

BYZANTINE-ROBUST LEARNING ON HETEROGENEOUS DATASETS VIA BUCKETING

Anonymous authors

Paper under double-blind review

ABSTRACT

In Byzantine robust distributed or federated learning, a central server wants to train a machine learning model over data distributed across multiple workers. However, a fraction of these workers may deviate from the prescribed algorithm and send arbitrary messages. While this problem has received significant attention recently, most current defenses assume that the workers have identical data. For realistic cases when the data across workers are heterogeneous (non-iid), we design new attacks which circumvent current defenses, leading to significant loss of performance. We then propose a simple bucketing scheme that adapts existing robust algorithms to heterogeneous datasets at a negligible computational cost. We also theoretically and experimentally validate our approach, showing that combining bucketing with existing robust algorithms is effective against challenging attacks. Our work is the first to establish guaranteed convergence for the non-iid Byzantine robust problem under realistic assumptions.

1 INTRODUCTION

Distributed or federated machine learning, where the data is distributed across multiple workers, has become an increasingly important learning paradigm both due to growing sizes of datasets, as well as data privacy concerns. In such a setting, the workers collaborate to train a single model without directly transmitting their training data (McMahan et al., 2016; Bonawitz et al., 2019; Kairouz et al., 2019). However, by decentralizing the training across a vast number of workers we potentially open ourselves to new security threats. Due to the presence of agents in the network which are actively malicious, or simply due to system and network failures, some workers may disobey the protocols and send arbitrary messages; such workers are also known as *Byzantine* workers (Lamport et al., 2019). Byzantine robust optimization algorithms attempt to combine the updates received from the workers using robust aggregation rules and ensure that the training is not impacted by the presence of a small number of malicious workers.

While this problem has received significant recent attention due to its importance, (Blanchard et al., 2017; Yin et al., 2018; Alistarh et al., 2018; Karimireddy et al., 2021), most of the current approaches assume that the data present on each different worker has identical distribution. This assumption is very unrealistic in practice and heterogeneity is inherent in distributed and federated learning (Kairouz et al., 2019). In this work, we show that existing Byzantine aggregation rules catastrophically fail with very simple attacks (or sometimes even with no attacks) in realistic settings. We carefully examine the causes of these failures, and propose a simple solution which provably solves the Byzantine resilient optimization problem under heterogeneous workers.

Concretely, our contributions in this work are summarized below

- We show that when the data across workers is heterogeneous, existing aggregation rules fail to converge, even when no Byzantine adversaries are present. We also propose a simple new attack, *mimic*, which explicitly takes advantage of data heterogeneity and circumvents median-based defenses. Together, these highlight the fragility of existing methods in real world settings.
- We then propose a simple fix - a new bucketing step which can be used before any existing aggregation rule. We introduce a formal notion of a robust aggregator (ARAGG) and prove that existing methods like KRUM, coordinate-wise median (CM), and geometric median aka robust federated averaging (RFA)—though insufficient on their own—become provably robust aggregators when augmented with our bucketing.

- We combine our notion of robust aggregator (ARAGG) with worker momentum to obtain optimal rates for Byzantine robust optimization with matching lower bounds. Unfortunately, our lower bounds imply that convergence to an exact optimum may not be possible due to heterogeneity. We then circumvent this lower bound and show that when heterogeneity is mild (or when the model is overparameterized), we can in fact converge to an exact optimum. This is the first result establishing convergence to the optimum for heterogeneous Byzantine robust optimization.
- Finally, we evaluate the effect of the proposed techniques (bucketing and worker momentum) against known and new attacks showcasing drastic improvement on realistic heterogeneously distributed datasets.

Setup and notations. Suppose that of the total n workers, the set of good workers is denoted by $\mathcal{G} \subseteq \{1, \dots, n\}$. Our objective is to minimize

$$f(\mathbf{x}) := \frac{1}{|\mathcal{G}|} \sum_{i \in \mathcal{G}} \{f_i(\mathbf{x}) := \mathbb{E}_{\xi_i}[F_i(\mathbf{x}; \xi_i)]\} \quad (1)$$

where f_i is the loss function on worker i defined over its own (heterogeneous) data distribution ξ_i . The (stochastic) gradient computed by a good worker $i \in \mathcal{G}$ over minibatch ξ_i is given as $\mathbf{g}_i(\mathbf{x}, \xi_i) := \nabla F_i(\mathbf{x}; \xi_i)$. The noise in every stochastic gradient is independent, unbiased with $\mathbb{E}_{\xi_i}[\mathbf{g}_i(\mathbf{x}, \xi_i)] = \nabla f_i(\mathbf{x})$, and has bounded variance $\mathbb{E}_{\xi_i} \|\mathbf{g}_i(\mathbf{x}, \xi_i) - \nabla f_i(\mathbf{x})\|^2 \leq \sigma^2$. Further, we assume that the data heterogeneity across the workers can be bounded as

$$\mathbb{E}_{j \sim \mathcal{G}} \|\nabla f_j(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq \zeta^2, \quad \forall \mathbf{x}.$$

We write \mathbf{g}_i^t or simply \mathbf{g}_i instead of $\mathbf{g}_i(\mathbf{x}^t, \xi_i^t)$ when there is no ambiguity.

Byzantine attack model. The set of Byzantine workers $\mathcal{B} \subset [n]$ is fixed over time, with the remaining workers \mathcal{G} being good, i.e. $[n] = \mathcal{B} \uplus \mathcal{G}$. We write δ for the fraction of Byzantine workers, $|\mathcal{B}| =: f \leq \delta n$. The Byzantine workers can deviate arbitrarily from our protocol, sending any update to the server. Further, they can collude and may even know the states of all other workers.

2 RELATED WORK

IID defenses. There has been a significant amount of recent work on the case when all workers have identical data distributions. [Blanchard et al. \(2017\)](#) initiated the study of Byzantine robust learning and proposed a distance-based aggregation approach Krum and extended to [\(Mhamdi et al., 2018; Damaskinos et al., 2019\)](#). [Yin et al. \(2018\)](#) propose to use and analyze the coordinate-wise median (CM), and [Pillutla et al. \(2019\)](#) use approximate geometric median. [Bernstein et al. \(2018\)](#) propose to use the signs of gradients and then aggregate them by majority vote, however, [Karimireddy et al. \(2019\)](#) show that it may fail to converge. Most recently, [Alistarh et al. \(2018\)](#); [Allen-Zhu et al. \(2021\)](#); [Mhamdi et al. \(2021\)](#); [Karimireddy et al. \(2021\)](#) showcase how to use past gradients to more accurately filter iid Byzantine workers and specifically *time-coupled* attacks. In particular, our work builds on top of [\(Karimireddy et al., 2021\)](#) and non-trivially extends to the non-iid setting.

IID vs. Non-IID attacks. For the iid setting, the state-of-the-art attacks are *time-coupled* attacks [\(Baruch et al., 2019; Xie et al., 2020\)](#). These attacks introduce a small but consistent bias at every step which is hard to detect in any particular round, but accumulates over time and eventually leads to divergence, breaking most prior robust methods. Our work focuses on developing attacks (and defenses) which specifically take advantages of the non-iid setting. The non-iid setting also enables targeted *backdoor* attacks which are designed to take advantage of heavy-tailed data [\(Bagdasaryan et al., 2020; Bhagoji et al., 2018\)](#). However, this is a challenging and open problem [\(Sun et al., 2019; Wang et al., 2020\)](#). Our focus is on the overall accuracy of the trained model, not on any subproblem.

Non-IID defenses. The non-iid defenses are relatively under-examined. [Li et al. \(2019\)](#) proposes an SGD variant (RSA) with additional ℓ_1 penalty. [Ghosh et al. \(2019\)](#); [Sattler et al. \(2020\)](#) use an outlier-robust clustering method. When the server has the entire training dataset, [Xie et al. \(2019\)](#); [Chen et al. \(2018\)](#); [Rajput et al. \(2019\)](#) show that non-iid-ness can be overcome. Typical examples are parallel training of neural networks on public cloud, or volunteer computing [\(Meeds et al., 2015; Miura & Harada, 2015\)](#). However, none of these methods are applicable to the standard federated learning. This is partially tackled in [\(Data & Diggavi, 2020; 2021\)](#) who analyze spectral methods for robust optimization. However, these methods require $\Omega(d^2)$ time, making them infeasible for large scale optimization. In an independent recent work, [Acharya et al. \(2021\)](#) analyze geometric median (GM) on non-iid data using sparsified gradients. However, they do not defend against time coupled

attacks, and their analysis neither proves convergence to the optimum nor recovers the standard rate of SGD when $\delta = 0$. In contrast, our analysis of GM addresses both issues and is more general. Finally, Yang & Li (2021) propose to use bucketing for asynchronous Byzantine learning which is very similar to the bucketing trick proposed in this paper for non-iid setup. We note that these two are parallel work with very different purposes, techniques, and analysis.¹

3 ATTACKS AGAINST EXISTING AGGREGATION SCHEMES

In this section we show that when the data across the workers is heterogeneous (non-iid), then we can design simple new attacks which take advantage of the heterogeneity, leading to the failure of existing aggregation schemes. We study three representative and widely used defenses:

Krum. For $i \neq j$, let $i \rightarrow j$ denote that x_j belongs to the $n - f - 2$ closest vectors to x_i . Then,

$$\text{KRUM}(x_1, \dots, x_n) := \operatorname{argmin}_{i \rightarrow j} \|x_i - x_j\|^2.$$

Krum is computationally expensive, requiring $\mathcal{O}(n^2)$ work by the server (Blanchard et al., 2017).

CM. Coordinate-wise median computes for the k th coordinate:

$$[\text{CM}(x_1, \dots, x_n)]_k := \text{median}([x_1]_k, \dots, [x_n]_k) = \operatorname{argmin}_{i \rightarrow j} \sum_{j=1}^n |[x_i]_k - [x_j]_k|.$$

Coordinate-wise median is fast to implement requiring only $\mathcal{O}(n)$ time (Chen et al., 2017).

RFA. Robust federated averaging (RFA) computes the geometric median

$$\text{RFA}(x_1, \dots, x_n) := \operatorname{argmin}_v \sum_{i=1}^n \|v - x_i\|_2.$$

While the geometric median has no closed form solution, (Pillutla et al., 2019) approximate it using multiple iterations of smoothed Weiszfeld algorithm, each of which requires $\mathcal{O}(n)$ computation.

3.1 FAILURE ON IMBALANCED DATA WITHOUT BYZANTINE WORKERS

We show that when the data amongst the workers is imbalanced, existing aggregation rules *fail* even in the *absence* of any Byzantine workers. Algorithms like KRUM select workers who are *representative* of a majority of the workers by relying on statistics such as pairwise differences between the various worker updates. Our key insight is that when the data across the workers is heterogeneous, there is no single worker who is representative of the whole dataset. This is because each worker computes their local gradient over vastly different local data.

Example. Suppose that there are $2n + 1$ workers with worker i holding $(-1)^i \in \{\pm 1\}$. This means that the true mean is ≈ 0 , but KRUM, CM, and RFA will output ± 1 . This motivates our next attack.

Hence, for convergence it is important to not only select a good (non-Byzantine) worker, but also ensure that each of the good workers is selected with roughly equal frequency. In Table 1, we demonstrate failures of such aggregators by training on MNIST with $n = 20$ and no attackers ($\delta = 0$). We construct an imbalanced dataset where each successive class has only a fraction of samples of the previous class. We defer details of the experiments to Appendix A. As we can see, KRUM, CM and RFA match the ideal performance of SGD in the iid case, but only attain less than 90% accuracy in the non-iid case. This corresponds to learning only the top 2–3 classes and ignoring the rest.

A similar phenomenon was observed when using batch-size 1 in the iid case by (Karimireddy et al., 2021). However, in the iid case this can be easily overcome by increasing the batch-size. In contrast, when the data across the works is non-iid (e.g. split by class), increasing the batch-size does *not* make the worker gradients any more similar and there remains a big drop in performance. Finally, note that in Table 1 a hitherto new algorithm (CCLIP) maintains its performance both in the iid and the non-iid setting. We will explore this in more detail in Section 4.

3.2 MIMIC ATTACK ON BALANCED DATA

Motivated by how data imbalance could lead to consistent errors in the aggregation rules and significant loss in accuracy, in this section, we will propose a new attack *mimic* which specifically tries to maximize the perceived data imbalance even if the original data is balanced.

¹The previous version of this work uses resampling which has identical performance as bucketing. The detailed comparison is listed in Appendix A.2.4.

Table 1: Test accuracy (%) with no Byzantine workers ($\delta=0$) on imbalanced data.

Aggr	iid	non-iid
AVG	98.79 \pm 0.10	98.75 \pm 0.02
KRUM	97.95 \pm 0.25	89.90 \pm 4.75
CM	97.72 \pm 0.22	80.36 \pm 0.05
RFA	98.62 \pm 0.08	82.60 \pm 0.84
CCLIP	98.78 \pm 0.10	98.78 \pm 0.06

Table 2: Test accuracy (%) under mimic attack with $\delta = 0.2$ fraction of Byzantine workers.

Aggr	iid	non-iid
AVG	93.20 \pm 0.21	92.73 \pm 0.32
KRUM	90.36 \pm 0.25	37.33 \pm 6.78
CM	90.80 \pm 0.12	64.27 \pm 3.70
RFA	92.92 \pm 0.25	78.93 \pm 9.27
CCLIP	93.16 \pm 0.22	91.53 \pm 0.06

Algorithm 1 Robust Aggregation (ARAGG) using bucketing

-
- 1: **input** $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, $s \in \mathbb{N}$, aggregation rule AGGR
 - 2: pick random permutation π of $[n]$
 - 3: compute $\mathbf{y}_i \leftarrow \frac{1}{s} \sum_{k=(i-1)s+1}^{\min(n, i \cdot s)} \mathbf{x}_{\pi(k)}$ for $i = \{1, \dots, \lceil n/s \rceil\}$
 - 4: **output** $\hat{\mathbf{x}} \leftarrow \text{AGGR}(\mathbf{y}_1, \dots, \mathbf{y}_{\lceil n/s \rceil})$ // aggregate after bucketing
-

Mimic attack. All Byzantine workers pick a good worker (say i_*) to mimic and copy its output ($\mathbf{x}_{i_*}^t$). This inserts a consistent bias towards over-emphasizing worker i_* and thus under-representing other workers. Since the attacker simply mimics a good worker, it is impossible to distinguish it from a real worker and hence it cannot be filtered out. Indeed, the target i_* can be any fixed good worker. In Appendix B, we present an empirical rule to choose i_* and include a simple example demonstrating how median based aggregators suffer from the heterogeneity under mimic attack.

Table 2 shows the effectiveness of mimic attack even when the fraction of Byzantine nodes is small (i.e. $n = 25$, $|\mathcal{B}| = 5$). Note that this attack specifically targets the non-iid nature of the data—all robust aggregators maintain their performance in the iid setting and only suffer in the non-iid setting. Their performance is in fact worse than even simply averaging. As predicted by our example, KRUM and CM have the worst performance and RFA performs slightly better. We will discuss the remarkable performance of CCLIP in the next section.

4 CONSTRUCTING AN AGNOSTIC ROBUST AGGREGATOR USING BUCKETING

In Section 3 we demonstrated how existing aggregation rules fail in realistic non-iid scenarios, with and without attackers. In this section, we show how using bucketing can provably fix such aggregation rules. The underlying reason for this failure, as we saw previously, is that the existing methods fixate on the contribution of only the most likely worker, and ignore the contributions from the rest. To overcome this issue, we propose to use bucketing which ‘mixes’ the data from all the workers thereby reducing the chance of any subset of the data being consistently ignored.

4.1 BUCKETING ALGORITHM

Given n inputs $\mathbf{x}_1, \dots, \mathbf{x}_n$, we perform s -bucketing which randomly partitions them into $\lceil n/s \rceil$ buckets with each bucket having no more than s elements. Then, the contents of each bucket are averaged to construct $\{\mathbf{y}_1, \dots, \mathbf{y}_{\lceil n/s \rceil}\}$ which are then input to an aggregator AGGR. The details are summarized in Algorithm 1. The key property of our approach is that after bucketing, the resulting set of averaged $\{\mathbf{y}_1, \dots, \mathbf{y}_{\lceil n/s \rceil}\}$ are much more homogeneous (lower variance) than the original inputs. Thus, when fed into existing aggregation schemes, the chance of success increases. We formalize this in the following simple lemma.

Lemma 1 (Bucketing reduces variance). *Suppose we are given n independent (but not identical) random vectors $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ such that a good subset $\mathcal{G} \subseteq [n]$ of size at least $|\mathcal{G}| \geq n(1 - \delta)$ satisfies:*

$$\mathbb{E}\|\mathbf{x}_i - \mathbf{x}_j\|^2 \leq \rho^2, \quad \text{for any fixed } i, j \in \mathcal{G}.$$

Define $\bar{\mathbf{x}} := \frac{1}{|\mathcal{G}|} \sum_{j \in \mathcal{G}} \mathbf{x}_j$. Let the outputs after s -bucketing be $\{\mathbf{y}_1, \dots, \mathbf{y}_{\lceil n/s \rceil}\}$. Then, a subset of the outputs $\tilde{\mathcal{G}} \subseteq \{1, \dots, \lceil n/s \rceil\}$ of size at least $|\tilde{\mathcal{G}}| \geq \lceil n/s \rceil(1 - \delta s)$ satisfies

$$\mathbb{E}[\mathbf{y}_i] = \mathbb{E}[\bar{\mathbf{x}}] \quad \text{and} \quad \mathbb{E}\|\mathbf{y}_i - \mathbf{y}_j\| \leq \rho^2/s \quad \text{for any fixed } i, j \in \tilde{\mathcal{G}}.$$

The expectation in the above lemma is taken both over the random vectors as well as over the randomness of the bucketing procedure.

Remark 2. Lemma 1 proves that after our bucketing procedure, we are left with outputs \mathbf{y}_i which have i) pairwise variance reduced by s , and ii) potentially s times more fraction of Byzantine vectors. Hence, bucketing trades off increasing influence of Byzantine inputs against having more homogeneous vectors. Using $s = 1$ simply shuffles the inputs and leaves them otherwise unchanged.

4.2 AGNOSTIC ROBUST AGGREGATION

We now define what it means for an agnostic robust aggregator to succeed.

Definition A ((δ_{\max}, c) -ARAGG). Suppose we are given input $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ of which a subset \mathcal{G} of size at least $|\mathcal{G}| > (1 - \delta)n$ for $\delta \leq \delta_{\max} < 0.5$ and satisfies $\mathbb{E}\|\mathbf{x}_i - \mathbf{x}_j\|^2 \leq \rho^2$. Then, the output $\hat{\mathbf{x}}$ of a Byzantine robust aggregator satisfies:

$$\mathbb{E}\|\hat{\mathbf{x}} - \bar{\mathbf{x}}\|^2 \leq c\delta\rho^2 \quad \text{where} \quad \hat{\mathbf{x}} = \text{ARAGG}_\delta(\mathbf{x}_1, \dots, \mathbf{x}_n).$$

Further, ARAGG does not need to know ρ^2 (only δ), and automatically adapts to any value ρ^2 .

Our robust aggregator is parameterized by $\delta_{\max} < 0.5$ which denotes the maximum amount of Byzantine inputs it can handle, and a constant c which determines its performance. If $\delta = 0$, i.e. when there are no Byzantine inputs, we are guaranteed to *exactly* recover the true average $\bar{\mathbf{x}}$. Exact recovery is also guaranteed when $\rho = 0$ since in that case it is easy to identify the good inputs since they are all equal and in majority. When both $\rho > 0$ and $\delta > 0$, we recover the average up to an additive error term. We also require that the robust aggregator is *agnostic* to the value of ρ^2 and automatically adjusts its output to the current ρ during training. The aggregator can take δ as an input though. This property is very useful in the context of Byzantine robust optimization since the variance ρ^2 keeps changing over the training period, whereas the fraction of Byzantine workers δ remains constant. This is a major difference from the definition used in (Karimireddy et al., 2021).

We next show that aggregators which we saw were not robust in Section 3, can be made to satisfy Definition A by combining with bucketing.

Theorem I. Suppose we are given n inputs $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ satisfying properties in Lemma 1 for some $\delta \leq \delta_{\max}$, with δ_{\max} to be defined. Then, running Algorithm 1 with $s = \lfloor \delta_{\max}/\delta \rfloor$ yields the following:

- *Krum*: $\mathbb{E}\|\text{KRUM} \circ \text{BUCKETING}(\mathbf{x}_1, \dots, \mathbf{x}_n) - \bar{\mathbf{x}}\|^2 \leq \mathcal{O}(\delta\rho^2)$ with $\delta_{\max} < 1/4$.
- *Geometric median*: $\mathbb{E}\|\text{RFA} \circ \text{BUCKETING}(\mathbf{x}_1, \dots, \mathbf{x}_n) - \bar{\mathbf{x}}\|^2 \leq \mathcal{O}(\delta\rho^2)$ with $\delta_{\max} < 1/2$.
- *Coordinate-wise median*: $\mathbb{E}\|\text{CM} \circ \text{BUCKETING}(\mathbf{x}_1, \dots, \mathbf{x}_n) - \bar{\mathbf{x}}\|^2 \leq \mathcal{O}(d\delta\rho^2)$ with $\delta_{\max} < 1/2$.

Note that all these methods satisfy our notion of an *agnostic* Byzantine robust aggregator (Definition A). This is because both our bucketing procedures as well as the underlying aggregators are independent of ρ^2 . Further, our error is $\mathcal{O}(\delta\rho^2)$ and is information theoretically optimal, unlike previous analyses (e.g. Acharya et al. (2021)) who had an error of $\mathcal{O}(\rho^2)$.

The error of CM depends on the dimension d which is problematic when $d \gg n$. However, we suspect this is because we measure stochasticity using Euclidean norms instead of coordinate-wise. In practice, we found that CM often outperforms KRUM, with RFA outperforming them both. Note that we select $s = \lfloor \delta_{\max}/\delta \rfloor$ to ensure that after bucketing, we have the maximum amount of Byzantine inputs tolerated by the method with $(s\delta) = \delta_{\max}$.

Remark 3 (Centered clipping). The centered clipping aggregator (CCLIP) given a clipping radius τ and an initial guess \mathbf{v} of the average $\bar{\mathbf{x}}$ performs

$$\text{CCLIP}(\mathbf{x}_1, \dots, \mathbf{x}_n) = \mathbf{v} + \frac{1}{n} \sum_{i \in [n]} (\mathbf{x}_i - \mathbf{v}) \min(1, \tau/\|\mathbf{x}_i - \mathbf{v}\|_2).$$

Karimireddy et al. (2021) prove that CCLIP even without bucketing satisfies Definition A with $\delta_{\max} = 0.1$, and $c = \mathcal{O}(1)$. This explains its good performance on non-iid data in Section 3. However, CCLIP is not agnostic since it requires clipping radius τ as an input which in turn depends on ρ^2 . Devising a version of CCLIP which automatically adapts its clipping radius is an important open question. Empirically however, we observe that simple rules for setting τ work quite well—we always use $\tau = \frac{10}{1-\beta}$ in our limited experiments where β is the coefficient of momentum.

While we have shown how to construct a robust aggregator which satisfies some notion of a robustness, we haven't yet seen how this affects the Byzantine robust optimization problem. We investigate this question theoretically in the next section and empirically in Section 6.

5 ROBUST NON-IID OPTIMIZATION USING A ROBUST AGGREGATOR

In this section, we study the problem of optimization in the presence of Byzantine workers and heterogeneity, given access to any robust aggregator satisfying Definition A. We then show that data heterogeneity makes Byzantine robust optimization especially challenging and prove lower bounds for the same. Finally, we see how mild heterogeneity, or sufficient overparameterization can circumvent these lower bounds, obtaining convergence to the optimum.

Algorithm 2 Robust Optimization using any Agnostic Robust Aggregator

Input: ARAGG, η, β
 1: **for** $t = 1, \dots$ **do**
 2: **for** worker $i \in [n]$ **in parallel**
 3: $\mathbf{g}_i \leftarrow \nabla F_i(\mathbf{x}, \xi_i)$ and $\mathbf{m}_i \leftarrow (1 - \beta)\mathbf{g}_i + \beta\mathbf{m}_i$ // worker momentum
 4: **send** \mathbf{m}_i if $i \in \mathcal{G}$, else send * if Byzantine
 5: $\hat{\mathbf{m}} = \text{ARAGG}(\mathbf{m}_1, \dots, \mathbf{m}_n)$ and $\mathbf{x} \leftarrow \mathbf{x} - \eta\hat{\mathbf{m}}$. // update params using robust aggregate

5.1 ALGORITHM DESCRIPTION

In Section 4 we saw that bucketing could tackle heterogeneity across the workers by reducing ζ^2 . However, there still remains variance σ^2 in the gradients within each worker since each worker uses stochastic gradients. To reduce the effect of this variance, we rely on worker momentum. Each worker sends their local worker momentum vector \mathbf{m}_i to be aggregated by ARAGG instead of \mathbf{g}_i directly:

$$\begin{aligned}\mathbf{m}_i^t &= \beta\mathbf{m}_i^{t-1} + (1 - \beta)\mathbf{g}_i(\mathbf{x}^{t-1}) \quad \text{for every } i \in \mathcal{G}, \\ \mathbf{x}^t &= \mathbf{x}^{t-1} - \eta\text{ARAGG}(\mathbf{m}_1^t, \dots, \mathbf{m}_n^t).\end{aligned}$$

This is equivalent to the usual momentum description up to a rescaling of step-size η . Intuitively, using worker momentum \mathbf{m}_i averages over $1/(1-\beta)$ independent stochastic gradients \mathbf{g}_i and thus reduces the effect of the within-worker-variance σ^2 (Karimireddy et al., 2021). Note that the resulting $\{\mathbf{m}_i\}$ are *still heterogeneous* across the workers. This heterogeneity is the key challenge we face.

5.2 CONVERGENCE RATES

We now turn towards proving convergence rates for our bucketing aggregation method Algorithm 1 based on any existing aggregator AGGR. We will assume that for any fixed $i \in \mathcal{G}$, the stochastic gradients computed satisfy

$$\mathbb{E}_{\xi_i} \|\mathbf{g}_i(\mathbf{x}) - \nabla f_i(\mathbf{x})\|^2 \leq \sigma^2 \text{ and } \mathbb{E}_{j \sim \mathcal{G}} \|\nabla f_j(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq \zeta^2, \quad \forall \mathbf{x}. \quad (2)$$

This first condition bounds the variance of the stochastic gradient within a worker whereas the latter is a standard measure of inter-client heterogeneity in federated learning (Yu et al., 2019; Khaled et al., 2020; Karimireddy et al., 2020). Under these conditions, we can prove the following.

Theorem II. *Suppose we are given a (δ_{\max}, c) -ARAGG satisfying Definition A, and n workers of which a subset \mathcal{G} of size at least $|\mathcal{G}| \geq n(1 - \delta)$ faithfully follow the algorithm for $\delta \leq \delta_{\max}$. Further, for any good worker $i \in \mathcal{G}$ let f_i be a possibly non-convex function with L -Lipschitz gradients, and the stochastic gradients on each worker be independent, unbiased and satisfy (2). Then, for $F^0 := f(\mathbf{x}^0) - f^*$, the output of Algorithm 2 satisfies*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\mathbf{x}^{t-1})\|^2 \leq \mathcal{O} \left(c\delta\zeta^2 + \sigma \sqrt{\frac{LF^0}{T}} (c\delta + 1/n) + \frac{LF^0}{T} \right).$$

Remark 4 (Unified proofs). Remark 3 shows that CCLIP is a robust aggregator, and Theorem I shows KRUM, RFA, and CM on combining with sufficient bucketing are all robust aggregators satisfying Definition A. Most of these methods had no end-to-end convergence guarantees prior to our results. Thus, Theorem II gives the first unified analysis in both the iid and non-iid settings.

When $\delta = 0$ i.e. there are no Byzantine workers, the above rate recovers the familiar $\mathcal{O}(\frac{\sigma}{\sqrt{Tn}})$ rate which is optimal for non-convex SGD and even has linear speed-up with respect to the n workers. In

contrast, all previous algorithms for non-iid data (e.g. (Data & Diggavi, 2021; Acharya et al., 2021)) do not recover SGD rates when $\delta = 0$. This is also empirically reflected in Section 3.1, where we showed how these algorithms can fail even in the absence of Byzantine workers ($\delta = 0$).

Further, in the iid case when $\zeta = 0$ the rate above simplifies to $\mathcal{O}(\frac{\sigma}{\sqrt{T}} \cdot \sqrt{c\delta + 1/n})$ which matches the optimal iid Byzantine robust rates of (Karimireddy et al., 2021). In both these cases we converge to the optimum and can make the gradient arbitrarily small. However, when $\delta > 0$ and $\zeta > 0$, Theorem II only shows convergence to a radius of $\mathcal{O}(\sqrt{\delta\zeta})$ and not to the actual optimum. We will next explore this limitation.

5.3 LOWER BOUNDS AND THE CHALLENGE OF HETEROGENEITY

Suppose worker j sends us an update which looks ‘weird’ and is very different from the updates from the rest of the workers. This may be because worker j might be malicious and their update represents an attempted attack. It may also be because worker j has highly *non-representative data*. In the former case the update should be ignored, whereas in the latter the update represents a valuable source of specialized data. However, it is impossible for the server to distinguish between the two situations. The above argument can in fact be formalized to prove the following lower bound.

Theorem III. *Given any optimization algorithm ALG, we can find n functions $\{f_1(x), \dots, f_n(x)\}$ of which at least $(1 - \delta)n$ are good (belong to \mathcal{G}), 1-smooth, μ -strongly convex functions, and satisfy $\mathbb{E}_{i \sim \mathcal{G}} \|\nabla f_i(x) - \nabla f(x)\| \leq \zeta^2$ such that the output of ALG given access to these n function has an error at least*

$$\mathbb{E}[f(\text{ALG}(f_1, \dots, f_n)) - f^*] \geq \Omega\left(\frac{\delta\zeta^2}{\mu}\right) \quad \text{and} \quad \mathbb{E}\|\nabla f(\text{ALG}(f_1, \dots, f_n))\|^2 \geq \Omega(\delta\zeta^2).$$

The expectation above is over the potential randomness of the algorithm. This theorem unfortunately implies that it is impossible to converge to the true optimum in the presence of Byzantine workers. Note that the above lower bound is information theoretic in nature and is independent of how many gradients are computed or how long the algorithm is run.

Remark 5 (Matches lower bound). *Suppose that we satisfy the heterogeneity condition (2) with $\zeta^2 > 0$ and $\sigma = 0$. Then, the rate in Theorem II can be simplified to $\mathcal{O}(\delta\zeta^2 + 1/T)$. While the second term in this decays to 0 with T , the first term remains, implying that we only converge to a radius of $\sqrt{\delta\zeta}$ around the optimum. However, this matches our lower bound result from Theorem III and hence is in general unimprovable.*

This is a very strong negative result and seems to indicate that Byzantine robustness might be impossible to achieve in real world federated learning. This would be major stumbling block for deployment since the system would provably be vulnerable to attackers. We will next carefully examine the lower bound and will attempt to circumvent it.

5.4 CIRCUMVENTING LOWER BOUNDS USING OVERPARAMETERIZATION

We previously saw some strong impossibility results posed by heterogeneity. In this section, we show that while indeed in the worst case being robust under heterogeneity is impossible, we may still converge to the true optimum under more realistic settings. Let us consider a different heterogeneity bound in place of (2):

$$\mathbb{E}_{j \sim \mathcal{G}} \|\nabla f_j(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq B^2 \|\nabla f(\mathbf{x})\|^2, \quad \forall \mathbf{x}. \quad (3)$$

Note that at the optimum \mathbf{x}^* we have $\nabla f(\mathbf{x}^*) = 0$, and hence this assumption implies that $\nabla f_j(\mathbf{x}^*) = 0$ for all $j \in \mathcal{G}$. This is satisfied if the model is *sufficiently over-parameterized* and typically holds in most realistic settings (Vaswani et al., 2019).

Theorem IV. *Suppose we are given a (δ_{\max}, c) -ARAGG and n workers with loss functions $\{f_1, \dots, f_n\}$ satisfying the conditions in Theorem II with $\delta \leq \delta_{\max}$ and (3) for some $B^2 < \frac{1}{3c\delta}$. Then, for $F^0 := f(\mathbf{x}^0) - f^*$, the output of Algorithm 2 satisfies*

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla f(\mathbf{x}^{t-1})\|^2 \leq \mathcal{O}\left(\frac{1}{1-3c\delta B^2} \cdot \left(\sigma \sqrt{\frac{LF^0}{T}} (c\delta + 1/n) + \frac{LF^0}{T}\right)\right).$$

Remark 6 (Overparameterization fixes convergence). *The rate in Theorem IV not only goes to 0 with T , but also matches that of the optimal iid rate of $\mathcal{O}(\frac{\sigma}{\sqrt{T}} \cdot \sqrt{c\delta + 1/n})$ (Karimireddy et al.,*

Table 3: Table 1 + Bucketing ($s=2$).

Aggr	iid	non-iid
AVG	98.80 \pm 0.10	98.74 \pm 0.02
KRUM	98.35 \pm 0.20	93.27 \pm 0.10
CM	98.26 \pm 0.22	95.59 \pm 0.89
RFA	98.75 \pm 0.14	97.34 \pm 0.58
CCLIP	98.79 \pm 0.10	98.75 \pm 0.02

Table 4: Table 2 + Bucketing ($s=2$).

Aggr	iid	non-iid
AVG	93.17 \pm 0.23	92.67 \pm 0.27
KRUM	91.64 \pm 0.30	53.15 \pm 3.96
CM	91.91 \pm 0.24	78.60 \pm 3.15
RFA	93.00 \pm 0.23	91.17 \pm 0.51
CCLIP	93.17 \pm 0.23	92.56 \pm 0.21

2021). Thus, using a stronger heterogeneity assumption allows us to circumvent lower bounds for the non-iid case and converge to a good solution even in the presence of Byzantine workers. This is the first result of its kind, and takes a major step towards realistic and practical robust algorithms.

In the overparameterized setting, we can be sure that we will be able to *simultaneously* optimize all worker’s losses. Hence, over time the agreement between all worker’s gradients increases. This in turn makes any attempts by the attackers to derail training stand out easily, especially towards the end of the training. To take advantage of this increasing closeness, we need an aggregator which automatically adapts the quality of its output as the good workers get closer. Thus, the *agnostic* robust aggregator is crucial to our overparameterized convergence result. We empirically demonstrate the effects of overparameterization in Appendix A.2.3.

6 EXPERIMENTS

In this section, we demonstrate the effects of bucketing on datasets distributed in a non-iid fashion. Throughout the section, we illustrate the tasks, attacks, and defenses by an example of training an MLP on a heterogeneous version of the MNIST dataset (LeCun et al., 1998). The dataset is sorted by labels and sequentially divided into equal parts among good workers; Byzantine workers have access to the entire dataset. Implementations are based on PyTorch (Paszke et al., 2019) and will be made publicly available.² We defer details of setup, implementation, and runtime to Appendix A.

Bucketing against the attacks on non-iid data. In Section 3 we have presented how heterogeneous data can lead to failure of existing robust aggregation rules. Here we apply our proposed bucketing with $s=2$ to the same aggregation rules, showing that bucketing overcomes the described failures. Results are presented in Table 3. Comparing Table 3 with Table 1, bucketing improves the aggregators’ top-1 test accuracy on long-tail and non-iid dataset by 4% to 14% and allows them to learn classes at the tail distribution. For non-iid balanced dataset, bucketing also greatly improves the performance of KRUM and CM and makes RFA and CCLIP close to ideal performance. Similarly, combining aggregators with bucketing also performs much better on non-iid dataset under mimic attack. In Table 4, RFA and CCLIP reach accuracy similar to the iid case, and the performance of KRUM, and CM are improved by around 15%.

Bucketing against general Byzantine attacks. In Figure 1, we present thorough experiments on non-iid data over 25 workers with 5 Byzantine workers, under different attacks. In each subfigure, we compare an aggregation rule with its variant with bucketing. The aggregation rules compared are KRUM, CM, RFA, CCLIP. 5 different kinds of attacks are applied (one per column in the figure): bit flipping (BF), label flipping (LF), *mimic* attack, as well as inner product manipulation (IPM) attack (Xie et al., 2020) and the “a little is enough” (ALIE) attack (Baruch et al., 2019).

- **Bit flipping:** A Byzantine worker sends $-\nabla f(x)$ instead of $\nabla f(x)$ due to hardware failures etc.
- **Label flipping:** Corrupt MNIST dataset by transforming labels by $\mathcal{T}(y) := 9 - y$.
- **Mimic:** Explained in Section 3.2.
- **IPM:** The attackers send $-\frac{\epsilon}{|\mathcal{G}|} \sum_{i \in \mathcal{G}} \nabla f(x_i)$ where ϵ controls the strength of the attack.
- **ALIE:** The attackers estimate the mean $\mu_{\mathcal{G}}$ and standard deviation $\sigma_{\mathcal{G}}$ of the good gradients, and send $\mu_{\mathcal{G}} - z\sigma_{\mathcal{G}}$ to the server where z is a small constant controlling the strength of the attack.

Both IPM and ALIE are the state-of-the-art attacks in the iid distributed learning setups which takes advantage of the variances among workers. These attacks are much stronger in the non-iid setup. In the last two columns of Figure 1 we show that worker momentum and bucketing reduce such variance while momentum alone is not enough. Overall, Figure 1 shows that bucketing improves

²The anonymous version of code is available at [this url](#).


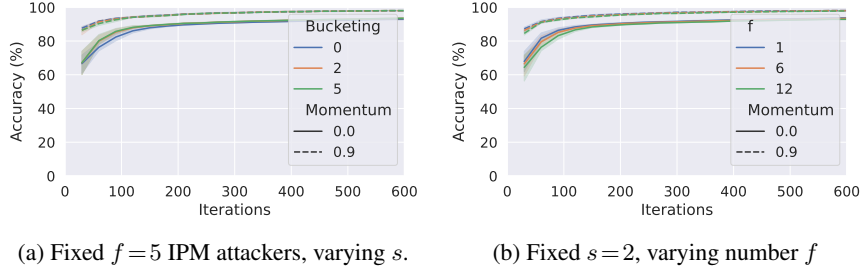


Figure 1: Top-1 test accuracies of KRUM, CM, CCLIP, RFA, under 5 attacks on non-iid datasets.

Figure 2: Top-1 accuracies of CCLIP with varying f and s when training on a cluster of $n=53$ nodes

the performances of almost all aggregators under all kinds of attacks. Note that τ of CCLIP is not finetuned for each attack but rather fixed to $\frac{10}{1-\beta}$ for all attacks. This scaling is required because CCLIP is *not agnostic*. We defer the discussion to Appendix A.2.1.

Bucketing hyperparameter. Finally we study the influence of s and f on the heterogeneous MNIST dataset. We use CCLIP as the base aggregator and apply IPM attack. The Figure 2a confirms that larger s gives faster convergence but $s=2$ can be good enough. The results in Figure 2b shows that $s=2$ still behaves well when increasing f close to 25%. The complete evaluation of the results are deferred to Appendix A.

Discussion. In all our experiments, we consistently observe: i) mild bucketing ($s=2$) improves performance, ii) worker momentum further stabilizes training, and finally iii) CCLIP recovers the ideal performance. Given its ease of implementation, this leads us to strongly recommend using CCLIP in practical federated learning to safeguard against actively malicious agents or passive failures. RFA combined with bucketing and worker momentum also nearly recovers ideal performance and can instead be used when a proper radius τ is hard to find. Ideally, one would want to *adaptively* and automatically set τ so that it works in all instances without any tuning. Designing such a clipping operator as well as its large scale empirical study is left for future work.

7 CONCLUSION

Heterogeneity poses unique challenges for Byzantine robust optimization. The first challenge is that existing defenses attempt to pick a “representative” update, which may not exist in the non-iid setting. This, we showed, can be overcome by using bucketing. A second more fundamental challenge is that it is difficult to distinguish between a “weird” but good worker from an actually Byzantine attacker. In fact, we proved strong impossibility results in such a setting. For this we showed how overparameterization (which is prevalent in real world deep learning) provides a solution, ensuring convergence to the optimum even in the presence of attackers. Together, our results represent a major breakthrough and yield the first practical provably Byzantine robust algorithms for the non-iid setting.

REFERENCES

- Anish Acharya, Abolfazl Hashemi, Prateek Jain, Sujay Sanghavi, Inderjit S Dhillon, and Ufuk Topcu. Robust training in high dimensions via block coordinate geometric median descent. *arXiv preprint arXiv:2106.08882*, 2021.
- Dan Alistarh, Zeyuan Allen-Zhu, and Jerry Li. Byzantine stochastic gradient descent. In *NeurIPS - Advances in Neural Information Processing Systems*, pp. 4613–4623, 2018.
- Zeyuan Allen-Zhu, Faeze Ebrahimi, Jerry Li, and Dan Alistarh. Byzantine-resilient non-convex stochastic gradient descent. In *ICLR 2021 - International Conference of Learning Representations*, 2021.
- Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *AISTATS - Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108, pp. 2938–2948, 2020.
- Moran Baruch, Gilad Baruch, and Yoav Goldberg. A Little Is Enough: Circumventing Defenses For Distributed Learning. *NeurIPS*, 2019.
- Jeremy Bernstein, Jiawei Zhao, Kamyar Azizzadenesheli, and Anima Anandkumar. signSGD with majority vote is communication efficient and fault tolerant. *arXiv preprint arXiv:1810.05291*, 2018.
- Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens, 2018.
- Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent. In *NeurIPS - Advances in Neural Information Processing Systems 30*, pp. 119–129, 2017.
- Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konecny, Stefano Mazzocchi, H Brendan McMahan, et al. Towards federated learning at scale: System design. In *SysML - Proceedings of the 2nd SysML Conference, Palo Alto, CA, USA*, 2019.
- Lingjiao Chen, Hongyi Wang, Zachary Charles, and Dimitris Papailiopoulos. Draco: Byzantine-resilient distributed training via redundant gradients. *arXiv preprint arXiv:1803.09877*, 2018.
- Yudong Chen, Lili Su, and Jiaming Xu. Distributed statistical machine learning in adversarial settings. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 1(2):1–25, Dec 2017. ISSN 2476-1249. doi: 10.1145/3154503.
- Michael B Cohen, Yin Tat Lee, Gary Miller, Jakub Pachocki, and Aaron Sidford. Geometric median in nearly linear time. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 9–21, 2016.
- Georgios Damaskinos, El Mahdi El Mhamdi, Rachid Guerraoui, Arsany Hany Abdelmessih Guirguis, and Sébastien Louis Alexandre Rouault. Aggregathor: Byzantine machine learning via robust gradient aggregation. *Conference on Systems and Machine Learning (SysML) 2019, Stanford, CA, USA*, pp. 19, 2019.
- Deepesh Data and Suhas Diggavi. Byzantine-resilient sgd in high dimensions on heterogeneous data. *arXiv preprint arXiv:2005.07866*, 2020.
- Deepesh Data and Suhas Diggavi. Byzantine-resilient high-dimensional sgd with local iterations on heterogeneous data. In *ICML 2021 - 37th International Conference on Machine Learning*, 2021.
- Avishek Ghosh, Justin Hong, Dong Yin, and Kannan Ramchandran. Robust federated learning in a heterogeneous environment. *arXiv preprint arXiv:1906.06629*, 2019.

- Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D'Oliveira, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badi Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- Sai Praneeth Karimireddy, Quentin Rebjock, Sebastian U Stich, and Martin Jaggi. Error Feedback Fixes SignSGD and other Gradient Compression Schemes. In *ICML 2019 - Proceedings of the 36th International Conference on Machine Learning*, 2019.
- Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. SCAFFOLD: Stochastic controlled averaging for federated learning. In *ICML 2020 - International Conference on Machine Learning*, 2020.
- Sai Praneeth Karimireddy, Li He, and Martin Jaggi. Learning from history for byzantine robust optimization. In *ICML 2021 - 37th International Conference on Machine Learning*, 2021.
- Ahmed Khaled, Konstantin Mishchenko, and Peter Richtárik. Tighter theory for local sgd on identical and heterogeneous data. In *AISTATS 2020 - International Conference on Artificial Intelligence and Statistics*, 2020.
- Leslie Lamport, Robert Shostak, and Marshall Pease. The byzantine generals problem. In *Concurrency: the Works of Leslie Lamport*, pp. 203–226. 2019.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Liping Li, Wei Xu, Tianyi Chen, Georgios B Giannakis, and Qing Ling. RSA: Byzantine-robust stochastic aggregation methods for distributed learning from heterogeneous datasets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1544–1551, 2019.
- H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629*, 2016.
- Edward Meeds, Remco Hendriks, Said Al Faraby, Magiel Bruntink, and Max Welling. Mlthb: machine learning in the browser. *PeerJ Computer Science*, 1:e11, Jul 2015. ISSN 2376-5992. doi: 10.7717/peerj-cs.11. URL <http://dx.doi.org/10.7717/peerj-cs.11>.
- El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. The hidden vulnerability of distributed learning in byzantium. *arXiv preprint arXiv:1802.07927*, 2018.
- El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. Distributed momentum for byzantine-resilient stochastic gradient descent. In *ICLR 2021 - International Conference on Learning Representations*, 2021.
- Stanislav Minsker et al. Geometric median and robust estimation in banach spaces. *Bernoulli*, 21(4): 2308–2335, 2015.
- Ken Miura and Tatsuya Harada. Implementation of a practical distributed calculation system with browsers and javascript, and application to distributed deep learning. *arXiv preprint arXiv:1503.05743*, 2015.
- Erkki Oja. Simplified neuron model as a principal component analyzer. *Journal of mathematical biology*, 15(3):267–273, 1982.

- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pp. 8024–8035, 2019.
- Krishna Pillutla, Sham M. Kakade, and Zaid Harchaoui. Robust Aggregation for Federated Learning. *arXiv preprint arXiv:1912.13445*, 2019.
- Shashank Rajput, Hongyi Wang, Zachary Charles, and Dimitris Papailiopoulos. Detox: A redundancy-based framework for faster and more robust gradient aggregation. *arXiv preprint arXiv:1907.12205*, 2019.
- F. Sattler, K. Müller, T. Wiegand, and W. Samek. On the byzantine robustness of clustered federated learning. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8861–8865, 2020.
- Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. Can you really backdoor federated learning? *arXiv preprint arXiv:1911.07963*, 2019.
- Sharan Vaswani, Francis Bach, and Mark Schmidt. Fast and faster convergence of sgd for over-parameterized models and an accelerated perceptron. In *AISTATS 2019 - The 22nd International Conference on Artificial Intelligence and Statistics*, 2019.
- Hongyi Wang, Kartik Sreenivasan, Shashank Rajput, Harit Vishwakarma, Saurabh Agarwal, Jy-yong Sohn, Kangwook Lee, and Dimitris Papailiopoulos. Attack of the tails: Yes, you really can backdoor federated learning. *arXiv preprint arXiv:2007.05084*, 2020.
- Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. Zeno: Distributed stochastic gradient descent with suspicion-based fault-tolerance. In *ICML 2019 - 35th International Conference on Machine Learning*, 2019.
- Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. Fall of Empires: Breaking Byzantine-tolerant SGD by Inner Product Manipulation. In *UAI - Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, 2020.
- Yi-Rui Yang and Wu-Jun Li. Basgd: Buffered asynchronous sgd for byzantine learning. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11751–11761. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/yang21e.html>.
- Dong Yin, Yudong Chen, Kannan Ramchandran, and Peter Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. *arXiv preprint arXiv:1803.01498*, 2018.
- Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *AAAI 2019 - Conference on Artificial Intelligence*, 2019.