to conduct random sampling without replacement. To this end, we randomly and independently generate another integer i_2^* from 1 to N . It is possible that i_2^* is an already sampled unit in $\mathcal{S}_0$, which leads to duplicated sampling. To avoid duplicated sampling, i_2^* needs to be compared with every already sampled unit in $\mathcal{S}_0$. If we find $i_2^* \in \mathcal{S}_0$ already, then i_2^* needs to be re-generated. Otherwise, $\mathcal{S}_0$ can be updated to be $\mathcal{S}_0 := \mathcal{S}_0 \cup \{i_2^*\}$ and $Y_{i_2^*}$ is read into memory.

- (4) Assume a total of K subsamples with size n need to be generated. Then, the size of the index set is about $|\mathcal{S}_0| = O(nK)$. To avoid duplicated sampling, every sampled unit needs to be compared with every unit in $\mathcal{S}_0$. This leads to a computation cost of order $O(nK)$ for every sampled unit on average. The total computation cost should be of order $O\{(nK)^2\}$ on average. This is an expensive cost. The whole

3.1 Why Sampling with Replacement

process is graphically illustrated in Figure 6.

- (5) Lastly, if we conduct random sampling with replacement, we **avoid** the need to: (a) keep updating $\mathcal{S}_0$ and compare whether $i_2^* \in \mathcal{S}_0$; or (b) keep updating $\mathcal{S}_1$. This makes our proposal computationally more efficient.

To summarize, compared with subsampling with replacement, subsampling without replacement with massive datasets is practically challenging. Therefore, for massive datasets on a hard drive we prefer sampling methods with replacement. In this case, no recording and comparison operations need to be conducted. Next, to further demonstrate this point, we develop an experiment to compare sampling with and without replacement on the hard drive. To this end, we generate independent and identical $X_i = (X_{i_1}, X_{i_2})$ from a standard bivariate normal distribution with $N = 10^9$. The interested parameter is the population mean μ . To estimate μ , the sample mean is calculated based on the two sampling strategies. We denote $\hat{\mu}_{rep}$ and $\hat{\mu}_{worep}$ to represent the estimator based on sampling with and without replacement, respectively. We repeat the experiment $R = 100$ times. Then, the average mean square error (MSE) and time cost (TC) for sampling are reported for both sampling strategies across R replications. All the results are summarized in Table 1.

3.2 An Algorithm for a GPU

Table 1: Comparison of the sampling with and without replacement on the hard drive based on $R = 100$ simulation replications for various (n, K) combinations.

n	K	TC		MSE ($\times 10^{-4}$)	
		$\hat{\mu}_{rep}$	$\hat{\mu}_{worep}$	$\hat{\mu}_{rep}$	$\hat{\mu}_{worep}$
100	50	0.18	0.49	4.04	4.04
	100	0.36	1.55	2.03	2.31
	200	0.73	5.41	0.87	0.90
500	50	0.43	7.64	0.81	0.79
	100	0.87	28.80	0.38	0.36
	200	1.73	110.97	0.18	0.19

From Table 1, we draw the following conclusions. First, the MSE values of the two sampling strategies are comparable, and they both decrease with increasing n or K . However, the TC values of the two strategies are quite different. Sampling with replacement method is much faster than the sampling without replacement method. As nK increases, the gap between the two strategies increases significantly. For instance, if $n = 500$ and $K = 200$, it takes only 1.73 seconds for the sampling with replacement method to complete the procedure, while the time required by the sampling without replacement method is almost 111 seconds.

3.2 An Algorithm for a GPU

We next develop a GPU-based algorithm for fast computation. Note

3.2 An Algorithm for a GPU

that the proposed method exhibits many theoretically and practically useful properties. Theoretically, it guarantees the statistical efficiency of subsample estimators with small subsample sizes. Practically, it is simple, automatic, and flexible. However, the associated computation cost is expensive because the new method requires not only subsampling K times but also jackknifing n times for each subsample. Consequently, it is computationally expensive. Accordingly, its implementation on a Central Processing Unit (CPU) might be inefficient because a standard CPU usually has a very limited number of computation cores. To ameliorate this issue, consider for example the MacBook Pro (13-inch, 2020). It uses the Intel Core i5 processor with only four cores. In contrast, a standard Graphical Processing Unit (GPU) may hold tens of thousands of cores. Accordingly, the GPU is an extremely powerful tool for parallel computation ([Krüger and Westermann, 2005](#); [Che et al., 2008](#)). Meanwhile, our method (particularly the jackknifing part) is extremely suitable for parallel computation. This inspires us to develop a GPU-based algorithm for the proposed method.

A standard GPU system should have two unique features. To make full use of its computational power, we need to take both features into consideration. The first unique feature of the GPU system is that it suffers from two types of communication cost; see Figure 1. The first type of commu-

3.2 An Algorithm for a GPU

Figure 1: Two types of communication cost for a GPU system.

communication cost refers to the time cost required for transferring data from the hard disk (HD) to the CPU memory (CM). This is a standard communication cost that is essentially required by any computation system. For our algorithm, this type of cost is primarily due to subsampling. The second type of communication cost refers to the time cost required for transferring data from the CM to the GPU memory (GM). The main purpose of transferring data from the CM to GM is to prepare data for parallel execution of jackknifing. Consequently, we consider that this part of the communication cost is mainly due to jackknifing. Note that the current GPU architecture does not allow the GPU to directly read the data from the HD. As a consequence, a good algorithm should simultaneously minimize both types of communication cost. Multiple communication between the HD, CM and GM should be avoided.

The second unique feature is that GPU systems are extremely suitable for tensor-type parallel computation. Through this type of computation,

3.2 An Algorithm for a GPU

Figure 2: A graphical illustration of the proposed GPU algorithm.

the parallel computation power of a GPU system can be fully utilized. This suggests that the jackknifing computation should be formulated into a tensor-type computation problem. Specifically, here, we develop a three-step algorithm to implement the proposed method. The process is shown in Figure 2. First, we obtain the k th subsample $\mathcal{S}_k \subset \mathbb{S}$ from the HD and place it in the CM. With slight abuse of notation, we assume that in this subsection, X_i is a p -dimensional vector for any $i \in \mathbb{S}$. Next, we can formulate the k th subsample into an $n \times p$ matrix format as $\mathbf{X}_k = [X_i, i \in \mathcal{S}_k] \in \mathbb{R}^{n \times p}$. We then pass $\mathbf{X}_k$ to the GM and replicate $\mathbf{X}_k$ n times so that a 3-dimensional tensor $\mathbb{X}_k = [\mathbf{X}_k, \dots, \mathbf{X}_k] \in \mathbb{R}^{n \times p \times n}$ can be constructed. Next, we define a function to compute the intended statistics with jackknifing. We then map this function to different channels of $\mathbb{X}_k$, where each $\mathbf{X}_k$ represents one channel of $\mathbb{X}_k$. By doing so, jackknifing computation can be executed by the GPU systems in a parallel fashion. We then collect the computation

3.3 The Communication and Computation Cost

results from each channel and reduce them into the desired statistics $\hat{\theta}_{JDS}^{(k)}$ and $\widehat{\text{SE}}_{(k)}^2 = \sum_{j \in \mathcal{S}_k} (\hat{\theta}_{-j}^{(k)} - \hat{\theta}^{(k)})^2$ for the k th subsample. Then we obtain the final estimators accordingly. This leads to the entire GPU algorithm. The details are provided below.

Input: Data $X_1, \dots, X_N$ on the HD, $X_i \in \mathbb{R}^p$;
 $g(\cdot)$: the function of interest of the moment;
 n : the subsample size K : the number of subsamples;
Output: A JDS estimator $\hat{\theta}_{JDS}$ and a JSE estimator $\widehat{\text{SE}}$.
for $k \leftarrow 1$ **to** K **do** subsampling
 Generate $\mathcal{S}_k \subset \mathbb{S}$, and place the $n \times p$ matrix $\mathbf{X}_k$ into the GM ;
 Compute $\hat{\theta}^{(k)} \leftarrow g(\hat{\mu}^{(k)})$ in the GM;
 Generate an $n \times p \times n$ tensor $\mathbb{X}_k$ in the GM ;
 Map the function $g(\cdot)$ to each channel of $\mathbf{X}_k$, which then leads to $\{\hat{\theta}_{-j}^{(k)}, j \in \mathcal{S}_k\}$;
 Compute $\hat{\theta}_{JDS}^{(k)} = \hat{\theta}^{(k)} - (n-1)\{n^{-1} \sum_{j=1}^n \hat{\theta}_{-j}^{(k)} - \hat{\theta}^{(k)}\}$ and
 $\widehat{\text{SE}}_{(k)}^2 = \sum_{j \in \mathcal{S}_k} (\hat{\theta}_{-j}^{(k)} - \hat{\theta}^{(k)})^2$ in the GM ;
end
 Compute $\hat{\theta}_{JDS} = K^{-1} \sum_{k=1}^K \hat{\theta}_{JDS}^{(k)}$ and
 $\widehat{\text{SE}}^2 = (1/K + n/K)K^{-1} \sum_{k=1}^K \widehat{\text{SE}}_{(k)}^2$ in the GM;
return $\hat{\theta}_{JDS}$ and $\widehat{\text{SE}}^2$.

Algorithm 1: The GPU algorithm

3.3 The Communication and Computation Cost

To evaluate the finite sample performance of the proposed method, we subsequently present a number of numerical experiments. We first consider how to generate the whole sample with a very large $N = 10^9$. For every

3.3 The Communication and Computation Cost

$1 \leq i \leq N$, we generate a 2-dimensional random variable $X_i = (X_{i1}, X_{i2})^\top$ independently and identically from a bivariate normal distribution with mean 0 and covariance $\Sigma = \{\sigma_{ij}\}_{2 \times 2}$, where $\sigma_{11} = 25$, $\sigma_{12} = \sigma_{21} = 10$, and $\sigma_{22} = 5$. We then define the parameter of interest to be the correlation coefficient $\text{Corr}(X_{i1}, X_{i2})$ as follows:

$$\theta = \text{Corr}(X_{i1}, X_{i2}) = \frac{\text{Cov}(X_{i1}, X_{i2})}{\sqrt{\text{var}(X_{i1}) \text{var}(X_{i2})}} = \frac{2}{\sqrt{5}}.$$

This parameter is a complex nonlinear function of various moments about X_i . Once the whole sample is generated, it is placed as a single file on the HD, requiring approximately 38.3 gigabytes. As one can see, this is a size that can hardly be read into a CM. Once the data are placed in the HD, they are fixed for the rest of the simulation experiments. In other words, we do not update the whole sample dataset on the HD across different simulation replications. For a reliable evaluation, we replicate the subsequent experiment a total of $M = 1000$ times. All computations are performed by using TensorFlow 2.2.0 on a single GPU device (NVIDIA Tesla P100).

In this subsection, we focus on the performance in terms of the time cost. We study both the communication cost and computation cost. The communication cost can be further divided into two parts. The first part is

3.3 The Communication and Computation Cost

the time cost required for transferring data from the HD to the CM. The second part is the time cost required for transferring data from the CM to GM. Next, we vary the subsample size n from 100 to 3,000 and K from 10 to 200. We then use $\mathcal{S}_k^{(m)} \subset \mathbb{S}$ to represent the k th subsample obtained in the m th simulation replication. The time cost used for obtaining $\mathcal{S}_k^{(m)}$ is recorded by $T_{1k}^{(m)}$. Based on $\mathcal{S}_k^{(m)}$, we can obtain matrix $\mathbf{X}_k^{(m)}$. We then transfer $\mathbf{X}_k^{(m)}$ from the HD to the CM, where the associated time cost is recorded as $T_{2k}^{(m)}$. The computation cost required for computing $\hat{\theta}_{JDS}$ and $\widehat{\text{SE}}$ is given by $T_{3k}^{(m)}$. Consequently, the total time cost is given by $T_k^{(m)} = T_{1k}^{(m)} + T_{2k}^{(m)} + T_{3k}^{(m)}$. Their averages are obtained as $T_1 = M^{-1} \sum_{k,m} T_{1k}^{(m)}$, $T_2 = M^{-1} \sum_{k,m} T_{2k}^{(m)}$, and $T_3 = M^{-1} \sum_{k,m} T_{3k}^{(m)}$. Their relationships with both K and n are investigated.

The detailed results are given in Figure 3. As one can see from Figure 3, all types of time cost increase as the number of subsamples K increases. In particular, the communication cost required by subsampling (i.e., T_1) is substantially larger than the other two types of time cost. Comparatively, the communication cost required by jackknifing (i.e., T_2) is the smallest. It is remarkable that T_3 is supposed to be very significant if a CPU-only system is used. However, due to the use of a GPU system, the corresponding time cost becomes practically ignorable. To understand this idea, considering one

3.3 The Communication and Computation Cost

Figure 3: The log-transformed time cost for different (n, K) combinations with $K = 10, 50$ and 200 . The communication cost due to subsampling T_1 is given in the left panel. The communication cost due to jackknifing T_2 is reported in the middle panel. The computation cost T_3 is presented in the right panel. The reported time costs (in log-scale) are averaged based on $M = 1000$ simulations.

special case with $K = 50$ and $n = 3000$, we have $T_1 = 1.882$ s, $T_2 = 0.017$ s and $T_3 = 0.055$ s. We also find in the middle and right panels of Figure 3 that for a fixed total subsample size K , T_2 and T_3 remain almost unchanged as n increases. This result demonstrates the excellent parallel capability of a GPU-based system and suggests that better computation efficiency can be achieved by setting the subsample size n to be as large as possible as long as the computer memory allows this.

Next, we demonstrate the computational advantage of a GPU system. To this end, we define $T_{GPU}^{(m)}$ as the total time cost required by the m th simulation replication except the communication cost due to subsampling

3.3 The Communication and Computation Cost

(such cost is required by any computation system). We then execute the same algorithm on a CPU-only system (in our case, TensorFlow 2.2.0 can also be executed on the CPU-only system). This leads to the total time cost except the communication cost due to subsampling required by the CPU-only system, which is recorded as $T_{CPU}^{(m)}$. We then compute their ratio for the m th replication as $R^{(m)} = T_{GPU}^{(m)} / T_{CPU}^{(m)}$. We define the averaged ratio as $AR = M^{-1} \sum_{m=1}^M R^{(m)}$. Then, the relationships of the log-transformed AR values for different (n, K) combinations are reported in Figure 4. As we can see from Figure 4, the $\log(AR)$ values are always smaller than 0. This suggests that the computational time cost required by a GPU-based system is always smaller than that required by a CPU-only system on average. In fact, the reported $\log(AR)$ values seem to be rather insensitive to the number of subsamples (i.e., K). Furthermore, for a fixed number of subsamples K , the $\log(AR)$ value decreases as the subsample size n increases. This is because a larger n requires a higher computation cost. Accordingly, the parallel computational power of a GPU system can be better demonstrated. For instance, considering the case with $K = 50$ and $n = 3000$, the averaged time cost of the GPU system is approximately 0.55 s, while that of the CPU system is approximately 4.83 s. The corresponding AR value is $AR = 0.011$. This suggests that the computation cost required by a GPU

3.4 Simulation Results of the JSE Estimator

system is only approximately 1.1% that of a CPU system on average.

Figure 4: Comparison of the computation efficiency between a GPU system and a CPU system. The AR is reported in log-scale based on $M = 1000$ simulation replications. The numbers of subsamples are fixed to $K = 10, 50$ and 200.

3.4 Simulation Results of the JSE Estimator

In this subsection, we focus on the finite sample performance of the JSE estimator $\widehat{\text{SE}}$. To this end, we follow the simulation setup in the previous subsection. Note that the data on the HD are generated only one time to conserve experimental time. Once the data are generated, we replicate experiments $M = 1000$ times based on the same whole sample dataset. Specifically, for the m th replication, we obtain an SOS estimator $\widehat{\theta}_{\text{SOS}}^{(m)}$, a JDS estimator $\widehat{\theta}_{\text{JDS}}^{(m)}$, and a JSE estimator $\widehat{\text{SE}}^{(m)}$. Define SE_{SOS} and

3.4 Simulation Results of the JSE Estimator

SE_{JDS} as the respective sample standard deviations of $\{\hat{\theta}_{SOS}^{(m)}, m = 1, \dots, M\}$ and $\{\hat{\theta}_{JDS}^{(m)}, m = 1, \dots, M\}$. Accordingly, SE_{SOS} and SE_{JDS} measure the variabilities of $\hat{\theta}_{SOS}$ and $\hat{\theta}_{JDS}$ conditional on the whole sample dataset on the HD. Because we have $N \gg nK$, they should be good approximations of the true variabilities of $\hat{\theta}_{SOS}$ and $\hat{\theta}_{JDS}$; see Theorems 2 and 3. Next, for the m th replication, we can define the relative absolute errors as $RAE_{SOS}^{(m)} = |\widehat{SE}^{(m)} / SE_{SOS} - 1|$ and $RAE_{JDS}^{(m)} = |\widehat{SE}^{(m)} / SE_{JDS} - 1|$. They are box plotted in Figure 5.

Figure 5: Boxplots of RAE values for the JDS (light box) and SOS (dark box) estimators. The left panel corresponds to the case with n fixed to $n = 100$. The right panel corresponds to the case with K fixed to $K = 100$. Each box is summarized based on $M = 1000$ simulation replications.

As one can see from the left panel of Figure 5, the RAE values of the SOS and JDS estimators are similar. They both decrease to 0 as K

3.5 Simulation Results of the JDS Estimator

increases. This suggests that a larger K leads to more accurate JSE estimators, under the condition that n is fixed. Qualitatively similar patterns are also observed for the right panel. We find that a larger n leads to a more accurate JSE estimation, under the condition that K is fixed. To summarize, both boxplots in Figure 5 suggest that the proposed JSE estimator is consistent as $nK \rightarrow \infty$.

3.5 Simulation Results of the JDS Estimator

Finally, we evaluate the finite sample performance of the point estimation $\hat{\theta}_{JDS}$ and its statistical inference in terms of the confidence interval. For comparison, that of the SOS estimator $\hat{\theta}_{SOS}$ is also evaluated. Specifically, following the simulation setup in the previous subsection, we replicate experiments $M = 1000$ times based on the same whole dataset on the HD. For the m th replication, we calculate the JDS estimator $\hat{\theta}_{JDS}^{(m)}$ and the corresponding JSE estimator $\widehat{SE}^{(m)}$. This leads to a total of M estimators $\{(\hat{\theta}_{JDS}^{(m)}, \widehat{SE}^{(m)}) : 1 \leq m \leq M\}$. Based on these estimators, the averaged bias can be computed as $\text{Bias} = M^{-1} \sum_{m=1}^M (\hat{\theta}_{JDS}^{(m)} - \theta)$, and the corresponding standard error (SE) can be obtained. In addition, for each estimator $\hat{\theta}_{JDS}^{(m)}$, a $(1 - \alpha)$ th level confidence interval for θ is constructed as $\text{CI}^{(m)} = [\hat{\theta}_{JDS}^{(m)} - \widehat{SE}^{(m)} Z_{1-\alpha/2}, \hat{\theta}_{JDS}^{(m)} + \widehat{SE}^{(m)} Z_{1-\alpha/2}]$, where $\alpha = 0.05$

3.6 Real Data Analysis

and Z_α represents the lower α quantile of the standard normal distribution. The empirical coverage probabilities are then also evaluated as $\text{ECP}_{JDS} = M^{-1} \sum_{m=1}^M I(\theta \in \text{CI}^{(m)})$, where $I(\cdot)$ is the indicator function. The SOS estimator $\hat{\theta}_{\text{SOS}}$ is evaluated similarly. The detailed results are given in Table 2.

From Table 2, we find that the two estimators perform similarly in terms of the standard error (SE) for various (n, K) combinations. However, they are very different in terms of bias. The bias of the SOS estimator $\hat{\theta}_{\text{SOS}}$ is much larger than that of $\hat{\theta}_{\text{JDS}}$. Considering for example the case with $n = 200$ and $K = 200$, the bias of $\hat{\theta}_{\text{SOS}}$ is 4.49×10^{-4} , while that of $\hat{\theta}_{\text{JDS}}$ is only 1.1×10^{-5} . As one can see, the former is approximately forty times larger than the latter. Moreover, the bias of $\hat{\theta}_{\text{SOS}}$ is quite comparable to its standard error. As a consequence, the confidence interval of $\hat{\theta}_{\text{SOS}}$ is poor, resulting from the fact that the corresponding ECP is significantly smaller than 95%. In contrast, the confidence interval of the JDS estimator is good since the corresponding ECP values of $\hat{\theta}_{\text{JDS}}$ are quite close to 95%.

3.6 Real Data Analysis

In this subsection, we study a real dataset: the U.S. Airline Dataset. The dataset is available on the official website of the American Statistical

3.6 Real Data Analysis

Table 2: Comparison of the SOS estimator $\hat{\theta}_{SOS}$ and the JDS estimator $\hat{\theta}_{JDS}$ based on $M = 1000$ simulation replications for various (n, K) combinations.

K	SE ($\times 10^{-3}$)		Bias ($\times 10^{-3}$)		ECP (%)	
	$\hat{\theta}_{SOS}$	$\hat{\theta}_{JDS}$	$\hat{\theta}_{SOS}$	$\hat{\theta}_{JDS}$	$\hat{\theta}_{SOS}$	$\hat{\theta}_{JDS}$
$n = 50$						
100	2.964	2.929	2.038	0.050	92.1	95.8
200	2.103	2.077	2.054	0.066	87.8	97.0
500	1.335	1.314	1.956	0.032	73.6	96.8
1000	0.972	0.959	1.907	0.078	50.8	95.4
$n = 100$						
100	2.068	2.057	0.910	0.029	93.3	96.1
200	1.446	1.437	0.960	0.019	91.7	95.3
500	0.938	0.932	0.877	0.065	84.8	95.6
1000	0.672	0.667	0.904	0.038	72.9	95.6
$n = 200$						
100	1.436	1.432	0.487	0.029	93.7	94.8
200	1.029	1.025	0.449	0.011	92.9	95.4
500	0.656	0.654	0.442	0.017	89.4	95.0
1000	0.441	0.440	0.477	0.018	83.3	96.4

Association (ASA). The airline dataset contains approximately 120 million records. It takes up approximately 12 gigabytes of space on a hard drive. Each record contains detailed information for one particular commercial flight in the USA from October 1987 to April 2008. The dataset contains 13 continuous variables and 16 categorical variables. For illustration, we focus on the 13 continuous variables. However, a significant portion of records are missing for many continuous variables. Only 5 of them have missing rates less than 10%: **ActualElapsedTime** (actual elapsed time), **CRSElapsedTime** (scheduled elapsed time), **Distance**, **DepDelay** (de-

3.6 Real Data Analysis

Table 3: Descriptive statistics for the 5 continuous variables based on the whole airline dataset after signed-log-transformation. The descriptive statistics are given by the sample mean (Mean), sample standard deviation (SD) and sample kurtosis (Kurt).

	Actual ElapsedTime	CRS ElapsedTime	Distance	DepDelay	ArrDelay
Mean	4.656	4.670	6.272	0.492	0.236
SD	0.525	0.513	0.777	1.905	2.463
Kurt	2.586	2.597	2.750	2.272	1.594

parture delay), and **ArrDelay** (arrival delay). As a consequence, only these 5 variables are subsequently illustrated. For more detailed variable information, refer to the ASA official website at <http://stat-computing.org/dataexpo/2009>.

For each variable, the signed-log-transformation is applied: $\log|x| \cdot \text{sign}(x)$ transformation. This transformation is conducted purely for illustration. Otherwise, many variables (e.g., **ArrDelay**) are so heavy-tailed that the existence of finite moments becomes questionable. For each transformed variable, the following parameters are studied: the mean, standard deviation, and kurtosis. Their WS estimators are given in Table 3. These WS estimators are then treated as if they were the true parameters. Accordingly, simulation experiments can be conducted as in the previous subsections. In this case, we fixed $nK = 6 \times 10^4$, with different (n, K) combinations, and replicated the experiments $M = 1000$ times. The detailed results are

3.6 Real Data Analysis

summarized in Table 4.

From Table 4, we can obtain the following interesting observations. First, note that the sample mean is an exactly unbiased estimator for the mean. Accordingly, both the SOS and JDS estimators are unbiased. In fact, they are identical to each other in this case. As a result, both estimators demonstrated identical simulation results, with ECP values both very close to their nominal level of 95% for all five variables. Second, for the other two parameters (i.e., standard deviation and kurtosis), the sample estimators are no longer unbiased. Accordingly, the SOS and JDS estimators are no longer identical. As we expect, both estimators are similar in terms of the standard error (SE). However, they are very different in terms of the empirical bias. Obviously, the bias of the SOS estimator is substantially larger than that of the JDS estimator for all reported cases. As a consequence, the ECP values of the SOS estimator significantly depart from their nominal level of 95%. In contrast, those of the JDS estimator remain very close to 95%. Consider for example the case of the kurtosis of `ArrDelay` with $n = 200$ and $K = 300$. The ECP value of the SOS estimator is only 35.6%. In contrast, that of the JDS is 95.7%.

4. CONCLUDING REMARKS

In this article, we develop a novel statistical method for datasets with large sizes. The new method is particularly designed for practitioners with limited computational resources. The proposed method combines the ideas of both subsampling and jackknifing. Subsampling allows our method to work with datasets with large sizes. Jackknifing further enhances this capability by significantly reducing the bias. To practically implement our method, a novel algorithm is developed for GPU systems. We theoretically show that the resulting estimator could be as good as the whole sample estimator under very mild regularity conditions. Extensive numerical studies built on both simulation and real datasets are presented to demonstrate its outstanding performance.

To conclude this work, we would like to discuss a few interesting topics for future study. First, the statistics considered in this work are relatively simple. They represent nonlinear transformation of various moments. It is then of great interest to develop similar methods for more general M estimators. Second, the data considered in this work are collected from independent samples. This makes the theoretical understanding of the resulting subsample estimator analytically simple. How to develop similar methods for data with a sophisticated dependence structure (e.g., spatial

REFERENCES

temporal data) is another interesting topic worth studying. Future research along this direction is definitely needed.

Supplementary Materials

The supplementary materials contain the proofs of all the theoretical results in the main text.

Acknowledgements

We thank the editor, the associate editor and two referees for insightful comments that have substantially improved this article. Xuening Zhu's research is supported by the National Natural Science Foundation of China (nos. 11901105, 71991472, U1811461), Shanghai Science and Technology Commission Grant (No. 17JC1420200), and the Shanghai Sailing Program for Youth Science and Technology Excellence (19YF1402700). Hansheng Wang's research is partially supported by National Natural Science Foundation of China (No. 11831008).

References

Cameron, A. C. and Trivedi, P. K. (2005), *Microeconometrics: methods and applications*, Cambridge university press.

REFERENCES

- Che, S., Michael, B., Meng, J., Tarjan, D., Sheaffer, J., and Skadron, K. (2008), “A performance study of general-purpose applications on graphics processors using CUDA,” *Journal of Parallel and Distributed Computing*, 68, 1370–1380.
- Drineas, P., Mahoney, M. W., Muthukrishnan, S., and Sarlós, T. (2011), “Faster least squares approximation,” *Numerische mathematik*, 117, 219–249.
- Efron, B. and Stein, C. (1981), “The jackknife estimate of variance,” *The Annals of Statistics*, 586–596.
- Huang, C. and Huo, X. (2015), “A distributed one-step estimator,” *arXiv preprint arXiv:1511.01443*.
- Jordan, M. I., Lee, J. D., and Yang, Y. (2019), “Communication-efficient distributed statistical inference,” *Journal of the American Statistical Association*, 114, 668–681.
- Krüger, J. and Westermann, R. (2005), “Linear algebra operators for GPU implementation of numerical algorithms,” in *ACM SIGGRAPH 2005 Courses*, pp. 234–es.
- Lehmann, E. L. and Casella, G. (2006), *Theory of point estimation*, Springer Science & Business Media.
- Ma, P., Mahoney, M. W., and Yu, B. (2015), “A statistical perspective on algorithmic leveraging,” *The Journal of Machine Learning Research*, 16, 861–911.
- Ma, P., Zhang, X., Xing, X., Ma, J., and Mahoney, M. (2020), “Asymptotic analysis of sampling estimators for randomized numerical linear algebra algorithms,” in *International Confer-*

REFERENCES

- ence on Artificial Intelligence and Statistics*, PMLR, pp. 1026–1035.
- Mahoney, M. W. (2011), “Randomized algorithms for matrices and data,” *Foundations and Trends® in Machine Learning*, 3, 123–224.
- Mcdonald, R., Mohri, M., Silberman, N., Walker, D., and Mann, G. S. (2009), “Efficient large-scale distributed training of conditional maximum entropy models,” in *Advances in neural information processing systems*, pp. 1231–1239.
- Quenouille, M. H. (1949), “Approximate Tests of Correlation in Time-Series,” *Journal of the Royal Statistical Society*, 11, 68–84.
- Shao, J. (2003), “Mathmetical Statistics,” *New York, Springer*.
- Suresh, A. T., Felix, X. Y., Kumar, S., and McMahan, H. B. (2017), “Distributed mean estimation with limited communication,” in *International Conference on Machine Learning*, PMLR, pp. 3329–3337.
- Wang, F., Huang, D., Zhu, Y., and Wang, H. (2020), “Efficient Estimation for Generalized Linear Models on a Distributed System with Nonrandomly Distributed Data,” *arXiv preprint arXiv:2004.02414*.
- Wang, H. (2019), “More Efficient Estimation for Logistic Regression with Optimal Subsamples,” *Journal of Machine Learning Research*, 20, 1–59.
- Wang, H., Zhu, R., and Ma, P. (2018), “Optimal subsampling for large sample logistic regression,” *Journal of the American Statistical Association*, 113, 829–844.

REFERENCES

Wu, C.-F. J. (1986), “Jackknife, bootstrap and other resampling methods in regression analysis,”
the Annals of Statistics, 14, 1261–1295.

Yu, J., Wang, H., Ai, M., and Zhang, H. (2020), “Optimal Distributed Subsampling for Maximum Quasi-Likelihood Estimators with Massive Data,” *Journal of the American Statistical Association*, Accepted on May 20, 2020.

Zhang, Y., Duchi, J. C., and Wainwright, M. J. (2013), “Communication-efficient algorithms for statistical optimization,” *The Journal of Machine Learning Research*, 14, 3321–3363.

Zhu, X., Pan, R., Wu, S., and Wang, H. (2021), “Feature Screening for Massive Data Analysis by Subsampling,” *Journal of Business & Economic Statistics*, 0, 1–31.

Zinkevich, M., Weimer, M., Smola, A. J., and Li, L. (2011), “Parallelized Stochastic Gradient Descent,” in *Advances in Neural Information Processing Systems 23: Conference on Neural Information Processing Systems A Meeting Held December*.

Guanghua School of Management, Peking University, Beijing, China

E-mail: shuyuan.w@pku.edu.cn

School of Data Science, Fudan University, Shanghai, China

E-mail: xueningzhu@fudan.edu.cn (Corresponding author)

Guanghua School of Management, Peking University, Beijing, China

E-mail: hansheng@pku.edu.cn

REFERENCES

Figure 6: The whole process for sampling without replacement on the hard drive.

Table 4: Simulation results for the airline dataset based on $M = 1000$ simulation replications. The parameters of interest include the mean (Mean), standard deviation (SD), and kurtosis (Kurt) for signed-log-transformed variables. Both the SOS estimator $\hat{\theta}_{SOS}$ and the JDS estimator $\hat{\theta}_{JDS}$ are compared in terms of Bias, SE and ECP. The nominal level of ECP is 95%.

	ActualElapsedTime			CRSElapsedTime			Distance			DepDelay			ArrDelay		
	JDS			$n = 300, K = 200$			SOS			JDS			SOS		
	SOS	Bias ($\times 10^{-2}$)	SE ($\times 10^{-2}$)	ECP (%)	SOS	Bias ($\times 10^{-2}$)	SE ($\times 10^{-2}$)	ECP (%)	SOS	Bias ($\times 10^{-2}$)	SE ($\times 10^{-2}$)	ECP (%)	SOS	Bias ($\times 10^{-2}$)	SE ($\times 10^{-2}$)
Mean	0.003	0.003	0.003	94.6	0.002	0.002	0.002	94.1	0.001	0.001	0.001	94.6	0.001	0.001	0.001
	0.219	0.219	0.216	94.6	0.216	0.216	0.216	94.1	0.326	0.326	0.326	94.6	0.771	0.771	0.771
	94.6	94.6	94.6	94.6	94.1	94.1	94.1	94.1	94.1	94.1	94.1	94.6	95.7	95.7	95.7
SD	0.130	0.008	0.134	95.7	0.127	0.007	0.130	95.6	0.194	0.008	0.424	96.0	0.476	0.002	0.371
	0.133	0.134	0.130	95.7	0.130	0.131	0.130	95.6	0.205	0.206	0.415	96.0	0.370	0.370	0.371
	83.9	95.7	84.3	95.7	84.3	95.6	84.7	96.5	84.7	96.5	86.2	96.0	78.6	96.5	96.5
Kurt	0.682	0.032	1.778	93.7	0.679	0.025	1.706	94.7	1.685	0.056	0.431	95.4	0.789	0.002	0.492
	1.576	1.778	1.706	93.7	1.706	2.103	1.979	94.7	1.979	2.119	1.197	95.4	0.487	0.487	0.492
	89.6	93.7	89.9	94.7	89.9	94.7	85.2	94.7	85.2	94.7	94.3	95.4	66.2	95.2	95.2
Mean	0.003	0.003	0.002	94.8	0.002	0.002	0.001	94.2	0.001	0.001	0.001	94.6	0.026	0.026	0.026
	0.219	0.219	0.216	94.8	0.216	0.216	0.326	94.2	0.326	0.326	0.771	94.6	0.980	0.980	0.980
	94.8	94.8	94.2	94.8	94.2	94.2	94.2	94.2	94.2	94.2	94.6	94.6	95.7	95.7	95.7
SD	0.192	0.008	0.187	95.3	0.187	0.007	0.288	95.2	0.288	0.008	0.634	96.2	0.711	0.000	0.371
	0.133	0.133	0.130	95.3	0.130	0.130	0.205	95.2	0.205	0.205	0.414	96.2	0.370	0.370	0.371
	71.4	95.3	72.6	95.3	72.6	95.2	72.5	96.2	72.5	96.2	70.2	96.2	55.7	96.6	96.6
Kurt	1.035	0.025	1.024	93.8	1.024	0.012	2.542	94.8	2.542	0.023	0.631	96.1	1.179	0.008	0.502
	1.511	1.745	1.571	93.8	1.571	1.984	1.952	94.8	1.952	2.158	1.203	96.1	0.494	0.494	0.502
	85.4	93.8	85.0	93.8	85.0	94.8	75.3	92.6	75.3	92.6	93.0	96.1	35.6	95.7	95.7

REFERENCES