ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) | 978-1-7281-6327-7/23/$31.00 ©2023 IEEE | DOI: 10.1109/ICASSP49357.2023.10095178

DISTRIBUTED ONLINE LEARNING WITH ADVERSARIAL PARTICIPANTS IN AN ADVERSARIAL ENVIRONMENT

*Xingrong Dong** *Zhaoxian Wu** *Qing Ling** *Zhi Tian*†

*Sun Yat-Sen University †George Mason University

ABSTRACT

This paper studies distributed online learning under Byzantine attacks. The performance of an online learning algorithm is characterized by (adversarial) regret, and a sublinear bound is preferred. But we prove that, even with a class of state-of-the-art robust aggregation rules, in an adversarial environment and with 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 a Byzantine-robust distributed online gradient descent algorithm with momentum to attain such a sublinear stochastic regret bound.

***Index Terms*—** Distributed optimization, Byzantine-robustness, Online learning

1. INTRODUCTION

Online learning is a powerful tool to process streaming data in a timely manner [1, 2, 3]. 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 adversarial regret increases sublinearly in time, which would lead to asymptotically vanishing performance loss. When the streaming data are separately collected by multiple participants and data privacy is a concern, distributed online learning becomes a natural choice. Each participant makes a local decision, and a server aggregates all the local decisions to a global one [4, 5]. Exemplary applications include online web ranking and online advertisement recommendation, to name a few [6, 7, 8, 9].

However, distributed online learning faces a new challenge in terms of robustness, since not all the participants are trustful. 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 [10]. 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 give 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 robust aggregation rules, we are able to control the constant of linear adversarial regret to a reasonable value. On the other hand, we further show that when the environment is not fully adversarial so that the losses of the honest participants are i.i.d. (independent and identically distributed), sublinear stochastic regret [11], in contrast to the aforementioned adversarial regret, is possible. We develop a Byzantine-robust distributed online gradient descent algorithm with momentum to attain such sublinear adversarial regret.

Related works. Similar to its centralized counterpart, the distributed online gradient descent algorithm has provable $\mathcal{O}(\sqrt{T})$ and $\mathcal{O}(logT)$ regret bounds for convex and strongly convex losses, respectively [12, 13]. However, its Byzantine-robust extension is rarely studied, and will be the focus of this paper. Another tightly related area is Byzantine-robust distributed stochastic optimization. Therein, the basic idea is to replace the vulnerable mean aggregation in distributed stochastic gradient descent with robust aggregation rules, including coordinate-wise median [14], trimmed mean [14, 15], geometric median [16], Krum [17], centered clipping [18], Phocas [19], FABA [20], etc. We will incorporate these robust aggregation rules with distributed online gradient descent to enable Byzantine-robustness.

Several recent works investigate distributed bandit 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 [21] that the proposed Byzantine-robust algorithms have linear adversarial regret bounds for multi-armed and linear-contextual problems. This is consistent with our result. Under the i.i.d. assumption, [22] proves $\mathcal{O}(T^{3/4})$ regret for linear bandit with high probability. Also under the i.i.d. assumption, [23] reaches $\mathcal{O}(\sqrt{T})$ regret but requires the action set to be finite. Our proposed algorithm, with the aid of momentum, attains the $\mathcal{O}(\sqrt{T})$ bound in terms of the stochastic regret. The work of [24] is free of the i.i.d. assumption, but the regret for multi-armed bandit is defined according to a suboptimal solution other than the optimal one. Therefore, the derived $\mathcal{O}(logT)$ sublinear regret bound is not comparable to others.

2. PROBLEM STATEMENT

Consider n participants in a set $\mathcal{N}$, among which h are honest and in $\mathscr{H}$, while b are Byzantine and in $\mathscr{B}$. We have $n = h + b$, but the identities and number of Byzantine participants are unknown. 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. 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 if it is from an honest or Byzantine participant. Upon receiving all z_t^j, the server aggregates them to yield a global decision $w_t \in \mathbb{R}^d$. The quality of the sequential decisions over T steps is evaluated by (adversarial) regret

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

where

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

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

For distributed online gradient descent, each honest participant j 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. 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 $j \in \mathscr{B}$ are arbitrarily malicious, such that w_{t+1} can be manipulated to reach infinite adversarial regret.

Motivated by the recent advances of Byzantine-robust distributed stochastic optimization, one may think of replacing the vulnerable mean aggregation with robust aggregation rules. Denoting AGG as a proper robust aggregation rule, the server makes the decision as

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

3. 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 $\frac{1}{2}$ [14, 16, 15, 17, 18, 19, 20]. Thus, one may wonder if 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. *An aggregation rule AGG is robust bounded aggregation, if the difference between its output and the mean of the honest messages is bounded by*

$$\|w_{t+1} - \bar{z}_t\|^2 = \|AGG(z_t^1, z_t^2, \cdots, z_t^n) - \bar{z}_t\|^2 \le C_\alpha^2 \zeta^2,$$

where $\bar{z}_t := \frac{1}{h}\sum_{j\in\mathscr{H}} z_t^j$ is the mean of the honest messages, ζ^2 is the largest deviation of the honest messages such that $\|\bar{z}_t - z_t^j\|^2 \le \zeta^2$ for all $j \in \mathscr{H}$, and C_α is an aggregation-specific constant influenced by the fraction of Byzantine participants $\alpha = \frac{b}{n}$.

We show that several state-of-the-art robust aggregation rules, including coordinate-wise median [14], trimmed mean [14, 15], geometric median [16], Krum [17], centered clipping [18], Phocas [19], and FABA [20], are all robust bounded aggregations. Their analysis and the corresponding constants C_α are left to an extended version of this paper.

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

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*

$$||\nabla f_t^j(x) - \nabla f_t^j(y)|| \le L||x - y||. \tag{6}$$

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*

$$\langle \nabla f_t^j(x), x - y\rangle \ge f_t^j(x) - f_t^j(y) + \frac{\mu}{2}\|x - y\|^2. \tag{7}$$

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

$$||\nabla f_t^j(w_t) - \nabla \bar{f}_t(w_t)||^2 \le \sigma^2. \tag{8}$$

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*

$$\|\frac{1}{h}\sum_{j\in\mathscr{H}} \nabla f_t^j(w^*)\|^2 \le \xi^2. \tag{9}$$

These assumptions are common in online learning. Some works make stronger assumptions [1, 2, 3], for example, bounded variable or bounded gradient that yields Assumptions 3 and 4.

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 robust bounded aggregation and constant step size $\eta_t = \eta \in (0, \frac{1}{8L}]$ have a linear adversarial regret bound*

$$\begin{aligned} R_T \le & \frac{1}{\eta}\|w_1 - w^*\|^2 + 4\eta\left(1 + \frac{8L^2\eta}{\mu}\right)\xi^2 T \\ & + 2\left(\eta + \frac{1}{\mu}\right) C_\alpha^2 \sigma^2 T. \end{aligned} \tag{10}$$

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

$$\begin{aligned} R_T \le & \frac{32L^2c^2}{\mu}\xi^2 + \left(\frac{\|w_1 - w^*\|^2}{c} + 2cC_\alpha^2\sigma^2 + 4c\xi^2\right)\sqrt{T} \\ & + \frac{2}{\mu}C_\alpha^2\sigma^2 T. \end{aligned} \tag{11}$$

We construct the following counter-example to demonstrate 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, $\mathscr{N} = \{1, 2, 3\}$, $\mathscr{H} = \{1, 2\}$ and $\mathscr{B} = \{3\}$. Suppose that at any step t, the losses of participants 1 and 2 are respectively given by*

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

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 1, participant 1 sends $z_1^1 = w_1^1 = w_1 - \eta(w_1 - \sigma) = \sigma$, while participant 2 sends $z_1^2 = w_1^2 = w_1 - \eta(w_1 + \sigma) = \sigma - 2\eta\sigma$. In this circumstance,

participant 3, who is Byzantine, can send $z_1^3 = \sigma$ so that the aggregation result is $w_2 = \sigma$. As such, for any step t, $f_t(w_t) = \sigma^2$ and $f_t(w^*) = \frac{1}{2}\sigma^2$, 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 $\bar{z}_t$ is $(1-\eta)\sigma$ and the largest deviation $\zeta^2 = \eta^2\sigma^2$. 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 adversarial regret. If the aggregation rule is majority-voting-based, such as coordinate-wise median and trimmed mean, sending $z_t^3 = \sigma$ is effective. For centered clipping, participant 3 can send $z_t^3 = \sigma + 2\eta\sigma$ instead.

Note that one can use a diminishing step size η_t in (3). We prove that it also yields linear adversarial regret, and leave the analysis to the extended version of this paper. In fact, Example 1 still holds true for a diminishing step size.

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 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 aggregation rule and the fraction of Byzantine participants, and the gradient deviation among honest participants.

4. SUBLINEAR STOCHASTIC REGRET BOUNDS OF BYZANTINE-ROBUST DISTRIBUTED ONLINE MOMENTUM GRADIENT DESCENT

According to Theorem 1, the linear adversarial regret is proportional to σ^2, the deviation between each honest gradient 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 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 $\mathscr{D}$ at all steps. Define the expected loss $F(w) := \mathbb{E}_{\mathscr{D}} f_t^j(w)$ for all $j \in \mathscr{H}$ and all t. In this setting, stochastic regret [11] is defined as

$$S_T := \mathbb{E}\sum_{t=1}^{T} F(w_t) - T \cdot \min_{w\in\mathbb{R}^d} F(w), \tag{12}$$

where the expectation is taken over the stochastic process. Note that the works of [22] and [23], which investigate Byzantine-robust distributed bandit, also make such an i.i.d. assumption.

However, naively applying robust aggregation rules (3) and (5) cannot guarantee sublinear stochastic gradient, 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 [25, 26, 27, 18, 28], we let each honest participant perform momentum gradient descent, instead of gradient descent, 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

$$m_t^j = \nu_t \nabla f_t^j(w_t) + (1-\nu_t) m_{t-1}^j, \tag{13}$$

where $\nu_t > 0$ is the momentum parameter. Then, it makes its local decision following

$$w_{t+1}^j = w_t - \eta_t m_t^j, \tag{14}$$

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

Corresponding to Assumptions 1, 2 and 3, the analysis needs the following assumptions.

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*

$$||\nabla F(x) - \nabla F(y)|| \le L||x - y||. \tag{15}$$

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*

$$\langle \nabla F(x), x - y\rangle \ge F(x) - F(y) + \frac{\mu}{2}\|x - y\|^2. \tag{16}$$

Assumption 7 (Bounded variance)**.** *The variance of each honest gradient is bounded by*

$$\mathbb{E}_{\mathscr{D}}||\nabla f_t^j(w_t) - \nabla F(w_t)||^2 \le \sigma^2. \tag{17}$$

In the 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 $\mathscr{D}$ with expectation $F := \mathbb{E}_{\mathscr{D}} f_t^j$. Under Assumptions 5, 6 and 7, the Byzantine-robust distributed online momentum gradient descent updates* (14) *and* (5) *with robust bounded aggregation, proper constant step size $\eta_t = \eta = \mathcal{O}(\frac{1}{\sqrt{T}})$ and proper constant momentum parameter $\nu_t = \nu = \mathcal{O}(\frac{1}{\sqrt{T}})$ have a sublinear stochastic regret bound*

$$S_T = \mathcal{O}\left(\frac{\sigma^2}{h}\left(1 + (h+1)C_\alpha^2\right)\frac{L^4}{\mu^4}\sqrt{T}\right). \tag{18}$$

In the sublinear stochastic regret bound (18), the constant $\frac{\sigma^2}{h}$ is inversely proportional to h, the number of honest participants, and highlights the benefit of collaboration. The constant $1 + (h+1)C_\alpha^2$ is determined by C_α that characterizes the defence 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 present. In this case, the derived stochastic regret bound degenerates to $\mathcal{O}(\frac{\sigma^2}{h}\sqrt{T})$.

The i.i.d. assumption is essential to the sublinear bound. Similar to the construction in Example 1, we can also show that Byzantine-robust online momentum gradient descent has a tight linear adversarial regret bound. But on the other hand, the momentum technique is important as we can show that Byzantine-robust online gradient descent without momentum has a linear stochastic regret bound. We omit the analysis due to the page limit.

5. NUMERICAL EXPERIMENTS

In this section, we demonstrate performance of the Byzantine-robust distributed online gradient descent and Byzantine-robust distributed online momentum gradient descent algorithms with experiments. We consider the softmax regression problem on the MNIST dataset, which contains 60,000 train samples and 10000 test samples. The batch size is set as 32 during training. We launch one server and 30 participants, and consider two data distributions. In the i.i.d. setting, all training samples are evenly distributed to all participants.

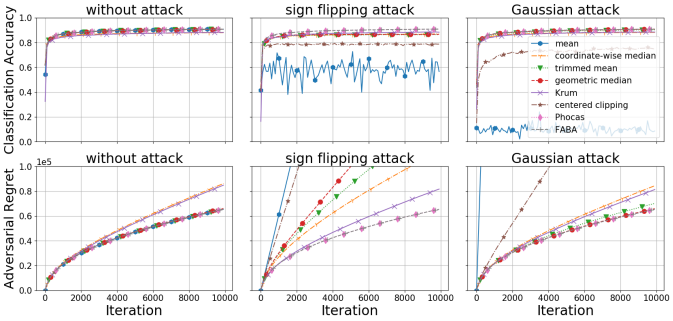

Fig. 1. Performance of Byzantine-robust distributed online gradient descent on i.i.d. data.

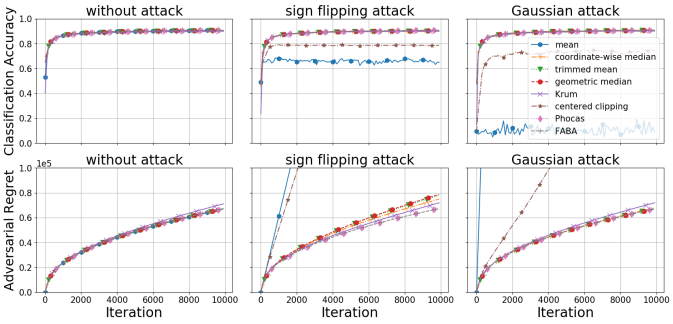

Fig. 2. Performance of Byzantine-robust distributed online momentum gradient descent on i.i.d. data.

In the non-i.i.d. setting, each class of training samples are evenly distributed to 3 participants. Under Byzantine attacks, 5 randomly chosen participants are adversarial. In the experiments the following two Byzantine attacks are considered:

Sign flipping attack. Each Byzantine participant sends a negative multiple of the honest message, and the multiple is -1.

Gaussian attack. Each Byzantine participant sends a random message, where each element obeys the Gaussian distribution $\mathbb{N}(0, 200)$.

We compare seven robust bounded aggregation rules with mean, including coordinate-wise median, trimmed mean, geometric median, Krum, centered clipping, Phocas and FABA.

The step size η and the momentum constant ν are set to 0.01. Other hyperparameters and the code are available online[1]. The performance metrics are classification accuracy and adversarial regret, because computing the stochastic regret is computationally demanding on the large training set. Note that on the i.i.d. data, adversarial regret is an approximation of the stochastic regret, but there is still a substantial gap between the two.

Experiments on i.i.d. data. As shown in Fig. 1, 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. Under both attacks, the algorithm with mean aggregation 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.

Also on the i.i.d. data, the Byzantine-robust distributed online momentum gradient descent algorithms improve over the ones without momentum in terms of classification accuracy and adversarial regret, as shown in Fig. 2. However, no sublinear adversarial regret bound is guaranteed, which confirms our theoretical prediction.

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

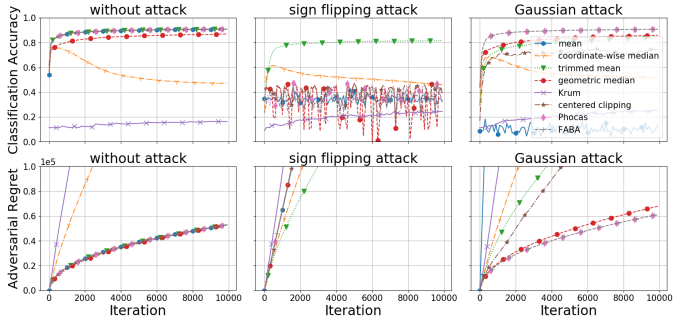

Fig. 3. Performance of Byzantine-robust distributed online gradient descent on non-i.i.d. data.

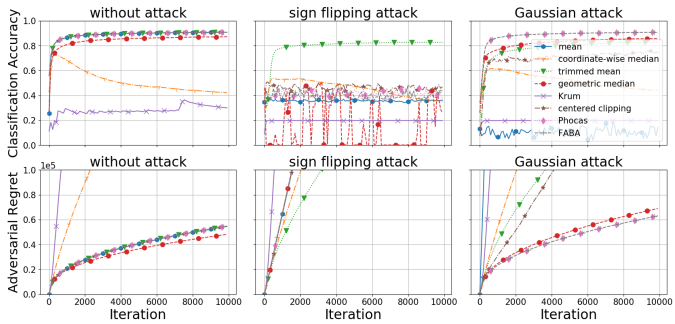

Fig. 4. Performance of Byzantine-robust distributed online momentum gradient descent on non-i.i.d. data.

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, the Byzantine-robust distributed online gradient descent algorithm, no matter with or without momentum, does not perform well, as in Figs. 3 and 4. This observation matches our conclusion on the hardness of handling adversarial participants in the adversarial environment.

6. CONCLUSIONS

This paper is the first 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 gradient descent 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.

Acknowledgement. The corresponding author, Qing Ling, is supported in part by NSF China grant 61973324, Guangdong Basic and Applied Basic Research Foundation grant 2021B1515020094, and Guangdong Provincial Key Laboratory of Computational Science grant 2020B1212060032.

7. REFERENCES

[1] Martin Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *International Conference on Machine Learning*, 2003, pp. 928–936.

[2] Elad Hazan, Amit Agarwal, and Satyen Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning*, vol. 69, no. 2, pp. 169–192, 2007.

[3] Elad Hazan, "Introduction to online convex optimization," *Foundations and Trends® in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2016.

[4] Konstantinos I Tsianos and Michael G Rabbat, "Distributed strongly convex optimization," in *Allerton Conference on Communication, Control, and Computing*, 2012, pp. 593–600.

[5] Saghar Hosseini, Airlie Chapman, and Mehran Mesbahi, "Online distributed convex optimization on dynamic networks," *IEEE Transactions on Automatic Control*, vol. 61, no. 11, pp. 3545–3550, 2016.

[6] Shai Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends® in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.

[7] Ofer Dekel, Philip M Long, and Yoram Singer, "Online multitask learning," in *International Conference on Computational Learning Theory*, 2006, pp. 453–467.

[8] Xin Jin, Ping Luo, Fuzhen Zhuang, Jia He, and Qing He, "Collaborating between local and global learning for distributed online multiple tasks," in *ACM International on Conference on Information and Knowledge Management*, 2015, pp. 113–122.

[9] Yujing Chen, Yue Ning, Martin Slawski, and Huzefa Rangwala, "Asynchronous online federated learning for edge devices with non-iid data," in *IEEE International Conference on Big Data*, 2020, pp. 15–24.

[10] Leslie Lamport, Robert Shostak, and Marshall Pease, "The Byzantine generals problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, 1982.

[11] Elad Hazan and Satyen Kale, "Beyond the regret minimization barrier: An optimal algorithm for stochastic strongly-convex optimization," in *Conference on Learning Theory*, 2011, pp. 421–436.

[12] Feng Yan, Shreyas Sundaram, SVN Vishwanathan, and Yuan Qi, "Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties," *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 11, pp. 2483–2493, 2012.

[13] Yuanyu Wan, Wei-Wei Tu, and Lijun Zhang, "Projection-free distributed online convex optimization with $O(\sqrt{T})$ communication complexity," in *International Conference on Machine Learning*, 2020, pp. 9818–9828.

[14] Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett, "Byzantine-robust distributed learning: Towards optimal statistical rates," in *International Conference on Machine Learning*, 2018, pp. 5650–5659.

[15] Lili Su and Nitin H Vaidya, "Byzantine-resilient multiagent optimization," *IEEE Transactions on Automatic Control*, vol. 66, no. 5, pp. 2227–2233, 2020.

[16] Yudong Chen, Lili Su, and Jiaming Xu, "Distributed statistical machine learning in adversarial settings: Byzantine gradient descent," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 1, no. 2, pp. 1–25, 2017.

[17] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," in *Advances in Neural Information Processing Systems*, 2017, pp. 118–128.

[18] Sai Praneeth Karimireddy, Lie He, and Martin Jaggi, "Learning from history for Byzantine robust optimization," in *International Conference on Machine Learning*, 2021, pp. 5311–5319.

[19] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta, "Phocas: Dimensional Byzantine-resilient stochastic gradient descent," *arXiv preprint arXiv:1805.09682*, 2018.

[20] Qi Xia, Zeyi Tao, Zijiang Hao, and Qun Li, "FABA: An algorithm for fast aggregation against Byzantine attacks in distributed neural networks," in *International Joint Conferences on Artificial Intelligence Organization*, 2019, pp. 4824–4830.

[21] Sayash Kapoor, Kumar Kshitij Patel, and Purushottam Kar, "Corruption-tolerant bandit learning," *Machine Learning*, vol. 108, no. 4, pp. 687–715, 2019.

[22] Ali Jadbabaie, Haochuan Li, Jian Qian, and Yi Tian, "Byzantine-robust federated linear bandits," *arXiv preprint arXiv:2204.01155*, 2022.

[23] Aritra Mitra, Arman Adibi, George J Pappas, and Hamed Hassani, "Collaborative linear bandits with adversarial agents: Near-optimal regret bounds," *arXiv preprint arXiv:2206.02834*, 2022.

[24] Ilker Demirel, Yigit Yildirim, and Cem Tekin, "Federated multi-armed bandits under Byzantine attacks," *arXiv preprint arXiv:2205.04134*, 2022.

[25] Zhaoxian Wu, Qing Ling, Tianyi Chen, and Georgios B Giannakis, "Federated variance-reduced stochastic gradient descent with robustness to Byzantine attacks," *IEEE Transactions on Signal Processing*, vol. 68, pp. 4583–4596, 2020.

[26] Prashant Khanduri, Saikiran Bulusu, Pranay Sharma, and Pramod K Varshney, "Byzantine resilient non-convex svrg with distributed batch gradient computations," *arXiv preprint arXiv:1912.04531*, 2019.

[27] El-Mahdi El-Mhamdi, Rachid Guerraoui, and Sébastien Rouault, "Distributed momentum for Byzantine-resilient learning," *arXiv preprint arXiv:2003.00010*, 2020.

[28] Eduard Gorbunov, Samuel Horváth, Peter Richtárik, and Gauthier Gidel, "Variance reduction is an antidote to Byzantines: Better rates, weaker assumptions and communication compression as a cherry on the top," *arXiv preprint arXiv:2206.00529*, 2022.