Byzantine-Robust Distributed Online Learning: Taming Adversarial Participants in An Adversarial Environment

Xingrong Dong, *Graduate Student Member, IEEE*, Zhaoxian Wu, Qing Ling, *Senior Member, IEEE*, and Zhi Tian, *Fellow, IEEE*

Abstract—**This paper studies distributed online learning under Byzantine attacks. The performance of an online learning algorithm is often characterized by (adversarial) regret, which evaluates the quality of one-step-ahead decision-making when an environment incurs adversarial losses, and a sublinear regret bound is preferred. But we prove that, even with a class of state-of-the-art robust aggregation rules, in an adversarial environment and in the presence of Byzantine participants, distributed online gradient descent can only achieve a linear adversarial regret bound, which is tight. This is the inevitable consequence of Byzantine attacks, even though we can control the constant of the linear adversarial regret to a reasonable level. Interestingly, when the environment is not fully adversarial so that the losses of the honest participants are i.i.d. (independent and identically distributed), we show that sublinear stochastic regret, in contrast to the aforementioned adversarial regret, is possible. We develop Byzantine-robust distributed online momentum algorithms to attain such sublinear stochastic regret bounds for a class of robust aggregation rules. Numerical experiments corroborate our theoretical analysis.**

Index Terms—**Distributed optimization, Byzantine-robustness, online learning.**

I. Introduction

ONLINE learning is a powerful tool to process streaming data in a timely manner [2], [3], [4]. In response to an environment that provides (adversarial) losses sequentially, an online learning algorithm makes one-step-ahead decisions. Its performance is characterized by (adversarial) regret, which measures the accumulative difference between the losses of the online decisions and those of the overall best solution in hindsight. It is preferred that the adversarial regret increases sublinearly in time, which would lead to asymptotically vanishing performance degradation. When the streaming data are separately collected by multiple participants and data privacy is a concern, distributed online learning becomes a natural choice [5], [6], [7], [8]. Each participant makes a local decision, and a server aggregates all the local decisions to a global one [9], [10]. Exemplary applications include online web ranking and online advertisement recommendation, etc. [11], [12], [13], [14].

In addition to the sequential losses caused by the adversarial environment, distributed online learning faces a new challenge in terms of robustness, because not all the participants are guaranteed to be trustworthy. Some participants may intentionally or unintentionally send wrong messages, instead of true local decisions, to the server. These adversarial participants are termed as Byzantine participants following the notion in distributed systems to describe the worst-case attacks [15]. Therefore, an interesting question arises: *Is it possible to develop a Byzantine-robust distributed online learning algorithm with provable sublinear adversarial regret, in an adversarial environment and in the presence of adversarial participants?*

In this paper, we provide a rather negative answer to this question. We show that even equipped with a class of state-of-the-art robust aggregation rules, distributed online gradient descent algorithms can only achieve linear adversarial regret bounds, which are tight. This rather negative result highlights the difficulty of Byzantine-robust distributed online learning. The joint impact from the adversarial environment and the adversarial participants leads the online decisions to deviate from the overall best solution in hindsight, no matter how long the learning time is. Nevertheless, we stress that it is the necessary price for handling arbitrarily malicious Byzantine attacks from the adversarial participants, and with the help of the state-of-the-art robust aggregation rules, we can control the constant of linear adversarial regret to a reasonable value.

On the other hand, we further show that if the environment is not fully adversarial so that the losses of the honest participants are i.i.d. (independent and identically distributed), then sublinear stochastic regret [16], in contrast to the aforementioned

Manuscript received 20 July 2023; revised 8 November 2023; accepted 3 December 2023. Date of publication 12 December 2023; date of current version 19 December 2023. The work of Qing Ling was supported in part by the China NSF under Grants 61973324, 12126610, and 62373388, and in part by Guangdong Basic and Applied Basic Research Foundation under Grant 2021B1515020094. The work of Zhi Tian was supported by the US NSF under Grants 1939553 and 2231209. An earlier version of this paper was presented at the ICASSP, 2023 [DOI: 10.1109/ICASSP49357.2023.10095178]. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Ketan Rajawat. *(Corresponding author: Qing Ling.)*

Xingrong Dong, Zhaoxian Wu, and Qing Ling are with the School of Computer Science and Engineering, Sun Yat-Sen University, Guangzhou, Guangdong 510006, China, and also with Peng Cheng Laboratory, Shenzhen, Guangdong 518066, China (e-mail: lingqing556@mail.sysu.edu.cn).

Zhi Tian is with the Department of Electrical and Computer Engineering, George Mason University, Fairfax, VA 22030 USA.

This article has supplementary downloadable material available at https://doi.org/10.1109/TSP.2023.3340028, provided by the authors.

Digital Object Identifier 10.1109/TSP.2023.3340028

adversarial regret, is possible. Accordingly, we develop a family of Byzantine-robust distributed online gradient descent algorithms enhanced with momentum to attain such sublinear stochastic regret bounds.

The rest of this paper is organized as follows. We briefly survey the related works in Section II, and give the problem statement in Section III. The linear adversarial regret bounds of Byzantine-robust distributed online gradient descent are established in Section IV, and the sublinear stochastic regret bounds of Byzantine-robust distributed online momentum are shown in Section V. We conduct numerical experiments in Section VI, followed by conclusions in Section VII.

II. Related Works

Online learning aims at sequentially making one-step-ahead decisions in an environment that provides (adversarial) losses. Classical online learning algorithms include but are not limited to online gradient descent [17], online conditional gradient [18], online mirror descent [19], adaptive gradient [20]. We focus on online gradient descent and its variants in this paper. Their performance is often characterized by (adversarial) regret, which measures the accumulative difference between the losses of the online decisions and those of the overall best solution in hindsight. These algorithms have provable adversarial regret bounds of $\mathcal{O}(\sqrt{T})$ and $\mathcal{O}(\log T)$ for convex and strongly convex losses, respectively, where T is the time horizon. When T goes to infinity, such sublinear adversarial regret bounds imply asymptotically vanishing performance degradation in the long run.

When the streaming data are separately collected by multiple participants, data privacy becomes a big concern. Therefore, distributed online learning, which avoids transmitting raw data from the participants to the server, has attracted extensive research attention [5], [6]. Similar to their centralized counterparts, the distributed online gradient descent algorithms have provable adversarial regret bounds of $\mathcal{O}(\sqrt{T})$ and $\mathcal{O}(\log T)$ for convex and strongly convex losses, respectively [21], [22].

However, in a distributed online learning system, some of the participants can be adversarial. They do not follow the prescribed algorithmic protocol but send arbitrarily malicious messages to the server. We characterize these adversarial participants with the classical Byzantine attacks model [15]. Interestingly, Byzantine-robust distributed online learning, which investigates reliable decision-making in an adversarial environment and in the presence of adversarial participants, is rarely studied. The work of [23] focuses on the case that the environment provides linear losses, which is different to ours. The proposed asynchronous distributed online learning algorithm in [23] also lacks regret bound analysis. The work of [24] considers online mean estimation over a decentralized network without a server. There is only one malicious participant, which has a limited budget to attack and only pollutes a faction of its messages to be transmitted. The performance metric is the Euclidean distance between the true mean and the estimate. In contrast, our work considers a general distributed online learning problem, the Byzantine participants have unlimited budgets to attack and can pollute all of their messages to be transmitted, and the performance metrics are adversarial and stochastic regrets. The work of [25] considers decentralized online learning, but relaxes the problem to minimizing a convex combination of the losses. Accordingly, an $\mathcal{O}(\log^2 T)$ relaxed adversarial regret bound is established. In our work, we do not introduce any relaxation and the $\mathcal{O}(\log^2 T)$ relaxed adversarial regret bound in [25] is not comparable to ours. The work of [26] also considers decentralized online learning, but confines the number of Byzantine participants to be small so as to establish a dynamic regret bound. In contrast, our analysis on the static regret bounds allows nearly up to half of the participants to be Byzantine. Byzantine-robust decentralized meta learning is investigated in [27], and a stochastic regret bound is established.

Several recent works investigate distributed bandits under Byzantine attacks. Different from online learning, participants receive values of losses, instead of gradients or functions, from an environment. It has been shown in [28] that the proposed Byzantine-robust algorithms have linear adversarial regret bounds for multi-armed and linear-contextual problems. This is consistent with our result. Some works make the i.i.d. assumption [29], [30], [31]. The work of [29] proves $\mathcal{O}(T^{3/4})$ regret for linear bandits with high probability. The work of [30] reaches $\mathcal{O}(\sqrt{T})$ regret but requires the action set to be finite. Our proposed algorithm, with the aid of momentum, attains the $\mathcal{O}(\sqrt{T})$ stochastic regret bound. The work of [31] considers multi-armed bandits, and uses historic information to reach $\mathcal{O}(\log T)$ regret, which is consistent with our stochastic regret bound established for Byzantine-robust distributed online momentum. The work of [32] is free of the i.i.d. assumption, but the regret for multi-armed bandits is defined with respect to a suboptimal solution other than the optimal one. Therefore, the derived $\mathcal{O}(\log T)$ sublinear regret bound is not comparable to others.

Another tightly related area is Byzantine-robust distributed stochastic optimization [33], [34], [35]. Therein, the basic idea is to replace the vulnerable mean aggregation rule in distributed stochastic gradient descent with robust aggregation rules, including coordinate-wise median [36], trimmed mean [36], [37], geometric median [38], Krum [39], centered clipping [40], Phocas [41], FABA [42], etc. Most of them belong to the category of robust bounded aggregation rules (see Definition 1). We will incorporate these robust bounded aggregation rules with distributed online gradient descent and momentum to enable Byzantine-robustness.

In Table I, we compare the adversarial regret bounds of distributed online gradient descent with the mean aggregation rule and without Byzantine attacks, the derived adversarial regret bounds of Byzantine-robust distributed online gradient descent with robust bounded aggregation rules, as well as the derived stochastic regret bounds of Byzantine-robust distributed online momentum with robust bounded aggregation rules.

III. Problem Statement

Consider n participants in a set $\mathcal{N}$, among which h are honest and in subset $\mathcal{H}$, while b are Byzantine and in subset $\mathcal{B}$.

TABLE I
REGRET BOUNDS OF DIFFERENT ALGORITHMS

	constant step size	diminishing step size
Byzantine-free mean[1]	$\mathcal{O}(\sqrt{T})$	$\mathcal{O}(\log T)$
Byzantine-robust[2]	$\mathcal{O}(T)$	$\mathcal{O}(T)$
Byzantine-robust momentum[3]	$\mathcal{O}(\sqrt{T})$	$\mathcal{O}(\log T)$

[1] Adversarial regret bounds of distributed online gradient descent with the mean aggregation rule and without Byzantine attacks.

[2] Adversarial regret bounds of Byzantine-robust distributed online gradient descent with robust bounded aggregation rules.

[3] Stochastic regret bounds of Byzantine-robust distributed online momentum with robust bounded aggregation rules.

We know $n = h + b$, but the identities of Byzantine participants are unknown and we can only roughly estimate an upper bound of b. At step t, each honest participant j makes its local decision of the model parameters $w_t^j \in \mathbb{R}^d$ and sends it to the server, while each Byzantine participant j sends an arbitrarily malicious message to the server. For notational convenience, denote $z_t^j \in \mathbb{R}^d$ as the message sent by participant j to the server at step t, no matter whether it is from an honest or Byzantine participant. Upon receiving all z_t^j, the server aggregates them to yield a global decision of the model parameters $w_t \in \mathbb{R}^d$. The quality of the sequential decisions over T steps is often evaluated by adversarial regret with respect to the overall best solution in hindsight, given by

$$R_T := \sum_{t=1}^{T} f_t(w_t) - \min_{w \in \mathbb{R}^d} \sum_{t=1}^{T} f_t(w), \tag{1}$$

where

$$f_t(w) := \frac{1}{h} \sum_{j \in \mathcal{H}} f_t^j(w), \tag{2}$$

and f_t^j is the loss revealed to $j \in \mathcal{H}$ at the end of step t.

For distributed online gradient descent, each honest participant $j \in \mathcal{H}$ makes its local decision following

$$w_{t+1}^j = w_t - \eta_t \nabla f_t^j(w_t), \tag{3}$$

where $\eta_t > 0$ is the step size, and sends $z_{t+1}^j = w_{t+1}^j$ to the server. The server aggregates the messages z_{t+1}^j to yield the mean value

$$w_{t+1} = \frac{1}{n} \sum_{j=1}^{n} z_{t+1}^j. \tag{4}$$

However, messages z_{t+1}^j from the Byzantine participants $j \in \mathcal{B}$ are arbitrarily malicious, such that w_{t+1} can be manipulated to reach arbitrarily large adversarial regret.

Motivated by the recent advances of Byzantine-robust distributed stochastic optimization, one may think of using robust aggregation rules to replace the vulnerable mean aggregation rule in (4). Denote AGG as a proper robust aggregation rule. Now, the server makes the decision as

$$w_{t+1} = AGG(z_{t+1}^1, z_{t+1}^2, \cdots, z_{t+1}^n). \tag{5}$$

Below we introduce two exemplary robust aggregation rules. More examples can be found in the extended version of this paper [43]. For notational convenience, we denote $\mathcal{Z}_{t+1} := \{z_{t+1}^1, z_{t+1}^2, \cdots, z_{t+1}^n\}$ as the set of the n received messages and $\mathcal{Z}_{t+1}[k] := \{z_{t+1}^1[k], z_{t+1}^2[k], \cdots, z_{t+1}^n[k]\}$ as the set of their k-th elements, where $k \in [d]$.

Coordinate-wise median. It yields the median for each dimension, given by

$$\begin{aligned} &\text{coomed}(\mathcal{Z}_{t+1}) \\ &\quad := [\text{med}(\mathcal{Z}_{t+1}[1]); \text{med}(\mathcal{Z}_{t+1}[2]); \cdots; \text{med}(\mathcal{Z}_{t+1}[d])], \end{aligned} \tag{6}$$

where $\text{med}(\cdot)$ calculates the median of the input scalars.

Trimmed mean. It is also coordinate-wise. Let $q < \frac{n}{2}$ be the estimated number of Byzantine participants. Given $\mathcal{Z}_{t+1}[k]$, trimmed mean removes the largest q inputs and the smallest q inputs, and then averages the rest to yield $\text{trimean}(\mathcal{Z}_{t+1}[k])$. The results of the d dimensions are stacked to yield

$$\begin{aligned} &\text{trimean}(\mathcal{Z}_{t+1}) \\ &\quad := [\text{trimean}(\mathcal{Z}_{t+1}[1]); \text{trimean}(\mathcal{Z}_{t+1}[2]); \\ &\qquad \cdots; \text{trimean}(\mathcal{Z}_{t+1}[d])]. \end{aligned} \tag{7}$$

IV. LINEAR ADVERSARIAL REGRET BOUNDS OF BYZANTINE-ROBUST DISTRIBUTED ONLINE GRADIENT DESCENT

Robust aggregation rules have been proven effective in distributed stochastic optimization, given that the fraction of Byzantine participants $\alpha = \frac{b}{n}$ is less than $\frac{1}{2}$ [36], [37], [38], [39], [40], [41], [42]. Thus, one may wonder whether the Byzantine-robust distributed online gradient descent updates (3) and (5) can achieve sublinear adversarial regret.

Our answer is negative. Even with a wide class of *robust bounded aggregation* rules, the tight adversarial regret bounds are linear.

Definition 1: (Robust bounded aggregation rule). Consider n messages $z_t^1, z_t^2, \cdots, z_t^n \in \mathbb{R}^d$ from h honest participants in $\mathcal{H}$ and b Byzantine participants in $\mathcal{B}$. The fraction of Byzantine participants $\alpha = \frac{b}{n} < \frac{1}{2}$. An aggregation rule AGG is a robust bounded aggregation rule, if the difference between its output and the mean of the honest messages is bounded by

$$\|w_t - \bar{z}_t\|^2 = \|AGG(z_t^1, z_t^2, \cdots, z_t^n) - \bar{z}_t\|^2 \leq C_\alpha^2 \zeta^2,$$

where $\bar{z}_t := \frac{1}{h} \sum_{j \in \mathcal{H}} z_t^j$ is the mean of the honest messages, ζ^2 is the largest deviation of the honest messages such that $\|\bar{z}_t - z_t^j\|^2 \leq \zeta^2$ for all $j \in \mathcal{H}$, and C_α is an aggregation-specific constant determined by α.

In Definition 1, ζ^2 characterizes the heterogeneity of the messages to be aggregated. For a robust bounded aggregation rule, the difference between its output and the mean of the honest messages is bounded by $C_\alpha^2 \zeta^2$. Therefore, a robust bounded aggregation rule with a smaller C_α^2 can better handle the heterogeneity of the messages to be aggregated.

We show that a number of state-of-the-art robust aggregation rules, including coordinate-wise median [36], trimmed mean

TABLE II
CONSTANTS C_α OF ROBUST BOUNDED AGGREGATION RULES, WITH α BEING THE FRACTION OF BYZANTINE PARTICIPANTS

	C_α
coordinate-wise median	$\mathcal{O}(\frac{1}{1-\alpha})$
trimmed mean	$\mathcal{O}(\frac{\sqrt{\alpha(1-\alpha)}}{(1-2\alpha)})$
geometric median	$\mathcal{O}(\frac{1-\alpha}{1-2\alpha})$
Krum	$\mathcal{O}(1+\sqrt{\frac{1-\alpha}{1-2\alpha}})$
centered clipping	$\mathcal{O}(\sqrt{\alpha})$
Phocas	$\mathcal{O}(\sqrt{1+\frac{\alpha(1-\alpha)}{(1-2\alpha)^2}})$
FABA	$\mathcal{O}(\frac{\alpha}{1-3\alpha})$

[36], [37], geometric median [38], Krum [39], centered clipping[1] [40], Phocas [41], and FABA[2] [42], all belong to robust bounded aggregation rules. Their corresponding constants C_α are listed in Table II and the derivations of these constants are left to the extended version of this paper [43].

Note that $\bar{z}_t$ is only used for the purpose of theoretical analysis. The server does not need to calculate the value of $\bar{z}_t$ during implementing a robust bounded aggregation rule.

To analyze the adversarial regret bounds, we make the following standard assumptions on the losses of any honest participant $j \in \mathcal{H}$.

Assumption 1: (L-smoothness). f_t^j is differentiable and has Lipschitz continuous gradients. For any $x, y \in \mathbb{R}^d$, there exists a constant $L > 0$ such that

$$||\nabla f_t^j(x) - \nabla f_t^j(y)|| \leq L||x-y||. \tag{8}$$

Assumption 2: (μ-strong convexity). f_t^j is strongly convex. For any $x, y \in \mathbb{R}^d$, there exists a constant $\mu > 0$ such that

$$\langle \nabla f_t^j(x), x-y \rangle \geq f_t^j(x) - f_t^j(y) + \frac{\mu}{2}\|x-y\|^2. \tag{9}$$

Assumption 3: (Bounded deviation). Define $\nabla f_t(w_t) := \frac{1}{h}\sum_{j\in\mathcal{H}} \nabla f_t^j(w_t)$. The deviation between each honest gradient $\nabla f_t^j(w_t)$ and the mean of the honest gradients is bounded by

$$||\nabla f_t^j(w_t) - \nabla f_t(w_t)||^2 \leq \sigma^2. \tag{10}$$

Assumption 4: (Bounded gradient at the overall best solution). Define $w^* = \arg\min_{w\in\mathbb{R}^d} \sum_{t=1}^T f_t(w)$ as the overall best solution. The mean of the honest gradients at this point is upper bounded by

$$\|\frac{1}{h}\sum_{j\in\mathcal{H}} \nabla f_t^j(w^*)\|^2 \leq \xi^2. \tag{11}$$

These assumptions are common in the analysis of online learning algorithms. Some works make stronger assumptions [2], [3], [4], for example, bounded variable or bounded gradient that yields Assumptions 3 and 4.

Next, we show that the Byzantine-robust distributed online gradient descent algorithm with a robust bounded aggregation rule can only reach a linear adversarial regret bound under Byzantine attacks. The proof is left to Appendix A. In contrast, the distributed online gradient descent algorithm with the mean aggregation rule can achieve a sublinear adversarial regret bound without Byzantine attacks, as shown in Appendix C of [43]. To distinguish the adversarial regret bounds with different step sizes, we denote $R_{T:\eta}, R_{T:\frac{1}{\sqrt{T}}}, R_{T:\frac{1}{t}}$ as the adversarial regrets with a constant step size η, a special constant step size $\mathcal{O}(\frac{1}{\sqrt{T}})$, and a diminishing step size $\mathcal{O}(\frac{1}{t})$, respectively.

Theorem 1: Suppose that the fraction of Byzantine participants $\alpha = \frac{b}{n} < \frac{1}{2}$. Under Assumptions 1, 2, 3, and 4, the Byzantine-robust distributed online gradient descent updates (3) and (5) with a robust bounded aggregation and a constant step size $\eta_t = \eta \in (0, \frac{1}{4L}]$ have an adversarial regret bound

$$R_{T:\eta} \leq \frac{1}{\eta}\|w_1 - w^*\|^2 + \left(2\eta + \frac{8L^2\eta^2}{\mu}\right)\xi^2 T + \frac{2}{\mu}C_\alpha^2\sigma^2 T. \tag{12}$$

In particular, if $\eta_t = \eta = \frac{c}{\sqrt{T}}$ where c is a sufficiently small positive constant, then the adversarial regret bound becomes

$$R_{T:\frac{1}{\sqrt{T}}} \leq \frac{8L^2c^2}{\mu}\xi^2 + \left(\frac{\|w_1 - w^*\|^2}{c} + 2c\xi^2\right)\sqrt{T} + \frac{2}{\mu}C_\alpha^2\sigma^2 T. \tag{13}$$

If we use a diminishing step size $\eta_t = \min\{\frac{1}{4L}, \frac{8}{\mu t}\}$, then the adversarial regret bound is

$$R_{T:\frac{1}{t}} \leq 4L\|w_1 - w^*\|^2 + \frac{48L}{\mu^2}\xi^2 \log T + \frac{2}{\mu}C_\alpha^2\sigma^2 T. \tag{14}$$

We construct the following counter-example to show that the derived $\mathcal{O}(\sigma^2 T)$ linear adversarial regret bound is tight.

Example 1: Consider a distributed online learning system with 3 participants, among which participant 3 is Byzantine. Thus, $\mathcal{N} = \{1,2,3\}$, $\mathcal{H} = \{1,2\}$ and $\mathcal{B} = \{3\}$. Suppose that at any step t, the losses of participants 1 and 2 are respectively given by

$$f_t^1(w) = \frac{1}{2}(w-\sigma)^2, \quad f_t^2(w) = \frac{1}{2}(w+\sigma)^2.$$

It is easy to check that these losses satisfy Assumptions 1, 2, 3, and 4. To be specific, the overall best solution $w^* = 0$, $L = 1$, $\mu = 1$, and $\xi^2 = 0$.

Take geometric median as an exemplary aggregation rule. Suppose that the algorithm is initialized by $w_1 = \sigma$. At step 2, participant 1 sends $z_2^1 = w_2^1 = w_1 - \eta_1(w_1 - \sigma) = \sigma$, while participant 2 sends $z_2^2 = w_2^2 = w_1 - \eta_1(w_1 + \sigma) = \sigma - 2\eta_1\sigma$. In this circumstance, participant 3, who is Byzantine, can send $z_2^3 = \sigma$ so that the aggregation result is $w_2 = \sigma$. As such, for any step t, $w_t = \sigma$, $f_t(w_t) = \sigma^2$, $f_t(w^*) = \frac{1}{2}\sigma^2$, and the adversarial regret is $\frac{1}{2}\sigma^2 T$.

For other robust bounded aggregation rules, we can observe that the mean of the honest messages is $\bar{z}_{t+1} = (1-\eta_t)\sigma$ and the largest deviation is $\zeta^2 = \eta_t^2\sigma^2$. According to Definition 1, participant 3 can always manipulate its message so that the

[1] Centered clipping requires $\alpha \leq 0.1$.

[2] FABA requires $\alpha < \frac{1}{3}$.

aggregation result is in the order of σ, which eventually yields linear adversarial regret. If the aggregation rule is majority-voting-based, such as coordinate-wise median and trimmed mean, sending $z_{t+1}^3 = \sigma$ is effective. For centered clipping, participant 3 can send $z_{t+1}^3 = \sigma + 2\eta_t\sigma$ instead.

Note that Example 1 holds for both constant and diminishing step sizes. Meanwhile, Example 1 can be extended to a larger number of participants.

The linear adversarial regret bound seems frustrating, but is the necessary price for handling arbitrarily malicious Byzantine attacks from the adversarial participants. With the help of robust bounded aggregation rules, we are able to control the constant of linear adversarial regret to a reasonable value $\frac{2}{\mu}C_\alpha^2\sigma^2$, which is determined by the property of losses, the robust bounded aggregation rule, the fraction of Byzantine participants, and the gradient deviation among honest participants.

V. Sublinear Stochastic Regret Bounds of Byzantine-Robust Distributed Online Momentum

According to Theorem 1, the established linear adversarial regret bounds are proportional to σ^2, the maximum between the honest gradients and the mean of the honest gradients. This makes sense as the disagreement among the honest participants is critical, especially in an adversarial environment. This observation motivates us to investigate whether it is possible to attain sublinear regret bounds when the disagreement among the honest participants is well-controlled.

To this end, suppose that the environment provides all the honest participants with independent losses from the same distribution $\mathcal{D}$ at all steps. Define the expected loss $F(w) := \mathbb{E}_{\mathcal{D}} f_t^j(w)$ for all $j \in \mathcal{H}$ and all t. Then, stochastic regret is defined as

$$S_T := \mathbb{E}\sum_{t=1}^{T} F(w_t) - T \cdot \min_{w\in\mathbb{R}^d} F(w), \tag{15}$$

where the expectation is taken over the stochastic process [16]. In such an i.i.d. setting, the notion of stochastic regret is natural and has been widely adopted [16], [44], [45]. Note that the works of [29] and [30], which investigate the problem of Byzantine-robust distributed bandits, also make a similar i.i.d. assumption.

However, naively applying robust bounded aggregation rules to (3) and (5) cannot guarantee sublinear stochastic regret, since the random perturbations of the honest losses still accumulate over time and the disagreement among the honest participants does not diminish. Motivated by the successful applications of variance reduction techniques in Byzantine-robust distributed stochastic optimization [40], [46], [47], [48], [49], we let each honest participant perform momentum steps, instead of gradient descent steps, to gradually eliminate the disagreement during the learning process.

In Byzantine-robust distributed online gradient descent with momentum, each honest participant j maintains a momentum vector

$$m_t^j = \nu_t \nabla f_t^j(w_t) + (1-\nu_t) m_{t-1}^j, \tag{16}$$

where $\nu_t \in (0,1)$ is the momentum parameter. Then, it makes the local decision following

$$w_{t+1}^j = w_t - \eta_t m_t^j, \tag{17}$$

instead of (3) and sends to the server. The server still aggregates the messages and makes the decision as (5).

The ensuing analysis needs the following assumptions on the expected loss, in lieu of Assumptions 1, 2 and 3 on the individual losses.

Assumption 5: (L-smoothness). F is differentiable and has Lipschitz continuous gradients. For any $x, y \in \mathbb{R}^d$, there exists a constant $L > 0$ such that

$$||\nabla F(x) - \nabla F(y)|| \le L||x-y||. \tag{18}$$

Assumption 6: (μ-strong convexity). F is strongly convex. For any $x, y \in \mathbb{R}^d$, there exists a constant $\mu > 0$ such that

$$\langle \nabla F(x), x-y\rangle \ge F(x) - F(y) + \frac{\mu}{2}\|x-y\|^2. \tag{19}$$

Assumption 7: (Bounded variance). The variance of each honest gradient $\nabla f_t^j(w_t)$ is bounded by

$$\mathbb{E}||\nabla f_t^j(w_t) - \nabla F(w_t)||^2 \le \sigma^2. \tag{20}$$

In the investigated i.i.d. setting, the overall best solution $w^* = \arg\min_{w\in\mathbb{R}^d} F(w)$ makes $\nabla F(w^*) = 0$, such that we no longer need to bound the gradient at the overall best solution as in Assumption 4.

Theorem 2: Suppose that the fraction of Byzantine participants $\alpha = \frac{b}{n} < \frac{1}{2}$ and that each honest participant j draws its loss f_t^j at step t from distribution $\mathcal{D}$ with expectation $F := \mathbb{E}_{\mathcal{D}} f_t^j$. Under Assumptions 5, 6 and 7, the Byzantine-robust distributed online momentum updates (16), (17) and (5) with a robust bounded aggregation rule, a constant step size $\eta_t = \eta \in (0, \frac{\mu}{16L^2})$ and a constant momentum parameter $\nu_t = \nu = \frac{8\sqrt{3}L^2}{\mu}\eta$ have a stochastic regret bound

$$S_{T:\eta} \le \mathcal{O}\left(\frac{1}{\eta} + \frac{\sigma^2}{h}\left(1 + h^2 C_\alpha^2\right)\frac{L^4}{\mu^4}\eta T\right). \tag{21}$$

In particular, if $\eta_t = \eta = \mathcal{O}(\frac{1}{\sqrt{T}})$ and $\nu_t = \nu = \mathcal{O}(\frac{1}{\sqrt{T}})$ are properly chosen, then the stochastic regret bound becomes

$$S_{T:\frac{1}{\sqrt{T}}} = \mathcal{O}\left(\sqrt{T} + \frac{\sigma^2}{h}\left(1 + h^2 C_\alpha^2\right)\frac{L^4}{\mu^4}\sqrt{T}\right). \tag{22}$$

If we use a proper diminishing step size $\eta_t = \mathcal{O}(\frac{1}{t})$ and a proper momentum parameter $\nu_t = \mathcal{O}(\frac{1}{t})$, then the stochastic regret bound is

$$S_{T:\frac{1}{t}} = \mathcal{O}\left(\frac{\sigma^2}{h}\left(1 + h^2 C_\alpha^2\right)\frac{L^4}{\mu^4}\log T\right). \tag{23}$$

The proof is left to Appendix B. With proper constant and diminishing step sizes, Theorem 2 establishes the $\mathcal{O}(\sqrt{T})$ and $\mathcal{O}(\log T)$ stochastic regret bounds of Byzantine-robust distributed online momentum in the i.i.d. setting. In the sublinear stochastic regret bounds (22) and (23), the coefficient $\frac{\sigma^2}{h}$ is inversely proportional to h, the number of honest participants, which highlights the benefit of collaboration. The constant is also determined by C_α that characterizes the defense ability of the robust bounded aggregation rule. Smaller C_α yields smaller stochastic regret. Besides, some robust bounded aggregation rules, including trimmed mean, centered clipping and FABA, have $C_\alpha = 0$ when $\alpha = 0$, namely, no Byzantine participants are present. In this case, the derived stochastic regret bounds respectively degenerate to $\mathcal{O}((\sigma^2/h)\sqrt{T})$ and $\mathcal{O}((\sigma^2/h)\log T)$.

The i.i.d. assumption is essential to the sublinear stochastic regret bound. Without the i.i.d. assumption, we show that Byzantine-robust distributed online momentum has a tight linear stochastic regret bound in Example 2, similar to the construction in Example 1.

Example 2: Consider a distributed online learning system with 3 participants, among which participant 3 is Byzantine. Thus, $\mathcal{N} = \{1, 2, 3\}$, $\mathcal{H} = \{1, 2\}$ and $\mathcal{B} = \{3\}$. Suppose that at any step t, the losses of participants 1 and 2 are respectively given by

$$f_t^1(w) = \frac{1}{2}(w - \sigma)^2, \quad f_t^2(w) = \frac{1}{2}(w + \sigma)^2.$$

The losses of participants 1 and 2 are non-i.i.d. and the expected loss of the honest participants is

$$F_t(w) = \frac{1}{2}\left(\frac{1}{2}(w - \sigma)^2 + \frac{1}{2}(w + \sigma)^2\right) = \frac{1}{2}(w^2 + \sigma^2).$$

It is easy to check that these losses satisfy Assumptions 5, 6 and 7. To be specific, the overall best solution $w^* = 0$, $L = 1$, and $\mu = 1$.

Take geometric median as an exemplary aggregation rule. Suppose that the algorithm is initialized by $w_1 = \sigma$, $m_0^1 = 0$ and $m_0^2 = 2\sigma$. Thus, $m_1^1 = \nu_1(w_1 - \sigma) + (1 - \nu_1)m_0^1 = 0$ and $m_1^2 = \nu_1(w_1 + \sigma) + (1 - \nu_1)m_0^2 = 2\sigma$. At step 2, participant 1 sends $z_2^1 = w_2^1 = w_1 - \eta_1 m_1^1 = \sigma$, while participant 2 sends $z_2^2 = w_2^2 = w_1 - \eta_1 m_1^2 = \sigma - 2\eta_1\sigma$. In this circumstance, participant 3, who is Byzantine, can send $z_2^3 = \sigma$ so that the aggregation result is $w_2 = \sigma$. As such, for any step t, $w_t = \sigma$, $F_t(w_t) = \sigma^2$, $F_t(w^*) = \frac{1}{2}\sigma^2$, and the stochastic regret is $\mathbb{E}\sum_{t=1}^T (F_t(w_t) - F_t(w^*)) = \frac{1}{2}\sigma^2 T$.

For other robust bounded aggregation rules, we can observe that the mean of the honest messages is $\bar{z}_{t+1} = (1 - \eta_t)\sigma$ and the largest deviation is $\zeta^2 = \eta_t^2\sigma^2$. According to Definition 1, participant 3 can always manipulate its message so that the aggregation result is in the order of σ, which eventually yields linear stochastic regret. If the aggregation rule is majority-voting-based, such as coordinate-wise median and trimmed mean, sending $z_{t+1}^3 = \sigma$ is effective. For centered clipping, participant 3 can send $z_{t+1}^3 = \sigma + 2\eta_t\sigma$ instead.

But on the other hand, the momentum technique is critical to the sublinear stochastic regret bound. In contrast, we can show that Byzantine-robust distributed online gradient descent without momentum has an undesired tight linear stochastic regret bound even with the i.i.d. assumption; see Example 3.

Example 3: Consider a distributed online learning system with 3 participants, among which participant 3 is Byzantine. Thus, $\mathcal{N} = \{1, 2, 3\}$, $\mathcal{H} = \{1, 2\}$ and $\mathcal{B} = \{3\}$. Suppose that at any step t, the losses of participants 1 and 2 are independently sampled from the following two functions with the same probability:

$$f_1(w) = \frac{1}{2}(w - \sigma)^2, \quad f_2(w) = \frac{1}{2}(w + \sigma)^2.$$

The losses of participants 1 and 2 are i.i.d. and the expected loss of honest participants is

$$F_t(w) = \frac{1}{2}\left(\frac{1}{2}(w - \sigma)^2 + \frac{1}{2}(w + \sigma)^2\right) = \frac{1}{2}(w^2 + \sigma^2).$$

It is easy to check that these losses satisfy Assumptions 5, 6 and 7. To be specific, the overall best solution $w^* = 0$, $L = 1$, and $\mu = 1$.

Take geometric median as an exemplary aggregation rule. Suppose that the algorithm is initialized by $w_1 = \frac{\sigma}{2}$. At step 2, participant 1 sends $z_2^1 = w_2^1 = w_1 - \eta_1(w_1 - \sigma) = (1 - \eta_1)w_1 + \eta_1\sigma$ or $z_2^1 = w_2^1 = w_1 - \eta_1(w_1 + \sigma) = (1 - \eta_1)w_1 - \eta_1\sigma$, each with a 50% probability. Participant 2 sends z_2^2 whose distribution is the same as that of z_2^1. In this circumstance, participant 3, who is Byzantine, can send $z_2^3 = (1 - \eta_1)w_1 + \eta_1\sigma$ so that the aggregation result is $w_2 = (1 - \eta_1)w_1 + \eta_1\sigma$ with a 75% probability or $w_2 = (1 - \eta_1)w_1 - \eta_1\sigma$ with a 25% probability. Thus, the expected aggregation result at step 2 is $\mathbb{E}w_2 = (1 - \eta_1)w_1 + \eta_1\frac{\sigma}{2}$.

At step 3, participant 1 sends $z_3^1 = w_3^1 = w_2 - \eta_2(w_2 - \sigma) = (1 - \eta_2)w_2 + \eta_2\sigma$ or $z_3^1 = w_3^1 = w_2 - \eta_2(w_2 + \sigma) = (1 - \eta_2)w_2 - \eta_2\sigma$, each with a 50% probability. Participant 2 sends z_3^2 whose distribution is the same as that of z_3^1. In this circumstance, participant 3, who is Byzantine, can send $z_3^3 = (1 - \eta_2)w_2 + \eta_2\sigma$ so that the aggregation result is $w_3 = (1 - \eta_2)w_2 + \eta_2\sigma$ with a 75% probability or $w_3 = (1 - \eta_2)w_2 - \eta_2\sigma$ with a 25% probability. Thus, the expected aggregation result at step 3 is $\mathbb{E}w_3 = (1 - \eta_2)\mathbb{E}w_2 + \eta_2\frac{\sigma}{2}$.

As such, for any step $t+1$, we have $\mathbb{E}w_{t+1} = (1 - \eta_t)\mathbb{E}w_t + \eta_t\frac{\sigma}{2}$. With the initialization $w_1 = \frac{\sigma}{2}$, for any step t, we get $\mathbb{E}w_t = \frac{\sigma}{2}$, $F_t(w_t) = \frac{1}{2}(w_t^2 + \sigma^2)$, $F_t(w^*) = \frac{1}{2}\sigma^2$, and the stochastic regret is at least

$$\begin{aligned}\mathbb{E}\sum_{t=1}^T (F_t(w_t) - F_t(w^*)) &= \frac{1}{2}\mathbb{E}\sum_{t=1}^T w_t^2 \\ &\geq \frac{1}{2}\sum_{t=1}^T (\mathbb{E}w_t)^2 = \frac{1}{8}\sigma^2 T.\end{aligned}$$

For other robust bounded aggregation rules, we can observe that the expected mean of the honest messages is $\mathbb{E}\bar{z}_{t+1} = (1 - \eta_t)\frac{\sigma}{2}$ and the largest deviation is $\zeta^2 = \eta_t^2\sigma^2$. According to Definition 1, participant 3 can always manipulate its message so that the expected aggregation result is in the order of σ, which eventually yields linear stochastic regret. If the aggregation rule is majority-voting-based, such as coordinate-wise median and trimmed mean, sending $z_{t+1}^3 = (1 - \eta_t)w_t + \eta_t\sigma$ is effective. For centered clipping, participant 3 can send $z_{t+1}^3 = (1 - \eta_t)w_t + 3\eta_t\sigma$ if $z_{t+1}^1 \neq z_{t+1}^2$ or $z_{t+1}^3 = z_{t+1}^1$ if $z_{t+1}^1 = z_{t+1}^2$ instead.

Remark 1: When the environment provides all the honest participants with independent losses from the same distribution

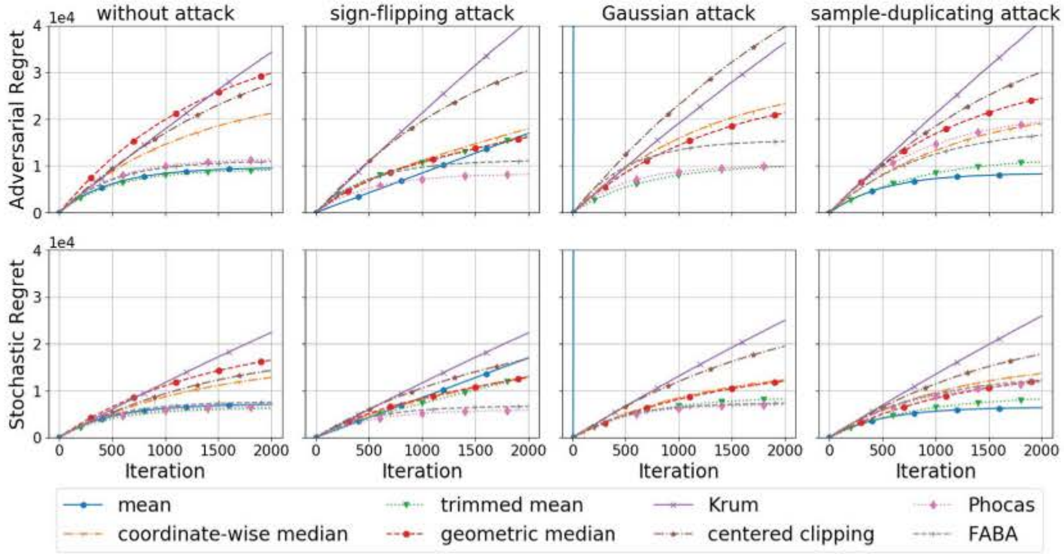

Fig. 1. Byzantine-robust distributed online gradient descent for least-squares regression on synthetic i.i.d. data with constant step size.

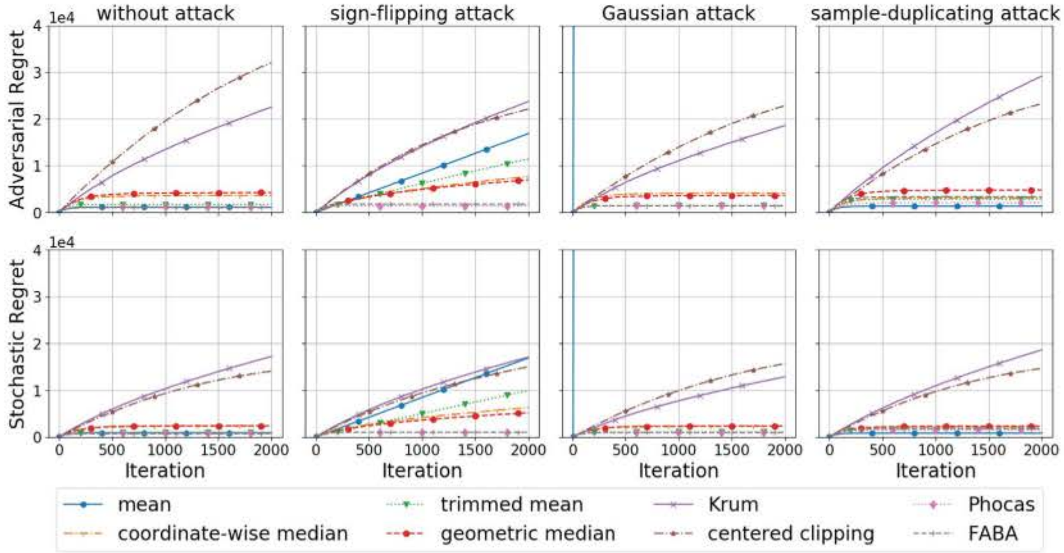

Fig. 2. Byzantine-robust distributed online gradient descent for least-squares regression on synthetic i.i.d. data with diminishing step size.

$\mathcal{D}$, the online learning and stochastic optimization formulations share similarities. However, they are from different perspectives, one is sequentially making decisions against a possibly adversarial environment with the objective of minimizing the regret, while another is actively sampling losses to approach the minimizer of the expected loss. In addition, from the online learning perspective, we can adopt different performance metrics, such as dynamic regret when the underlying distribution is time-varying [2]. Our results can also be extended to new online learning algorithms, such as online conditional gradient [18], online mirror descent [19], adaptive gradient [20], etc.

VI. Numerical Experiments

In this section, we show the performance of the Byzantine-robust distributed online gradient descent and momentum algorithms through numerical experiments, including least-squares regression on synthetic datasets, softmax regression on the MNIST dataset and Resnet18 training on the CIFAR10 dataset. Due to the page limit, we left Resnet18 training on the CIFAR10 dataset to the extended version of this paper [43]. The source code is available online[3].

In addition to the non-robust mean aggregation rule, we test seven robust bounded aggregation rules, including coordinate-wise median, trimmed mean, geometric median, Krum, centered clipping, Phocas, and FABA. We consider the following three commonly-used Byzantine attacks.

Sign-flipping attack. Each Byzantine participant sends a negative multiple of its true message, and the coefficient is set as -3, -1 and -1 for the three numerical experiments, respectively.

[3] https://github.com/wanger521/OGD

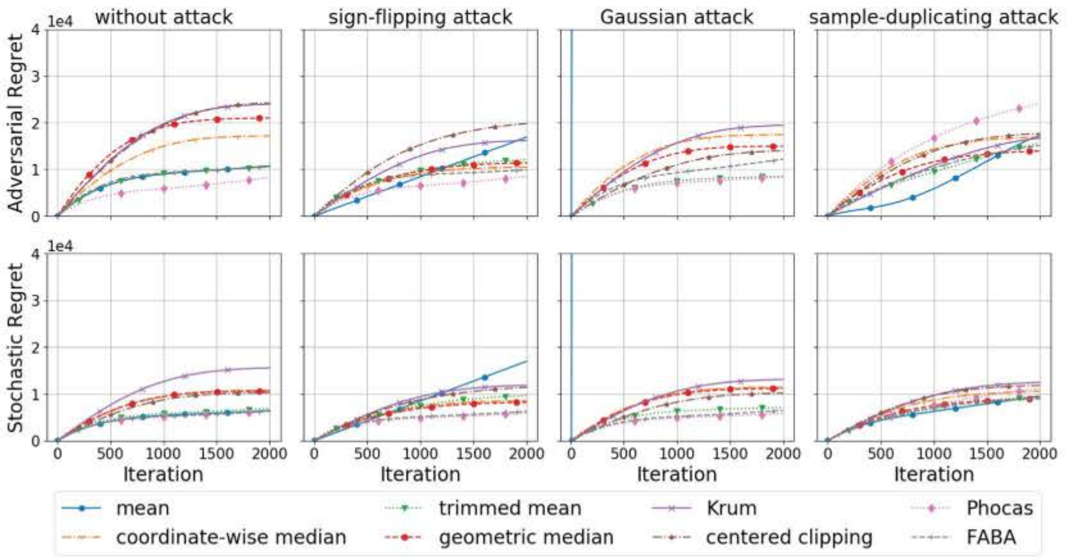

Fig. 3. Byzantine-robust distributed online momentum for least-squares regression on synthetic i.i.d. data with constant step size.

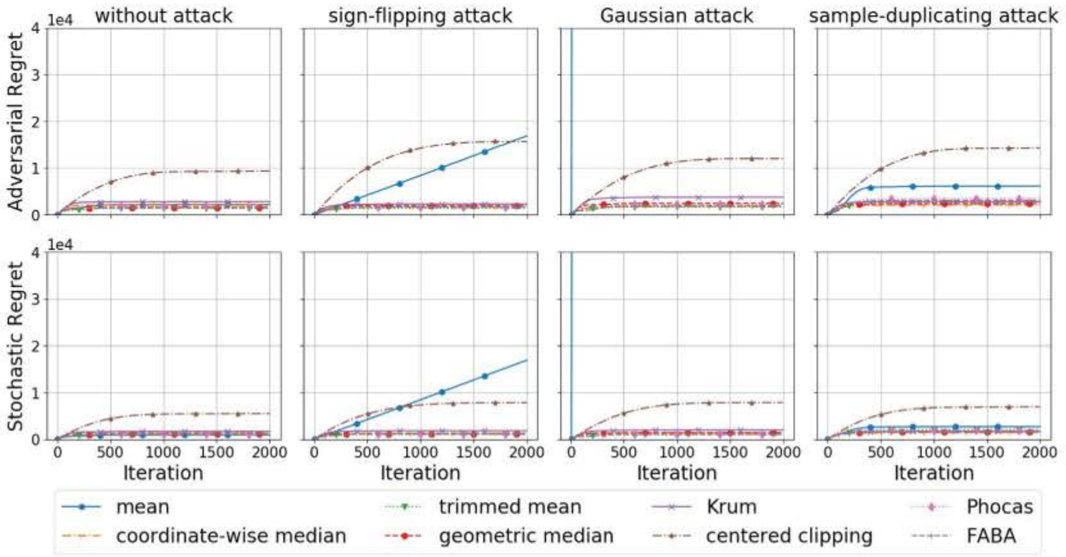

Fig. 4. Byzantine-robust distributed online momentum for least-squares regression on synthetic i.i.d. data with diminishing step size.

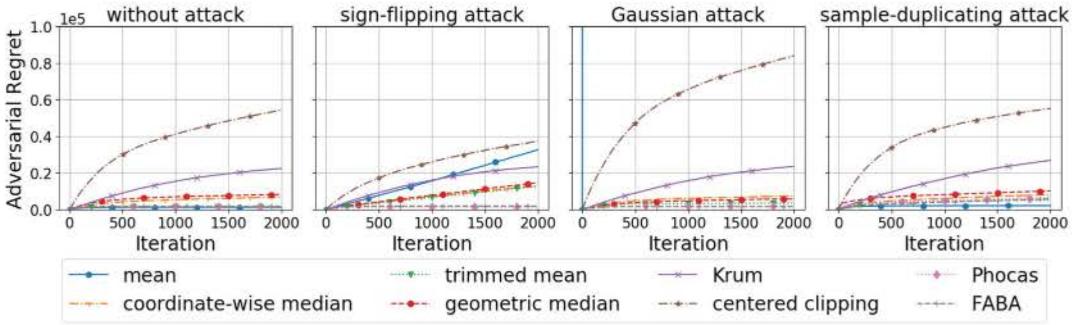

Fig. 5. Byzantine-robust distributed online gradient descent for least-squares regression on synthetic non-i.i.d. data with constant step size.

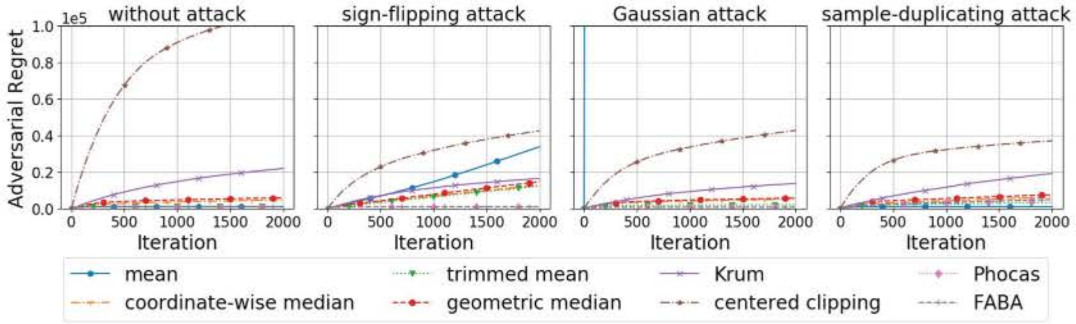

Fig. 6. Byzantine-robust distributed online gradient descent for least-squares regression on synthetic non-i.i.d. data with diminishing step size.

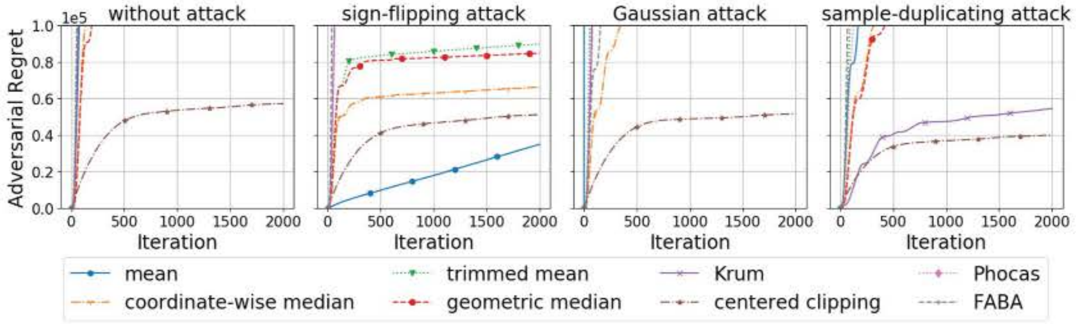

Fig. 7. Byzantine-robust distributed online momentum for least-squares regression on synthetic non-i.i.d. data with constant step size.

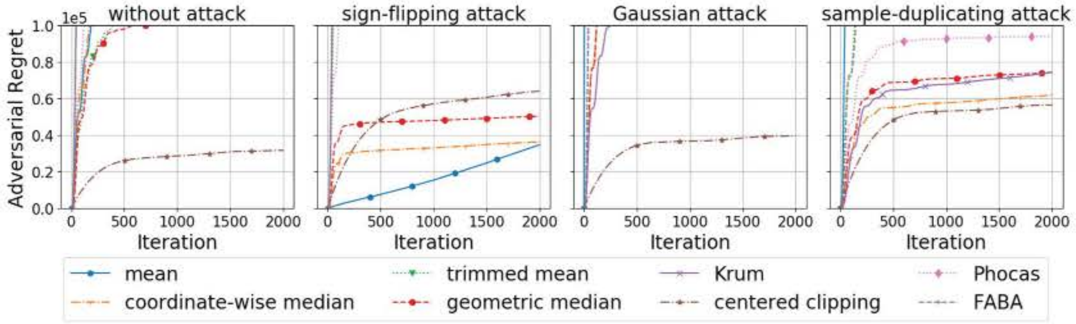

Fig. 8. Byzantine-robust distributed online momentum for least-squares regression on synthetic non-i.i.d. data with diminishing step size.

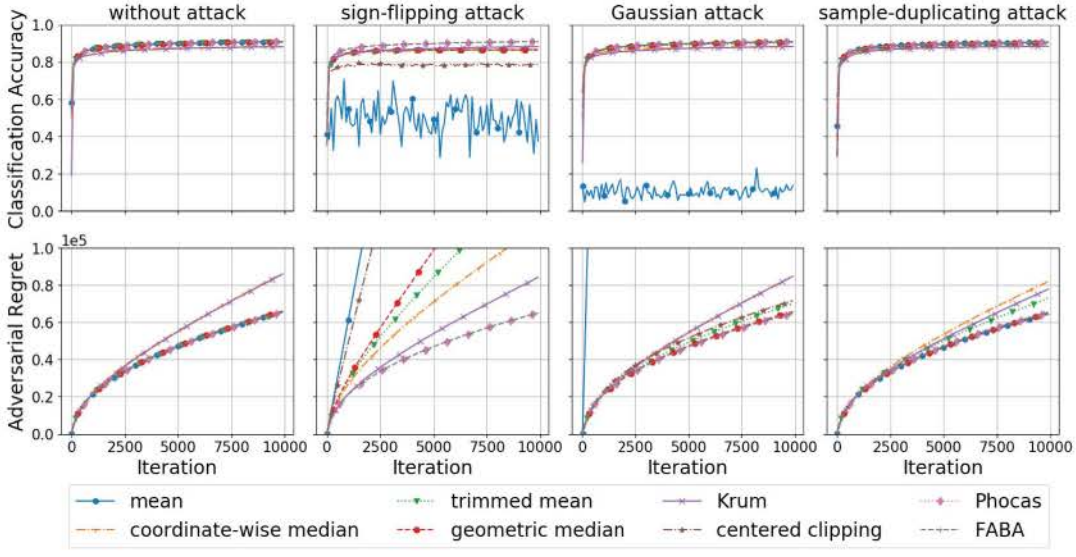

Fig. 9. Byzantine-robust distributed online gradient descent for softmax regression on MNIST i.i.d. data with constant step size.

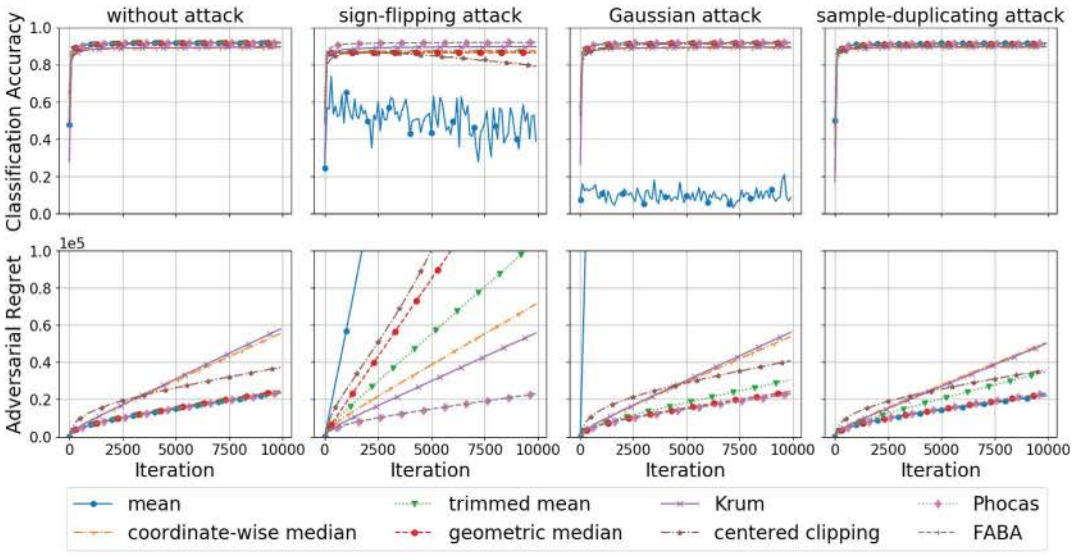

Fig. 10. Byzantine-robust distributed online gradient descent for softmax regression on MNIST i.i.d. data with diminishing step size.

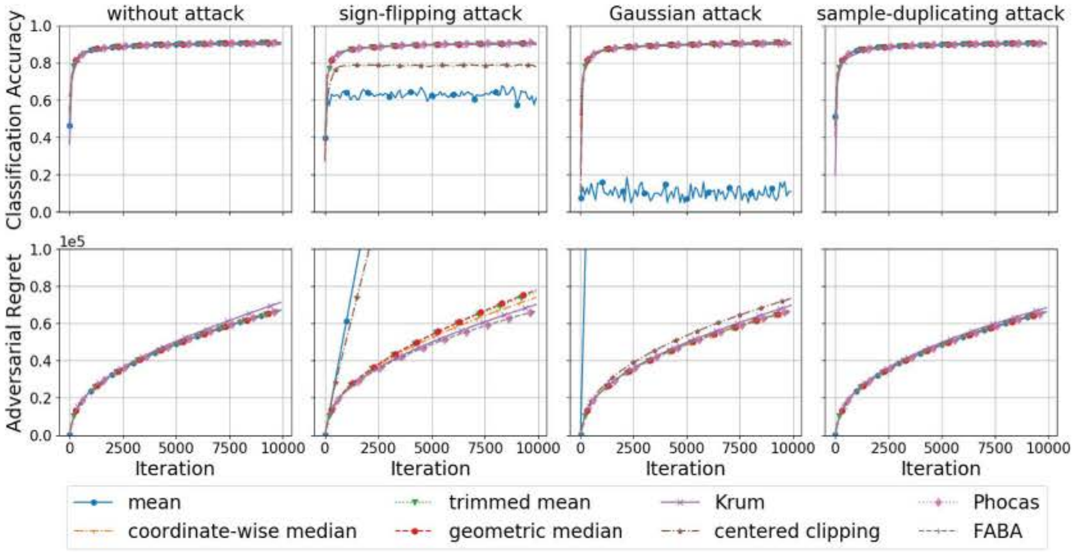

Fig. 11. Byzantine-robust distributed online momentum for softmax regression on MNIST i.i.d. data with constant step size.

Gaussian attack. Each Byzantine participant sends a random message, where each element follows the Gaussian distribution $\mathcal{N}(0,500)$, $\mathcal{N}(0,200)$ and $\mathcal{N}(0,200)$ for the three numerical experiments, respectively.

Sample-duplicating attack. The Byzantine participants jointly choose one honest participant, and duplicate its message to send. This amounts to that the Byzantine participants duplicate the samples of the chosen honest participant.

A. Least-Squares Regression on Synthetic Datasets

We start with least-squares regression on synthetic datasets, each of which contains 60,000 training samples. The dimensionality of decision variable is $d=10$. During training, the batch size is 1. We launch one server and 30 participants. Under Byzantine attacks, 5 randomly chosen participants are adversarial.

We take into account two data distributions. In the i.i.d. setting, each element of the regressors is drawn from the Gaussian distribution $\mathcal{N}(0,1)$. We also randomly generate each dimension of the ground-truth solution from the Gaussian distribution $\mathcal{N}(0,1)$. Then, the labels are obtained via multiplying the regressors by the ground-truth solution, followed by adding Gaussian noise $\mathcal{N}(0,0.1)$. These training samples are evenly distributed to all participants. In the non-i.i.d. setting, each element of the regressors and the ground-truth solution is evenly drawn from three pairs of Gaussian distributions: $(\mathcal{N}(0,1), g+\mathcal{N}(0,0.5))$, $(\mathcal{N}(1,1), g+\mathcal{N}(0.2,0.5))$, and $(\mathcal{N}(2,1), g+\mathcal{N}(0.4,0.5))$, where $g\sim\mathcal{N}(0,1)$. The added Gaussian noise is still from $\mathcal{N}(0,0.1)$. For each of the

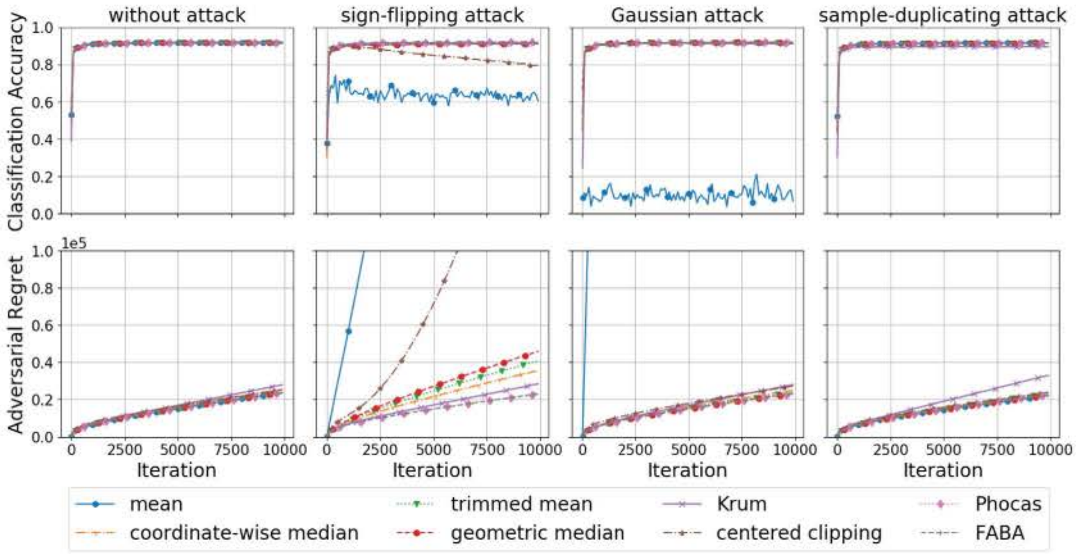

Fig. 12. Byzantine-robust distributed online momentum for softmax regression on MNIST i.i.d. data with diminishing step size.

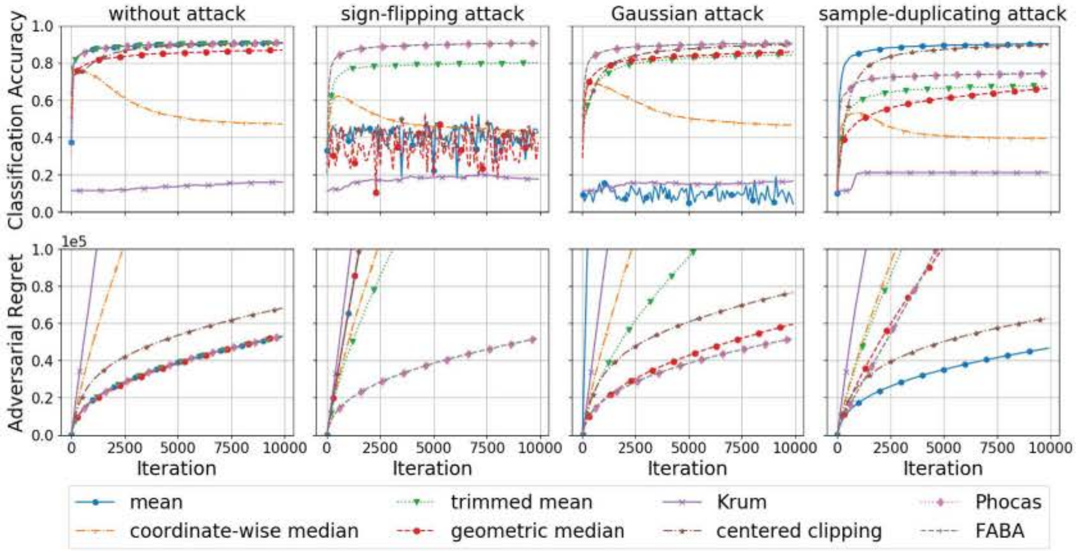

Fig. 13. Byzantine-robust distributed online gradient descent for softmax regression on MNIST non-i.i.d. data with constant step size.

three classes, the training samples are evenly distributed to 10 participants.

The performance metrics are adversarial regret and stochastic regret for the i.i.d. setting, and adversarial regret for the non-i.i.d. setting. We repeat generating the datasets and conducing the experiments for 10 times to calculate the regrets. This way, taking the average approximates the stochastic regret bound, while choosing the worst approximates the adversarial regret bound.

When the step size η and the momentum parameter ν are constant, they are set to 0.01 for the i.i.d. setting and 0.005 for the non-i.i.d. setting. For the diminishing step size η_t and momentum parameter ν_t, they are set to 0.008 in the first 500 iterations, and $\frac{4}{t}$ afterwards.

Numerical experiments on i.i.d. data. As shown in Figs. 1 and 2, the Byzantine-robust distributed online gradient descent algorithms equipped with robust bounded aggregation rules demonstrate trends of linear regret bounds, no matter using constant or diminishing step size. Take Fig. 2 as an example. Although trimmed mean, Phocas, coordinate-wise median, geometric median and FABA show sublinear regret bounds under the Gaussian and sample-duplicating attacks, they yield to linear regret bounds under the sign-flipping attack. This validates the tightness of Theorem 1 even on i.i.d. data. The Byzantine-robust distributed online momentum algorithms significantly improves the regret bounds, as shown in Figs. 3 and 4. Their regret bounds are all sublinear, which corroborate with Theorem 2.

Numerical experiments on non-i.i.d. data. On the non-i.i.d. data, the environment is more adversarial than that on the i.i.d. data. As shown in Figs. 5 and 6, the Byzantine-robust distributed online gradient descent algorithms, whether under attack or not, exhibit linear adversarial regret bounds. The Byzantine-robust distributed online momentum algorithms,

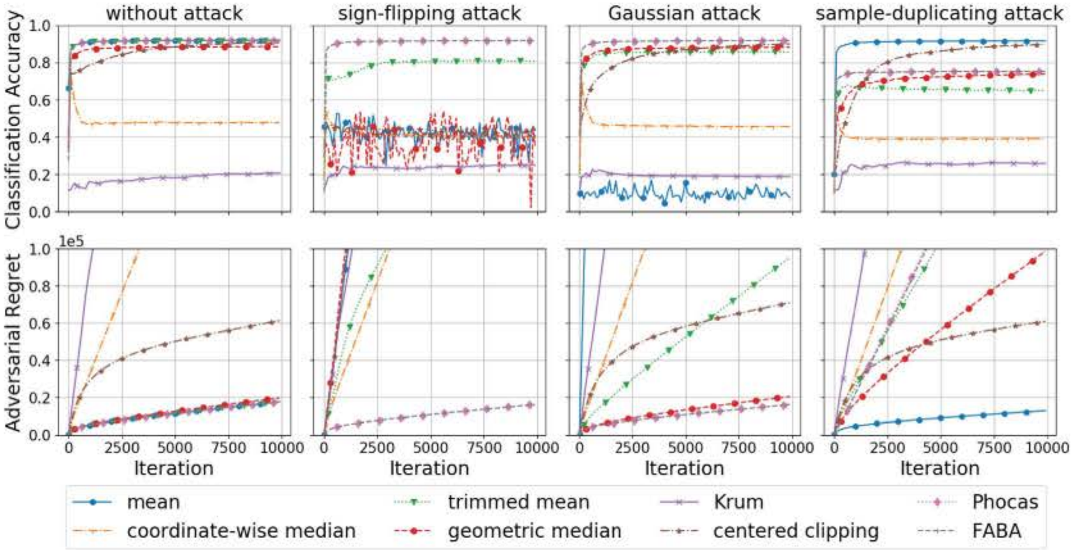

Fig. 14. Byzantine-robust distributed online gradient descent for softmax regression on MNIST non-i.i.d. data with diminishing step size.

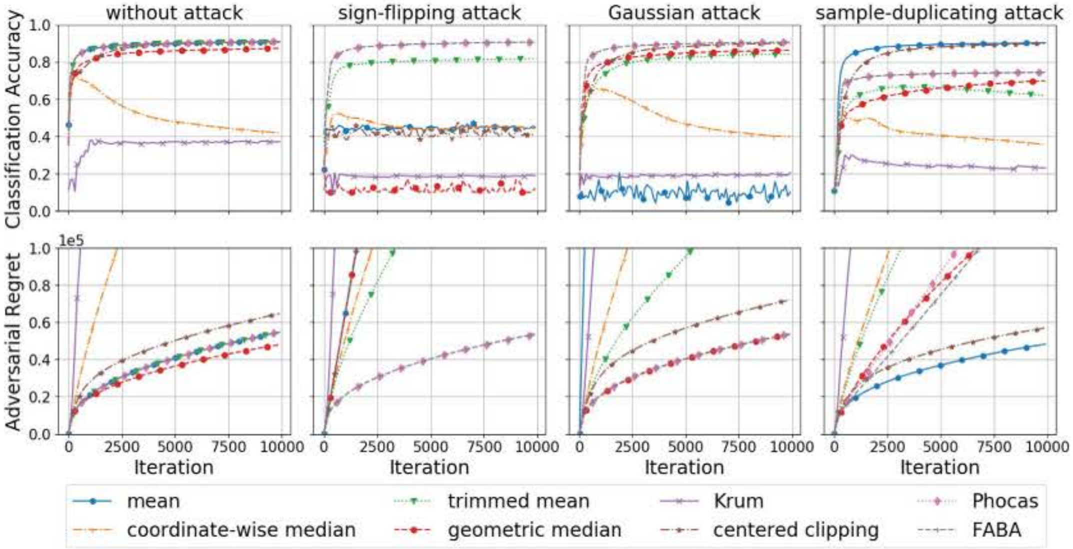

Fig. 15. Byzantine-robust distributed online momentum for softmax regression on MNIST non-i.i.d. data with constant step size.

as shown in Figs. 7 and 8, may have even larger regrets than those without momentum. This phenomenon underscores the importance of i.i.d. data distribution to Byzantine-robustness.

B. Softmax Regression on the MNIST Dataset

We next consider softmax regression on the MNIST dataset, which contains 60,000 training samples and 10,000 testing samples. The batch size is set to 32 during training. We launch one server and 30 participants, and consider two data distributions. In the i.i.d. setting, all the training samples are randomly and evenly allocated to all participants. In the non-i.i.d. setting, each class of the training samples are randomly and evenly distributed to 3 participants. Under Byzantine attacks, 5 randomly chosen participants are adversarial.

The performance metrics are classification accuracy on the testing samples and adversarial regret on the training samples. Since accurately calculating the adversarial and stochastic regret bounds is computationally demanding on such a large dataset, we only conduct the numerical experiments once, and calculate the adversarial regret to approximate its bound. Note that in the i.i.d. setting, adversarial regret is an approximation of stochastic regret, but there is still a substantial gap between the two.

When the step size η is constant, it is set to 0.01 and the momentum parameter ν is also set to 0.01. For the diminishing step size η_t and momentum parameter ν_t, they are set to 0.1 in the first 500 steps and $\frac{1}{t}$ afterwards.

Numerical experiments on i.i.d. data. As shown in Figs. 9 and 10, on the i.i.d. data, Byzantine-robust distributed online gradient descent equipped with robust bounded aggregation rules all perform well when no attack presents or under the sample-duplicating attack. Under the sign-flipping and Gaussian attacks, the algorithm with mean aggregation

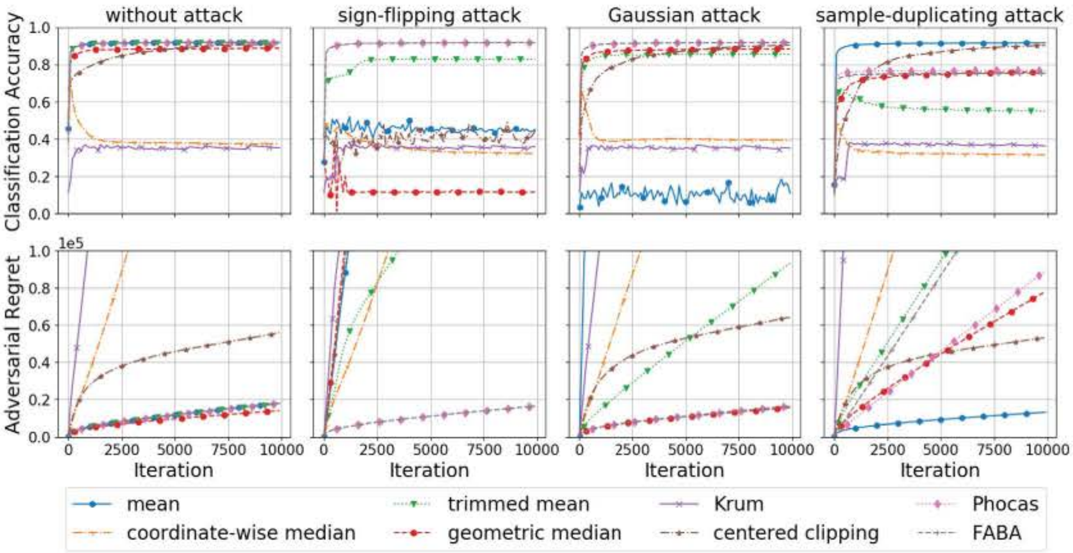

Fig. 16. Byzantine-robust distributed online momentum for softmax regression on MNIST non-i.i.d. data with diminishing step size.

fails, and the others demonstrate satisfactory robustness. The sign-flipping attack turns to be slightly stronger than the Gaussian attack; under the former, the algorithm with centered clipping performs worse, but is still much better than the one with mean aggregation.

The Byzantine-robust distributed online gradient descent algorithms with momentum improve over the ones without momentum in terms of classification accuracy and adversarial regret, as shown in Figs. 11 and 12. However, no sublinear adversarial regret bound is guaranteed, which confirms our theoretical prediction.

Numerical experiments on non-i.i.d. data. On the non-i.i.d. data, the environment is more adversarial than on the i.i.d. data. In this case, Byzantine-robust distributed online gradient descent, no matter with or without momentum, does not perform well, as in Figs. 13, 14, 15, and 16. This observation matches our conclusion on the hardness of handling adversarial participants in the adversarial environment.

VII. Conclusion

This paper is among the first efforts to investigate the Byzantine-robustness of distributed online learning. We show that Byzantine-robust distributed online gradient descent has linear adversarial regret, and the constant of the linear term is determined by the robust aggregation rule. On the other hand, we also establish the sublinear stochastic regret bound for Byzantine-robust distributed online momentum under the i.i.d. assumption.

Our future focus is to improve the Byzantine-robustness of distributed online learning algorithms in the non-i.i.d. setting, which is of practical importance in processing streaming data.

References

[1] X. Dong, Z. Wu, Q. Ling, and Z. Tian, “Distributed online learning with adversarial participants in an adversarial environment,” in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2023, pp. 1–5.

[2] M. Zinkevich, “Online convex programming and generalized infinitesimal gradient ascent,” in *Proc. Int. Conf. Mach. Learn.*, 2003, pp. 928–936.

[3] E. Hazan, A. Agarwal, and S. Kale, “Logarithmic regret algorithms for online convex optimization,” *Mach. Learn.*, vol. 69, no. 2, pp. 169–192, 2007.

[4] E. Hazan, “Introduction to online convex optimization,” *Found. Trends® Optim.*, vol. 2, no. 3–4, pp. 157–325, 2016.

[5] S. M. Fosson, “Centralized and distributed online learning for sparse time-varying optimization,” *IEEE Trans. Autom. Control*, vol. 66, no. 6, pp. 2542–2557, Jun. 2021.

[6] R. Li, F. Ma, W. Jiang, and J. Gao, “Online federated multitask learning,” in *Proc. IEEE Int. Conf. Big Data*, 2019, pp. 215–220.

[7] S. Paternain, S. Lee, M. M. Zavlanos, and A. Ribeiro, “Distributed constrained online learning,” *IEEE Trans. Signal Process.*, vol. 68, pp. 3486–3499, 2020.

[8] X. Yi, X. Li, L. Xie, and K. H. Johansson, “Distributed online convex optimization with time-varying coupled inequality constraints,” *IEEE Trans. Signal Process.*, vol. 68, pp. 731–746, 2020.

[9] K. I. Tsianos and M. G. Rabbat, “Distributed strongly convex optimization,” in *Proc. Allerton Conf. Commun. Control Comput.*, 2012, pp. 593–600.

[10] S. Hosseini, A. Chapman, and M. Mesbahi, “Online distributed convex optimization on dynamic networks,” *IEEE Trans. Autom. Control*, vol. 61, no. 11, pp. 3545–3550, Nov. 2016.

[11] S. Shalev-Shwartz, “Online learning and online convex optimization,” *Found. Trends® Mach. Learn.*, vol. 4, no. 2, pp. 107–194, 2012.

[12] O. Dekel, P. M. Long, and Y. Singer, “Online multitask learning,” in *Proc. Int. Conf. Comput. Learn. Theory*, 2006, pp. 453–467.

[13] X. Jin, P. Luo, F. Zhuang, J. He, and Q. He, “Collaborating between local and global learning for distributed online multiple tasks,” in *Proc. ACM Int. Conf. Inf. Knowl. Manage.*, 2015, pp. 113–122.

[14] Y. Chen, Y. Ning, M. Slawski, and H. Rangwala, “Asynchronous online federated learning for edge devices with non-iid data,” in *Proc. IEEE Int. Conf. Big Data*, 2020, pp. 15–24.

[15] L. Lamport, R. Shostak, and M. Pease, “The Byzantine generals problem,” *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.

[16] E. Hazan and S. Kale, “Beyond the regret minimization barrier: An optimal algorithm for stochastic strongly-convex optimization,” in *Proc. Conf. Learn. Theory*, 2011, pp. 421–436.

[17] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, “Online optimization in dynamic environments: Improved regret rates for strongly convex problems,” in *Proc. IEEE Conf. Decis. Control*, 2016, pp. 7195–7201.

[18] D. Garber and E. Hazan, “A linearly convergent variant of the conditional gradient algorithm under strong convexity, with applications to online and stochastic optimization,” *SIAM J. Optim.*, vol. 26, no. 3, pp. 1493–1528, 2016.

[19] E. C. Hall and R. M. Willett, "Online convex optimization in dynamic environments," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 4, pp. 647–662, Jun. 2015.

[20] Y. Ding, P. Zhao, S. Hoi, and Y.-S. Ong, "An adaptive gradient method for online AUC maximization," in *Proc. Assoc. Adv. Artif. Intell.*, 2015, pp. 2568–2574.

[21] Y. Wan, W.-W. Tu, and L. Zhang, "Projection-free distributed online convex optimization with $O(\sqrt{T})$ communication complexity," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 9818–9828.

[22] F. Yan, S. Sundaram, S. Vishwanathan, and Y. Qi, "Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties," *IEEE Trans. Knowl. Data Eng.*, vol. 25, no. 11, pp. 2483–2493, Nov. 2013.

[23] S. Ganesh, A. Reiffers-Masson, and G. Thoppe, "Online learning with adversaries: A differential inclusion analysis," 2023, *arXiv:2304.01525*.

[24] T. Yao and S. Sundaram, "Robust online and distributed mean estimation under adversarial data corruption," in *Proc. IEEE Conf. Decis. Control*, 2022, pp. 4193–4198.

[25] S. Sahoo, A. Gokhale, and R. K. Kalaimani, "Distributed online optimization with Byzantine adversarial agents," in *Proc. Amer. Control Conf.*, 2022, pp. 222–227.

[26] O. T. Odeyomi and G. Zaruba, "Byzantine-resilient federated learning with differential privacy using online mirror descent," in *Proc. Int. Conf. Comput., Netw. Commun.*, 2023, pp. 66–70.

[27] O. T. Odeyomi, B. Ude, and K. Roy, "Online decentralized multiagents meta-learning with Byzantine resiliency," *IEEE Access*, vol. 11, pp. 68286–68300, 2023.

[28] S. Kapoor, K. K. Patel, and P. Kar, "Corruption-tolerant bandit learning," *Mach. Learn.*, vol. 108, no. 4, pp. 687–715, 2019.

[29] A. Jadbabaie, H. Li, J. Qian, and Y. Tian, "Byzantine-robust federated linear bandits," in *Proc. IEEE Conf. Decis. Control*, 2022, pp. 5206–5213.

[30] A. Mitra, A. Adibi, G. J. Pappas, and H. Hassani, "Collaborative linear bandits with adversarial agents: Near-optimal regret bounds," in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, pp. 22602–22616.

[31] J. Zhu, A. Koppel, A. Velasquez, and J. Liu, "Byzantine-resilient decentralized multi-armed bandits," 2023, *arXiv:2310.07320*.

[32] I. Demirel, Y. Yildirim, and C. Tekin, "Federated multi-armed bandits under Byzantine attacks," 2022, *arXiv:2205.04134*.

[33] K. Pillutla, S. M. Kakade, and Z. Harchaoui, "Robust aggregation for federated learning," *IEEE Trans. Signal Process.*, vol. 70, pp. 1142–1154, 2022.

[34] S. Yu and S. Kar, "Secure distributed optimization under gradient attacks," *IEEE Trans. Signal Process.*, vol. 71, pp. 1802–1816, 2023.

[35] J. Xu, S.-L. Huang, L. Song, and T. Lan, "Byzantine-robust federated learning through collaborative malicious gradient filtering," in *Proc. Int. Conf. Distrib. Comput. Syst.*, 2022, pp. 1223–1235.

[36] D. Yin, Y. Chen, R. Kannan, and P. Bartlett, "Byzantine-robust distributed learning: Towards optimal statistical rates," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 5650–5659.

[37] L. Su and N. H. Vaidya, "Byzantine-resilient multiagent optimization," *IEEE Trans. Autom. Control*, vol. 66, no. 5, pp. 2227–2233, May 2021.

[38] Y. Chen, L. Su, and J. Xu, "Distributed statistical machine learning in adversarial settings: Byzantine gradient descent," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 1, no. 2, pp. 1–25, 2017.

[39] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 118–128.

[40] S. P. Karimireddy, L. He, and M. Jaggi, "Learning from history for Byzantine robust optimization," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 5311–5319.

[41] C. Xie, O. Koyejo, and I. Gupta, "Phocas: Dimensional Byzantine-resilient stochastic gradient descent," 2018, *arXiv:1805.09682*.

[42] Q. Xia, Z. Tao, Z. Hao, and Q. Li, "FABA: An algorithm for fast aggregation against Byzantine attacks in distributed neural networks," in *Proc. Int. Joint Conf. Artif. Intell.*, 2019, pp. 4824–4830.

[43] X. Dong, Z. Wu, Q. Ling, and Z. Tian, "Byzantine-robust distributed online learning: Taming adversarial participants in an adversarial environment," 2023, *arXiv:2307:07980*.

[44] S. Chen, W.-W. Tu, P. Zhao, and L. Zhang, "Optimistic online mirror descent for bridging stochastic and adversarial online convex optimization," 2023, *arXiv:2302.04552*.

[45] X. Li, L. Xie, and N. Li, "A survey on distributed online optimization and online games," *Annu. Rev. Control*, vol. 56, 2023, Art. no. 100904.

[46] Z. Wu, Q. Ling, T. Chen, and G. B. Giannakis, "Federated variance-reduced stochastic gradient descent with robustness to Byzantine attacks," *IEEE Trans. Signal Process.*, vol. 68, pp. 4583–4596, 2020.

[47] S. Bulusu, P. Khanduri, S. Kafle, P. Sharma, and P. K. Varshney, "Byzantine resilient non-convex SCSG with distributed batch gradient computations," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 7, pp. 754–766, 2021.

[48] E.-M. El-Mhamdi, R. Guerraoui, and S. Rouault, "Distributed momentum for Byzantine-resilient learning," 2020, *arXiv:2003.00010*.

[49] E. Gorbunov, S. Horváth, P. Richtárik, and G. Gidel, "Variance reduction is an antidote to Byzantines: Better rates, weaker assumptions and communication compression as a cherry on the top," in *Proc. Int. Conf. Learn. Represent.*, 2023, pp. 1–42.

Xingrong Dong (Graduate Student Member, IEEE) received the B.Phil. degree in logic from the Department of Philosophy, Sun Yat-Sen University, Guangzhou, China, in 2021. She is pursuing the M.E. degree with the School of Computer Science and Engineering, Sun Yat-Sen University, Guangzhou, China. Her research interest is online distributed optimization.

Zhaoxian Wu received the B.E. degree in software engineering and the M.E. degree in computer science and technology from the School of Computer Science and Engineering, Sun Yat-Sen University, Guangzhou, China, in 2020 and 2023, respectively. He is pursuing the Ph.D. degree with the Department of Electrical, Computer, and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY, USA. He was a recipient of the China National Scholarship in 2016. His research interest is distributed learning and optimization.

Qing Ling (Senior Member, IEEE) received the B.E. degree in automation and the Ph.D. degree in control theory and control engineering from the University of Science and Technology of China, Hefei, China, in 2001 and 2006, respectively. He was a Postdoctoral Research Fellow with the Department of Electrical and Computer Engineering, Michigan Technological University, Houghton, MI, USA, from 2006 to 2009, and an Associate Professor with the Department of Automation, University of Science and Technology of China, from 2009 to 2017. He is currently a Professor with the School of Computer Science and Engineering, Sun Yat-Sen University, Guangzhou, China. His current research interest includes distributed and decentralized optimization and its application in machine learning. He received the 2017 IEEE Signal Processing Society Young Author Best Paper Award as a supervisor. He served as a TPC Chair of the 2023 IEEE SPAWC Workshop. He was an Associate Editor and a Senior Area Editor of IEEE SIGNAL PROCESSING LETTERS, and an Associate Editor of IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT. He is now an Associate Editor of IEEE TRANSACTIONS ON SIGNAL PROCESSING.

Zhi Tian (Fellow, IEEE) has been a Professor with the Electrical and Computer Engineering Department, George Mason University since 2015. Prior to that, she was with the Faculty of Michigan Technological University from 2000 to 2014. She served as a Program Director with the US National Science Foundation from 2012 to 2014. Her research interest lies in the areas of statistical signal processing, wireless communications, and distributed machine learning. Her current research focuses on distributed network optimization and learning, wireless spectrum sensing, and millimeter-wave systems. She was an IEEE Distinguished Lecturer for both the IEEE Communications Society and the IEEE Vehicular Technology Society. She served as Associate Editor for IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS and IEEE TRANSACTIONS ON SIGNAL PROCESSING. She was a General Co-Chair of the 2016 IEEE GlobalSIP Conference and the 2023 IEEE SPAWC Workshop. She was a member-at-large of the Board of Governors of the IEEE Signal Processing Society for the term of 2019–2021. She received the IEEE Communications Society TCCN Publication Award in 2018.