



A selective review on statistical methods for massive data computation: distributed computing, subsampling, and minibatch techniques

Xuetong Li, Yuan Gao, Hong Chang, Danyang Huang, Yingying Ma, Rui Pan, Haobo Qi, Feifei Wang, Shuyuan Wu, Ke Xu, Jing Zhou, Xuening Zhu, Yingqiu Zhu & Hansheng Wang

To cite this article: Xuetong Li, Yuan Gao, Hong Chang, Danyang Huang, Yingying Ma, Rui Pan, Haobo Qi, Feifei Wang, Shuyuan Wu, Ke Xu, Jing Zhou, Xuening Zhu, Yingqiu Zhu & Hansheng Wang (2024) A selective review on statistical methods for massive data computation: distributed computing, subsampling, and minibatch techniques, *Statistical Theory and Related Fields*, 8:3, 163-185, DOI: [10.1080/24754269.2024.2343151](https://doi.org/10.1080/24754269.2024.2343151)

To link to this article: <https://doi.org/10.1080/24754269.2024.2343151>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 23 Apr 2024.



[Submit your article to this journal](#)



Article views: 1015



[View related articles](#)



[View Crossmark data](#)



Citing articles: 2 [View citing articles](#)



A selective review on statistical methods for massive data computation: distributed computing, subsampling, and minibatch techniques

Xuetong Li^a, Yuan Gao^{id}^a, Hong Chang^a, Danyang Huang^b, Yingying Ma^c, Rui Pan^d, Haobo Qi^e, Feifei Wang^b, Shuyuan Wu^f, Ke Xu^g, Jing Zhou^b, Xuening Zhu^h, Yingqiu Zhu^g and Hansheng Wang^a

^aGuanghua School of Management, Peking University, Beijing, People's Republic of China; ^bCenter for Applied Statistics and School of Statistics, Renmin University of China, Beijing, People's Republic of China; ^cSchool of Economics and Management, Beihang University, Beijing, People's Republic of China; ^dSchool of Statistics and Mathematics, Central University of Finance and Economics, Beijing, People's Republic of China; ^eSchool of Statistics, Beijing Normal University, Beijing, People's Republic of China; ^fSchool of Statistics and Management, Shanghai University of Finance and Economics, Shanghai, People's Republic of China; ^gSchool of Statistics, University of International Business and Economics, Beijing, People's Republic of China; ^hSchool of Data Science and MOE Laboratory for National Development and Intelligent Governance, Fudan University, Shanghai, People's Republic of China

ABSTRACT

This paper presents a selective review of statistical computation methods for massive data analysis. A huge amount of statistical methods for massive data computation have been rapidly developed in the past decades. In this work, we focus on three categories of statistical computation methods: (1) distributed computing, (2) subsampling methods, and (3) minibatch gradient techniques. The first class of literature is about distributed computing and focuses on the situation, where the dataset size is too huge to be comfortably handled by one single computer. In this case, a distributed computation system with multiple computers has to be utilized. The second class of literature is about subsampling methods and concerns about the situation, where the blacksample size of dataset is small enough to be placed on one single computer but too large to be easily processed by its memory as a whole. The last class of literature studies those minibatch gradient related optimization techniques, which have been extensively used for optimizing various deep learning models.

ARTICLE HISTORY

Received 28 January 2024
Revised 18 March 2024
Accepted 5 April 2024

KEYWORDS

Distributed computing;
massive data analysis;
minibatch techniques;
stochastic optimization;
subsampling methods

1. Introduction

Modern statistical analysis often involves datasets of massive size (Fan et al., 2020), for which effective computation methods are indispensable. On one side, the huge demand for computation methods for massive data analysis places serious challenges on the traditional statistical methods, which have been developed for datasets of regular size. On the other side, it also stimulates new research efforts, which try to conquer computation challenges by statistical wisdom. The research along this direction not only benefits the real practices with massive datasets but also inspires new statistical theory. The objective of this work is to provide a selective review about this exciting research area, which has been rapidly developed during the past many years. Then the objective here is not to provide a complete list for all the research works related to massive data computation. This is obviously a mission impossible. Instead, we should focus on three categories of statistical computing methods. They are, respectively, distributed computing methods, subsampling methods, and minibatch gradient descent methods. We try to organize the most important and relevant ones associated with those three categories in a structured way, so that follow-up researchers might benefit. Due to our limited understanding about the past literature and also the space constraint of a regular manuscript, we might miss some important references therein. If that happens, our sincere apology in advance and we should be more than happy to hear the feedback.

As the title suggests, this review is a selective review about statistical methods for massive data analysis. Then, the meaning of 'massive data' needs to be precisely defined. We argue that whether a dataset is massive or not is relative to the computation resource. In the most ideal situation, if one is given a super computer with an unlimited amount of hard drive and memory (both CPU and GPU memory) together with a super powerful CPU, then no dataset can be considered as massive. In this ideal situation, any dataset of any size can be easily placed on one hard drive, loaded into the memory as a whole, and then processed in no time. Unfortunately, such an ideal situation never happens in reality. In real practice, most researchers are given only a limited amount of computation resources, which put various constraints on the computation power. The constraints could be the hard drive. If the size of the

CONTACT Yuan Gao ✉ yuan_gao@pku.edu.cn, ygao_stat@outlook.com 📧 Guanghua School of Management, Peking University, 5 Yiheyuan Road (Summer Palace Road), Haidian District, Beijing

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

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. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

data exceeds one single computer's hard drive capacity, then a distributed system has to be used to place the data. A natural question arises immediately: should we also compute it in a distributed way? This inspires a large amount of research for distributed computing, if a powerful distributed computation system is indeed available. With such a system, we find that it might remain to be practically appealing to distribute a large dataset on a powerful distributed system, even if the data size is not strictly larger than one single computer's hard drive. This makes the subsequent computation more convenient. This constitutes the first part of the selective literature to be reviewed in this work.

Distributed computing is a powerful solution for problems with extremely large scale datasets. However, this seems not the most typical situation. The most typical one in real practice is an embarrassing situation, where the dataset sizes are not extremely large but large enough to cause a lot of computation challenges. To fix the idea, consider for example a dataset of a size (for example) 100GB. Note that this is a size substantially smaller than that of a hard drive of a modern computer (e.g., 2TB), but much larger than the size of the typical memory (e.g., 32GB). For a dataset of this embarrassing size, one straightforward solution remains to be distributed computing, as mentioned in the previous paragraph. However, this straightforward solution seems not the only best one for at least two reasons. First, to implement a distributed computing algorithm, one needs a powerful distributed computation system. Depending on its size, the distributed computer system could be very expensive, if not completely not affordable. Second, to utilize a distributed computing system, appropriate programming techniques are necessarily needed. Popularly used programming frameworks (e.g., Spark, Hadoop) need to be learned. This is unfortunately a painful learning process for most field practitioners (e.g., a medical scientist), who are not professional statistical programmers. Therefore, it seems that there is a practical need for a handy computation method, which can deal with datasets of any size on a given hard drive and can be easily implemented on one single computer. The key challenge here is how to accomplish a massive data computation task with limited memory (both CPU and GPU memories) constraints. This leads to a huge body of literature about subsampling, and/or streaming data analysis. This constitutes the second part of the literature to be reviewed this work.

Both the problems of distributed computing and computing with memory constraint concern about the computation problems of datasets with massive sizes. However, we often encounter situations where not only the dataset sizes are large, but also the model sizes are extremely large. The most typical example in this regard is various deep learning methods. To fix the idea, consider for example the famous *ImageNet* dataset of Deng et al. (2009), which contains a total of over 1.3 million colour images belonging to 1000 classes. The total amount of images is about 150GB in size. Next, consider for example a classical deep learning model VGG16 (Simonyan & Zisserman, 2015). This is a convolution neural network (CNN) model with a total of over 130 million parameters. To train this VGG16 model on the *ImageNet* dataset, all the model parameters need to be fully placed in the GPU memory for fast tensor computation. Unfortunately, once this sophisticated VGG16 model is fully loaded into the GPU memory, the space left for data processing is inevitably significantly reduced. Consequently, the *ImageNet* dataset has to be processed in a minibatch-by-minibatch manner. Here, a minibatch refers to a small or even tiny subset of the whole sample. Depending on the way this subsample is generated, we might have streaming data based minibatches (Chen, Lee et al., 2020), subsampling based stochastic minibatches (Gower et al., 2019), and random partition based minibatches (Gao et al., 2023; Qi, Wang et al., 2023). This constitutes the third part of the literature to be reviewed in this work.

The rest of this article is organized as follows. Section 2 reviews the literature about distributed computing. Section 3 studies various subsampling methods. Section 4 discusses minibatch gradient related techniques. The article is then concluded with a brief discussion in Section 5.

2. Distributed computing

2.1. Theoretical framework of distributed computing

Consider a standard statistical learning problem. Assume a total of N observations, where N is notably large. For each observation i , we collect a response variable $Y_i \in \mathbb{R}^1$ and its corresponding feature vector $X_i \in \mathbb{R}^p$. The primary goal here is to accurately estimate an unknown parameter $\theta_0 \in \mathbb{R}^p$ through a suitably defined loss function. To be specific, define the empirical loss function as $\mathcal{L}(\theta) = \sum_{i=1}^N \ell(X_i, Y_i; \theta)$. Here, $\ell(X_i, Y_i; \theta)$ represents the loss function for the i -th observation.

In conventional scenarios with relatively small N , this learning problem could be easily solved using various standard optimization algorithms (e.g., Newton-Raphson method and gradient descent method). Nevertheless, for datasets of massive size, implementation of these standard algorithms becomes practically challenging or even infeasible. Consider, for example, the classical Newton-Raphson algorithm. Let $\widehat{\theta}^{(t)}$ be the estimator derived in the t -th iteration. Then, the $(t + 1)$ -th step estimator is updated as follows:

$$\widehat{\theta}^{(t+1)} = \widehat{\theta}^{(t)} - \{\ddot{\mathcal{L}}(\widehat{\theta}^{(t)})\}^{-1} \dot{\mathcal{L}}(\widehat{\theta}^{(t)}), \quad (1)$$

where $\dot{\mathcal{L}}(\theta)$ and $\ddot{\mathcal{L}}(\theta)$ represent the 1st- and 2nd-order derivatives of the loss function $\mathcal{L}(\cdot)$ with respect to θ . With a fixed feature dimension p , the computational complexity of the Newton-Raphson algorithm is at least of the order $O(N)$ in each iteration. In the case of large datasets with an exceedingly large N , such computation costs could be practically challenging or even infeasible. To address this issue, various distributed computing methods have been developed. The key idea of distributed computing is to divide a massive dataset into smaller pieces, which can be processed simultaneously across many multiple computer machines.

Assume the N samples are distributed across a total of M different local machines and each machine is assigned n_m observations for $1 \leq m \leq M$. It follows that $\sum_{m=1}^M n_m = N$. We then denote the whole sample as $\mathcal{S}_F = \{1, 2, \dots, N\}$ and the sample assigned to the m -th local computer as $\mathcal{S}_{(m)} \subset \mathcal{S}_F$. Thus, we have $\cup_{m=1}^M \mathcal{S}_{(m)} = \mathcal{S}_F$, and $\mathcal{S}_{(m_1)} \cap \mathcal{S}_{(m_2)} = \emptyset$ for any $m_1 \neq m_2$, and $|\mathcal{S}_{(m)}| = n_m$. Recall that the global loss function is defined as $\mathcal{L}(\theta) = N^{-1} \sum_{i=1}^N \ell(X_i, Y_i; \theta)$. The averaged local sample size is denoted as $n = M^{-1} \sum_{m=1}^M n_m$. Define $\hat{\theta} = \operatorname{argmin}_{\theta} \mathcal{L}(\theta)$ and $\theta_0 = \operatorname{argmin}_{\theta} E\{\ell(X_i, Y_i; \theta)\}$ as the global estimator and true parameter, respectively. Subsequently, define the local loss function on the m -th local computer as $\mathcal{L}_{(m)}(\theta) = n_m^{-1} \sum_{i \in \mathcal{S}_{(m)}} \ell(X_i, Y_i; \theta)$. Let $\hat{\theta}_{(m)} = \operatorname{argmin}_{\theta} \mathcal{L}_{(m)}(\theta)$ be the estimator locally obtained on the m -th local computer. Moreover, denote $\dot{\ell}(X_i, Y_i; \theta) = \partial \ell(X_i, Y_i; \theta) / \partial \theta \in \mathbb{R}^p$ and $\ddot{\ell}(X_i, Y_i; \theta) = \partial^2 \ell(X_i, Y_i; \theta) / \partial \theta \theta^\top \in \mathbb{R}^{p \times p}$ as the 1st- and 2nd-order derivatives of $\ell(X_i, Y_i; \theta)$ with respect to θ , respectively.

2.2. One-shot methods

For distributed statistical learning, various one-shot (OS) methods have been developed (Hector & Song, 2020; Lee et al., 2017; McDonald et al., 2009; Rosenblatt & Nadler, 2016; Zhang et al., 2012; Zinkevich et al., 2010). The basic idea is to calculate some important statistics on each local machine based on the data stored in each local machine in a fully parallel way. Subsequently, they are sent to the central machine, where these statistics are then assembled into one final estimator. Specifically, each local machine $1 \leq m \leq M$ uses local sample $\mathcal{S}_{(m)}$ to compute the local estimator $\hat{\theta}_{(m)} = \operatorname{argmin}_{\theta} \mathcal{L}_{(m)}(\theta)$. Subsequently, the central server collects these local estimates and aggregates them to obtain the final estimator $M^{-1} \sum_{m=1}^M \hat{\theta}_{(m)}$, which is denoted as the OS estimator $\hat{\theta}_{\text{os}}$.

Extensive research has been proposed in this field. For example, Chen and Xie (2014) studied the properties of one-shot estimator based on penalized generalized linear regression with smoothly clipped absolute deviation (SCAD) penalty. Battey et al. (2018) proposed high-dimensional one-shot estimators based on Wald tests and Rao's score tests. They extended the classical one-shot estimator from low-dimensional generalized linear regression to high-dimensional sparse scenarios. Lian and Fan (2018) developed a debiased form of one-shot estimator for support vector machines for ultra-high-dimensional data. Tang et al. (2020) used the confidence distribution approach to combine bias-corrected lasso-type estimates computed in each local machine in the generalized linear model setting. The one-shot strategy for correlated outcomes is also discussed in the previous literature. One notable work is the distributed and integrated method of moments (DIMM) proposed by Hector and Song (2021), which addresses the estimation problem in a regression setting with high-dimensional correlated outcomes. The key idea is to split all outcomes into blocks of low-dimensional response subvectors, then analyse these blocks in a distributed scheme, and finally combine the block-specific results using a closed-form meta-estimator. By this way, the computational challenges associated with high-dimensional correlated outcomes are alleviated. A generalization of DIMM is further developed in Hector and Song (2020), which doubly divides the data at both the outcome and subject levels to speed up computation. Recently, a distributed empirical likelihood (DEL) method has been proposed to solve the estimation problem for imbalanced local datasets in the framework of integrative analysis (Zhou et al., 2023).

As one can see, the OS method is easy to implement. It is also communicationally efficient because it requires only one round of communication between the local computers and the central computer (i.e., transferring the local estimates $\hat{\theta}_{(m)}$ s). However, many researchers (Wang et al., 2021; Wu et al., 2023c; Zhang et al., 2012) have pointed out that a number of critical conditions are necessarily needed by various OS methods to achieve the same asymptotic efficiency as the global estimator $\hat{\theta}$. The first condition is *uniformity*, implying that the massive data should be distributed across local computers in a relatively uniform manner so that the local sample sizes across different local machines should be approximately equal. The second condition is *randomness*, indicating that the massive data should be distributed across local computers as randomly as possible. The third condition is *sufficiency*, signifying that the sample size on each local machine should not be too small. To be more precise, it typically requires $n^2/N \rightarrow \infty$ unless some important biased reduction techniques (e.g., jackknifing) have been used (Wu et al., 2023c).

However, in real practice, these conditions are often violated to some extent. For example, practitioners rarely distribute large datasets in a completely uniform and random manner. Consequently, understanding how the violation of these conditions affects the statistical performance of the OS estimators becomes a topic of significant interest. Intuitively, when the uniformity condition is violated, at least one local computer ends up with a relatively

tiny sample size. Consequently, the local estimates generated by these local computers may exhibit significantly larger variability or bias than others. When the randomness condition is violated, the local estimates from different local computers could be seriously biased. When the sufficiency condition is violated, the bias of each local estimator $\hat{\theta}_{(m)}$ with an order $O(1/n)$ becomes non-negligible as compared with $O(1/\sqrt{N})$, leading to a noticeable bias in the resulting OS estimator $\hat{\theta}_{os}$. In each scenario, $\hat{\theta}_{os}$ becomes statistically inefficient or even inconsistent, as rigorously demonstrated by Wang et al. (2021).

To address the challenges posed by the lack of distribution uniformity and randomness, a one-step upgraded pilot (OSUP) estimator was proposed by Wang et al. (2021). The OSUP method comprises several steps and is well-suited for a broad class of models with a likelihood specification. Specifically, to compute the OSUP estimator, a number of n_0 pilot samples should be randomly selected from different local computers and then transferred to the central computer. Then the central computer computes a pilot estimator $\hat{\theta}_p$ by maximizing the log-likelihood function $\ell(X_i, Y_i; \theta)$ of all the pilot samples, i.e., $\hat{\theta}_p = \operatorname{argmax}_{\theta} \sum_{i \in \mathcal{P}} \ell(X_i, Y_i; \theta)$, where \mathcal{P} is the set of pilot samples. The pilot estimator $\hat{\theta}_p$ is $\sqrt{n_0}$ -consistent for the target parameter θ . However, due to its smaller sample size (i.e., $n_0 \ll N$), $\hat{\theta}_p$ is not statistically as efficient as the global estimator. To further improve the statistical efficiency, the central computer broadcasts $\hat{\theta}_p$ back to all local computers. Then each local computer considers $\hat{\theta}_p$ as an initial point and computes the 1st and 2nd order derivatives for its local log-likelihood function as $\dot{\ell}_m(\hat{\theta}_p) = \sum_{i \in \mathcal{S}_{(m)}} \dot{\ell}(X_i, Y_i; \hat{\theta}_p)$ and $\ddot{\ell}_m(\hat{\theta}_p) = \sum_{i \in \mathcal{S}_{(m)}} \ddot{\ell}(X_i, Y_i; \hat{\theta}_p)$. These derivatives are then communicated to the central computer for summation, i.e., $\dot{\ell}(\hat{\theta}_p) = \sum_m \dot{\ell}_m(\hat{\theta}_p)$ and $\ddot{\ell}(\hat{\theta}_p) = \sum_m \ddot{\ell}_m(\hat{\theta}_p)$. Based on the summarized derivative information along with the pilot estimate, a novel one-step upgrading is performed by the central computer. This leads to the final OSUP estimator as follows,

$$\hat{\theta}_{OSUP} = \hat{\theta}_p - \left\{ \frac{1}{N} \ddot{\ell}(\hat{\theta}_p) \right\}^{-1} \left\{ \frac{1}{N} \dot{\ell}(\hat{\theta}_p) \right\}.$$

Compared with a standard OS estimator, the OSUP estimator incurs an extra computational cost for obtaining the pilot sample. However, the benefits derived from the OSUP method are significant, leading to an estimator with the same statistical efficiency as the global estimator under very mild conditions (Wang et al., 2021). A similar one-step estimator is also studied by Huang and Huo (2019). However, the key difference is that the initial estimator used in Huang and Huo (2019) is the simple average of all local estimators.

To address the challenges posed by the lack of local sufficient sample size, Wu et al. (2023c) developed a jackknife debiased (JDS) estimator to reduce the estimation bias based on the moment estimator. It should be noted that this method was originally proposed for subsampling. However, the key idea is also readily applicable to distributed computing. To be specific, they first define a jackknife estimator $\hat{\theta}_{-j}^{(m)}$ for the m -th machine as

$$\hat{\theta}_{-j}^{(m)} = \operatorname{argmin}_{\theta} \frac{1}{n-1} \sum_{i \in \mathcal{S}_{(m)}^{i \neq j}} \ell(X_i, Y_i; \theta).$$

It could be verified that $\operatorname{Bias}(\hat{\theta}_{-j}^{(m)})$ approximately equals $\tau/(n-1)$ for some constant τ (Shao & Tu, 1995). Then, $n^{-1} \sum_{j \in \mathcal{S}_m} \operatorname{Bias}(\hat{\theta}_{-j}^{(m)}) \approx \tau/(n-1)$ and $E(n^{-1} \sum_{j \in \mathcal{S}_m} \hat{\theta}_{-j}^{(m)} - \hat{\theta}^{(m)}) \approx \tau/\{n(n-1)\}$. This inspires an estimator for the bias, which is given by $\widehat{\operatorname{Bias}}^{(m)} = (n-1)n^{-1} \sum_{j \in \mathcal{S}_m} \hat{\theta}_{-j}^{(m)} - (n-1)\hat{\theta}^{(m)}$. Accordingly, a bias-corrected estimator for the m -th machine could be proposed as $\hat{\theta}_{JDS}^{(m)} = \hat{\theta}^{(m)} - \widehat{\operatorname{Bias}}^{(m)}$. Thereafter, $\hat{\theta}_{JDS}^{(m)}$ s can be further averaged across different m . As a consequence, the final JDS estimator could be obtained as $\hat{\theta}_{JDS} = M^{-1} \sum_{m=1}^M \hat{\theta}_{JDS}^{(m)}$. Subsequently, Wu et al. (2023c) rigorously verified that $\operatorname{Bias}(\hat{\theta}_{JDS}) = O(1/n^2) + O(1/N)$ and the asymptotic variance of $\hat{\theta}_{JDS}$ remains the same as that of $\hat{\theta}$. As a consequence, excellent statistical efficiency can be achieved by $\hat{\theta}_{JDS}$ with a very small size n . This bias correction method has been theoretically studied for moment estimator; however, the theoretical properties for the estimator computed from a general loss function remain unknown.

2.3. Efficient iterative approach

To improve the statistical efficiency of one-shot estimators, various distributed iterative methods can be considered. Since distributed computing requires passing messages among multiple computers, naively applying traditional iterative methods to a distributed system often incurs expensive communication costs. Therefore, how to achieve excellent statistical efficiency with well-controlled communication costs becomes the key issue (Jordan et al., 2019).

To illustrate this point, consider for example extending the iterative Newton-Raphson algorithm (1) to the distributed scenario. Recall that $\mathcal{S}_{(m)}$ collects the indices of samples allocated to the m -th local computer. Given the

current estimator $\widehat{\theta}^{(t)}$, we can compute the 1st and 2nd order derivatives of the loss function as

$$\dot{\mathcal{L}}(\widehat{\theta}^{(t)}) = M^{-1} \sum_{m=1}^M \dot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t)}) \quad \text{and} \quad \ddot{\mathcal{L}}(\widehat{\theta}^{(t)}) = M^{-1} \sum_{m=1}^M \ddot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t)}),$$

where $\dot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t)}) = \sum_{i \in \mathcal{S}_{(m)}} \dot{\ell}(X_i, Y_i; \widehat{\theta}^{(t)})$ and $\ddot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t)}) = \sum_{i \in \mathcal{S}_{(m)}} \ddot{\ell}(X_i, Y_i; \widehat{\theta}^{(t)})$. Note that $\dot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t)})$ and $\ddot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t)})$ are computed on the m -th local computer. They are then transferred to the central computer to update $\widehat{\theta}^{(t+1)}$ according to (1). As one can see, this is a solution easy to implement but suffers several serious limitations. First, inverting the $(p \times p)$ -dimensional Hessian matrix $\ddot{\mathcal{L}}(\widehat{\theta}^{(t)})$ in the central computer incurs a computation cost with the order $O(p^3)$ for each iteration. Second, transferring the local Hessian matrix $\ddot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t)})$ from each local computer to the central computer incurs a communication cost of order $O(p^2)$ for each local computer in each iteration. Thus, this approach leads to high computation and communication costs for high-dimensional data.

To address this issue, various communication-efficient Newton-type methods have been proposed to alleviate high communication costs. One of the underlying key ideas is to avoid Hessian matrix transmission (Crane & Roosta, 2019; Jordan et al., 2019; Luo & Song, 2020; Shamir et al., 2014; Wang et al., 2017; Wang, Roosta et al., 2018; Zhang & Lin, 2015). For example, the entire sample Hessian matrix can be approximated using some appropriate local estimators, which are computed on one single computer (e.g., the central computer). Consequently, the communication cost due to transferring the whole sample Hessian matrix between computers can be avoided. One notable work in this regard is Jordan et al. (2019). Motivated from the Taylor series expansion of $\mathcal{L}(\theta)$, Jordan et al. (2019) defined a surrogate loss function as $\widetilde{\mathcal{L}}(\theta) = \mathcal{L}_{(1)}(\theta) - \theta^\top \{\dot{\mathcal{L}}_{(1)}(\bar{\theta}) - \dot{\mathcal{L}}(\bar{\theta})\}$, where $\bar{\theta}$ denotes any initial estimator of θ . Then, based on the surrogate loss function, the updating formula of Newton's method is modified as $\widehat{\theta}^{(t+1)} = \widehat{\theta}^{(t)} - \{\ddot{\mathcal{L}}_{(1)}(\widehat{\theta}^{(t)})\}^{-1} \dot{\mathcal{L}}(\widehat{\theta}^{(t)})$, where $\ddot{\mathcal{L}}_{(1)}(\widehat{\theta}^{(t)})$ is the local Hessian matrix computed on the 1st local computer. Thus no Hessian matrix communication is needed for each iteration.

The implementation of communication-efficient Newton-type methods significantly reduces communication costs. However, when dealing with high-dimensional data, computing the inverse of the Hessian matrix remains to be a computationally expensive problem. To further improve computation efficiency, various methods avoiding matrix inverse calculation have been proposed. The first type is distributed (stochastic) gradient descent algorithms (Chen, Liu et al., 2022; Goyal et al., 2017; Li et al., 2022; Lin & Zhou, 2018; Qu & Li, 2019; Su & Xu, 2019), which compute only the 1st order derivatives of the loss function (i.e., gradients). To be specific, in distributed gradient descent methods, each local computer first receives the current parameter estimator from the central computer. Once the current parameter estimator is received, each local computer calculates its own gradient and sends it back to the central computer. Lastly, the central computer aggregates the local gradients and updates the parameter estimator as

$$\widehat{\theta}^{(t+1)} = \widehat{\theta}^{(t)} - \alpha_t M^{-1} \sum_{m=1}^M \dot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t)}), \quad (2)$$

where $\alpha_t > 0$ represents the learning rate. To further reduce the communication cost in (2), a local (stochastic) gradient descent algorithm is proposed (Stich, 2019; Woodworth et al., 2020). The key idea is to run the (stochastic) gradient descent algorithm independently and locally on different local computers in a fully parallel way. Subsequently, the local estimators are transferred to the central computer and then iteratively updated to form the final estimator.

However, for those distributed gradient descent algorithms, a large number of iterations are typically required for numerical convergence, and the choice of hyperparameters (e.g., α_t) is critically important and also difficult (Zhu, Huang et al., 2021). To address this problem, various quasi-Newton methods in a distributed manner have been developed (Chen et al., 2014; Eisen et al., 2017; Lee et al., 2018; Soori et al., 2020; Wu et al., 2023b). The key idea of distributed quasi-Newton methods is to approximate the Hessian inverse in each iteration without actually inverting the matrix (Davidon, 1991; Goldfarb, 1970). The communication cost of these methods could have orders as low as $O(p)$ in each iteration. In the meanwhile, the convergence rate of distributed quasi-Newton methods is superlinear, surpassing the linear convergence of distributed gradient descent methods (Broyden et al., 1973).

As an important method along this direction, Wu et al. (2023b) developed a K -stage distributed quasi-Newton method. Specifically, Wu et al. (2023b) started with the following one-stage distributed quasi-Newton method

$$\widehat{\theta}_{\text{stage},1} = \widehat{\theta}_{\text{stage},0} - M^{-1} \sum_{m=1}^M \left\{ H_{(m,0)} \dot{\mathcal{L}}(\widehat{\theta}_{\text{stage},0}) \right\}. \quad (3)$$

Here $\widehat{\theta}_{\text{stage},0}$ is a suitable initial estimator, such as the OS estimator $\widehat{\theta}_{\text{os}}$. $H_{(m,0)}$ represents the local inverse Hessian estimator obtained by each local computer after implementing the quasi-Newton algorithm. Next, the local

computer sends $H_{(m,0)} \dot{\mathcal{L}}(\hat{\theta}_{\text{stage},0})$ as a whole to the central computer. Therefore, the communication cost is of the order $O(p)$. Wu et al. (2023b) has verified that the optimal statistical efficiency can be achieved by the one-stage distributed quasi-Newton estimator as long as $N(\log p)^4/n^4 \rightarrow 0$. This condition can be further relaxed by allowing for multi-stage updating. This leads to the K -stage distributed quasi-Newton estimator, which is extremely efficient both computationally and communicationally.

2.4. Distributed quantile regression

Quantile regression is an important class of regression methods for its robustness against heavy-tailed distributed responses and outliers (Koenker, 2005; Koenker & Bassett, 1978). Its applications span various disciplines, including agriculture (Kostov & Davidova, 2013), climate change (Reich et al., 2012), health studies (Alhamzawi & Ali, 2018), house pricing (Chen et al., 2013), and others (Xu et al., 2017; Zhong et al., 2022). With the availability of large-scale datasets, extensive research has been dedicated to distributed estimation and inference for quantile regression. For instance, Yang et al. (2013) proposed a subspace preserving sampling technique for quantile regression on massive datasets. However, their method suffers from the problem of statistical inefficiency. Later, Xu et al. (2020) developed a block average approach, employing the one-shot strategy by averaging estimators derived from each local computer.

To guarantee statistical efficiency, researchers address the challenges of distributed quantile regression by proposing various loss functions and iterative algorithms. For instance, Volgushev et al. (2019) proposed a two-step quantile projection algorithm, incorporating valid statistical inference. In the initial step, conditional quantile functions are estimated at different levels. Subsequently, a quantile regression process is constructed through projection. Chen et al. (2019) presented a computationally efficient method, which involves multiple rounds of aggregations. After limited q iterations, the authors show that the statistical efficiency of the final estimator becomes the same as the one computed on the whole data. In a related work, Chen, Liu et al. (2020) studied the linear regression problem with heavy-tailed noises. Since the quantile regression loss function is a non-smooth function, the authors established a connection between quantile regression and ordinary linear regression by transforming the response. This results in a distributed estimator that is efficient in both computation and communication. Instead of dealing with conventional smoothing functions, Hu et al. (2021) extended the communication-efficient surrogate likelihood method proposed by Jordan et al. (2019). The authors constructed a surrogate loss function and established the consistency and asymptotic normality of the proposed methodology. In a recent development, Tan et al. (2022) developed a double-smoothing approach to the local and global objective functions of quantile regression.

In a recent work of Pan et al. (2022), the authors proposed a one-step approach that is efficient both communicationally and statistically. Notably, the derived estimator is robust against data distribution heterogeneity across local computers. Specifically, assume there are N observations indexed by $i = 1, \dots, N$. The response Y_i and the p -dimensional predictor X_i follow the standard τ -th quantile regression model $Y_i = X_i^\top \beta_\tau + \varepsilon_i$, where β_τ is the associated regression coefficient vector and $\tau \in (0, 1)$. Additionally, ε_i is the error term satisfying $P(\varepsilon_i \leq 0 | X_i) = \tau$. The standard check loss function can be constructed as $\mathcal{L}(\beta) = \sum_{i=1}^N \rho_\tau(Y_i - X_i^\top \beta)$, where $\rho_\tau(\mu) = \mu\{\tau - I(\mu \leq 0)\}$ represents the check function and $I(\cdot)$ the indicator function. Consequently, the standard estimator of β_τ can be obtained by $\hat{\beta}_\tau = \arg \min_\beta \mathcal{L}(\beta)$. Assume that a pilot sample (i.e., indexed by \mathcal{Q}) is derived across K local computers. The sample size of the pilot sample is n satisfying $n/N \rightarrow 0$. As a result, a pilot estimator can be obtained by $\hat{\beta}_\tau^{\mathcal{Q}} = \arg \min_\beta \sum_{i \in \mathcal{Q}} \rho_\tau(Y_i - X_i^\top \beta)$, which is \sqrt{n} -consistent. The one-step estimator proposed by Pan et al. (2022) is derived as

$$\hat{\beta}_\tau^{(1)} = \hat{\beta}_\tau^{\mathcal{Q}} + \frac{1}{\hat{f}(0)} \left(\sum_{i=1}^N X_i X_i^\top \right)^{-1} \left[\sum_{i=1}^N X_i \left\{ \tau - I(\hat{\varepsilon}_i \leq 0) \right\} \right], \quad (4)$$

where $\hat{f}(\cdot)$ is a kernel density estimator and $\hat{\varepsilon}_i$ is the residual. It can be proved that $\hat{\beta}_\tau^{(1)}$ is \sqrt{N} -consistent and asymptotically normal, regardless of how the raw data are distributed across local computers.

2.5. Distributed logistic regression with rare events data

The rare events problem in this subsection refers to a binary data classification problem, where one class (often assumed to be the negative class) has a much greater number of instances than the other class (often assumed to be the positive class). Rare events data are prevalent in scientific fields and applications. The rare events data examples include but are not limited to drug discovery (Korkmaz, 2020; Zhu et al., 2006), software defects (Richardson & Lidbury, 2013), and rare disease diagnosis (Zhao et al., 2018; Zhuang et al., 2019). In traditional statistical theory, one often assumes that the probability of any type of event to happen is fixed. However, for rare events, it is more

appropriate to assume that the positive class probability should decay towards zero at an appropriate rate as the total sample size increases (Wang, 2020). In this regard, Wang (2020) constructed a novel theoretical framework to accurately describe rare events data. Under this theoretical framework, it was demonstrated that the convergence rate of the global maximum likelihood estimator (MLE) is mainly determined by the sample size of the positive class instead of the total sample size. This implies a considerably slower convergence rate than that of the classical cases. It seems that limited attempts have been made for the rigorous asymptotic theory for distributed classification problems with rare event data. This motivates (Li et al., 2023) to develop a novel distributed logistic regression method with solid statistical theory support for massive rare event data.

More specifically, assume there are a total of N observations indexed by $1 \leq i \leq N$. The i -th observation is denoted as (X_i, Y_i) , where $X_i \in \mathbb{R}^p$ is a p -dimensional covariate and $Y_i \in \{0, 1\}$ is the binary response. Let $N_1 = \sum_{i=1}^N Y_i$ be the total number of positive instances. To model their regression relationship, the following logistic regression model is considered

$$P(Y_i = 1 | X_i) = p_i(\alpha, \beta) = \frac{\exp(\alpha + X_i^\top \beta)}{1 + \exp(\alpha + X_i^\top \beta)},$$

where $\alpha \in \mathbb{R}$ is the intercept and $\beta \in \mathbb{R}^p$ is the slope parameter. To reflect the asymptotic behaviour, two important assumptions should be imposed (Wang, 2020). First, the percentage of positive instances should be extremely small. Statistically, the positive response probability is specified to converge towards 0 as the total sample size $N \rightarrow \infty$. Rewrite α as α_N . This leads to $\alpha_N \rightarrow -\infty$ as $N \rightarrow \infty$. Second, the total number of positive instances should diverge to infinity. Otherwise, the parameters of interest cannot be estimated consistently. Mathematically, it follows that $E(N_1) \approx N \exp(\alpha_N) E\{\exp(X_i^\top \beta)\}$ as $N \rightarrow \infty$. This suggests that $\alpha_N \rightarrow -\infty$ and $\alpha_N + \log N \rightarrow \infty$ as $N \rightarrow \infty$ (Li et al., 2023; Wang, 2020).

Assume a distributed computation system with one central computer and a total of K local computers indexed by $1 \leq k \leq K$. Li et al. (2023) first discussed two different data distribution strategies. They are RANDOM and COPY strategies, respectively. Specifically, the RANDOM strategy is used to randomly distribute the full data to each local computer with approximately equal sizes. For the COPY strategy, all the positive instances are copied to every local computer. In contrast, the negative instances are randomly distributed on different local computers. Next, Li et al. (2023) studied three types of objective functions. They are, respectively,

$$\begin{aligned} \mathcal{L}_{R,k}(\theta) &= \sum_{i=1}^N a_i^{(k)} \left\{ Y_i \log p_i(\alpha_N, \beta) + (1 - Y_i) \log (1 - p_i(\alpha_N, \beta)) \right\}, \\ \mathcal{L}_{US,k}(\theta) &= \sum_{i=1}^N \left\{ Y_i \log p_i(\alpha_N, \beta) + (1 - Y_i) a_i^{(k)} \log (1 - p_i(\alpha_N, \beta)) \right\}, \\ \mathcal{L}_{IPW,k}(\theta) &= \sum_{i=1}^N \left\{ Y_i \log p_i(\alpha_N, \beta) + K(1 - Y_i) a_i^{(k)} \log (1 - p_i(\alpha_N, \beta)) \right\}. \end{aligned}$$

Here $P(a_i^{(k)} = 1) = 1/K$ and $a_i^{(k)} = 1$ if the i -th observation is randomly distributed to k -th local computer. Local estimators can be computed based on the local data according to different loss functions. The local estimators are then averaged by the central computer in the last step (Chang et al., 2017; Zhang et al., 2012). This leads to three distributed estimators (i.e., $\hat{\theta}_{RMLE}$, $\hat{\theta}_{US}$ and $\hat{\theta}_{IPW}$). Under some regularity conditions, Li et al. (2023) found that the COPY strategy together with the inverse probability weighted objective function $\mathcal{L}_{IPW,k}(\theta)$ seems to be the best choice.

2.6. Decentralized distributed computing

The distributed computing methods outlined earlier share a common characteristic. That is the requirement of a central computer, which is responsible for communicating with every local computer. Such type of architecture is easy to implement. However, it suffers from several serious limitations. First, this centralized network structure is extremely fragile. If the central computer stops working, the entire network stops. Second, there exists the issue of privacy disclosure for the centralized structure. This is because if the central machine is attacked, the attacker is given the chance to communicate with every local computer. Third, a centralized network structure has a high requirement for network bandwidth, since the central machine should communicate with numerous local computers (Bellet et al., 2018; Li et al., 2021).

To fix these problems, a number of researchers advocate the idea of fully decentralized distributed computing, which is also called decentralized federated learning (DFL) (Yuan et al., 2016). The key feature of DFL is that there

is no central computer involved for model training and communication. Different local computers are directly connected through a sophisticated communication network. All computation-related communications should occur only between network-connected individual local computers. To be specific, define $\hat{\theta}^{(t,m)}$ to be the t -th estimator obtained on the m -th local computer, and the updating formula of the decentralized federated learning algorithm is given by

$$\hat{\theta}^{(t+1,m)} = \tilde{\theta}^{(t,m)} - \alpha \dot{\mathcal{L}}_{(m)} \left(\tilde{\theta}^{(t,m)} \right). \quad (5)$$

Here, $\tilde{\theta}^{(t,m)}$ is the neighbourhood-averaged estimator obtained in the t -th iteration for the m -th local computer. Numerous studies have investigated the numerical convergence properties of the method (Blot et al., 2016; Lalitha et al., 2018; Lian et al., 2017; Nedic et al., 2017; Tang et al., 2018; Vanhaesebrouck et al., 2017). It has been demonstrated that the algorithm can achieve a linear convergence rate even with data heterogeneity (Richards et al., 2020; Savazzi et al., 2020). To achieve this nice theoretical property, several stringent conditions have been assumed about the network structure in the past literature. The most typical assumption is that the transition matrix determined by the network structure should be doubly stochastic (Tang et al., 2018; Yuan et al., 2016). Unfortunately, this is an assumption that can hardly be satisfied in real practice.

To relax this stringent condition, Wu et al. (2023a) developed a novel methodology for DFL, which only requires the network structure to be weakly balanced. To be specific, take a linear regression as an example. Assume there are a total of M local computers, which are connected by a communication network. The adjacency matrix of the network is defined as $A = (a_{ij}) \in \mathbb{R}^{M \times M}$, and the corresponding weighting matrix is defined as $W = (w_{ij}) \in \mathbb{R}^{M \times M}$ with $w_{ij} = a_{ij}/d_m$, where $d_m = \sum_j a_{ij}$ represents the in-degree of A . Then algorithm (5) could be rewritten as

$$\hat{\theta}^{(t+1,m)} = \left(I_p - \alpha \hat{\Sigma}_{xx}^{(m)} \right) \left(\sum_{k=1}^M w_{mk} \hat{\theta}^{(t,k)} \right) + \alpha \hat{\Sigma}_{xy}^{(m)}, \quad (6)$$

where $I_p \in \mathbb{R}^{p \times p}$ is an identity matrix, $\hat{\Sigma}_{xx}^{(m)} = \sum_{i \in \mathcal{S}(m)} X_i X_i^\top / n \in \mathbb{R}^{p \times p}$, and $\hat{\Sigma}_{xy}^{(m)} = \sum_{i \in \mathcal{S}(m)} X_i Y_i / n \in \mathbb{R}^p$. Next, define $\hat{\theta}^{*(t)} = (\hat{\theta}^{(t,1)\top}, \dots, \hat{\theta}^{(t,M)\top})^\top \in \mathbb{R}^{Mp}$, $\hat{\Sigma}_{xy}^* = (\hat{\Sigma}_{xy}^{(1)\top}, \dots, \hat{\Sigma}_{xy}^{(M)\top})^\top \in \mathbb{R}^{Mp}$, and $\Delta^* = \text{diag}\{I_p - \alpha \hat{\Sigma}_{xx}^{(1)}, \dots, I_p - \alpha \hat{\Sigma}_{xx}^{(M)}\}$. This leads to a matrix form of (6) as

$$\hat{\theta}^{*(t+1)} = \Delta^* (W \otimes I_p) \hat{\theta}^{*(t)} + \alpha \hat{\Sigma}_{xy}^*,$$

where \otimes stands for the Kronecker product. Assume the stable solution of this system (denoted by $\hat{\theta}^*$) exists, it follows then $\hat{\theta}^* = \alpha \{I_p - \Delta^* (W \otimes I_p)\}^{-1} \hat{\Sigma}_{xy}^*$. Theoretically, Wu et al. (2023a) proved that the statistical efficiency of the DFL is determined by three factors: (1) the learning rate, (2) the network structure, and (3) the data distribution pattern. The optimal statistical efficiency can be guaranteed if the learning rate is relatively small and the network structure is relatively balanced, even if data are distributed heterogeneously.

There has been extensive research to further extend the classical DFL algorithm from different perspectives. To address the issue of insufficient labels, Gao et al. (2019) developed a heterogeneous horizontal federated learning framework. To enhance network security, Chen, Yu et al. (2022) introduced a decentralized federated learning algorithm that meets differential privacy requirements. The key idea is that each client adds random noise to their parameter estimators before communication. The classical DFL algorithm suffers from high communication costs and low convergence rates in non-convex situations. To fix this problem, Nadiradze et al. (2021) proposed an asynchronous decentralized federated learning method. A novel DFL framework was developed by Liu, Chen et al. (2022), which optimally balances communication efficiency and statistical efficiency. Liu, Mao et al. (2022) further proposed a decentralized surrogate median regression method for non-smooth sparse problems. For valid statistical inference in DFL, Gu and Chen (2023) studied communication-efficient M -estimation. To achieve optimal efficiency, they proposed a one-step DFL estimation method that allows a relatively large number of clients.

2.7. Distributed statistical inference

In addition to estimation, other statistical inference tools, such as hypothesis testing and confidence intervals, also play a crucial role in scientific research and data analysis. These inference tools allow researchers to quantify the uncertainty of the estimators obtained from the data, and then help practitioners interpret the results appropriately (Casella & Berger, 2002). In the above discussion of this section, various distributed estimation methods have been introduced. Most of them have proven that the distributed estimator can be statistically as efficient as the global

estimator under certain conditions (Fan et al., 2023; Jordan et al., 2019; Pan et al., 2022; Volgushev et al., 2019; Wang et al., 2017, 2019a; Zhu, Li et al., 2021). Consequently, the distributed estimators generally share the same asymptotic distribution as the global estimator. This means that if one can consistently estimate the asymptotic covariance matrix, then asymptotically valid statistical inference can be directly conducted. Then the key issue becomes how to estimate the asymptotic covariance matrix consistently and distributedly. To this end, various plug-in approaches have been widely adopted for various models. For example, general M -estimation problems with smoothed loss (Jordan et al., 2019), quantile regression models (Pan et al., 2022), support vector machine (Wang et al., 2019a), and various debiased estimators for high-dimensional models (Fan et al., 2019; Tu et al., 2023).

However, if the asymptotic covariance of an estimator is too complex, it can be very challenging to construct the corresponding estimator analytically. In this case, bootstrap provides a more directed inference approach (Efron & Stein, 1981; Shao & Tu, 1995). Nevertheless, bootstrap usually requires multiple resampling procedures over the whole dataset. This is computationally too expensive to be acceptable, especially for massive datasets. To solve this problem, researchers developed some more computationally feasible bootstrap methods. Among them, Kleiner et al. (2014) introduced a method called the bag of little bootstraps (BLB). The BLB method first divides the whole sample into multiple subsets. It then computes multiple repeated estimates of the estimator (or related statistics) based on the inflated resamples. Finally, an averaging step is applied to aggregate these estimates. Taking a similar approach, Sengupta et al. (2016) further proposed the method of subsampled double bootstrap (SDB). Instead of directly dividing the whole sample into the disjoint subsets, the SDB method selects multiple subsets from the whole dataset by sampling with replacement. The computational efficiency and inferential reliability of these subsample-based bootstrap methods depend on various hyperparameters, such as the subsample size and the number of replicates. To address this issue, Ma et al. (2024) developed an interesting approach for selecting the optimal hyperparameters for the subsample-based bootstrap methods, including the BLB and the SDB.

However, the above bootstrap variants generally require that the involved estimator (or statistic) has a weighted subsample representation. This may not be true for some general statistics. For example, the class of symmetric statistics considered in Chen and Peng (2021), which includes the U -statistics as an important example. Chen and Peng (2021) investigated the theoretical properties of the one-shot averaging type distributed statistics for both the degenerate and non-degenerate cases. For inference purposes, they developed a distributed bootstrap procedure, where no technique such as inflated resampling technique in the BLB method of Kleiner et al. (2014) is required. To further ease the computational burden, Chen and Peng (2021) proposed a pseudo-distributed bootstrap (PDB) procedure. The consistency of the PDB procedure has also been theoretically proved and numerically validated.

3. Subsampling models

3.1. Sequential addressing subsampling

Note that subsampling technique is closely related to the idea of bootstrap (Bickel & Freedman, 1981; Efron, 1979). However, the classical full size bootstrap is often computationally too expensive for massive data analysis. A practical solution in this regard is to repeatedly generate subsamples of small sizes for parameter estimation and statistical inference. Consequently, various subsampling methods are proposed. These methods include, but are not limited to, the m out of n bootstrap (Bickel et al., 1997), the bag of little bootstrap (Kleiner et al., 2014), the subsampling double bootstrap (Sengupta et al., 2016), the distributed bootstrap (Chen & Peng, 2021), the optimal subsampling bootstrap (Ma et al., 2024), and possibly others. These methods are particularly useful for the situation, where the data is small enough to be comfortably placed on one single hard drive but large enough so that it cannot be fully loaded into the computer memory as a whole.

When computational resources are limited, an alternative way of subsampling is to comprehensively utilize both the computer memory and the hard drive. Some early literature has proposed out-of-core sampling methods, which obtain samples by randomly accessing data points on the hard drive without loading the whole data file in advance (Li, 1994; Vitter, 1985). However, due to the hardware limitations at that time, such out-of-core sampling methods have not been tested on massive datasets. When dealing with massive datasets, the time cost is of the most critical concern. The time required to sample a single data point from the hard drive often exceeds that of in-memory sampling (Suwandarathna & Koggalage, 2007). This time cost comprises two main components. They are, respectively, the *addressing cost* associated with identifying the target data point on the hard drive and the *I/O cost* associated with reading the target data point into memory. It is then referred to as the *hard drive sampling cost* (HDSC), representing the time needed to fetch a specific data point from the hard disk into computer memory (Pan et al., 2023). To reduce the HDSC for massive data, Pan et al. (2023) developed a computationally efficient method known as *sequential addressing subsampling* (SAS). This method involves a two-step process. They are, respectively, a random shuffling operation aiming at randomly sorting the raw data and a sequential sampling step for obtaining the desired

subsamples. It is noteworthy that the random addressing operation, a crucial component of obtaining a subsample, is only performed once. The subsample obtained through this process is referred to as the SAS subsample, and various statistics can be constructed and theoretically studied using these SAS subsamples.

The subsampling method in Pan et al. (2023) provides a promising solution to accelerate the subsampling process on the hard drive. Moreover, their SAS method has been proven to be a robust tool for making statistical inference on massive datasets. Consider the sample mean as a concrete example to illustrate the theoretical findings based on the SAS method. Assume a set of N samples represented by X_1, \dots, X_N with mean μ and variance σ^2 . Additionally, assume that $E(X_i - \mu)^4 = \gamma \sigma^4$. Let $\{X_k, X_{k+1}, \dots, X_{k+n-1}\}$ be the k -th subsample with a sample size of n , where the sample mean is defined as $\bar{X}_k = n^{-1} \sum_{i=k}^{k+n-1} X_i$. In practice, assume that B sequential subsamples are obtained using the SAS method. Denote the corresponding sample means of these B sequential subsamples as $\{\bar{X}_{(1)}, \dots, \bar{X}_{(b)}, \dots, \bar{X}_{(B)}\}$. Define the sample mean of interest as $\bar{\bar{X}}_B = B^{-1} \sum_{b=1}^B \bar{X}_{(b)}$. It is easy to see that $E(\bar{\bar{X}}_B) = \mu$. Given the assumptions that (1) $n \rightarrow \infty$ and $n/N \rightarrow 0$ as $N \rightarrow \infty$; (2) $B/N \rightarrow 0$ and $nB = O(N)$ as $N \rightarrow \infty$, the variance of $\bar{\bar{X}}_B$ can be presented as

$$\text{var}(\bar{\bar{X}}_B) = \sigma^2 \left(\frac{1}{nB} + \frac{1}{N} \right) \{1 + o(1)\}.$$

On one hand, the term σ^2/N is associated with the overall sample and cannot be eliminated by subsampling. On the other hand, the term $\sigma^2/(nB)$ can be reduced by increasing the subsample size n or the number of subsamples B . To perform automatic inference, the standard error of $\bar{\bar{X}}_B$ is proposed as

$$\widehat{\text{SE}}^2(\bar{\bar{X}}_B) = \frac{n}{B-1} \left(\frac{1}{nB} + \frac{1}{N} \right) \sum_{b=1}^B (\bar{X}_{(b)} - \bar{\bar{X}}_B)^2,$$

where the theoretical results can be found in Pan et al. (2023). In summary, the SAS method is time-saving in terms of HDSC as well as useful for automatic statistical inferences.

3.2. Subsampling-based estimation methods

In the previous subsection, we have thoroughly reviewed various subsampling methods and discussed them from the perspective of computational efficiency. In the meanwhile, how to estimate the parameters of interest with the best statistical efficiency is also a crucial concern. To formulate this problem, let $\mathcal{F} = \{i : 1 \leq i \leq N\}$ represent an index set for an extremely large dataset with sample size N . Let Y_i be the response associated with the i -th subject and $X_i = (X_{ij}) \in \mathbb{R}^p$ be the corresponding p -dimensional feature vector. Define π_i to be the sampling probability for each sample $1 \leq i \leq N$. A random subsample of size n can be drawn (with replacement) with $\pi = (\pi_1, \dots, \pi_N)^\top \in \mathbb{R}^N$. Then the key research question here is how to define π_i s appropriately, so that a small but representative subsample can be obtained for downstream statistical models.

Various subsampling methods have been proposed to address this concern. For the linear regression model, the algorithm leveraging methods have been extensively discussed, utilizing the empirical statistical leverage scores of the input covariate matrix to define the sampling probabilities (Drineas et al., 2006, 2011; Mahoney, 2011). Ma et al. (2014) further provided an effective framework to evaluate the statistical properties of parameter estimation in these algorithmic leveraging methods. Wang et al. (2019b) introduced the information-based optimal subdata selection (IBOSS) method, which can deterministically identify a subsample with the maximum information matrix under the D -optimality criterion. The IBOSS approach is further extended to a divide-and-conquer setting by Wang (2019a). For binary logistic regression, Wang, Zhu et al. (2018) proposed the optimal subsampling method motivated by the A -optimality criterion (OSMAC) by minimizing the asymptotic mean squared error (MSE) of the subsample estimator to design the subsampling probability. The OSMAC method can be enhanced by incorporating unweighted objective functions and Poisson subsampling, resulting in improved efficiency (Wang, 2019b). Furthermore, the applicability of OSMAC is extended to various classes of models, including multi-class logistic regression (Yao & Wang, 2019), generalized linear models (Ai et al., 2021), quantile regression (Wang & Ma, 2021) and quasi-likelihood (Yu et al., 2022).

Although the above methods are demonstrated to be statistically efficient, computing π_i for the entire dataset poses a substantial computational challenge, especially when the data size N is extremely large. Take the OSMAC method (Wang, Zhu et al., 2018) as an example. For the OSMAC method, determining the optimal subsampling probabilities involves a computational complexity of $O(Np)$. As a result, this optimal subsampling algorithm becomes computationally expensive when dealing with a very large sample size N . To address this challenge, the repeated subsampling method can be adopted. The key idea of repeated subsampling is to draw a subsample with

uniform probability (i.e., $\pi_i = 1/N$) while operating the subsampling step repeatedly. Using the uniform probability in the subsampling process eliminates the need to compute probabilities for the entire dataset in advance. Therefore the computation cost is significantly reduced. Through repeated subsampling, the selected data approximates the whole dataset. Note that the cost associated with subsampling cannot be negligible, particularly true for the repeated subsampling methods. Nevertheless, thanks to the SAS method of blackPan et al. (2023), the hard drive sampling cost can be significantly reduced.

Based on repeated subsampling, a variety of statistical models have been developed. Here we introduce the sequential one-step (SOS) estimator for generalized linear models (Wang et al., 2022) for example. Assume the whole dataset has been randomly distributed on the hard drive and the SAS method is used to obtain each sub-data. Assume the subsampling is repeated for K times. In the k -th subsampling with $1 \leq k \leq K$, denote \mathcal{S}_k to be the indices of selected observations in the whole dataset. Based on \mathcal{S}_k , the SOS estimator is computed as follows. First, we need to calculate an initial estimator $\bar{\beta}_1$ based on \mathcal{S}_1 . This initial estimator could, for instance, be a maximum likelihood estimator (MLE) of the generalized linear regression models. Assume $\bar{\beta}_k$ to be the current estimator in the k -th step. Subsequently, in the $(k+1)$ -th subsampling step, a new SAS subsample \mathcal{S}_{k+1} is obtained. Then a one-step update is performed based on $\bar{\beta}_k$ to obtain the one-step updated estimator, i.e., $\hat{\beta}_{k+1} = \bar{\beta}_k - \{\dot{\ell}_{\mathcal{S}_{k+1}}(\bar{\beta}_k)\}^{-1} \dot{\ell}_{\mathcal{S}_{k+1}}(\bar{\beta}_k)$, where $\dot{\ell}_{\mathcal{S}_{k+1}}(\bar{\beta}_k)$ and $\ddot{\ell}_{\mathcal{S}_{k+1}}(\bar{\beta}_k)$ denote the 1st and 2nd order derivatives of the likelihood function based on the $(k+1)$ -th subsample, respectively. Next, the SOS subsampling estimator for the $(k+1)$ -th step is computed as $\bar{\beta}_{k+1} = \{k\bar{\beta}_k + \hat{\beta}_{k+1}\}/(k+1) = \sum_{l=1}^{k+1} \hat{\beta}_l/(k+1)$. The corresponding estimator obtained in the last K -th step is the final SOS estimator, i.e., $\hat{\beta}^{\text{SOS}} = \bar{\beta}_K$. It is noteworthy that the SOS method represents an extension of the classical one-step estimator black(Shao, 2003; Zou & Li, 2008) but within the context of subsampling. The theoretical properties of the SOS estimator are also established in Wang et al. (2022), demonstrating that both the bias and variance of the SOS estimator decrease as the number of sampling iterations K increases.

3.3. Subsample feature screening

Feature screening plays a critically important role for ultrahigh dimensional data analysis. Extensive literature has been developed along this direction. Since Fan and Lv (2008) introduced the seminal work of sure independence screening (SIS), a large amount of follow-up research has been inspired. The key idea of SIS is to rank and then select important features by certain appropriately defined correlation measures. For example, under a linear regression model setup and assuming some appropriate regularity conditions, Fan and Lv (2008) showed that the top features selected according to marginal sample correlation coefficients are screening consistent. In other words, the selected top features are assured to asymptotically cover the underlying low dimensional true model with probability tending to one. Wang (2009) further improved SIS by the method of forward regression for a significantly improved finite-sample performance. Li et al. (2012) proposed a distance correlation based independent screening method (DC-SIS) so that variable screening can be conducted in a model free manner. Recently, a distributed feature selection method has been developed by Li et al. (2020) for massive data analysis.

Zhu et al. (2022) proposed a novel subsampling-based feature selection method for large datasets with ultrahigh dimensional features. They consider a classical linear regression model as $Y_i = X_i^\top \beta + \varepsilon_i$ (Fan & Lv, 2008), where $\beta \in \mathbb{R}^p$ is regression parameters, ε_i is the independent noise term with $\text{var}(\varepsilon_i) = \sigma^2$. Assume a total of B subsamples with size n , which are denoted by $\mathcal{S}_{(b)} \subset \mathcal{S}_F = \{1, 2, \dots, N\}$ with $|\mathcal{S}_{(b)}| = n$. Define $\mathbb{X}_{(b)} = (X_i : i \in \mathcal{S}_{(b)}) \in \mathbb{R}^{n \times p}$ as the subsampled design matrix and $\mathbb{Y}_{(b)} = (Y_i : i \in \mathcal{S}_{(b)}) \in \mathbb{R}^n$ as the associate response vector. Define a candidate model as $\mathcal{M} = \{j_1, \dots, j_m\}$ with $1 \leq j_k \leq p$ for every $1 \leq k \leq K$. Define the design matrix associate with \mathcal{M} as $\mathbb{X}_{(b)}^{(\mathcal{M})} \in \mathbb{R}^{n \times |\mathcal{M}|}$. Then, the R -Squared statistic to be computed from the b -th subsample for the model \mathcal{M} is given by

$$R_{(b)}^2(\mathcal{M}) = \left(\mathbb{X}_{(b)}^{(\mathcal{M})\top} \mathbb{Y}_{(b)} \right)^\top \left(\mathbb{X}_{(b)}^{(\mathcal{M})\top} \mathbb{X}_{(b)}^{(\mathcal{M})} \right)^{-1} \left(\mathbb{X}_{(b)}^{(\mathcal{M})\top} \mathbb{Y}_{(b)} \right) \left\| \mathbb{Y}_{(b)} - \bar{\mathbb{Y}}_{(b)} \right\|^{-2},$$

where $\bar{\mathbb{Y}}_{(b)} = n^{-1} \mathbf{1}^\top \mathbb{Y}_{(b)}$. Thereafter, a one-shot type statistic can be assembled as $R_{\text{OS}}^2(\mathcal{M}) = B^{-1} \sum_{b=1}^B R_{(b)}^2(\mathcal{M})$.

The one-shot estimator $R_{\text{OS}}^2(\mathcal{M})$ is easy to compute. However, the drawback is that it might suffer from non-ignorable estimation bias if the subsample size is relatively small. To address this issue, Zhu et al. (2022) developed two improved methods for bias reduction. The first method is a jackknife-based bias-correction method. To be specific, define a delete-one estimator as $R_{(b)i}^2(\mathcal{M})$, which is the $R_{(b)}^2(\mathcal{M})$ statistic computed without the i -th observation. This leads to the jackknife estimator for the bias as $\hat{\Delta}_{(b)} = n^{-1}(n-1) \sum_i R_{(b)}^2(\mathcal{M}) - (n-1) R_{(b)i}^2(\mathcal{M})$. Then the jackknife-based bias-correction metric is defined as $R_{\text{JBC}}^2(\mathcal{M}) = B^{-1} \sum_{k=1}^B \{R_{(b)}^2(\mathcal{M}) - \hat{\Delta}_{(b)}\}$. The second method is an aggregated moment method. This method

first decomposes $R_{(b)}^2(\mathcal{M})$ into several moment components, and then aggregates the components separately. To be more precise, note that R -Squared statistic is composed of three components, namely $\widehat{\Sigma}_{\mathbb{X}(b)}^{\mathcal{M}} = n^{-1}(\mathbb{X}_{(b)}^{\mathcal{M}})^{\top} \mathbb{X}_{(b)}^{\mathcal{M}}$, $\widehat{\Sigma}_{\mathbb{X}\mathbb{Y}(b)}^{\mathcal{M}} = n^{-1}(\mathbb{X}_{(b)}^{\mathcal{M}})^{\top} \mathbb{Y}_{(b)}$, and $\widehat{\sigma}_{y(b)}^2 = n^{-1} \|\mathbb{Y}_{(b)} - \bar{\mathbb{Y}}_{(b)}\|^2$. Averaging over all subsamples leads to $\widehat{\Sigma}_{\mathbb{X}}^{\mathcal{M}} = B^{-1} \sum_{b=1}^B \widehat{\Sigma}_{\mathbb{X}(b)}^{\mathcal{M}}$, $\widehat{\Sigma}_{\mathbb{X}\mathbb{Y}}^{\mathcal{M}} = B^{-1} \sum_{k=1}^B \widehat{\Sigma}_{\mathbb{X}\mathbb{Y}(k)}^{\mathcal{M}}$ and $\widehat{\sigma}_{\mathbb{Y}}^2 = B^{-1} \sum_{b=1}^B \widehat{\sigma}_{y(b)}^2$. Then, an aggregated moment estimator is defined as $R_{\text{AM}}^2(\mathcal{M}) = \widehat{\sigma}_{\mathbb{Y}}^{-2} (\widehat{\Sigma}_{\mathbb{X}\mathbb{Y}}^{\mathcal{M}})^{\top} (\widehat{\Sigma}_{\mathbb{X}}^{\mathcal{M}})^{-1} (\widehat{\Sigma}_{\mathbb{X}\mathbb{Y}}^{\mathcal{M}})$. Both two methods can reduce the bias of $R_{\text{OS}}^2(\mathcal{M})$ significantly without inflecting the asymptotic variance.

4. Minibatch related techniques

4.1. A selective review on statistical optimization

In statistical research, many estimation problems can ultimately be transformed into optimization problems. For example, for the generalized linear models (GLMs), one usually estimates the model parameters by maximizing the likelihood function (Nelder & Wedderburn, 1972). The likelihood function is generally a sufficiently smooth and strongly convex function. Consequently, the Newton's method or Fisher's score method can be easily implemented. Often the numerical convergence can be achieved in a few iterations black(Shao, 2003). Therefore, researchers usually do not concern much about the specific optimization process but directly study the statistical properties of the optimizer (Van der Vaart, 2000).

However, with the rapid development of information technology, not only datasets are becoming increasingly large, but models are also becoming even more complex (Fan et al., 2020). As mentioned before, this poses two challenges to traditional optimization methods. First, the dimension of the model parameters p can be very high. This would make it difficult to invert the $p \times p$ Hessian matrix for Newton's type methods. Second, the dataset may be too large to be read into computer memory as a whole. Then the optimization methods based on whole datasets become no longer feasible. For the former challenge, one can consider gradient-based first-order optimization methods, such as gradient descent (GD) method, quasi-Newton method, and conjugate gradient method (Beck, 2017). For the latter challenge, one can load the data into the memory in a minibatch-wise manner, and then implement the algorithms based on these small minibatches. This leads to various minibatch-based methods. In particular, when the minibatches are generated randomly, they are also referred to as stochastic optimization methods (Lan, 2020), such as stochastic gradient descent (SGD). Due to the scalability of these minibatch-based first-order optimization methods, they are now widely used in large-scale learning tasks such as deep learning (Bottou et al., 2018; He et al., 2016; Simonyan & Zisserman, 2015).

In addition to the popularity in practice, the theoretical properties of minibatch related algorithms have also attracted increasing attention from researchers. The early literature on solving optimization problems using the idea of stochastic approximation can be traced back to Robbins and Monro (1951) and Kiefer and Wolfowitz (1952). In order to improve the efficiency of approximation, Polyak and Juditsky (1992) further proposed to perform averaging over the iterates, also known as PJR-averaging operation. More recently, Moulines and Bach (2011) and Bach and Moulines (2013) investigated the SGD algorithm for objective functions with and without strong convexity, establishing the non-asymptotic convergence upper bounds. For objective functions of finite sums, several variance reduction techniques have been found useful in achieving a faster convergence rate compared to the classical SGD (Defazio et al., 2014; Johnson & Zhang, 2013; Roux et al., 2012). To further accelerate the convergence rate for ill-conditioned problems, momentum based methods have also attracted great attention (Assran & Rabbat, 2020; Gitman et al., 2019; Liu et al., 2020). Another line of research considered how to generate minibatches to improve the minibatch based GD (Gürbüzbalaban et al., 2021; Gower et al., 2019; Mishchenko et al., 2020; Needell & Ward, 2017). Apart from research on minibatch related methods from an optimization perspective, there are also many studies conducted from a statistical perspective, paying more attention to characterizing the statistical properties of the resulting estimators. These works include but are not limited to Toulis and Airolidi (2017), Chen, Lee et al. (2020), Luo and Song (2020), Zhu et al. (2023), and Tang et al. (2023).

4.2. Minibatch gradient descent algorithms

For high dimensional data analysis, various stochastic minibatch gradient descent (SMGD) methods have received increasing attention in recent literature due to their outstanding performances and relatively easier theoretical properties (Duchi et al., 2011; Kingma & Ba, 2014). The SMGD algorithms can be mainly categorized into two groups according to the generation of minibatch data. The first group assumes the minibatches are independently generated from the given sample with replacement (Gao et al., 2021). Then the noise introduced by minibatch data can be viewed as conditionally independent. The second category assumes that the noises introduced by minibatch data form a martingale difference sequence. This assumption is particularly true for streaming data analysis

(Chen, Mou et al., 2022; Mou et al., 2020; Yu et al., 2021). Taking the SMGD studied in Chen, Lee et al. (2020) as a concrete example, they assume that the gradient noise from different minibatch data forms a martingale difference. Following the previous work of Polyak and Juditsky (1992), the asymptotic distribution of the averaged SMGD estimator can be established. To conduct statistical inference of the averaged SMGD estimator, different inference procedures are proposed for both fixed dimension case and diverged dimension case. In the fixed dimension case, Chen, Lee et al. (2020) proposed a novel batch-means covariance estimator which can avoid computing the inverse of the Hessian matrix as compared with the naive plug-in estimator. In the high dimension case, Chen, Lee et al. (2020) proposed an online debiased lasso procedure to construct the confidence interval of element-wise regression coefficient.

Different from the SMGD algorithm, another way to generate minibatch data is the random partition, which is arguably the most popularly used minibatch method in offline real practice, since it has been well implemented by many standard deep learning programs such as TensorFlow and PyTorch. Since the minibatches form a partition of the whole sample data, they are no longer independent or conditionally independent with each other. As a result, it does not match the model assumption in the past literature (Bottou et al., 2018; Chen, Mou et al., 2022; Dieuleveut et al., 2020; Lan, 2020; Mou et al., 2020; Yu et al., 2021), which calls for theoretical investigation. To fill this theoretical gap, Qi, Wang et al. (2023) studied the properties of fixed minibatch gradient descent (FMGD) algorithm and the resulting estimator. Let $\mathcal{S} = \{1, 2, \dots, N\}$ be the index set of the whole sample. Let $Y_i \in \mathbb{R}^1$ be the response of interest and $X_i = (X_{i1}, X_{i2}, \dots, X_{ip})^\top \in \mathbb{R}^p$ be the associated p -dimensional predictor. Define the loss function evaluated at sample i as $\ell(X_i, Y_i; \theta)$, where $\theta \in \mathbb{R}^q$ denotes the parameter. Then the global loss function can be constructed as $\mathcal{L}(\theta) = N^{-1} \sum_{i=1}^N \ell(X_i, Y_i; \theta)$. The global estimator can be defined as $\hat{\theta} = \operatorname{argmin} \mathcal{L}(\theta)$. Denote $\{\mathcal{S}^{(t,m)}\}_{m=1}^M$ as the minibatch index sets in the t -th epoch. Then one should have $\mathcal{S} = \bigcup_m \mathcal{S}^{(t,m)}$ and $\mathcal{S}^{(t,m_1)} \cap \mathcal{S}^{(t,m_2)} = \emptyset$ for any $t \geq 1$ and $m_1 \neq m_2$. For convenience, assume N and M are particularly designed so that $n = N/M$ is an integer and all minibatches have the same sample size as $|\mathcal{S}^{(t,m)}| = n$. Then the updating formula of MGD can be expressed as

$$\begin{aligned} \hat{\theta}^{(t,1)} &= \hat{\theta}^{(t-1,M)} - \alpha \dot{\mathcal{L}}^{(t,1)}(\hat{\theta}^{(t-1,M)}), \\ \hat{\theta}^{(t,m)} &= \hat{\theta}^{(t,m-1)} - \alpha \dot{\mathcal{L}}^{(t,m)}(\hat{\theta}^{(t,m-1)}) \text{ for } 2 \leq m \leq M, \end{aligned} \quad (7)$$

where $\alpha > 0$ is the learning rate, $\mathcal{L}^{(t,m)}(\theta) = n^{-1} \sum_{i \in \mathcal{S}^{(t,m)}} \ell(X_i, Y_i; \theta)$ is the loss function for the m -th minibatch in the t -th epoch, and $\dot{\mathcal{L}}^{(t,m)}(\theta)$ is the first-order derivatives of $\mathcal{L}^{(t,m)}(\theta)$ with respect to θ .

To study the algorithm, Qi, Wang et al. (2023) first consider FMGD algorithm under the linear regression model with a fixed sample partition. Then the above updating formulas (7) naturally form a linear system. Under appropriate technical assumptions, Qi, Wang et al. (2023) show that the FMGD estimator converges linearly to the stable solution of the linear system as

$$\left\| \hat{\theta}^{(t,m)} - \hat{\theta}^{(m)} \right\| \leq \rho_{\alpha, M}^{t-1} \left\| \hat{\theta}^{(0,m)} - \hat{\theta}^{(m)} \right\|,$$

where $\rho_{\alpha, M} \in (0, 1)$ is a contraction factor depending on α and M . The asymptotic normality result is also established by Qi, Wang et al. (2023). However, Qi, Wang et al. (2023) find that the FMGD estimator is biased for any constant learning $\alpha > 0$. To further reduce the error upper bound, Qi, Wang et al. (2023) consider the diminishing learning rate scheduling. As long as $\sum_{t=1}^{\infty} \alpha_t = \infty$ and $\sum_{t=1}^{\infty} \alpha_t^2 < \infty$, the FMGD estimator should converge to the OLS estimator as $t \rightarrow \infty$. Qi, Wang et al. (2023) then extend their theoretical investigations to random partition with shuffling and general loss function, and similar results are established.

4.3. Minibatch gradient descent with momentum

Despite the practical usefulness of the MGD algorithm, it still can be extremely time-consuming for large-scale statistical analysis with high dimensional parameters, particularly in the research field of deep learning (Devlin et al., 2019; Goodfellow et al., 2016; He et al., 2016; Kingma & Ba, 2014). To address this issue, various improved algorithms have been proposed. One direction is to investigate the accelerated gradient descent algorithm. As a first-order optimization method, gradient descent (GD) does not require the computation of the second derivative (i.e., the Hessian matrix) of the objective function. However, due to the neglect of the second-order information of the objective function, the numerical convergence rate of the standard GD algorithm is often much slower than that of second-order optimization algorithms such as Newton's method. This is particularly true when the objective function is severely ill-conditioned. To address this issue, Polyak (1964) proposed a so-called "heavy-ball" method. This method utilizes not only the gradient of the current step but also the information from the previous step (i.e.,

momentum). Specifically, it updates the estimates as

$$\widehat{\theta}^{(t)} = \widehat{\theta}^{(t-1)} - \alpha \dot{\mathcal{L}}(\widehat{\theta}^{(t-1)}) + \gamma (\widehat{\theta}^{(t-1)} - \widehat{\theta}^{(t-2)}),$$

where $\widehat{\theta}^{(t)}$ is the t -th estimate, $\dot{\mathcal{L}}(\theta)$ is the gradient, $\alpha > 0$ is the learning rate, and $\gamma > 0$ is the momentum parameter. Further theoretical analysis by Polyak (1964) showed that the convergence rate could be much improved compared to the standard GD algorithm by this modification. This method is now commonly referred to as the gradient descent with momentum (GDM).

The success of the momentum idea has attracted considerable attention from both theoretical and practical perspectives (Bottou et al., 2018; Goodfellow et al., 2016; Sutskever et al., 2013). For example, Nesterov (1983) proposed a method that uses the momentum and predicted gradient to update the parameters. By a similar idea, Beck and Teboulle (2009) developed an accelerated optimization algorithm for non-smooth objective functions. Kingma and Ba (2014) proposed an adaptive momentum method called ADAM, which is widely used in the fields of deep learning. Cyrus et al. (2018) described a robust momentum method that generalizes the triple momentum method proposed in Van Scoy et al. (2017). Ma and Yarats (2018) developed a more general variant called the quasi-hyperbolic momentum (QHM) algorithm, whose theoretical properties were further investigated by Gitman et al. (2019). In practice, when the whole dataset is too large to be loaded into the memory, one has to process the data in a minibatch-by-minibatch manner. This is particularly true for many deep learning tasks (He et al., 2016; Krizhevsky et al., 2012; Simonyan & Zisserman, 2015). This leads to various minibatch-based GDM (MGDM) methods. In fact, the MGDM methods have been incorporated into many important software libraries, including TensorFlow (Abadi et al., 2016) and PyTorch (Paszke et al., 2019).

In recent years, many efforts have been devoted to the theoretical analysis of the MGDM methods. Most of these studies investigated various MGDM methods from an optimization perspective, and have shown that the momentum term can effectively improve the numerical convergence (Assran & Rabbat, 2020; Gitman et al., 2019; Liu et al., 2020; Loizou & Richtárik, 2020). A recent work of Tang et al. (2023) analysed a PJR-averaging version of the MGDM method for general statistical optimization problems and established the asymptotic distribution of the resulting estimator. It is worth noting that most of the theoretical studies mentioned above typically require that the minibatches are sampled independently and identically from the whole dataset (or the population distribution). However, in practice, such as in TensorFlow or PyTorch, minibatches are often obtained through random partition. As mentioned before, random partition means that the whole dataset is randomly partitioned into several non-overlapping minibatches. Unfortunately, minibatches generated in this way no longer satisfy the aforementioned requirements. To bridge the gap between theory and practice, Gao et al. (2023) considered the random partition based MGDM algorithm as a linear dynamical system:

$$\widehat{\theta}^{(t,m)} = \widehat{\theta}^{(t,m-1)} - \alpha \dot{\mathcal{L}}_{(m)}(\widehat{\theta}^{(t,m-1)}) + \gamma (\widehat{\theta}^{(t,m-1)} - \widehat{\theta}^{(t,m-2)}), \quad \text{for } 1 \leq m \leq M,$$

where $\widehat{\theta}^{(t,m)}$ is the estimate from the m -th minibatch in the t -th epoch, and $\dot{\mathcal{L}}_{(m)}(\theta)$ is the gradient computed on the m -th minibatch ($1 \leq m \leq M$).

Based on the linear regression model, a closed form of the stable solution to the above linear dynamical system can be obtained. Specifically, let $\widehat{\theta}^{(m)}$ be the stable solution corresponding to the m -th minibatch. Under appropriate conditions, one can obtain the following linear convergence result:

$$\left\| \widehat{\theta}^{(t,m)} - \widehat{\theta}^{(m)} \right\| \leq \left(\rho_{\alpha,\gamma}^M + \varepsilon_{n,t} \right)^t \left(\left\| \widehat{\theta}^{(0,m)} - \widehat{\theta}^{(m)} \right\| + \left\| \widehat{\theta}^{(0,m-1)} - \widehat{\theta}^{(m-1)} \right\| \right),$$

for each $1 \leq m \leq M$, where $\rho_{\alpha,\gamma} \in (0, 1)$ is the contraction factor controlling the convergence rate, and $\varepsilon_{n,t}$ is some small number. By choosing appropriate tuning parameters α and γ , one can achieve the minimal (and thus the optimal) contraction factor $\rho_{\min} = (\sqrt{\kappa} - 1)/(\sqrt{\kappa} + 1)$, where κ is the condition number of the least squares problem. In addition, Gao et al. (2023) established the asymptotic normality for the stable solution. The results showed that the stable solution can be statistically as efficient as the whole sample OLS estimator, as long as learning rate α is sufficiently small. However, this interesting result relies on the least squares loss function. Investigating the MGDM algorithm based on a randomly partition strategy under general loss functions is a problem worth further exploration.

4.4. Communication reduction for minibatch gradient descent algorithm

Another research direction focuses on how to reduce communication costs during the training process of MGD algorithms. Massive datasets are usually stored on hard drives due to a limited storage capacity of the computation

devices. However, the data need to be transmitted into the system RAM and/or graphical memory before gradient computation. Therefore, the communication cost during the training process of the MGD algorithm might also lead to expensive time consumption (Chien et al., 2018; Pumma et al., 2017). The communication mechanism of MGD algorithm also causes another problem when computing in a CPU-GPU system, that is, the GPU idling problem. This problem refers to the phenomenon where the GPU waits for the CPU to perform data communication before gradient computation during the training process (Choi et al., 2019; Zhang & Ré, 2016; Zinkevich et al., 2010). It could become even worse when the number of updates is extremely large. Then how to reduce the communication cost during the training of MGD algorithms for faster computation becomes a problem of great interest (Bauer et al., 2011; Chien et al., 2018; Ofeidis et al., 2022; Zhu et al., 2019).

This problem can be solved by either advanced hardware technology or algorithmic innovation. There has been remarkable technological advancement in the first direction. For example, the GPUDirect Storage technology developed by NVIDIA enables direct communication between the hard drive and the GPU graphical memory. Furthermore, their Remote Direct Memory Access (RDMA) technology provides a direct path among different GPUs within a distributed system. As for the second direction, various methods have been proposed to either reduce the idling time or improve the communication efficiency. For example, Choi et al. (2019) proposed a data echoing method, which repeatedly computes the gradient on the currently loaded minibatch data before the next minibatch to be prepared. Hoffer et al. (2019) further considered applying different data augmentation method to the currently loaded minibatch data. To reduce the idling time of GPU and enhance computation efficiency, the idea of pre-loading and buffering has also received a lot of attention; see for example the data processing pipelines (Bauer et al., 2011; Ma et al., 2003; Nitzberg & Lo, 1997) and asynchronous data loading (Ofeidis et al., 2022; Zhu et al., 2019). Despite their usefulness, the existing methods still suffer from several challenges. First, novel hardware technologies often require substantial engineering labours and financial expenses. As a result, these methods have not been extensively adopted by general practitioners and researchers for now. Second, since solving the GPU idling problem is quite engineering-oriented, most existing algorithms and training strategies lack theoretical analysis. Consequently, it is of great importance to investigate the theoretical properties of those methods under certain framework.

To address this issue, Qi, Huang et al. (2023) consider a buffered minibatch gradient descent (BMGD) algorithm. The proposed BMGD algorithm consists of two steps, that is the buffering step and the computation step. In the buffering step, a large amount of data are loaded into the CPU memory. To this end, assume that the entire sample can be divided into K non-overlapping blocks, which are referred to as the buffered data. Their indices are collected by $\mathcal{S}_{r,k}$ ($1 \leq r \leq R, 1 \leq k \leq K$). For all $1 \leq r \leq R$, one can obtain $\mathcal{S} = \bigcup_k \mathcal{S}_{r,k}$ and $\mathcal{S}_{r,k_1} \cap \mathcal{S}_{r,k_2} = \emptyset$ for $k_1 \neq k_2$. Then in the computation step, a standard MGD updating procedure is applied on the buffered data. Assume that each buffered data can be further decomposed into M minibatches. Let $\mathcal{S}_{r,k}^{(t,m)}$ be the index set of the sample calculated on the m -th minibatch in the t -th epoch for the k -th buffer in the r -th iteration and assume that $|\mathcal{S}_{r,k}^{(t,m)}| = n$ for all r, k, t, m . By inheriting the model assumptions in Section 4.1, the updating formula of the BMGD algorithm can be written as

$$\begin{aligned}\hat{\theta}_{r,k}^{(t,1)} &= \hat{\theta}_{r,k}^{(t-1,M)} - \alpha \dot{\mathcal{L}}_{r,k}^{(t,1)} \left(\hat{\theta}_{r,k}^{(t-1,M)} \right), \\ \hat{\theta}_{r,k}^{(t,m)} &= \hat{\theta}_{r,k}^{(t,m-1)} - \alpha \dot{\mathcal{L}}_{r,k}^{(t,m)} \left(\hat{\theta}_{r,k}^{(t,m-1)} \right) \quad \text{for } 2 \leq m \leq M,\end{aligned}\tag{8}$$

where $\alpha > 0$ is the learning rate, and $\dot{\mathcal{L}}_{r,k}^{(t,m)}(\theta) = n^{-1} \sum_{i \in \mathcal{S}_{r,k}^{(t,m)}} \dot{\ell}(X_i, Y_i; \theta)$ is the gradient computed on minibatch $\mathcal{S}_{r,k}^{(t,m)}$. The buffering idea of the BMGD algorithm can help reduce the GPU idling time and improve communication efficiency. A rigorous asymptotic theory was developed by Qi, Huang et al. (2023) to support the BMGD method. Its applicability is also extended to the Polyak-Lojasiewicz (PL) function class, which not only contains a wide range of statistical models (e.g., the generalized linear model) but also some non-convex loss functions.

4.5. Learning rate scheduling

One crucial component to the minibatch gradient descent algorithm and its variants is an appropriate learning rate scheduling. Despite previous efforts showing that the algorithm should converge under certain conditions, determining the correct learning rate in real practice largely relies on subjective judgment (Ruder, 2016). If the learning rate is set inappropriately, it can lead to training failure or slow convergence. To address this issue, various scheduling approaches have been proposed to achieve adaptive adjustment of the learning rate. Examples include rule-based scheduling, such as the step decay scheduler (Ge et al., 2019) and the ‘reduce learning rate on plateau’

method (Nakamura et al., 2021). These methods automatically reduce the learning rate when optimization encounters bottlenecks. Besides, Duchi et al. (2011) proposed AdaGrad, which iteratively decreases the step size based on a pre-specified function. However, AdaGrad requires setting an additional parameter related to the learning rate, which needs to be subjectively chosen. Extensions of AdaGrad, such as RMSProp black (Mukkamala & Hein, 2017) and AdaDelta (Zeiler, 2012), have been proposed. RMSProp introduces a decay factor to adjust the weights of previous sample gradients. Moreover, Adam (Kingma & Ba, 2014) combines RMSProp with a momentum-based method called adaptive moment estimation. In Adam, both the step size and update direction are adjusted during each iteration. However, because the step sizes are adjusted without considering the loss function, the loss reduction obtained for each update step is suboptimal. Therefore, the convergence rate can still be further improved. New techniques have been introduced to adjust step sizes in gradient-based optimization methods. For instance, Baydin et al. (2017) introduced a ‘learning rate for the learning rate’ hyperparameter, which is updated using gradient descent to adjust the learning rate. Shu et al. (2022) have built an additional network to predict learning rate values in different iterations. However, those works bring more hyperparameters and thus result in more parameter tuning as well as more uncertainties.

From the perspective of optimization, the key to improving training algorithms lies in how to appropriately exploit 1st-order information (i.e., the gradient) and the 2nd-order information (i.e., the Hessian matrix). When the scale of parameters to be estimated is relatively small, the most common approach has been to apply the Newton’s method. For training neural networks with a large number of parameters, several generalized optimization methods (Agarwal et al., 2017; Bergou et al., 2022; Gargiani et al., 2020; Tan et al., 2016) have been proposed, inspired by the Newton–Raphson iteration. Due to the high computational cost, existing studies have tried to approximate the Hessian matrix, including the Barzilai–Borwein method (Tan et al., 2016), subsampling methods (Agarwal et al., 2017; Bergou et al., 2022), and generalized Gauss-Newton method (Gargiani et al., 2020). However, most of those methods still involve the computation and storage of 1st- and 2nd-order derivatives. Thus, they may be less efficient or even not feasible practically, when the dimension of parameter estimation is extremely high (Sutskever, 2013).

In order to achieve automatic and nearly optimal optimization, Zhu, Huang et al. (2021) propose a novel optimization method based on local quadratic approximation (LQA). Viewing the learning rate as a time-varying parameter, they treat the loss reduction as a function of the temporal learning rate. Then, the learning rate is dynamically adjusted through the maximization of the loss reduction. Their aim is to make use of 2nd-order derivative information to accelerate the optimization while avoiding calculating the 2nd-order derivatives directly. To this end, they combine techniques of Taylor expansion and quadratic approximation to propose an improved optimization algorithm with low computational costs.

Denote the time-varying learning rate as $\alpha_{t,k}$, where t and k are indices of iteration and minibatch, respectively. Given a loss function $\mathcal{L}(X; \theta)$ with the model parameter θ and input sample X , let $\widehat{\theta}^{(t,k)}$ and $\Delta \mathcal{L}(\alpha_{t,k})$ denote the estimate and the loss reduction at the k -th minibatch of the t -th iteration of the training, respectively. For simplicity, let $g_{t,k} = |\mathcal{S}_k|^{-1} \sum_{i \in \mathcal{S}_k} \dot{\mathcal{L}}(\widehat{\theta}^{(t,k)})$ denote the current gradient, where \mathcal{S}_k is the index set of the k -th minibatch. Then, the LQA method adopts Taylor expansion to explore the loss reduction. From the Taylor expansion of $\mathcal{L}(\theta)$ around $\widehat{\theta}^{(t,k)}$, one can obtain the following approximation,

$$\begin{aligned} \Delta \mathcal{L}(\alpha_{t,k}) &= \frac{1}{|\mathcal{S}_k|} \sum_{i \in \mathcal{S}_k} \left\{ \mathcal{L} \left(X_i; \widehat{\theta}^{(t,k)} - \alpha_{t,k} g_{t,k} \right) - \mathcal{L} \left(X_i; \widehat{\theta}^{(t,k)} \right) \right\} \\ &= -a_{t,k} \alpha_{t,k} + b_{t,k} \alpha_{t,k}^2 + o(\alpha_{t,k}^2 g_{t,k}^\top g_{t,k}), \end{aligned} \quad (9)$$

where the two constants are given by $a_{t,k} = |\mathcal{S}_k|^{-1} \sum_{i \in \mathcal{S}_k} \dot{\mathcal{L}}(X_i; \widehat{\theta}^{(t,k)}) g_{t,k}$ and $b_{t,k} = (2|\mathcal{S}_k|)^{-1} \sum_{i \in \mathcal{S}_k} g_{t,k}^\top \ddot{\mathcal{L}}(X_i; \widehat{\theta}^{(t,k)}) g_{t,k}$ with $\dot{\mathcal{L}}(\widehat{\theta}^{(t,k)})$ and $\ddot{\mathcal{L}}(\widehat{\theta}^{(t,k)})$ denoting the 1st- and 2nd-order derivatives of the local loss function, respectively. Since the higher order terms here are negligible, Equation (9) can be simply written as $\Delta \mathcal{L}(\alpha_{t,k}) \approx -a_{t,k} \alpha_{t,k} + b_{t,k} \alpha_{t,k}^2$. To optimize $\Delta \mathcal{L}(\alpha_{t,k})$ with respect to $\alpha_{t,k}$, one can take the corresponding derivative of the loss reduction, which leads to the approximated optimal learning rate $\alpha_{t,k}^* = (2b_{t,k})^{-1} a_{t,k}$. This suggests that once the coefficients $a_{t,k}$ and $b_{t,k}$ are estimated, a nearly optimal choice for the learning rate can thus be determined. In this way, the reduction in the loss function is nearly optimal for each batch step. As a consequence, the total number of iterations required for convergence can be significantly reduced, allowing for the algorithm to converge much faster than usual.

5. Concluding remarks

This paper gives a selective review about three categories of statistical computation methods for massive data analysis. Firstly, we considered the distributed computing methods, which provide a feasible solution when the data size

is too large to be comfortably accommodated by one single computer. Secondly, we considered the subsampling methods, which are practically useful, when limited computing resources are available for a massive data analysis task. Finally, we considered the minibatch gradient techniques, which have been widely used for the training of complicated models with a large number of parameters, such as deep neural network models.

To conclude this paper, we discuss here some potential future research directions. Firstly, we remark that all the methods reviewed in this work are developed for independent data. On the other side, datasets with sophisticated dependence structure (e.g., spatial-temporal data) are commonly encountered in real practice. Then how to conduct distributed computation for data with sophisticated dependent structure is an important research direction. Second, all the methods reviewed in this work are all developed for models with relatively limited dimension and simple structure. On the other side, models with extremely high dimension and extremely complicated structure are increasingly available. This is particularly true for various well-celebrated deep learning models. Then how to develop statistical computing theory for models with extremely high dimension and extremely sophisticated structure is another important research direction. Third, all the methods reviewed in this work assume a global model universally for all the data points with a common set of model parameters. However, this seems an obviously unrealistic assumption for dataset with a massive size. Then how to conduct effective statistical learning with more flexible model parameters is also an important research direction. Lastly, we remark that the three categories of methods reviewed in this paper can be combined together for even better performance. See for example (Yu et al., 2022) and Chen, Liu et al. (2022). This is the last research direction worth pursuing.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by the National Natural Science Foundation of China [grant numbers 72301070, 72171226, 12271012, 12171020, 12071477, 72371241, 72222009, 71991472 and 12331009]; the National Statistical Science Research Project [grant number 2023LD008]; the Fundamental Research Funds for the Central Universities in UIBE [grant number CXTD13-04]; the MOE Project of Key Research Institute of Humanities and Social Sciences [grant number 22JJD110001]; the Program for Innovation Research, the disciplinary funding and the Emerging Interdisciplinary Project of Central University of Finance and Economics; and the Postdoctoral Fellowship Program of CPSF [grant numbers GZC20231522, GZC20230111 and GZB20230070].

ORCID

Yuan Gao  <http://orcid.org/0000-0001-6389-0192>

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I. J., Harp, A., Irving, G., Isard, M., Jia, Y., Józefowicz, R., Kaiser, L., Kudlur, M., ... Zheng, X. (2016). *TensorFlow: Large-scale machine learning on heterogeneous distributed systems*. arXiv: 1603.04467.
- Agarwal, N., Bullins, B., & Hazan, E. (2017). Second-order stochastic optimization for machine learning in linear time. *Journal of Machine Learning Research*, 18(1), 4148–4187.
- Ai, M., Yu, J., Zhang, H., & Wang, H. (2021). Optimal subsampling algorithms for big data regressions. *Statistica Sinica*, 31(2), 749–772.
- Alhamzawi, R., & Ali, H. T. M. (2018). Bayesian quantile regression for ordinal longitudinal data. *Journal of Applied Statistics*, 45(5), 815–828. <https://doi.org/10.1080/02664763.2017.1315059>
- Assran, M., & Rabbat, M. (2020). On the convergence of Nesterov's accelerated gradient method in stochastic settings. In *International Conference on Machine Learning*. PMLR.
- Bach, F., & Moulines, E. (2013). Non-strongly-convex smooth stochastic approximation with convergence rate $O(1/n)$. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Bathey, H., Fan, J., Liu, H., Lu, J., & Zhu, Z. (2018). Distributed testing and estimation under sparse high dimensional models. *The Annals of Statistics*, 46(3), 1352–1382. <https://doi.org/10.1214/17-AOS1587>
- Bauer, M., Cook, H., & Khailany, B. (2011). CudaDMA: Optimizing GPU memory bandwidth via warp specialization. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*. Association for Computing Machinery.
- Baydin, A. G., Cornish, R., Rubio, D. M., Schmidt, M., & Wood, F. (2017). *Online learning rate adaptation with hypergradient descent*. arXiv: 1703.04782.
- Beck, A. (2017). *First-order methods in optimization*. Society for Industrial and Applied Mathematics, SIAM.
- Beck, A., & Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1), 183–202. <https://doi.org/10.1137/080716542>
- Bellet, A., Guerraoui, R., Taziki, M., & Tommasi, M. (2018). Personalized and private peer-to-peer machine learning. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*. PMLR.

- Bergou, E. H., Diouane, Y., Kunc, V., Kungurtsev, V., & Royer, C. W. (2022). A subsampling line-search method with second-order results. *INFORMS Journal on Optimization*, 4(4), 403–425. <https://doi.org/10.1287/ijoo.2022.0072>
- Bickel, P. J., & Freedman, D. A. (1981). Some asymptotic theory for the bootstrap. *The Annals of Statistics*, 9(6), 1196–1217.
- Bickel, P. J., Götze, F., & van Zwet, W. R. (1997). Resampling fewer than n observations: Gains, losses, and remedies for losses. *Statistica Sinica*, 7(1), 1–31.
- Blot, M., Picard, D., Cord, M., & Thome, N. (2016). *Gossip training for deep learning*. arXiv: 1611.09726.
- Bottou, L., Curtis, F. E., & Nocedal, J. (2018). Optimization methods for large-scale machine learning. *SIAM Review*, 60(2), 223–311. <https://doi.org/10.1137/16M1080173>
- Broyden, C. G., Dennis Jr, J. E., & Moré, J. J. (1973). On the local and superlinear convergence of quasi-Newton methods. *IMA Journal of Applied Mathematics*, 12(3), 223–245. <https://doi.org/10.1093/imat/12.3.223>
- Casella, G., & Berger, R. L. (2002). *Statistical inference*. Duxbury Pacific Grove.
- Chang, X., Lin, S., & Wang, Y. (2017). Divide and conquer local average regression. *Electronic Journal of Statistics*, 11(1), 1326–1350. <https://doi.org/10.1214/17-EJS1265>
- Chen, C. W., Dunson, D. B., Reed, C., & Yu, K. (2013). Bayesian variable selection in quantile regression. *Statistics and Its Interface*, 6(2), 261–274. <https://doi.org/10.4310/SII.2013.v6.n2.a9>
- Chen, S., Yu, D., Zou, Y., Yu, J., & Cheng, X. (2022). Decentralized wireless federated learning with differential privacy. *IEEE Transactions on Industrial Informatics*, 18(9), 6273–6282. <https://doi.org/10.1109/TII.2022.3145010>
- Chen, S. X., & Peng, L. (2021). Distributed statistical inference for massive data. *The Annals of Statistics*, 49(5), 2851–2869.
- Chen, W., Wang, Z., & Zhou, J. (2014). Large-scale L-BFGS using mapreduce. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Chen, X., Lee, J. D., Tong, X. T., & Zhang, Y. (2020). Statistical inference for model parameters in stochastic gradient descent. *The Annals of Statistics*, 48(1), 251–273. <https://doi.org/10.1214/18-AOS1801>
- Chen, X., Liu, W., Mao, X., & Yang, Z. (2020). Distributed high-dimensional regression under a quantile loss function. *Journal of Machine Learning Research*, 21(1), 7432–7474.
- Chen, X., Liu, W., & Zhang, Y. (2019). Quantile regression under memory constraint. *The Annals of Statistics*, 47(6), 3244–3273.
- Chen, X., Liu, W., & Zhang, Y. (2022). First-order newton-type estimator for distributed estimation and inference. *Journal of the American Statistical Association*, 117(540), 1858–1874. <https://doi.org/10.1080/01621459.2021.1891925>
- Chen, X., & Xie, M. (2014). A split-and-conquer approach for analysis of extraordinarily large data. *Statistica Sinica*, 24(4), 1655–1684.
- Chen, Z., Mou, S., & Maguluri, S. T. (2022). Stationary behavior of constant stepsize SGD type algorithms: An asymptotic characterization. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 6(1), 1–24. <https://doi.org/10.1145/3508039>
- Chien, S. W. D., Markidis, S., Sishtla, C. P., Santos, L., Herman, P., Narasimhamurthy, S., & Laure, E. (2018). Characterizing deep-learning I/O workloads in TensorFlow. In *International Workshop on Parallel Data Storage and Data Intensive Scalable Computing Systems*. IEEE.
- Choi, D., Passos, A., Shallue, C. J., & Dahl, G. E. (2019). *Faster neural network training with data echoing*. arXiv: 1907.05550.
- Crane, R., & Roosta, F. (2019). DINGO: Distributed Newton-type method for gradient-norm optimization. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Cyrus, S., Hu, B., Van Scoy, B., & Lessard, L. (2018). A robust accelerated optimization algorithm for strongly convex functions. In *2018 Annual American Control Conference* (pp. 1376–1381). IEEE.
- Davidon, W. C. (1991). Variable metric method for minimization. *SIAM Journal on Optimization*, 1(1), 1–17. <https://doi.org/10.1137/0801001>
- Defazio, A., Bach, F., & Lacoste Julien, S. (2014). SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Deng, J., Dong, W., Socher, R., Li, Li, & Li, F. (2009). ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 248–255). IEEE.
- Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics.
- Dieuleveut, A., Durmus, A., & Bach, F. (2020). Bridging the gap between constant step size stochastic gradient descent and Markov chains. *The Annals of Statistics*, 48(3), 1348–1382. <https://doi.org/10.1214/19-AOS1850>
- Drineas, P., Mahoney, M. W., & Muthukrishnan, S. (2006). Sampling algorithms for ℓ_2 regression and applications. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithm*. Society for Industrial and Applied Mathematics.
- Drineas, P., Mahoney, M. W., Muthukrishnan, S., & Sarlós, T. (2011). Faster least squares approximation. *Numerische Mathematik*, 117(2), 219–249. <https://doi.org/10.1007/s00211-010-0331-6>
- Duchi, J., Hazan, E., & Singer, Y. (2011). Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7), 257–269.
- Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1), 1–26. <https://doi.org/10.1214/aos/1176344552>
- Efron, B., & Stein, C. (1981). The jackknife estimate of variance. *The Annals of Statistics*, 9(3), 586–596. <https://doi.org/10.1214/aos/1176345462>
- Eisen, M., Mokhtari, A., & Ribeiro, A. (2017). Decentralized quasi-Newton methods. *IEEE Transactions on Signal Processing*, 65(10), 2613–2628. <https://doi.org/10.1109/TSP.2017.2666776>
- Fan, J., Guo, Y., & Wang, K. (2023). Communication-efficient accurate statistical estimation. *Journal of the American Statistical Association*, 118(542), 1000–1010. <https://doi.org/10.1080/01621459.2021.1969238>
- Fan, J., Li, R., Zhang, C., & Zou, H. (2020). *Statistical foundations of data science*. Chapman and Hall/CRC.
- Fan, J., & Lv, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(5), 849–911. <https://doi.org/10.1111/j.1467-9868.2008.00674.x>

- Fan, J., Wang, D., Wang, K., & Zhu, Z. (2019). Distributed estimation of principal eigenspaces. *The Annals of Statistics*, 47(6), 3009–3031. <https://doi.org/10.1214/18-AOS1713>
- Gao, D., Ju, C., Wei, X., Liu, Y., Chen, T., & Yang, Q. (2019). *HHHFL: Hierarchical heterogeneous horizontal federated learning for electroencephalography*. arXiv: 1909.05784.
- Gao, Y., Li, J., Zhou, Y., Xiao, F., & Liu, H. (2021). Optimization methods for large-scale machine learning. In *2021 18th International Computer Conference on Wavelet Active Media Technology and Information Processing*. IEEE.
- Gao, Y., Zhu, X., Qi, H., Li, G., Zhang, R., & Wang, H. (2023). An asymptotic analysis of random partition based minibatch momentum methods for linear regression models. *Journal of Computational and Graphical Statistics*, 32(3), 1083–1096. <https://doi.org/10.1080/10618600.2022.2143786>
- Gargiani, M., Zanelli, A., Diehl, M., & Hutter, F. (2020). *On the promise of the stochastic generalized Gauss-Newton method for training DNNs*. arXiv: 2006.02409.
- Ge, R., Kakade, S. M., Kidambi, R., & Netrapalli, P. (2019). The step decay schedule: A near optimal, geometrically decaying learning rate procedure for least squares. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Gitman, I., Lang, H., Zhang, P., & Xiao, L. (2019). Understanding the role of momentum in stochastic gradient methods. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Goldfarb, D. (1970). A family of variable-metric methods derived by variational means. *Mathematics of Computation*, 24(109), 23–26. <https://doi.org/10.1090/mcom/1970-24-109>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.
- Gower, R. M., Loizou, N., Qian, X., Sailanbayev, A., Shulgin, E., & Richtárik, P. (2019). SGD: General analysis and improved rates. In *International Conference on Machine Learning*. PMLR.
- Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., & He, K. (2017). *Accurate, large minibatch SGD: Training ImageNet in 1 hour*. arXiv: 1706.02677.
- Gu, J., & Chen, S. (2023). Statistical inference for decentralized federated learning. *Working Paper*.
- Gürbüzbalaban, M., Ozdaglar, A., & Parrilo, P. A. (2021). Why random reshuffling beats stochastic gradient descent. *Mathematical Programming*, 186, 49–84. <https://doi.org/10.1007/s10107-019-01440-w>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE.
- Hector, E. C., & Song, P. X. (2020). Doubly distributed supervised learning and inference with high-dimensional correlated outcomes. *Journal of Machine Learning Research*, 21(1), 6983–7017.
- Hector, E. C., & Song, P. X. (2021). A distributed and integrated method of moments for high-dimensional correlated data analysis. *Journal of the American Statistical Association*, 116(534), 805–818. <https://doi.org/10.1080/01621459.2020.1736082>
- Hoffer, E., Nun, T. B., Hubara, I., Giladi, N., Hoefler, T., & Soudry, D. (2019). *Augment your batch: Better training with larger batches*. arXiv: 1901.09335.
- Hu, A., Jiao, Y., Liu, Y., Shi, Y., & Wu, Y. (2021). Distributed quantile regression for massive heterogeneous data. *Neurocomputing*, 448, 249–262. <https://doi.org/10.1016/j.neucom.2021.03.041>
- Huang, C., & Huo, X. (2019). A distributed one-step estimator. *Mathematical Programming*, 174(1-2), 41–76. <https://doi.org/10.1007/s10107-019-01369-0>
- Johnson, R., & Zhang, T. (2013). Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Jordan, M. I., Lee, J. D., & Yang, Y. (2019). Communication-efficient distributed statistical inference. *Journal of the American Statistical Association*, 114(526), 668–681. <https://doi.org/10.1080/01621459.2018.1429274>
- Kiefer, J., & Wolfowitz, J. (1952). Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23, 462–466. <https://doi.org/10.1214/aoms/1177729392>
- Kingma, D. P., & Ba, J. (2014). *Adam: A method for stochastic optimization*. arXiv: 1412.6980.
- Kleiner, A., Talwalkar, A., Sarkar, P., & Jordan, M. I. (2014). A scalable bootstrap for massive data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(4), 795–816. <https://doi.org/10.1111/rssb.12050>
- Koenker, R. (2005). *Quantile regression*. Cambridge University Press.
- Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, 46(1), 33–50. <https://doi.org/10.2307/1913643>
- Korkmaz, S. (2020). Deep learning-based imbalanced data classification for drug discovery. *Journal of Chemical Information and Modeling*, 60(9), 4180–4190. <https://doi.org/10.1021/acs.jcim.9b01162>
- Kostov, P., & Davidova, S. (2013). A quantile regression analysis of the effect of farmers' attitudes and perceptions on market participation. *Journal of Agricultural Economics*, 64(1), 112–132. <https://doi.org/10.1111/jage.2013.64.issue-1>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Lalitha, A., Shekhar, S., Javidi, T., & Koushanfar, F. (2018). Fully decentralized federated learning. In *3rd Workshop on Bayesian Deep Learning (NeurIPS)*. Curran Associates Inc.
- Lan, G. (2020). *First-order and stochastic optimization methods for machine learning*. Springer.
- Lee, C., Lim, C. H., & Wright, S. J. (2018). A distributed quasi-Newton algorithm for empirical risk minimization with nonsmooth regularization. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Association for Computing Machinery.
- Lee, J. D., Liu, Q., Sun, Y., & Taylor, J. E. (2017). Communication-efficient sparse regression. *Journal of Machine Learning Research*, 18(1), 115–144.
- Li, K. H. (1994). Reservoir-sampling algorithms of time complexity $O(n(1 + \log(N/n)))$. *ACM Transactions on Mathematical Software*, 20(4), 481–493. <https://doi.org/10.1145/198429.198435>
- Li, R., Zhong, W., & Zhu, L. (2012). Feature screening via distance correlation learning. *Journal of the American Statistical Association*, 107(499), 1129–1139. <https://doi.org/10.1080/01621459.2012.695654>
- Li, X., Li, R., Xia, Z., & Xu, C. (2020). Distributed feature screening via componentwise debiasing. *Journal of Machine Learning Research*, 21(24), 1–32.

- Li, X., Liang, J., Chang, X., & Zhang, Z. (2022). Statistical estimation and online inference via local SGD. In *Conference on Learning Theory*. PMLR.
- Li, X., Zhu, X., & Wang, H. (2023). *Distributed logistic regression for massive data with rare events*. arXiv: 2304.02269.
- Li, Y., Chen, C., Liu, N., Huang, H., Zheng, Z., & Yan, Q. (2021). A blockchain-based decentralized federated learning framework with committee consensus. *IEEE Network*, 35(1), 234–241. <https://doi.org/10.1109/MNET.65>
- Lian, H., & Fan, Z. (2018). Divide-and-conquer for debiased ℓ_1 -norm support vector machine in ultra-high dimensions. *Journal of Machine Learning Research*, 18(182), 1–26.
- Lian, X., Zhang, C., Zhang, H., Hsieh, C. J., Zhang, W., & Liu, J. (2017). Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent. In *Advances in neural information processing systems*. PMLR.
- Lin, S., & Zhou, D. (2018). Distributed kernel-based gradient descent algorithms. *Constructive Approximation*, 47(2), 249–276. <https://doi.org/10.1007/s00365-017-9379-1>
- Liu, W., Chen, L., & Wang, W. (2022). General decentralized federated learning for communication-computation tradeoff. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications Workshops*. IEEE.
- Liu, W., Mao, X., & Zhang, X. (2022). Fast and robust sparsity learning over networks: A decentralized surrogate median regression approach. *IEEE Transactions on Signal Processing*, 70, 797–809. <https://doi.org/10.1109/TSP.2022.3146785>
- Liu, Y., Gao, Y., & Yin, W. (2020). An improved analysis of stochastic gradient descent with momentum. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Loizou, N., & Richtárik, P. (2020). Momentum and stochastic momentum for stochastic gradient, newton, proximal point and subspace descent methods. *Computational Optimization and Applications*, 77(3), 653–710. <https://doi.org/10.1007/s10589-020-00220-z>
- Luo, L., & Song, P. X. (2020). Renewable estimation and incremental inference in generalized linear models with streaming data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82(1), 69–97. <https://doi.org/10.1111/rssb.12352>
- Ma, J., & Yarats, D. (2018). *Quasi-hyperbolic momentum and Adam for deep learning*. arXiv: 1810.06801.
- Ma, P., Mahoney, M., & Yu, B. (2014). A statistical perspective on algorithmic leveraging. In *International Conference on Machine Learning*. PMLR.
- Ma, X., Winslett, M., Lee, J., & Yu, S. (2003). Improving MPI-IO output performance with active buffering plus threads. In *International Parallel and Distributed Processing Symposium*. IEEE.
- Ma, Y., Leng, C., & Wang, H. (2024). Optimal subsampling bootstrap for massive data. *Journal of Business and Economic Statistics*, 42(1), 174–186. <https://doi.org/10.1080/07350015.2023.2166514>
- Mahoney, M. W. (2011). Randomized algorithms for matrices and data. *Foundations and Trends in Machine Learning*, 3(2), 123–224.
- McDonald, R., Mohri, M., Silberman, N., Walker, D., & Mann, G. S. (2009). Efficient large-scale distributed training of conditional maximum entropy models. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Mishchenko, K., Khaled, A., & Richtárik, P. (2020). Random reshuffling: Simple analysis with vast improvements. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Mou, W., Li, C. J., Wainwright, M. J., Bartlett, P. L., & Jordan, M. I. (2020). On linear stochastic approximation: Fine-grained Polyak-Ruppert and non-asymptotic concentration. In *Conference on Learning Theory*. PMLR.
- Moulines, E., & Bach, F. (2011). Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Mukkamala, M. C., & Hein, M. (2017). Variants of RMSProp and adagrad with logarithmic regret bounds. In *International Conference on Machine Learning*. PMLR.
- Nadiradze, G., Sabour, A., Davies, P., Li, S., & Alistarh, D. (2021). Asynchronous decentralized SGD with quantized and local updates. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Nakamura, K., Derbel, B., Won, K. J., & Hong, B. W. (2021). Learning-rate annealing methods for deep neural networks. *Electronics*, 10(16), 2029. <https://doi.org/10.3390/electronics10162029>
- Nedic, A., Olshevsky, A., & Shi, W. (2017). Achieving geometric convergence for distributed optimization over time-varying graphs. *SIAM Journal on Optimization*, 27(4), 2597–2633. <https://doi.org/10.1137/16M1084316>
- Needell, D., & Ward, R. (2017). Batched stochastic gradient descent with weighted sampling. In *Approximation theory XV: San Antonio 2016 15*. Springer.
- Nelder, J. A., & Wedderburn, R. W. (1972). Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3), 370–384. <https://doi.org/10.2307/2344614>
- Nesterov, Y. E. (1983). A method for solving the convex programming problem with convergence rate $O(1/k^2)$. In *Doklady Akademii Nauk*. Russian Academy of Sciences.
- Nitzberg, B., & Lo, V. (1997). Collective buffering: Improving parallel I/O performance. In *IEEE International Symposium on High Performance Distributed Computing*. IEEE.
- Ofeidis, I., Kiedanski, D., & Tassioulas, L. (2022). *An overview of the data-loader landscape: Comparative performance analysis*. arXiv: 2209.13705.
- Pan, R., Ren, T., Guo, B., Li, F., Li, G., & Wang, H. (2022). A note on distributed quantile regression by pilot sampling and one-step updating. *Journal of Business and Economic Statistics*, 40(4), 1691–1700. <https://doi.org/10.1080/07350015.2021.1961789>
- Pan, R., Zhu, Y., Guo, B., Zhu, X., & Wang, H. (2023). A sequential addressing subsampling method for massive data analysis under memory constraint. *IEEE Transactions on Knowledge and Data Engineering*, 35(9), 9502–9513. <https://doi.org/10.1109/TKDE.2023.3241075>
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., ... Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*. Curran Associates, Inc.

- Polyak, B. T. (1964). Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5), 1–17. [https://doi.org/10.1016/0041-5553\(64\)90137-5](https://doi.org/10.1016/0041-5553(64)90137-5)
- Polyak, B. T., & Juditsky, A. B. (1992). Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4), 838–855. <https://doi.org/10.1137/0330046>
- Pumma, S., Si, M., Feng, W., & Balaji, P. (2017). Parallel I/O optimizations for scalable deep learning. In *IEEE International Conference on Parallel and Distributed Systems*. IEEE.
- Qi, H., Huang, D., Zhu, Y., Huang, D., & Wang, H. (2023). Mini-batch gradient descent with buffer. arXiv: 2312.08728.
- Qi, H., Wang, F., & Wang, H. (2023). Statistical analysis of fixed mini-batch gradient descent estimator. *Journal of Computational and Graphical Statistics*, 32(4), 1348–1360. <https://doi.org/10.1080/10618600.2023.2204130>
- Qu, G., & Li, N. (2019). Accelerated distributed Nesterov gradient descent. *IEEE Transactions on Automatic Control*, 65(6), 2566–2581. <https://doi.org/10.1109/TAC.9>
- Reich, B. J., Fuentes, M., & Dunson, D. B. (2012). Bayesian spatial quantile regression. *Journal of the American Statistical Association*, 106(493), 6–20. <https://doi.org/10.1198/jasa.2010.ap09237>
- Richards, D., Rebeschini, P., & Rosasco, L. (2020). Decentralised learning with random features and distributed gradient descent. In *International Conference on Machine Learning*. PMLR.
- Richardson, A. M., & Lidbury, B. A. (2013). Infection status outcome, machine learning method and virus type interact to affect the optimised prediction of hepatitis virus immunoassay results from routine pathology laboratory assays in unbalanced data. *BMC Bioinformatics*, 14, 206. <https://doi.org/10.1186/1471-2105-14-206>
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3), 400–407. <https://doi.org/10.1214/aoms/1177729586>
- Rosenblatt, J. D., & Nadler, B. (2016). On the optimality of averaging in distributed statistical learning. *Journal of the IMA*, 5(4), 379–404.
- Roux, N., Schmidt, M., & Bach, F. (2012). A stochastic gradient method with an exponential convergence rate for finite training sets. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Ruder, S. (2016). *An overview of gradient descent optimization algorithms*. arXiv: 1609.04747.
- Savazzi, S., Nicoli, M., & Rampa, V. (2020). Federated learning with cooperating devices: A consensus approach for massive IoT networks. *IEEE Internet of Things Journal*, 7(5), 4641–4654. <https://doi.org/10.1109/JIoT.6488907>
- Sengupta, S., Volgushev, S., & Shao, X. (2016). A subsampled double bootstrap for massive data. *Journal of the American Statistical Association*, 111(515), 1222–1232. <https://doi.org/10.1080/01621459.2015.1080709>
- Shamir, O., Srebro, N., & Zhang, T. (2014). Communication-efficient distributed optimization using an approximate Newton-type method. In *International Conference on Machine Learning*. PMLR.
- Shao, J. (2003). *Mathematical statistics*. Springer.
- Shao, J., & Tu, D. (1995). *The jackknife and bootstrap*. Springer.
- Shu, J., Zhu, Y., Zhao, Q., Meng, D., & Xu, Z. (2022). Mlr-snet: Transferable LR schedules for heterogeneous tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3), 3505–3521.
- Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*. OpenReview.net.
- Soori, S., Mishchenko, K., Mokhtari, A., Dehnavi, M. M., & Gurbuzbalaban, M. (2020). DAVE-QN: A distributed averaged quasi-Newton method with local superlinear convergence rate. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*. PMLR.
- Stich, S. U. (2019). Local SGD converges fast and communicates little. In *2019 International Conference on Learning Representations*.
- Su, L., & Xu, J. (2019). Securing distributed gradient descent in high dimensional statistical learning. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(1), 1–41.
- Sutskever, I. (2013). *Training recurrent neural networks*. University of Toronto Toronto.
- Sutskever, I., Martens, J., Dahl, G., & Hinton, G. (2013). On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning*. PMLR.
- Suwandarathna, S., & Koggalage, R. (2007). Increasing hard drive performance – from a thermal perspective. In *International Conference on Industrial and Information Systems*.
- Tan, C., Ma, S., Dai, Y., & Qian, Y. (2016). Barzilai-Borwein step size for stochastic gradient descent. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Tan, K. M., Battey, H., & Zhou, W. (2022). Communication-constrained distributed quantile regression with optimal statistical guarantees. *Journal of Machine Learning Research*, 23(1), 12456–12516.
- Tang, H., Lian, X., Yan, M., Zhang, C., & Liu, J. (2018). Decentralized training over decentralized data. In *International Conference on Machine Learning*. PMLR.
- Tang, K., Liu, W., & Zhang, Y. (2023). Acceleration of stochastic gradient descent with momentum by averaging: Finite-sample rates and asymptotic normality. arXiv: 2305.17665.
- Tang, L., Zhou, L., & Song, P. X. K. (2020). Distributed simultaneous inference in generalized linear models via confidence distribution. *Journal of Multivariate Analysis*, 176, Article 104567. <https://doi.org/10.1016/j.jmva.2019.104567>
- Toulis, P., & Airoldi, E. M. (2017). Asymptotic and finite-sample properties of estimators based on stochastic gradients. *The Annals of Statistics*, 45(4), 1694–1727. <https://doi.org/10.1214/16-AOS1506>
- Tu, J., Liu, W., Mao, X., & Xu, M. (2023). Distributed semi-supervised sparse statistical inference. arXiv: 2306.10395.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*. Cambridge University Press.
- Vanhaesebrouck, P., Bellet, A., & Tommasi, M. (2017). Decentralized collaborative learning of personalized models over networks. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*. PMLR.
- Van Scoy, B., Freeman, R. A., & Lynch, K. M. (2017). The fastest known globally convergent first-order method for minimizing strongly convex functions. *IEEE Control Systems Letters*, 2(1), 49–54. <https://doi.org/10.1109/LCSYS.2017.2722406>

- Vitter, J. S. (1985). Random sampling with a reservoir. *ACM Transactions on Mathematical Software*, 1(1), 37–57. <https://doi.org/10.1145/3147.3165>
- Volgushev, S., Chao, S., & Cheng, G. (2019). Distributed inference for quantile regression processes. *The Annals of Statistics*, 47(3), 1634–1662. <https://doi.org/10.1214/18-AOS1730>
- Wang, F., Huang, D., Gao, T., Wu, S., & Wang, H. (2022). Sequential one-step estimator by sub-sampling for customer churn analysis with massive data sets. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 71(5), 1753–1786.
- Wang, F., Zhu, Y., Huang, D., Qi, H., & Wang, H. (2021). Distributed one-step upgraded estimation for non-uniformly and non-randomly distributed data. *Computational Statistics and Data Analysis*, 162, Article 107265. <https://doi.org/10.1016/j.csda.2021.107265>
- Wang, H. (2009). Forward regression for ultra-high dimensional variable screening. *Journal of the American Statistical Association*, 104(488), 1512–1524. <https://doi.org/10.1198/jasa.2008.tm08516>
- Wang, H. (2019a). Divide-and-conquer information-based optimal subdata selection algorithm. *Journal of Statistical Theory and Practice*, 13(3), 46. <https://doi.org/10.1007/s42519-019-0048-5>
- Wang, H. (2019b). More efficient estimation for logistic regression with optimal subsamples. *Journal of Machine Learning Research*, 20(132), 1–59.
- Wang, H. (2020). Logistic regression for massive data with rare events. In *International Conference on Machine Learning*. PMLR.
- Wang, H., & Ma, Y. (2021). Optimal subsampling for quantile regression in big data. *Biometrika*, 108(1), 99–112. <https://doi.org/10.1093/biomet/asaa043>
- Wang, H., Yang, M., & Stufken, J. (2019b). Information-based optimal subdata selection for big data linear regression. *Journal of the American Statistical Association*, 114(525), 393–405. <https://doi.org/10.1080/01621459.2017.1408468>
- Wang, H., Zhu, R., & Ma, P. (2018). Optimal subsampling for large sample logistic regression. *Journal of the American Statistical Association*, 113(522), 829–844. <https://doi.org/10.1080/01621459.2017.1292914>
- Wang, J., Kolar, M., Srebro, N., & Zhang, T. (2017). Efficient distributed learning with sparsity. In *International Conference on Machine Learning*. PMLR.
- Wang, S., Roosta, F., Xu, P., & Mahoney, M. W. (2018). Giant: Globally improved approximate Newton method for distributed optimization. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Wang, X., Yang, Z., Chen, X., & Liu, W. (2019a). Distributed inference for linear support vector machine. *Journal of Machine Learning Research*, 20(113), 1–41.
- Woodworth, B., Patel, K. K., Stich, S., Dai, Z., Bullins, B., McMahan, B., Shamir, O., & Srebro, N. (2020). Is local SGD better than minibatch SGD? In *International Conference on Machine Learning*. PMLR.
- Wu, S., Huang, D., & Wang, H. (2023a). Network gradient descent algorithm for decentralized federated learning. *Journal of Business and Economic Statistics*, 41(3), 806–818. <https://doi.org/10.1080/07350015.2022.2074426>
- Wu, S., Huang, D., & Wang, H. (2023b). Quasi-Newton updating for large-scale distributed learning. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 85(4), 1326–1354. <https://doi.org/10.1093/jrsssb/qkad059>
- Wu, S., Zhu, X., & Wang, H. (2023c). Subsampling and jackknifing: A practically convenient solution for large data analysis with limited computational resources. *Statistica Sinica*, 33(3), 2041–2064.
- Xu, G., Sit, T., Wang, L., & Huang, C. Y. (2017). Estimation and inference of quantile regression for survival data under biased sampling. *Journal of the American Statistical Association*, 112(520), 1571–1586. <https://doi.org/10.1080/01621459.2016.1222286>
- Xu, Q., Cai, C., Jiang, C., Sun, F., & Huang, X. (2020). Block average quantile regression for massive dataset. *Statistical Papers*, 61, 141–165. <https://doi.org/10.1007/s00362-017-0932-6>
- Yang, J., Meng, X., & Mahoney, M. (2013). Quantile regression for large-scale applications. In *International Conference on Machine Learning*. PMLR.
- Yao, Y., & Wang, H. (2019). Optimal subsampling for softmax regression. *Statistical Papers*, 60(2), 585–599. <https://doi.org/10.1007/s00362-018-01068-6>
- Yu, J., Wang, H., Ai, M., & Zhang, H. (2022). Optimal distributed subsampling for maximum quasi-likelihood estimators with massive data. *Journal of the American Statistical Association*, 117(537), 265–276. <https://doi.org/10.1080/01621459.2020.1773832>
- Yu, L., Balasubramanian, K., Volgushev, S., & Erdogdu, M. (2021). An analysis of constant step size SGD in the non-convex regime: Asymptotic normality and bias. In *Advanced in neural information processing systems*. Curran Associates, Inc.
- Yuan, K., Ling, Q., & Yin, W. (2016). On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, 26(3), 1835–1854. <https://doi.org/10.1137/130943170>
- Zeiler, M. D. (2012). *AdaDelta: An adaptive learning rate method*. arXiv: 1212.5701.
- Zhang, J., & Ré, C. (2016). ParallelSGD: When does averaging help? In *International Conference on Machine Learning Workshop on Optimization in Machine Learning*. International Conference on Machine Learning (ICML).
- Zhang, Y., & Lin, X. (2015). DiSCO: Distributed optimization for self-concordant empirical loss. In *International Conference on Machine Learning*. PMLR.
- Zhang, Y., Wainwright, M. J., & Duchi, J. C. (2012). Communication-efficient algorithms for statistical optimization. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Zhao, Y., Wong, Z. S. Y., & Tsui, K. L. (2018). A framework of rebalancing imbalanced healthcare data for rare events' classification: A case of look-alike sound-alike mix-up incident detection. *Journal of Healthcare Engineering*, 2018, Article 6275435–
- Zhong, W., Wan, C., & Zhang, W. (2022). Estimation and inference for multi-kink quantile regression. *Journal of Business and Economic Statistics*, 40(3), 1123–1139. <https://doi.org/10.1080/07350015.2021.1901720>
- Zhou, L., She, X., & Song, P. X. (2023). Distributed empirical likelihood approach to integrating unbalanced datasets. *Statistica Sinica*, 33(3), 2209–2231.
- Zhu, M., Su, W., & Chipman, H. A. (2006). LAGO: A computationally efficient approach for statistical detection. *Technometrics*, 48(2), 193–205. <https://doi.org/10.1198/004017005000000643>

- Zhu, W., Chen, X., & Wu, W. B. (2023). Online covariance matrix estimation in stochastic gradient descent. *Journal of the American Statistical Association*, 118(541), 393–404. <https://doi.org/10.1080/01621459.2021.1933498>
- Zhu, X., Li, F., & Wang, H. (2021). Least-square approximation for a distributed system. *Journal of Computational and Graphical Statistics*, 30(4), 1004–1018. <https://doi.org/10.1080/10618600.2021.1923517>
- Zhu, X., Pan, R., Wu, S., & Wang, H. (2022). Feature screening for massive data analysis by subsampling. *Journal of Business and Economic Statistics*, 40(4), 1892–1903. <https://doi.org/10.1080/07350015.2021.1990771>
- Zhu, Y., Huang, D., Gao, Y., Wu, R., Chen, Y., Zhang, B., & Wang, H. (2021). Automatic, dynamic, and nearly optimal learning rate specification via local quadratic approximation. *Neural Networks*, 141, 11–29. <https://doi.org/10.1016/j.neunet.2021.03.025>
- Zhu, Y., Yu, W., Jiao, B., Mohror, K., Moody, A., & Chowdhury, F. (2019). Efficient user-level storage disaggregation for deep learning. In *2019 IEEE International Conference on Cluster Computing*.
- Zhuang, J., Cai, J., Wang, R., Zhang, J., & Zheng, W. (2019). CARE: Class attention to regions of lesion for classification on imbalanced data. In *International Conference on Medical Imaging with Deep Learning*. PMLR.
- Zinkevich, M., Weimer, M., Li, L., & Smola, A. (2010). Parallelized stochastic gradient descent. In *Advances in neural information processing systems*. Curran Associates, Inc.
- Zou, H., & Li, R. (2008). One-step sparse estimates in nonconcave penalized likelihood models. *The Annals of Statistics*, 36(4), 1509–1533.