Distributed Statistical Machine Learning in Adversarial Settings: Byzantine Gradient Descent

Extended Abstract*

Yudong Chen
Cornell University
Ithaca, NY, USA
yudong.chen@cornell.edu

Lili Su[†]
Massachusetts Institute of Technology
Cambridge, MA, USA
lilisu@mit.edu

Jiaming Xu
Purdue University
West Lafayette, IN, USA
xu972@purdue.edu

ABSTRACT

We consider the distributed statistical learning problem over decentralized systems that are prone to adversarial attacks. This setup arises in many practical applications, including Google's *Federated Learning*. Formally, we focus on a decentralized system that consists of a parameter server and m working machines; each working machine keeps N/m data samples, where N is the total number of samples. In each iteration, up to q of the m working machines suffer Byzantine faults – a faulty machine in the given iteration behaves arbitrarily badly against the system and has complete knowledge of the system. Additionally, the sets of faulty machines may be different across iterations. Our goal is to design robust algorithms such that the system can learn the underlying true parameter, which is of dimension d , despite the interruption of the Byzantine attacks.

In this paper, based on the *geometric median of means* of the gradients, we propose a simple variant of the classical gradient descent method. We show that our method can tolerate q Byzantine failures up to $2(1 + \epsilon)q \leq m$ for an arbitrarily small but fixed constant $\epsilon > 0$. The parameter estimate converges in $O(\log N)$ rounds with an estimation error on the order of $\max\{\sqrt{dq/N}, \sqrt{d/N}\}$, which is larger than the minimax-optimal error rate $\sqrt{d/N}$ in the centralized and failure-free setting by at most a factor of $\sqrt{q}$. The total computational complexity of our algorithm is of $O((Nd/m) \log N)$ at each working machine and $O(md + kd \log^3 N)$ at the central server, and the total communication cost is of $O(md \log N)$. We further provide an application of our general results to the linear regression problem.

A key challenge arises in the above problem is that Byzantine failures create arbitrary and unspecified dependency among the iterations and the aggregated gradients. To handle this issue in the analysis, we prove that the aggregated gradient, as a function of model parameter, converges *uniformly* to the true gradient function.

CCS CONCEPTS

• **Security and privacy** → **Distributed systems security; Mobile and wireless security**; • **Computing methodologies** → **Batch learning; MapReduce algorithms**;

KEYWORDS

Distributed Systems, Learning, Security, Byzantine Adversaries

ACM Reference Format:

Yudong Chen, Lili Su, and Jiaming Xu. 2018. Distributed Statistical Machine Learning in Adversarial Settings: Byzantine Gradient Descent: Extended Abstract. In *SIGMETRICS '18 Abstracts: ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems Abstracts, June 18–22, 2018, Irvine, CA, USA*. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3219617.3219655>

REFERENCES

- [1] Yudong Chen, Lili Su, and Jiaming Xu. 2017. Distributed Statistical Machine Learning in Adversarial Settings: Byzantine Gradient Descent. *Proc. ACM Meas. Anal. Comput. Syst.* 1, 2, Article 44 (Dec. 2017), 25 pages. <https://doi.org/10.1145/3154503>

*The full version of this paper is available as [1]

[†]This is the corresponding author

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS '18 Abstracts, June 18–22, 2018, Irvine, CA, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5846-0/18/06.

<https://doi.org/10.1145/3219617.3219655>