

---

# Fixing by Mixing: A Recipe for Optimal Byzantine ML under Heterogeneity

---

Youssef Allouah<sup>†\*</sup>

Nirupam Gupta<sup>†</sup>

Sadegh Farhadkhani<sup>†</sup>

Rafaël Pinot<sup>†</sup>

Rachid Guerraoui<sup>†</sup>

John Stephan<sup>†</sup>

<sup>†</sup>Ecole Polytechnique Fédérale de Lausanne (EPFL)

\*Correspondence to: [youssef.allouah@epfl.ch](mailto:youssef.allouah@epfl.ch)

## Abstract

Byzantine machine learning (ML) aims to ensure the resilience of distributed learning algorithms to misbehaving (or *Byzantine*) machines. Although this problem received significant attention, prior works often assume the data held by the machines to be *homogeneous*, which is seldom true in practical settings. Data *heterogeneity* makes Byzantine ML considerably more challenging, since a Byzantine machine can hardly be distinguished from a non-Byzantine outlier. A few solutions have been proposed to tackle this issue, but these provide suboptimal probabilistic guarantees and fare poorly in practice.

This paper closes the theoretical gap, achieving optimality and inducing good empirical results. In fact, we show how to automatically adapt existing solutions for (homogeneous) Byzantine ML to the heterogeneous setting through a powerful mechanism, we call *nearest neighbor mixing* (NNM), which boosts any standard robust distributed gradient descent variant to yield optimal Byzantine resilience under heterogeneity. We obtain similar guarantees (in expectation) by plugging NNM in the distributed *stochastic* heavy ball method, a practical substitute to distributed gradient descent. We obtain empirical results that significantly outperform state-of-the-art Byzantine ML solutions.

## 1 INTRODUCTION

Distributed machine learning (ML), i.e., distributing the training process amongst several machines (or *workers*), has been pivotal to the development of large complex models

with high accuracy guarantees (Abadi et al., 2015; Konečný et al., 2016; Kairouz et al., 2021). In the now standard *master-worker* architecture, distributed ML essentially consists in the workers sharing their local actions with the help of a master machine (a.k.a., *server*) to compute an accurate *global* model. Despite its rising popularity, distributed ML is arguably very fragile and not yet ready for real-world deployment. In particular, a handful of misbehaving (a.k.a., *Byzantine* (Lamport et al., 1982)) workers can highly compromise the efficacy of standard distributed ML algorithms such as distributed gradient descent (DGD) (Bertsekas and Tsitsiklis, 2015), by sending erroneous information to the server, e.g., see (Baruch et al., 2019; Xie et al., 2019). Such behaviors can be caused by software and hardware bugs, by poisoned data or by malicious players controlling part of the system. Besides, this vulnerability can lead to severe societal repercussions if the resulting ML models are deployed in sensitive data-oriented applications such as medicine.

The problem of Byzantine resilience in distributed ML (or *Byzantine ML*) has received significant attention in recent years (Blanchard et al., 2017; Chen et al., 2017; Yin et al., 2018; Allen-Zhu et al., 2020; El Mhamdi et al., 2021a; Karimireddy et al., 2021; Karimireddy et al., 2022; Farhadkhani et al., 2022b). It consists in designing a distributed algorithm that delivers an accurate model despite the presence of a subset of Byzantine workers. The standard approach (in a *master-worker* architecture) consists in having the server compute a robust aggregation to merge the information sent by the workers, to discard outliers. Most prior works relies however upon the strong assumption of *homogeneity*; i.e., the data sampled by the workers during the training process are assumed to be identically distributed. This assumption, although justifiable in a centralized setting (Bottou et al., 2018), is impractical in a distributed environment (Kairouz et al., 2021). Indeed, as each worker only holds a small part of the entire training dataset, the data samples at the workers are usually *heterogeneous* and need not be an accurate representation of the entire population. These differences in the workers' data samples can camouflage disruptive deviations of Byzantine machines from the prescribed algorithm,

making the problem of Byzantine ML under heterogeneity significantly more challenging than its homogeneous counterpart (El Mhamdi et al., 2021a; Karimireddy et al., 2022).

Recent works have shown that heterogeneity in workers’ data inevitably prevents any distributed ML algorithm from delivering an arbitrarily accurate model in the presence of Byzantine workers (Charikar et al., 2017; Liu et al., 2021). More specifically, in the context of smooth loss functions, a Byzantine ML solution can only deliver an approximate stationary point with an error lower bounded by the fraction of Byzantine workers times the maximal dispersion in workers’ gradients (stemming from data heterogeneity) (Karimireddy et al., 2022). Despite a few attempts (Gupta and Vaidya, 2020; El Mhamdi et al., 2021a; Karimireddy et al., 2022), no solution has so far matched the above lower bound deterministically, i.e., achieved *optimal Byzantine resilience*. The best solutions so far (Karimireddy et al., 2022) are randomized and could only match the lower bound in *expectation* (see Section 2).

**Main result.** We close the theoretical gap and reach optimal Byzantine resilience under heterogeneity. Specifically, we show how to automatically adapt existing solutions for homogeneous Byzantine ML to the heterogeneous setting, while ensuring optimality. We do so through *nearest neighbor mixing* (NNM), a pre-aggregation method that averages each input with a subset of their nearest neighbors. We show that enhancing D-GD using a composition of NNM and a standard robust aggregation directly yields the first Byzantine ML solution to achieve optimal resilience under heterogeneity. Our guarantee holds as long as less than half of the workers are Byzantine, which is optimal.

**Technical contributions.** To prove our guarantees, we introduce a novel robustness criterion called  $(f, \kappa)$ -robustness. This notion quantifies the ability of an aggregation rule to estimate the average of honest workers’ inputs despite  $f$  out of  $n$  workers being Byzantine. Crucially, our criterion characterizes a class of aggregations rules (for which  $\kappa = \mathcal{O}(f/n)$ ) that grant optimal Byzantine resilience to D-GD under heterogeneity. While many notable aggregation rules (e.g., *geometric median* (Small, 1990), *coordinate-wise median* (Yin et al., 2018), and *Krum* (Blanchard et al., 2017)) satisfy  $(f, \kappa)$ -robustness, they fall short of optimal robustness, as their  $\kappa$  is in  $\Theta(1)$ . Our main technical contribution is showing that NNM overcomes this shortcoming. Particularly, we prove that NNM deterministically reduces the variance of the honest inputs by a factor  $\mathcal{O}(f/n)$ , while sufficiently limiting their drift from the true average. Consequently, we show that composing a  $(f, \mathcal{O}(1))$ -robust aggregation with NNM enables the larger class of  $(f, \mathcal{O}(1))$ -robust aggregation rules to confer optimal Byzantine resilience to D-GD. As a byproduct, we provide novel and tight robustness analyses of aggregation rules, notably *coordinate-wise trimmed mean* (Yin et al., 2018), which

may be of independent interest to the Byzantine ML and robust statistics communities.

**Empirical evaluation.** Although we prove D-GD enhanced by NNM to be optimal, workers still need to compute the *full* gradient of their local loss function at every step, and risk getting stuck at saddle points (Du et al., 2017; Bottou et al., 2018). We go one step further and show that applying NNM to a stochastic variant of D-GD, namely the distributed stochastic heavy ball method (D-SHB) (Polyak, 1964), matches the lower bound in expectation<sup>1</sup>. We empirically show that the resulting scheme significantly improves over the state-of-the-art when tested on standard classification tasks such as MNIST and CIFAR-10. In short, our approach enables a modular practice of Byzantine resilient ML under heterogeneity, by first patching existing solutions with NNM and then deploying them with D-SHB.

**Paper outline.** Section 2 formally presents the problem of Byzantine ML and discusses the related work. Section 3 introduces the analysis framework for robust variants of D-GD. Section 4 presents NNM, our solution to render existing methods optimal. Section 5 presents a practical stochastic extension of our approach. Section 6 presents our experimental results. We defer all proofs to Appendix 8-13.

## 2 BACKGROUND & RELATED WORK

We consider a master-worker distributed architecture with  $n$  workers  $w_1, \dots, w_n$  and a central server. The workers hold local datasets  $\mathcal{D}_1, \dots, \mathcal{D}_n$ , composed each of  $m$  data points from an input space  $\mathcal{Z}$ . Specifically, for any  $i \in [n]$ ,  $\mathcal{D}_i := \{z_1^{(i)}, \dots, z_m^{(i)}\} \subset \mathcal{Z}^m$ . For a given model parameterized by vector  $\theta \in \mathbb{R}^d$ , each worker  $w_i$  has a loss function

$$\mathcal{L}_i(\theta) := \frac{1}{m} \sum_{k=1}^m \ell(\theta, z_k^{(i)}),$$

where  $\ell : \mathbb{R}^d \times \mathcal{Z} \rightarrow \mathbb{R}$  is a point-wise loss function, which we assume to be differentiable with respect to  $\theta$ . Furthermore, we assume each loss function  $\mathcal{L}_i$  to be  $L$ -smooth; i.e., for all  $\theta, \theta' \in \mathbb{R}^d$ ,

$$\|\nabla \mathcal{L}_i(\theta) - \nabla \mathcal{L}_i(\theta')\| \leq L \|\theta - \theta'\|.$$

**Objective in a Byzantine-free setting.** When all the workers are honest, i.e., they follow the prescribed algorithm correctly, the goal of the server is to compute a stationary point of the global loss function<sup>2</sup>  $\mathcal{L}(\theta) := \frac{1}{n} \sum_{i=1}^n \mathcal{L}_i(\theta)$ .

<sup>1</sup>The expectation is on the randomness due to data subsampling, unlike (Karimireddy et al., 2022) which features additional sources of randomness that hinder empirical performance. See Section 5.2 for details.

<sup>2</sup>Ideally, the server seeks to find a global minimizer of  $\mathcal{L}$ . However, as the loss could be non-convex (e.g., for neural networks), global minimization is NP-hard in general (Boyd and Vandenberghe, 2004). Hence, we usually aim to find stationary points.

Specifically, the server seeks to compute a stationary point  $\theta^*$ , i.e.  $\nabla \mathcal{L}(\theta^*) = 0$ , where  $\nabla \mathcal{L}$  denotes the *gradient* of  $\mathcal{L}$ . Assuming that  $\mathcal{L}$  admits a stationary point, this goal can be achieved using the classical distributed (stochastic) gradient descent (D-GD/SGD) method (Bertsekas and Tsitsiklis, 2015), up to an approximation error. In this method, the server maintains a model which is updated iteratively upon averaging the (stochastic) gradients computed by the workers on their loss functions.

**Objective with Byzantine workers.** We consider an adversarial setting where  $f$  workers, of a priori unknown identities, are Byzantine (Lamport et al., 1982). These workers need not follow the prescribed protocol and may send arbitrary messages to the server. Here, the goal of the server is finding a stationary point of the average loss function of the honest (i.e., non-Byzantine) workers. We formally define Byzantine resilience in this context as follows.

**Definition 1** ( $(f, \varepsilon)$ -Byzantine resilience). *A learning algorithm is said  $(f, \varepsilon)$ -Byzantine resilient if, despite the presence of  $f$  Byzantine workers, it outputs  $\hat{\theta}$  such that*

$$\left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \leq \varepsilon,$$

where  $\mathcal{H}$  denotes the set of indices of honest workers and  $\mathcal{L}_{\mathcal{H}}(\theta) := \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \mathcal{L}_i(\theta)$ .

In words, an  $(f, \varepsilon)$ -Byzantine resilient algorithm finds an  $\varepsilon$ -approximate stationary point for the honest empirical loss even in the presence of  $f$  Byzantine workers. Note that  $(f, \varepsilon)$ -Byzantine resilience is impossible in general (for any  $\varepsilon$ ) when  $f \geq n/2$ , see (Liu et al., 2021). Therefore, throughout the paper, we assume that  $f < n/2$ .

**Standard Byzantine ML solutions.** A typical strategy to obtain  $(f, \varepsilon)$ -Byzantine resilience is to enhance the D-GD (or D-SGD) method by replacing the simple averaging of the workers' gradients at the server by a *robust aggregation rule*. Basically, such a scheme aims to mitigate the negative impact of Byzantine gradients by accurately estimating the average of honest workers' gradients. Prominent aggregation rules include Krum (Blanchard et al., 2017), geometric median (GM) (Small, 1990; Pillutla et al., 2022; Acharya et al., 2022), coordinate-wise median (CWMed) (Yin et al., 2018), coordinate-wise trimmed mean (CWTM) (Yin et al., 2018), and minimum diameter averaging (MDA) (Rousseeuw, 1985; El Mhamdi et al., 2021a). Recent studies have also explored the idea of using the *history* of workers' gradients to strengthen the resilience guarantees of these solutions, e.g., by using distributed momentum (Karimireddy et al., 2022; Farhadkhani et al., 2022b) or by tracking the history of the workers' gradients (Allen-Zhu et al., 2020).

Nevertheless, all these works rely heavily on the assumption that the honest workers have *homogeneous* data, i.e., there exists a *ground-truth distribution*  $\mathcal{D}$  such that  $D_i \sim \mathcal{D}^m$  for

all  $i \in \mathcal{H}$ . Crucially, the robustness of these methods deteriorates drastically when this assumption is violated (Karimireddy et al., 2022).

**The challenge of heterogeneity in Byzantine ML.** The data across the honest workers is arguably *heterogeneous* in real-world distributed settings. In the context of non-convex optimization, data heterogeneity can be modeled as stated per the following assumption (Karimireddy et al., 2020; Karimireddy et al., 2022).

**Assumption 1** (Bounded heterogeneity). *There exists a real value  $G$  such that for all  $\theta \in \mathbb{R}^d$ ,*

$$\frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \|\nabla \mathcal{L}_i(\theta) - \nabla \mathcal{L}_{\mathcal{H}}(\theta)\|^2 \leq G^2.$$

Such heterogeneity makes the problem of Byzantine ML much more challenging, as the server can confuse incorrect gradients (from a Byzantine worker) with correct gradients from an honest worker holding outlier data points. Indeed, recent works show that there exists a lower bound on the error of any distributed algorithm in the presence of Byzantine workers (El Mhamdi et al., 2021a; Karimireddy et al., 2022). Specifically, we have the following result, owing to Theorem III in (Karimireddy et al., 2022).

**Proposition 1.** *If a learning algorithm  $\mathcal{A}$  is  $(f, \varepsilon)$ -Byzantine resilient for every collection of smooth loss functions  $\mathcal{L}_1, \dots, \mathcal{L}_n$  satisfying Assumption 1, then  $\varepsilon = \Omega(f/n G^2)$ .*

**Brittleness of existing solutions for heterogeneity.** Only a handful of prior works have studied Byzantine ML under heterogeneity. The problem was first formally addressed in (Li et al., 2019) by proposing *RSA*, a variant of D-SGD built upon  $\ell_p$  regularization. However, the analysis of *RSA* relies on the assumption of strong convexity, which is rarely satisfied in modern-day ML. A subsequent work (Gupta and Vaidya, 2020) introduced a novel aggregation rule, namely comparative gradient elimination (CGE). While this work provides some valuable insight on the general problem of heterogeneous Byzantine ML (notably on the need for redundancy), CGE fails to guarantee convergence even in the homogeneous setting (see (Farhadkhani et al., 2022b)). The work (El Mhamdi et al., 2021a) considers a peer-to-peer setting with asynchronous communication and heterogeneous data. However, the proposed algorithm is only analyzed asymptotically and provides suboptimal probabilistic robustness. In this peer-to-peer setting, another work used a nearest neighbor scheme similar to ours (Farhadkhani et al., 2022a), with honest nodes using their own honest vectors as pivots. However, in our master-worker setting, their technique cannot be used as the server does not have access to any reliable vector.

The closest relevant work to ours is (Karimireddy et al., 2022), which also proposes a pre-aggregation step called *Bucketing*, which is reminiscent of the *median-of-means*

estimator (Nemirovsky et al., 1983; Jerrum et al., 1986; Alon et al., 1996; Tu et al., 2021). Essentially, Bucketing consists in randomly partitioning the inputs into *buckets*, and then feeding the average of the inputs of each bucket to the robust aggregation rule. The randomness of the partition process reduces the *empirical variance* of the honest inputs in expectation, but it highly compromises the worst-case robustness of the scheme (see Appendix 10 for a detailed explanation). Furthermore, our experiments expose the inability of these methods to defend against state-of-the-art Byzantine attacks (Xie et al., 2019; Baruch et al., 2019; Allen-Zhu et al., 2020; Karimireddy et al., 2022).

### 3 ROBUST DISTRIBUTED GRADIENT DESCENT

In this section, we introduce a general framework for analyzing the convergence of robust variants of D-GD (or *robust D-GD*). We use this framework later in Section 4 to showcase the benefits of our proposed pre-aggregation step NNM. We start by presenting the skeleton of robust D-GD in Section 3.1 below. Then, we present the convergence analysis of robust D-GD in Section 3.2 upon introducing the notion of  $(f, \kappa)$ -robustness.

#### 3.1 Description of Robust D-GD

Robust D-GD, summarized below in Algorithm 1, is an iterative algorithm that proceeds in  $T$  steps.

---

#### Algorithm 1: Robust D-GD

---

**Input:** Initial model  $\theta_0$ , robust aggregation  $F$ , learning rate  $\gamma$ , and number of steps  $T$ .

**for**  $t = 1 \dots T$  **do**

**Server** broadcasts  $\theta_{t-1}$  to all workers;  
**for every honest worker,  $i$  in parallel do**  
└ Compute and send gradient  $g_t^{(i)} = \nabla \mathcal{L}_i(\theta_{t-1})$ ;  
**Server** aggregates the gradients:  
 $R_t = F(g_t^{(1)}, \dots, g_t^{(n)})$ ;  
**Server** updates the model:  $\theta_t = \theta_{t-1} - \gamma R_t$ ;

**Server** finds  $\tau \in \arg \min_{1 \leq t \leq T} \|R_t\|$  and sets  $\hat{\theta} = \theta_{\tau-1}$ .

**return**  $\hat{\theta}$ ;

---

Essentially, the server starts by initializing a parameter vector  $\theta_0$ . At each step  $t \in [T]$ , the server broadcasts the model  $\theta_{t-1}$  to all the workers. After receiving  $\theta_{t-1}$ , each honest worker sends back the gradient  $g_t^{(i)} = \nabla \mathcal{L}_i(\theta_{t-1})$ , computed by evaluating  $\theta_{t-1}$  on their local dataset  $\mathcal{D}_i$ . While the honest workers follow the algorithm correctly, a Byzantine worker  $w_i$  may send any arbitrary value for  $g_t^{(i)}$ . Upon receiving the gradients from all the workers, the server aggregates them using a robust aggregation rule  $F: \mathbb{R}^{d \times n} \rightarrow \mathbb{R}^d$ .

Specifically, the server computes

$$R_t = F(g_t^{(1)}, \dots, g_t^{(n)}).$$

Finally, the server updates the model to  $\theta_t = \theta_{t-1} - \gamma R_t$ , where  $\gamma > 0$  is referred to as the learning rate. After  $T$  steps, the server outputs the model  $\theta_t$  for which the associated aggregate  $R_t$  has the smallest norm. That is, the algorithm outputs  $\hat{\theta} = \theta_{\tau-1}$ , where  $\tau \in \arg \min_{1 \leq t \leq T} \|R_t\|$ .

#### 3.2 Analysis of Robust D-GD

At the core of robust D-GD lies the aggregation rule  $F$ . To tightly analyze the utility of  $F$ , we introduce the notion of  $(f, \kappa)$ -robustness which unifies previous robustness criteria and is sufficiently fine-grained to obtain tight convergence guarantees. In words,  $(f, \kappa)$ -robustness ensures that the error of an aggregation rule, in estimating the average of the honest inputs, is *uniformly* bounded by  $\kappa$  times the variance of honest inputs. Formally, it is defined as follows.

**Definition 2 ( $(f, \kappa)$ -robustness).** *Let  $f < n/2$  and  $\kappa \geq 0$ . An aggregation rule  $F$  is said to be  $(f, \kappa)$ -robust if for any vectors  $x_1, \dots, x_n \in \mathbb{R}^d$ , and any set  $S \subseteq [n]$  of size  $n - f$ ,*

$$\|F(x_1, \dots, x_n) - \bar{x}_S\|^2 \leq \frac{\kappa}{|S|} \sum_{i \in S} \|x_i - \bar{x}_S\|^2$$

where  $\bar{x}_S = \frac{1}{|S|} \sum_{i \in S} x_i$ . We refer to  $\kappa$  as the robustness coefficient.

Our criterion unifies the existing robustness definitions including  $(f, \lambda)$ -resilient averaging (Farhadkhani et al., 2022b) and  $(\delta_{max}, c)$ -ARAgg (Karimireddy et al., 2022). Specifically, when  $\kappa = \mathcal{O}(f/n)$ , our definition implies both  $(f, \lambda)$ -resilient averaging and  $(\delta_{max}, c)$ -ARAgg<sup>3</sup>. We prove this claim in Appendix 8.3. Furthermore,  $(f, \kappa)$ -robustness allows us to devise a general convergence analysis for robust D-GD when up to  $f$  workers are Byzantine, under the standard heterogeneity assumption. We present our general convergence analysis of robust D-GD in Theorem 1 below, assuming  $F$  to be  $(f, \kappa)$ -robust. Recall that  $\mathcal{H}$  denotes the set of indices for honest workers.

**Theorem 1.** *Let Assumption 1 hold and recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 1 with  $T \geq 1$  and learning rate  $\gamma = 1/L$ . If  $F$  is  $(f, \kappa)$ -robust then*

$$\|\nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta})\|^2 \leq 4\kappa G^2 + \frac{4L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{T},$$

where  $\mathcal{L}^* := \inf_{\theta \in \mathbb{R}^d} \mathcal{L}_{\mathcal{H}}(\theta)$ .

According to Theorem 1, the asymptotic error for robust D-GD (when  $T \rightarrow \infty$ ) is optimal, i.e., it matches the

<sup>3</sup>Although for satisfying  $(f, \lambda)$ -resilient averaging condition having  $\kappa = \sqrt{\lambda}$  suffices, i.e., it need not be in  $\mathcal{O}(f/n)$ .

Aggregation	GM	CWTM	CWMed	Krum	Lower bound
$\kappa$	$\left(1 + \frac{f}{n-2f}\right)^2$	$\frac{f}{n-2f} \left(1 + \frac{f}{n-2f}\right)$	$\left(1 + \frac{f}{n-2f}\right)^2$	$1 + \frac{f}{n-2f}$	$\frac{f}{n-2f}$

Table 1: Robustness coefficient  $\kappa$  for different  $(f, \kappa)$ -robust aggregation rules, ignoring numerical constants. The exact robustness coefficients are deferred to Appendix 8.1. Note also that our analysis is tight and significantly improves over previous works (Farhadkhani et al., 2022b; Karimireddy et al., 2022), e.g., eliminating the dependence on dimension for CWMed and CWTM.

lower bound from Proposition 1 if  $F$  is  $(f, \kappa)$ -robust with  $\kappa = \mathcal{O}(f/n)$ . In the homogeneous case, i.e., when  $G = 0$ , robust D-GD can asymptotically reach a stationary point of the average honest loss function despite the presence of  $f$  Byzantine workers, as long as  $F$  is  $(f, \kappa)$ -robust with a bounded  $\kappa$ . Note that the convergence rate of  $1/T$  is standard for smooth non-convex loss functions when analyzing first-order methods such as gradient descent (Ghadimi and Lan, 2016).

**Suboptimality of existing aggregations.** Several existing aggregation rules such as CWTM, Krum, GM, and CWMed can be shown to be  $(f, \kappa)$ -robust. The robustness coefficients  $\kappa$  for these rules are listed in Table 1, and the formal derivations are deferred to Appendix 8.1. Note also that for any  $f < n/2$ , an aggregation rule cannot be  $(f, \kappa)$ -robust for  $\kappa < \frac{f}{n-2f}$  (see Appendix 8.2 for details). This lower bound means that, in general, a robust aggregation rule cannot provide an estimate that is arbitrarily close to the average of honest inputs. This also indicates that the values in Table 1 for Krum, GM, and CWMed are suboptimal, as they do not match the lower bound. We show in the next section that NNM solves this issue by boosting the robustness of these aggregation rules and provides optimal convergence for robust D-GD.

## 4 FIXING BY NEAREST NEIGHBOR MIXING

In this section, we present a principled way of fixing the suboptimality of existing solutions in terms of Byzantine resilience. Specifically, we introduce a pre-aggregation algorithm called nearest neighbor mixing (NNM), and prove optimal Byzantine robustness when it is embedded in D-GD. We describe the NNM procedure in Section 4.1 and demonstrate in Section 4.2 that NNM amplifies robustness when applied prior to an aggregation rule.

### 4.1 Description of NNM

Given a set of  $n$  input vectors  $x_1, \dots, x_n \in \mathbb{R}^d$ , NNM replaces every vector with the average of its  $n - f$  nearest neighbors (including itself). Formally, NNM outputs

$(y_1, \dots, y_n) = \text{NNM}(x_1, \dots, x_n)$  where for each  $i \in [n]$ ,

$$y_i = \frac{1}{n-f} \sum_{j=1}^{n-f} x_{i:j}; \quad (1)$$

where  $x_{i:j}$  is the  $j^{\text{th}}$  nearest vector to  $x_i$  in  $(x_1, \dots, x_n)$ . Intuitively, in the context of robust D-GD, applying NNM mixes the gradients artificially, hence making every *mixed gradient* a better representation of  $\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})$ . The overall procedure for NNM can be found in Algorithm 2.

---

#### Algorithm 2: Nearest Neighbor Mixing (NNM)

---

**Input:** number of inputs  $n$ , number of Byzantine inputs  $f < n/2$ , vectors  $x_1, \dots, x_n \in \mathbb{R}^d$ .

**for**  $i = 1 \dots n$  **do**

Sort inputs to get  $(x_{i:1}, \dots, x_{i:n})$  such that

$$\|x_{i:1} - x_i\| \leq \dots \leq \|x_{i:n} - x_i\|;$$

Average the  $n - f$  nearest neighbors of  $x_i$ , i.e.,

$$y_i = \frac{1}{n-f} \sum_{j=1}^{n-f} x_{i:j};$$

**return**  $y_1, \dots, y_n$ ;

---

**Remark 1.** *The computational cost of NNM is  $\mathcal{O}(dn^2)$  in the worst case, which is due to the search of the  $n - f$  nearest neighbors of each input. Faster algorithms for approximate nearest neighbor search (Hajebi et al., 2011; Muja and Lowe, 2014) could be used for efficiency. Nevertheless, we argue that the cost of NNM is comparable to (or even smaller than) several aggregation rules including Krum (Blanchard et al., 2017), Multi-Krum (Blanchard et al., 2017), and MDA (Rousseeuw, 1985; El Mhamdi et al., 2018). Below, we list the cost of prominent aggregation rules: Krum and Multi-Krum:  $\mathcal{O}(dn^2)$ , CWMed (Yin et al., 2018) and MeaMed (Xie et al., 2018):  $\mathcal{O}(dn)$ , CWTM (Yin et al., 2018):  $\mathcal{O}(dn \log n)$ ,  $\epsilon$ -approximate GM (Acharya et al., 2022):  $\mathcal{O}(dn + d\epsilon^{-2})$ , MDA:  $\mathcal{O}(dn^2 + \binom{n}{f}n^2)$ . Finally, unlike spectral methods (Shejwalkar and Houmansadr, 2021), we stress that our algorithm preserves linear dependency in  $d$ , which may be extremely large in modern-day ML (i.e.,  $d \gg n$ ).*

## 4.2 Analysis of NNM

We now present in Lemma 1 the robustness amplification that NNM brings to a  $(f, \kappa)$ -robust aggregation rule. We then show in Corollary 1 how it leads to optimal Byzantine resilience guarantees under heterogeneity.

**Lemma 1.** *Let  $f < n/2$  and  $F: \mathbb{R}^{d \times n} \rightarrow \mathbb{R}^d$ . If  $F$  is  $(f, \kappa)$ -robust, then  $F \circ \text{NNM}$  is  $(f, \kappa')$ -robust with*

$$\kappa' \leq \frac{8f}{n-f}(\kappa + 1).$$

Essentially, Lemma 1 means that the composition of NNM with any  $(f, \kappa)$ -robust aggregation rule improves the order of magnitude of the robustness coefficient to  $\mathcal{O}(f/n(\kappa + 1))$ . Specifically, this means that if  $F$  has a robustness coefficient  $\kappa = \mathcal{O}(1)$ , NNM renders the new robustness coefficient optimal. We believe that the condition  $\kappa = \mathcal{O}(1)$  is general enough to hold for any standard robust aggregation rule. In fact, assuming that there exists  $\nu > 0$  such that  $n \geq (2 + \nu)f$ , all the aforementioned aggregation rules are  $(f, \kappa)$ -robust with  $\kappa = \mathcal{O}(1)$ . Thus, from a theoretical point of view, any aggregation rule from Table 1 becomes a good candidate in Algorithm 1 when combined with NNM. Then, as stated in Corollary 1 below, we obtain optimal Byzantine resilience under heterogeneity.

**Corollary 1.** *Let Assumption 1 hold and recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 1 with  $\gamma = 1/L$  and aggregation  $F \circ \text{NNM}$ . If  $F$  is  $(f, \kappa)$ -robust with  $\kappa = \mathcal{O}(1)$ , then Algorithm 1 is  $(f, \varepsilon)$ -Byzantine resilient with*

$$\varepsilon = \mathcal{O}(f/nG^2 + 1/T).$$

**Remark 2.** *Note that CWTM achieves an order-optimal robustness coefficient  $\kappa = \mathcal{O}(f/n)$  without the use of NNM, whenever  $n \geq (2 + \nu)f$  for some constant  $\nu > 0$ . However, using NNM prior to CWTM significantly improves its empirical performance (see Section 6). We believe that an average-case analysis, instead of our worst-case approach, could capture this improvement in theory.*

## 5 STOCHASTIC EXTENSION

Despite its optimality, our solution for robust D-GD is computationally demanding, as it requires the honest workers to compute each gradient on the whole dataset. In practice, it is more common to consider stochastic variants of D-GD, where workers compute gradients on random mini-batches of datasets. To accommodate this, we show that NNM can also be used to enhance the performance of robust variants of distributed stochastic heavy ball (D-SHB) that has been recently proven to perform well in Byzantine ML in the homogeneous setting (El Mhamdi et al., 2021b; Karimireddy et al., 2021; Farhadkhani et al., 2022b).

## 5.1 Description of Robust D-SHB

Similarly to robust D-GD, the algorithm proceeds in  $T$  iterations as follows. At every step  $t \in [T]$ , the server holds a model  $\theta_{t-1}$  and each honest worker holds a local momentum  $m_{t-1}^{(i)}$ <sup>4</sup>. The server broadcasts the current model  $\theta_{t-1}$  to all the workers for them to update their local momentum. To do so, each honest worker samples a mini-batch of data  $S_t^{(i)}$  at random from  $D_i$  and computes a stochastic estimate  $g_t^{(i)}$  of its gradient  $\nabla \mathcal{L}_i(\theta_{t-1})$ , defined as

$$g_t^{(i)} = \frac{1}{|S_t^{(i)}|} \sum_{z \in S_t^{(i)}} \nabla \ell(\theta_{t-1}, z). \quad (2)$$

Then, each honest worker updates and sends to the server its local momentum

$$m_t^{(i)} = \beta m_{t-1}^{(i)} + (1 - \beta) g_t^{(i)}, \quad (3)$$

where  $\beta \in (0, 1)$  is the *momentum parameter* and is shared by all the honest users. Similarly to D-SGD, the server computes an aggregate of the momentums it receives as  $R_t = F(m_t^{(1)}, \dots, m_t^{(n)})$ . Finally, the server updates the model to  $\theta_t = \theta_{t-1} - \gamma R_t$ , where  $\gamma > 0$  is the learning rate. After the  $T$  iterations, the server outputs  $\hat{\theta}$  by sampling uniformly from  $(\theta_0, \dots, \theta_{T-1})$ . The overall procedure for robust D-SHB is summarized in Algorithm 3.

---

### Algorithm 3: Robust D-SHB

---

**Input:** Initial model  $\theta_0$ , initial momentum  $m_0^{(i)} = 0$  for honest workers, robust aggregation  $F$ , learning rate  $\gamma$ , momentum coefficient  $\beta$ , and number of steps  $T$ .

**for**  $t = 1 \dots T$  **do**

**Server** broadcasts  $\theta_{t-1}$  to all workers;

**for every honest worker  $i$ , in parallel do**

Compute a stochastic gradient  $g_t^{(i)}$  as per (2);

Update local momentum  $m_t^{(i)}$  as per (3);

Send  $m_t^{(i)}$  to the server;

**Server** aggregates the momentums:

$R_t = F(m_t^{(1)}, \dots, m_t^{(n)})$ ;

**Server** updates the model:  $\theta_t = \theta_{t-1} - \gamma R_t$ ;

Sample  $\hat{\theta}$  uniformly from  $(\theta_0, \dots, \theta_{T-1})$ ;

**return**  $\hat{\theta}$ ;

---

## 5.2 Analysis of Robust D-SHB with NNM

We now provide convergence guarantees for robust D-SHB with NNM. To do so, we make an additional (standard) assumption on the variance of the stochastic gradients.

<sup>4</sup> $\theta_0$  is set by the server and  $m_0^{(i)} = 0$  for all honest workers.

**Assumption 2** (Bounded variance). *For each honest worker  $w_i, i \in \mathcal{H}$ , and all  $\theta \in \mathbb{R}^d$ , it holds that*

$$\frac{1}{m} \sum_{z \in \mathcal{D}_i} \|\nabla_{\theta} \ell(\theta, z) - \nabla \mathcal{L}_i(\theta)\|^2 \leq \sigma^2.$$

We now present in Theorem 2 the extension of our results to D-SHB. Essentially, we analyze Algorithm 3 upon assuming a constant learning rate, that assumptions 1 and 2 hold, and that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. For convenience, we introduce the following numerical constants before stating the theorem:

$$a_1 := 36, a_2 := 6\sqrt{\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*}, a_3 := 1728L, \\ a_4 := 288L, a_5 := 6L a_2^2, \text{ and } \mathcal{L}^* = \inf_{\theta \in \mathbb{R}^d} \mathcal{L}_{\mathcal{H}}(\theta).$$

**Theorem 2.** *Let assumption 1 and 2 hold and recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Let  $F$  be a  $(f, \kappa)$ -robust aggregation rule. Consider Algorithm 3 with momentum coefficient  $\beta = \sqrt{1 - 24\gamma L}$ , and learning rate*

$$\gamma = \min \left\{ \frac{1}{24L}, \frac{a_2}{2a_{\kappa}\sigma\sqrt{T}} \right\},$$

with  $a_{\kappa}^2 := a_3\kappa + \frac{a_4}{n-f}$ . For all  $T \geq 1$ , we have

$$\mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \right] \leq a_1\kappa G^2 + \frac{a_2 a_{\kappa} \sigma}{\sqrt{T}} + \frac{a_5}{T} + \frac{a_2 a_4 \sigma}{n a_{\kappa} T^{3/2}},$$

where the expectation is over the algorithm's randomness.

**Tight probabilistic guarantee.** The non-vanishing error in Theorem 2 is in  $\mathcal{O}(\kappa G^2)$ . Hence, this error is tight when  $\kappa = \mathcal{O}(f/n)$ . However, the main difference with robust D-GD (Theorem 1) is that the inequality only holds in expectation, and therefore may not verify  $(f, \varepsilon)$ -Byzantine resilience. Nevertheless, this result is consistent with the state-of-the-art convergence guarantees (Karimireddy et al., 2022) (Theorem II). Note that a subtle difference remains between our result and the one from (Karimireddy et al., 2022). The randomness of our result in Theorem 2 only depends on the random subsampling of data and the final choice of the model. These are natural sources of randomness that are usually considered when studying stochastic gradient descent (Bottou et al., 2018). On the other hand, the results from (Karimireddy et al., 2022) also incorporate an additional (exogenous) randomness introduced by the shuffling operation of Bucketing. This source of randomness cannot be canceled even if true gradients are computed by the workers, and it may amplify the uncertainty in the computations (see Section 6 and Appendix 10 for more details). As a result, we believe the probabilistic convergence guarantees from Theorem 2 to be strictly stronger than those obtained in (Karimireddy et al., 2022).

Once again, using Lemma 1, we can show that the tight (probabilistic) resilience guarantee implied by Theorem 2 holds for the larger class of  $(f, \kappa)$ -robust aggregation rules with  $\kappa = \mathcal{O}(1)$ . We formalize this in Corollary 2 below.

**Corollary 2.** *Let Assumption 1 hold and recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 3 with aggregation  $F \circ \text{NNM}$ , under the same setting as Theorem 2. If  $F$  is  $(f, \kappa)$ -robust with  $\kappa = \mathcal{O}(1)$ , then we have*

$$\mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \right] = \mathcal{O}(f/n G^2 + 1/\sqrt{T}),$$

where the expectation is over the algorithm's randomness.

## 6 EXPERIMENTAL EVALUATION

In this section, we investigate the practical performances of NNM. We report on a comprehensive set of experiments evaluating our solution against the state-of-the-art on three benchmark image classification tasks and under five different Byzantine attacks.

### 6.1 Experimental Setup

**Datasets, models, and hyperparameters.** We consider three image classification datasets, namely MNIST (LeCun and Cortes, 2010), Fashion-MNIST (Xiao et al., 2017), and CIFAR-10 (Krizhevsky et al., 2014); and we implement NNM on top of robust D-SHB. Due to space limitation, we only present results on MNIST and CIFAR-10, and defer the remaining results to Appendix 15, including experiments on D-GD with NNM.

On MNIST, we train a convolutional neural network (CNN) for  $T = 800$  steps using a batch size  $b = 25$ , with a decaying learning rate starting at 0.75, and a momentum parameter  $\beta = 0.9$ . On CIFAR-10, we use a CNN with  $b = 50$ ,  $T = 2000$ ,  $\gamma = 0.25$  that decays at step 1500, and  $\beta = 0.9$ . Furthermore, we implement our solution with four aggregation rules namely Krum, CWTM, CWMed, and GM<sup>5</sup>, and compare its performance against the Bucketing (Karimireddy et al., 2022) method and vanilla aggregation rules. As a benchmark, we also implement vanilla D-SHB (i.e., robust D-SHB with  $F = \text{average}$ ) in a setting where there are no faults ( $f = 0$ ).

**Heterogeneity.** We simulate heterogeneity in honest workers' data by sampling from the original dataset using a Dirichlet distribution of parameter  $\alpha$  (as done in (Hsu et al., 2019)). We consider three heterogeneity regimes: *extreme* ( $\alpha = 0.1$ ), *moderate* ( $\alpha = 1$ ), and *low* ( $\alpha = 10$ ). A pictorial representation of the resulting heterogeneity as a function of  $\alpha$  can be found in Appendix 14. We run our algorithm on MNIST and Fashion-MNIST over the whole spectrum of heterogeneity defined above. However, as CIFAR-10 is considerably more challenging, we restrict the heterogeneity to  $\alpha \in \{1, 10\}$ .

**Distributed system and Byzantine attacks.** We consider a distributed system of  $n = 17$  workers, among which

<sup>5</sup>We implement GM using the approximation from (Pillutla et al., 2022).

	Aggregation	ALIE	FOE	LF	SF	Mimic	Worst Case
①	Krum	10.19 ± 00.61	10.28 ± 00.57	12.76 ± 06.25	51.17 ± 06.14	86.78 ± 07.04	10.19 ± 00.61
	Bucketing + Krum	44.78 ± 06.18	36.10 ± 15.70	56.21 ± 19.49	47.54 ± 11.74	94.92 ± 03.93	36.10 ± 15.70
	NNM + Krum	<b>78.30 ± 07.78</b>	<b>70.07 ± 04.39</b>	<b>93.07 ± 05.24</b>	<b>82.44 ± 02.86</b>	<b>97.69 ± 00.77</b>	70.07 ± 04.39
②	GM	<b>92.01 ± 04.35</b>	65.61 ± 12.17	93.94 ± 03.70	57.86 ± 10.42	96.85 ± 01.57	57.86 ± 10.42
	Bucketing + GM	39.83 ± 11.35	44.73 ± 16.47	<b>96.22 ± 02.83</b>	<b>91.30 ± 03.91</b>	<b>97.68 ± 00.91</b>	39.83 ± 11.35
	NNM + GM	81.26 ± 08.91	<b>75.27 ± 02.69</b>	94.23 ± 03.20	86.33 ± 03.73	97.17 ± 01.09	75.27 ± 02.69
③	CWMed	68.74 ± 06.99	19.48 ± 10.97	33.34 ± 17.02	27.96 ± 09.97	64.01 ± 12.77	19.48 ± 10.97
	Bucketing + CWMed	55.86 ± 10.00	42.80 ± 21.25	70.16 ± 11.65	50.96 ± 16.52	94.43 ± 03.48	42.80 ± 21.25
	NNM + CWMed	<b>80.52 ± 07.45</b>	<b>75.20 ± 08.80</b>	<b>93.42 ± 02.98</b>	<b>85.10 ± 06.05</b>	<b>97.38 ± 00.70</b>	75.20 ± 08.80
④	CWTM	76.16 ± 07.68	69.96 ± 16.57	36.87 ± 21.43	27.45 ± 08.83	89.83 ± 02.83	27.45 ± 08.83
	Bucketing + CWTM	55.86 ± 10.00	42.80 ± 21.25	70.16 ± 11.65	50.96 ± 16.52	94.43 ± 03.48	42.80 ± 21.25
	NNM + CWTM	<b>79.04 ± 09.19</b>	<b>79.91 ± 03.94</b>	<b>94.75 ± 02.22</b>	<b>84.78 ± 05.78</b>	<b>96.02 ± 03.25</b>	79.04 ± 09.19

Table 2: Maximum test accuracies (%) across  $T = 800$  learning steps on MNIST, under extreme heterogeneity ( $\alpha = 0.1$ ) and five Byzantine attacks. There are  $f = 4$  Byzantine workers among  $n = 17$ . The baseline accuracy (D-SHB) is  $98.03 \pm 0.70\%$ . In each of the four horizontal blocks and under each attack, we highlight in **bold** the best accuracy. For every method, we also show the worst-case accuracy across attacks, and highlight the **best** and **worst** one in each block.

$f < n/2$  can be Byzantine. We vary  $f \in \{4, 6, 8\}$  on the MNIST dataset, and  $f \in \{2, 3, 4\}$  on CIFAR-10. The Byzantine workers execute five state-of-the-art gradient attacks, namely Fall of Empires (FOE) (Xie et al., 2019), A Little is Enough (ALIE) (Baruch et al., 2019), Sign Flipping (SF) (Allen-Zhu et al., 2020), Label Flipping (Allen-Zhu et al., 2020), and Mimic (Karimireddy et al., 2022). Note that for ALIE and FOE, we design optimized versions of the attacks, as done in (Shejwalkar and Houmansadr, 2021). We explain this further in Appendix 14.

**Reproducibility and reusability.** All experiments are run with five seeds from 1 to 5 for reproducibility purposes. The code will also be made available for reusability. Additional details on the setup can be found in Appendix 14.

## 6.2 Empirical Results on MNIST

In Table 2, we carefully examine the performance of NNM on MNIST under extreme heterogeneity ( $\alpha = 0.1$ ), in comparison with Bucketing and vanilla aggregation rules. For every block (i.e., every aggregation rule) and under every attack, we highlight in bold the algorithm resulting in the highest accuracy in the considered scenario.

**NNM improves robustness.** We clearly see that our algorithm boosts the resilience of aggregation rules in Byzantine settings, and provides the most consistent behavior across attacks. In fact, for Krum, CWMed, and CWTM (i.e., blocks 1, 3, and 4 in Table 2), NNM outputs the maximal accuracy under all attacks. The two other techniques (Bucketing and vanilla) showcase much weaker performances and are sometimes on par with a random classifier (e.g., 10.19% under

ALIE and 10.28% under FOE for vanilla Krum).

The case of GM is less evident since the other techniques can outperform our method in some settings (e.g., ALIE, SF). However, a crucial observation is that there always exists at least one attack that considerably deteriorates the performance of Bucketing+GM and vanilla GM, whereas NNM+GM consistently yields desirable accuracies in all attack scenarios. Indeed, the minimum accuracy across attacks achieved by NNM+GM is 75.27%, which is considerably better than Bucketing+GM and vanilla GM with 39.83% and 57.86% minimum accuracies, respectively. Moreover, although vanilla GM scores the highest under ALIE, it showcases low accuracies of 65.61% and 57.86% under FOE and SF, respectively. Bucketing+GM also fails considerably under ALIE and FOE with accuracies far below 50%. Finally, even though Bucketing+GM scores the highest under LF and Mimic, NNM+GM is also excellent under these two attacks with accuracies greater than 96%.

**Worst-case performance.** We show in the last column of Table 2 the *worst-case* performance (across attacks) of every aggregation technique (i.e., every row), and rank them within each block from worst (in red) to best (in green). We argue that this is a critical metric to correctly evaluate Byzantine resilience, as the same algorithm can simultaneously greatly defend against some attacks but perform poorly against others. Accordingly, we see from the last column that our method always displays the “best” worst case. In fact, the lowest accuracy NNM yields across attacks and aggregation rules is 70.07% with Krum. Furthermore, the worst case performances of the other techniques are much worse than ours, yielding values as low as 10.19%






Figure 1: Experiments on CIFAR-10 with  $f = 2$  Byzantine workers among  $n = 17$ . The Byzantine workers execute the ALIE (column 1) and LF (column 2) attacks. 1<sup>st</sup> row: moderate heterogeneity ( $\alpha = 1$ ). 2<sup>nd</sup> row: low heterogeneity ( $\alpha = 10$ ).

with Krum and 39.83% with Bucketing+GM.

Additionally, the worst case performance of Bucketing is almost equally poor for all aggregation rules (36.1%, 39.83%, and 42.80%). However, all worst case accuracies of NNM are within the range 70-79%. This highly suggests the unreliability of Bucketing, due to its subpar worst case behavior independently of the aggregation rule used.

### 6.3 Empirical Results on CIFAR-10

In Figure 1, we plot the performance of NNM and Bucketing on CIFAR-10 during  $T = 2000$  steps of learning, with GM, CWMed, CWTM, and Krum as base aggregation rules<sup>6</sup>. We consider two heterogeneity levels ( $\alpha = 1$  in row 1, and  $\alpha = 10$  in row 2), with two Byzantine workers executing the ALIE and LF attacks.

**Heterogeneity impairs learning.** Our first observation is that stronger heterogeneity regimes hinder the learning, as expected. As can be seen in Figure 1, increasing  $\alpha$  from 1 to 10 significantly improves all NNM aggregations (especially under ALIE), whereas Bucketing+CWMed still stagnates at 10% and Bucketing+GM barely reaches 30%.

**Empirical superiority of NNM.** Under ALIE, we clearly see that NNM greatly outperforms Bucketing in both heterogeneity regimes. In particular, ALIE deteriorates the learning when using Bucketing with GM and CWMed, with both methods reaching a final accuracy close to 10% when  $\alpha = 1$ . Although Bucketing+Krum has a better performance, it still displays a lower accuracy and a much larger variance than any NNM aggregation. On the other hand, we observe that LF is a much weaker attack. Although all

aggregation methods are able to converge to a desirable accuracy, our algorithm still portrays better convergence rates than Bucketing. This is particularly apparent for  $\alpha = 10$  where all NNM aggregations almost exactly match D-SHB’s convergence rate, whereas Bucketing converges slower, especially in the case of Krum. The remaining three attacks have a similar behavior to LF and are in Appendix 15.

## 7 CONCLUSION AND FUTURE WORK

We show that robust D-GD enhanced by NNM is an optimal solution for Byzantine ML under heterogeneity. We also derive similar results for robust D-SHB (with NNM) in expectation. We believe that an interesting future direction would be to tighten these guarantees by investigating almost sure convergence rates for our stochastic solution. Such guarantees have been recently introduced in the Byzantine-free setting for SHB (Sebbouh et al., 2021; Liu and Yuan, 2022; Gadat et al., 2018). Yet, adapting them to our setting remains an open question. In general, we believe subsequent works on Byzantine ML should strive to obtain almost sure convergence guarantees (instead of expectation), precluding the use of unnecessary randomness and certifying the robustness of these algorithms.

### Acknowledgements

This work has been supported in part by SNSF grants 200021\_200477 and 200021\_182542, and an EPFL-Ecocloud postdoctoral grant. The authors are thankful to the anonymous reviewers for their constructive comments.

<sup>6</sup>In the considered setting, Bucketing with CWMed and CWTM are exactly equivalent.

## References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. (2015). TensorFlow: Large-scale machine learning on heterogeneous distributed systems.
- Acharya, A., Hashemi, A., Jain, P., Sanghavi, S., Dhillon, I. S., and Topcu, U. (2022). Robust training in high dimensions via block coordinate geometric median descent. In Camps-Valls, G., Ruiz, F. J. R., and Valera, I., editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 11145–11168. PMLR.
- Allen-Zhu, Z., Ebrahimiaghazani, F., Li, J., and Alistarh, D. (2020). Byzantine-resilient non-convex stochastic gradient descent. In *International Conference on Learning Representations*.
- Alon, N., Matias, Y., and Szegedy, M. (1996). The space complexity of approximating the frequency moments. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 20–29.
- Baruch, M., Baruch, G., and Goldberg, Y. (2019). A little is enough: Circumventing defenses for distributed learning. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, 8-14 December 2019, Long Beach, CA, USA*.
- Bertsekas, D. and Tsitsiklis, J. (2015). *Parallel and distributed computation: numerical methods*. Athena Scientific.
- Blanchard, P., El Mhamdi, E. M., Guerraoui, R., and Stainer, J. (2017). Machine learning with adversaries: Byzantine tolerant gradient descent. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., editors, *Advances in Neural Information Processing Systems 30*, pages 119–129. Curran Associates, Inc.
- Bottou, L., Curtis, F. E., and Nocedal, J. (2018). Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311.
- Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.
- Charikar, M., Steinhardt, J., and Valiant, G. (2017). Learning from untrusted data. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 47–60.
- Chen, Y., Su, L., and Xu, J. (2017). Distributed statistical machine learning in adversarial settings: Byzantine gradient descent. *Proc. ACM Meas. Anal. Comput. Syst.*, 1(2).
- Du, S. S., Jin, C., Lee, J. D., Jordan, M. I., Singh, A., and Póczos, B. (2017). Gradient descent can take exponential time to escape saddle points. *Advances in neural information processing systems*, 30.
- El Mhamdi, E. M., Farhadkhani, S., Guerraoui, R., Guirguis, A., Hoang, L. N., and Rouault, S. (2021a). Collaborative learning in the jungle (decentralized, Byzantine, heterogeneous, asynchronous and nonconvex learning). In *Thirty-Fifth Conference on Neural Information Processing Systems*.
- El Mhamdi, E. M., Guerraoui, R., and Rouault, S. (2018). The hidden vulnerability of distributed learning in Byzantium. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3521–3530. PMLR.
- El Mhamdi, E. M., Guerraoui, R., and Rouault, S. L. A. (2021b). Distributed momentum for Byzantine-resilient stochastic gradient descent. In *9th International Conference on Learning Representations (ICLR)*.
- Farhadkhani, S., Guerraoui, R., Gupta, N., Hoang, L. N., Pinot, R., and Stephan, J. (2022a). Making Byzantine decentralized learning efficient. *arXiv preprint arXiv:2209.10931*.
- Farhadkhani, S., Guerraoui, R., Gupta, N., Pinot, R., and Stephan, J. (2022b). Byzantine machine learning made easy by resilient averaging of momentums. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S., editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 6246–6283. PMLR.
- Gadat, S., Panloup, F., and Saadane, S. (2018). Stochastic heavy ball. *Electronic Journal of Statistics*, 12(1):461–529.
- Ghadimi, S. and Lan, G. (2016). Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Mathematical Programming*, 156(1):59–99.
- Gupta, N. and Vaidya, N. H. (2020). Fault-tolerance in distributed optimization: The case of redundancy. In

- Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 365–374.
- Hajebi, K., Abbasi-Yadkori, Y., Shahbazi, H., and Zhang, H. (2011). Fast approximate nearest-neighbor search with k-nearest neighbor graph. In *Twenty-Second International Joint Conference on Artificial Intelligence*.
- Hsu, T.-M. H., Qi, H., and Brown, M. (2019). Measuring the effects of non-identical data distribution for federated visual classification.
- Jerrum, M. R., Valiant, L. G., and Vazirani, V. V. (1986). Random generation of combinatorial structures from a uniform distribution. *Theoretical computer science*, 43:169–188.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D’Oliveira, R. G. L., Eichner, H., Rouayheb, S. E., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., Gruteser, M., Harchaoui, Z., He, C., He, L., Huo, Z., Hutchinson, B., Hsu, J., Jaggi, M., Javidi, T., Joshi, G., Khodak, M., Konečný, J., Korolova, A., Koushanfar, F., Koyejo, S., Lepoint, T., Liu, Y., Mittal, P., Mohri, M., Nock, R., Özgür, A., Pagh, R., Qi, H., Ramage, D., Raskar, R., Raykova, M., Song, D., Song, W., Stich, S. U., Sun, Z., Suresh, A. T., Tramèr, F., Vepakomma, P., Wang, J., Xiong, L., Xu, Z., Yang, Q., Yu, F. X., Yu, H., and Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210.
- Karimireddy, S. P., He, L., and Jaggi, M. (2021). Learning from history for Byzantine robust optimization. *International Conference On Machine Learning, Vol 139*, 139.
- Karimireddy, S. P., He, L., and Jaggi, M. (2022). Byzantine-robust learning on heterogeneous datasets via bucketing. In *International Conference on Learning Representations*.
- Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S., Stich, S., and Suresh, A. T. (2020). Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pages 5132–5143. PMLR.
- Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., and Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*.
- Krizhevsky, A., Nair, V., and Hinton, G. (2014). The CIFAR-10 dataset. *online: <http://www.cs.toronto.edu/kriz/cifar.html>*, 55(5).
- Lamport, L., Shostak, R., and Pease, M. (1982). The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401.
- LeCun, Y. and Cortes, C. (2010). MNIST handwritten digit database.
- Li, L., Xu, W., Chen, T., Giannakis, G. B., and Ling, Q. (2019). RSA: Byzantine-robust stochastic aggregation methods for distributed learning from heterogeneous datasets. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI’19/IAAI’19/EAAI’19*. AAAI Press.
- Liu, J. and Yuan, Y. (2022). On almost sure convergence rates of stochastic gradient methods. In Loh, P.-L. and Raginsky, M., editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 2963–2983. PMLR.
- Liu, S., Gupta, N., and Vaidya, N. H. (2021). Approximate Byzantine fault-tolerance in distributed optimization. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing, PODC’21*, page 379–389, New York, NY, USA. Association for Computing Machinery.
- Muja, M. and Lowe, D. G. (2014). Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on pattern analysis and machine intelligence*, 36(11):2227–2240.
- Nemirovsky, A., dın, D., and Yudin, D. (1983). *Problem Complexity and Method Efficiency in Optimization*. A Wiley-Interscience publication. Wiley.
- Nesterov, Y. et al. (2018). *Lectures on convex optimization*, volume 137. Springer.
- Pillutla, K., Kakade, S. M., and Harchaoui, Z. (2022). Robust aggregation for federated learning. *IEEE Transactions on Signal Processing*, 70:1142–1154.
- Polyak, B. T. (1964). Some methods of speeding up the convergence of iteration methods. *USSR computational mathematics and mathematical physics*, 4(5):1–17.
- Rousseeuw, P. J. (1985). Multivariate estimation with high breakdown point. *Mathematical statistics and applications*, 8(37):283–297.
- Sebbouh, O., Gower, R. M., and Defazio, A. (2021). Almost sure convergence rates for stochastic gradient descent and stochastic heavy ball. In *Conference on Learning Theory*, pages 3935–3971. PMLR.

- Shejwalkar, V. and Houmansadr, A. (2021). Manipulating the Byzantine: Optimizing model poisoning attacks and defenses for federated learning. In *NDSS*.
- Small, C. G. (1990). A survey of multidimensional medians. *International Statistical Review/Revue Internationale de Statistique*, pages 263–277.
- Tu, J., Liu, W., Mao, X., and Chen, X. (2021). Variance reduced median-of-means estimator for Byzantine-robust distributed inference. *J. Mach. Learn. Res.*, 22:84–1.
- Xiao, H., Rasul, K., and Vollgraf, R. (2017). Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Xie, C., Koyejo, O., and Gupta, I. (2018). Generalized Byzantine-tolerant sgd.
- Xie, C., Koyejo, O., and Gupta, I. (2019). Fall of empires: Breaking Byzantine-tolerant SGD by inner product manipulation. In *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI 2019, Tel Aviv, Israel, July 22-25, 2019*, page 83.
- Yin, D., Chen, Y., Kannan, R., and Bartlett, P. (2018). Byzantine-robust distributed learning: Towards optimal statistical rates. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5650–5659. PMLR.

---

# Supplementary Materials:

## Fixing by Mixing: A Recipe for Optimal Byzantine ML under Heterogeneity

---

### ORGANIZATION OF THE APPENDIX

Appendix 8 contains proofs related to  $(f, \kappa)$ -robustness. Appendix 9 contains the analysis of NNM (proof of Lemma 1). Appendix 10 contains a formal discussion on Bucketing (Karimireddy et al., 2022). Appendix 11 contains the analysis of robust D-GD (proof of Theorem 1). Appendix 12 contains the analysis of robust D-GD with NNM (proof of Corollary 1), and the Byzantine resilience lower bound (Proposition 1). Appendix 13 contains the analysis of robust D-SHB with NNM (proofs of Theorem 2 and Corollary 2). Appendix 14 contains details on the experimental setup. Appendix 15 contains all experimental results.

### 8 ROBUSTNESS ANALYSIS

In this section, we prove all our claims related to  $(f, \kappa)$ -robustness. In Section 8.1, we prove several existing aggregation rules to be  $(f, \kappa)$ -robust and give their exact robustness coefficients. We then establish the tightness of our analysis in Section 8.2 by proving a universal lower bound on  $\kappa$ , and an aggregation-specific lower bound. Finally, we prove that  $(f, \kappa)$ -robustness unifies existing robustness definitions in Section 8.3.

We first recall the definition of  $(f, \kappa)$ -robustness:

**Definition 2.** Let  $f < n/2$  and  $\kappa \geq 0$ . An aggregation rule  $F$  is said to be  $(f, \kappa)$ -robust if for any vectors  $x_1, \dots, x_n \in \mathbb{R}^d$ , and any set  $S \subseteq [n]$  of size  $n - f$ ,

$$\|F(x_1, \dots, x_n) - \bar{x}_S\|^2 \leq \frac{\kappa}{|S|} \sum_{i \in S} \|x_i - \bar{x}_S\|^2$$

where  $\bar{x}_S = \frac{1}{|S|} \sum_{i \in S} x_i$ . We refer to  $\kappa$  as the robustness coefficient.

#### 8.1 $(f, \kappa)$ -robust Aggregation Rules

In this section, we prove that coordinate-wise Trimmed Mean (CWTM) (Yin et al., 2018), Krum (Blanchard et al., 2017), Geometric Median (GM) (Small, 1990), and coordinate-wise Median (CWMed) (Yin et al., 2018) all satisfy  $(f, \kappa)$ -robustness.

##### 8.1.1 Trimmed Mean

Let  $x \in \mathbb{R}^d$ , we denote by  $[x]_k$ , the  $k$ -th coordinate of  $x$ . Given the input vectors  $x_1, \dots, x_n \in \mathbb{R}^d$ , we let  $\tau_k$  denote a permutation on  $[n]$  that sorts the  $k$ -th coordinate of the input vectors in non-decreasing order, i.e.,  $[x_{\tau_k(1)}]_k \leq [x_{\tau_k(2)}]_k \leq \dots \leq [x_{\tau_k(n)}]_k$ . Then, the coordinate-wise trimmed mean of  $x_1, \dots, x_n$ , denoted by  $\text{CWTM}(x_1, \dots, x_n)$ , is a vector in  $\mathbb{R}^d$  whose  $k$ -th coordinate is defined as follows,

$$[\text{CWTM}(x_1, \dots, x_n)]_k := \frac{1}{n - 2f} \sum_{j=f+1}^{n-f} [x_{\tau_k(j)}]_k.$$

We first show a general lemma simplifying the analysis of  $(f, \kappa)$ -robustness for coordinate-wise aggregations, by reducing the analysis to scalars without loss of generality. Specifically, we show that if  $F$  is a coordinate-wise function, i.e., each  $k$ -th coordinate of  $F$  denoted by  $F_k$  only depends on the respective  $k$ -th coordinates of the inputs  $[x_1]_k, \dots, [x_n]_k$ , then coordinate-wise robustness implies overall robustness.

**Lemma 2.** Assume that  $F : \mathbb{R}^{d \times n} \rightarrow \mathbb{R}^d$  is a coordinate-wise aggregation function, i.e., there exist  $d$  real-valued functions  $F_1, \dots, F_d : \mathbb{R}^n \rightarrow \mathbb{R}$  such that for all  $x_1, \dots, x_n \in \mathbb{R}^d$ ,  $[F(x_1, \dots, x_n)]_k = F_k([x_1]_k, \dots, [x_n]_k)$ . If for each  $k \in [d]$ ,  $F_k$  is  $(f, \kappa)$ -robust then  $F$  is  $(f, \kappa)$ -robust.

*Proof.* Consider  $n$  arbitrary vectors in  $\mathbb{R}^d$ ,  $x_1, \dots, x_n$ , and an arbitrary set  $S \subseteq [n]$  such that  $|S| = n - f$ . Assume that for each  $k \in [d]$ ,  $F_k$  is  $(f, \kappa)$ -robust. Since  $F$  is assume to be a coordinate-wise aggregation function, we have

$$\|F(x_1, \dots, x_n) - \bar{x}_S\|^2 = \sum_{k=1}^d |F_k([x_1]_k, \dots, [x_n]_k) - [\bar{x}_S]_k|^2. \quad (4)$$

Since for each  $k \in [d]$ ,  $F_k$  is  $(f, \kappa)$ -robust, we have

$$\sum_{k=1}^d |F_k([x_1]_k, \dots, [x_n]_k) - [\bar{x}_S]_k|^2 \leq \sum_{k=1}^d \frac{\kappa}{|S|} \sum_{i \in S} |[x_i]_k - [\bar{x}_S]_k|^2 = \frac{\kappa}{|S|} \sum_{i \in S} \sum_{k=1}^d |[x_i]_k - [\bar{x}_S]_k|^2 = \frac{\kappa}{|S|} \sum_{i \in S} \|x_i - \bar{x}_S\|^2.$$

Substituting from above in (4) concludes the proof.  $\square$

We now show an important property in Lemma 3 below on the sorting of real values that proves essential in obtaining a tight  $(f, \kappa)$ -robustness guarantee for trimmed mean.

**Lemma 3.** Consider  $n$  real values  $x_1, \dots, x_n$  be such that  $x_1 \leq \dots \leq x_n$ . Let  $S \subseteq [n]$  of size  $|S| = n - f$ , and  $I := \{f + 1, \dots, n - f\}$ . We obtain that

$$\sum_{i \in I \setminus S} |x_i - \bar{x}_S|^2 \leq \sum_{i \in S \setminus I} |x_i - \bar{x}_S|^2.$$

*Proof.* First, note that, as  $|I| = n - 2f$  and  $|S| = n - f$ ,  $|I \setminus S| \leq f$  and  $|S \setminus I| \geq f$ . Thus,  $|I \setminus S| \leq |S \setminus I|$ . To prove the lemma, we show that there exists an injection  $\phi : I \setminus S \rightarrow S \setminus I$  such that

$$\forall i \in I \setminus S, |x_i - \bar{x}_S| \leq |x_{\phi(i)} - \bar{x}_S|. \quad (5)$$

As  $|I \setminus S| \leq |S \setminus I|$ , we have  $\sum_{i \in I \setminus S} |x_{\phi(i)} - \bar{x}_S|^2 \leq \sum_{i \in S \setminus I} |x_i - \bar{x}_S|^2$ . Hence, (5) proves the lemma.

We denote by  $B$  the complement of  $S$  in  $[n]$ , i.e.,  $B = [n] \setminus S$ . Therefore,  $|B| = f$  and  $I \setminus S = I \cap B$ . We denote  $I^+ := \{n - f + 1, \dots, n\}$ , and  $I^- := \{1, \dots, f\}$  to be the indices of values that are larger (or equal to) and smaller (or equal to) the values in  $I$ , respectively. Let  $|I \cap B| = q$ . As  $B = [n] \setminus S$ , we obtain that

$$|I^+ \cap S| = f - |I^+ \cap B| \geq f - |B \setminus (I \cap B)| = f - (f - q) = q. \quad (6)$$

Similarly, we obtain that

$$|I^- \cap S| \geq q. \quad (7)$$

Recall that  $q = |I \cap B|$  where  $I \setminus S = I \cap B$ . Therefore, due to (6) and (7), there exists an injection from  $I \setminus S$  to  $(I^- \cap S) \times (I^+ \cap S)$ . Let  $\psi$  be such an injection. For each  $i \in I \setminus S$ ,  $\psi(i)$  is a pair, denoted by  $(\psi^-(i), \psi^+(i))$ , in  $(I^- \cap S) \times (I^+ \cap S)$ . Consider an arbitrary  $i \in I \setminus S$ . By definition of  $I^-$  and  $I^+$ , we have

$$x_{\psi^-(i)} \leq x_i \leq x_{\psi^+(i)}.$$

Therefore, for any real value  $y$ ,

$$|x_i - y| \leq \max\{|x_{\psi^+(i)} - y|, |x_{\psi^-(i)} - y|\}.$$

The above proves (5), where the injection  $\phi$  is defined as  $\phi(i) = \arg \max_{j \in \{\psi^+(i), \psi^-(i)\}} |x_j - \bar{x}_S|$  for all  $i \in I \setminus S$ .  $\square$

We now prove the  $(f, \kappa)$ -robustness property of trimmed mean in the proposition below.

**Proposition 2.** Let  $n \in \mathbb{N}^*$  and  $f < n/2$ . CWTM is  $(f, \kappa)$ -robust with  $\kappa = \frac{6f}{n-2f} \left(1 + \frac{f}{n-2f}\right)$ .

*Proof.* First, note that trimmed mean Trimmed Mean is a coordinate-wise aggregation, defined in Lemma 2. Thus, due to Lemma 2, it suffices to show that Trimmed Mean is  $(f, \kappa)$ -robust in the scalar domain, i.e., when  $d = 1$ .

Let  $x_1, \dots, x_n \in \mathbb{R}$  and, without loss of generality, let us assume that  $x_1 \leq \dots \leq x_n$ . We denote by  $I := \{f+1, \dots, n-f\}$ , and let  $S$  be an arbitrary subset of  $[n]$  of size  $n - f$ . Recall that the set of indices selected by TM  $(x_1, \dots, x_n) = \frac{1}{n-2f} \sum_{i \in I} x_i$ . We obtain that

$$\begin{aligned} |\text{CWTM}(x_1, \dots, x_n) - \bar{x}_S|^2 &= \left| \frac{1}{n-2f} \sum_{i \in I} x_i - \bar{x}_S \right|^2 = \left| \frac{1}{n-2f} \sum_{i \in I} (x_i - \bar{x}_S) \right|^2 \\ &= \left| \frac{1}{n-2f} \sum_{i \in I} (x_i - \bar{x}_S) - \frac{1}{n-2f} \sum_{i \in S} (x_i - \bar{x}_S) \right|^2 \quad \left( \text{as } \sum_{i \in S} (x_i - \bar{x}_S) = 0 \right) \\ &= \frac{1}{(n-2f)^2} \left| \sum_{i \in I \setminus S} (x_i - \bar{x}_S) - \sum_{i \in S \setminus I} (x_i - \bar{x}_S) \right|^2. \end{aligned}$$

Using Jensen's inequality above, we obtain that

$$|\text{CWTM}(x_1, \dots, x_n) - \bar{x}_S|^2 \leq \frac{|I \setminus S| + |S \setminus I|}{(n-2f)^2} \left( \sum_{i \in I \setminus S} |x_i - \bar{x}_S|^2 + \sum_{i \in S \setminus I} |x_i - \bar{x}_S|^2 \right).$$

Note that  $|I \setminus S| = |I \cup S| - |S| \leq n - (n - f) = f$ . Similarly,  $|S \setminus I| = |I \cup S| - |I| \leq n - (n - 2f) = 2f$ . Accordingly,  $|I \setminus S| + |S \setminus I| \leq 3f$ . Therefore, we have

$$|\text{CWTM}(x_1, \dots, x_n) - \bar{x}_S|^2 \leq \frac{3f}{(n-2f)^2} \left( \sum_{i \in I \setminus S} |x_i - \bar{x}_S|^2 + \sum_{i \in S \setminus I} |x_i - \bar{x}_S|^2 \right).$$

Recall from Lemma 3 that  $\sum_{i \in I \setminus S} |x_i - \bar{x}_S|^2 \leq \sum_{i \in S \setminus I} |x_i - \bar{x}_S|^2$ . Using this fact above we obtain that

$$|\text{CWTM}(x_1, \dots, x_n) - \bar{x}_S|^2 \leq \frac{6f}{(n-2f)^2} \sum_{i \in S \setminus I} |x_i - \bar{x}_S|^2 \leq \frac{6f}{(n-2f)^2} \sum_{i \in S} |x_i - \bar{x}_S|^2.$$

Finally, as  $\frac{f}{(n-2f)^2} = \frac{f(n-f)}{(n-2f)^2} \left( \frac{1}{n-f} \right) = \left( \frac{f}{n-2f} + \left( \frac{f}{n-2f} \right)^2 \right) \frac{1}{n-f}$ , we obtain that

$$|\text{CWTM}(x_1, \dots, x_n) - \bar{x}_S|^2 \leq 6 \left( \frac{f}{n-2f} + \left( \frac{f}{n-2f} \right)^2 \right) \frac{1}{|S|} \sum_{i \in S} |x_i - \bar{x}_S|^2. \quad (8)$$

The above concludes the proof.  $\square$

### 8.1.2 Krum

In this section, we study a slight adaptation of the Krum algorithm first introduced in (Blanchard et al., 2017). Essentially, given the input vectors  $x_1, \dots, x_n$ , Krum outputs the vector that is the nearest to its neighbors upon discarding  $f$  (as opposed to  $f + 1$  in the original version) furthest vectors. Specifically, we denote by  $\mathcal{N}_j$  the set of indices of the  $n - f$  nearest neighbors of  $x_j$  in  $\{x_1, \dots, x_n\}$ , with ties arbitrarily broken. Krum outputs the vector  $x_{k^*}$  such that

$$k^* \in \arg \min_{j \in [n]} \sum_{i \in \mathcal{N}_j} \|x_j - x_i\|^2,$$

with ties arbitrarily broken if the set of minimizers above includes more than one element.

**Proposition 3.** Let  $n \in \mathbb{N}^*$  and  $f < n/2$ . Krum is  $(f, \kappa)$ -robust with  $\kappa = 6(1 + \frac{f}{n-2f})$ .

*Proof.* Let  $n \in \mathbb{N}$ ,  $f < n/2$ , and  $x_1, \dots, x_n \in \mathbb{R}^d$ . Consider any subset  $S \subseteq [n]$  of size  $|S| = n - f$ . In the following, for every  $j \in [n]$ , we denote by  $\mathcal{N}_j$  the set the of indices of the  $n - f$  nearest neighbors of  $x_j$  in  $\{x_1, \dots, x_n\}$ , with ties arbitrarily broken. Observe that this implies, for every  $j \in [n]$ , that

$$\sum_{i \in \mathcal{N}_j} \|x_j - x_i\|^2 \leq \sum_{i \in S} \|x_j - x_i\|^2. \quad (9)$$

Let  $k^* \in [n]$  be the index selected by Krum. By definition, it holds that  $k^* \in \arg \min_{j \in [n]} \sum_{i \in \mathcal{N}_j} \|x_j - x_i\|^2$ . Therefore, leveraging (9) we have

$$\begin{aligned} \sum_{i \in \mathcal{N}_{k^*}} \|x_{k^*} - x_i\|^2 &= \min_{j \in [n]} \sum_{i \in \mathcal{N}_j} \|x_j - x_i\|^2 \leq \min_{j \in S} \sum_{i \in \mathcal{N}_j} \|x_j - x_i\|^2 \leq \frac{1}{|S|} \sum_{j \in S} \sum_{i \in \mathcal{N}_j} \|x_j - x_i\|^2 \\ &\leq \frac{1}{|S|} \sum_{j \in S} \sum_{i \in S} \|x_j - x_i\|^2 = \frac{1}{|S|} \sum_{i, j \in S} \|x_j - \bar{x}_S - (x_i - \bar{x}_S)\|^2 \\ &= \frac{1}{|S|} \sum_{i, j \in S} \left( \|x_j - \bar{x}_S\|^2 + \|x_i - \bar{x}_S\|^2 - 2 \langle x_j - \bar{x}_S, x_i - \bar{x}_S \rangle \right) \\ &= \frac{1}{|S|} \left[ \sum_{i, j \in S} \|x_j - \bar{x}_S\|^2 + \sum_{i, j \in S} \|x_i - \bar{x}_S\|^2 - 2 \sum_{i, j \in S} \langle x_j - \bar{x}_S, x_i - \bar{x}_S \rangle \right] \\ &= \frac{1}{|S|} \left[ 2|S| \sum_{i \in S} \|x_i - \bar{x}_S\|^2 - 2 \sum_{i, j \in S} \langle x_j - \bar{x}_S, x_i - \bar{x}_S \rangle \right] \\ &= \frac{1}{|S|} \left[ 2|S| \sum_{i \in S} \|x_i - \bar{x}_S\|^2 - 2 \underbrace{\sum_{i \in S} \left\langle \sum_{j \in S} (x_j - \bar{x}_S), x_i - \bar{x}_S \right\rangle}_{=0} \right] = 2 \sum_{i \in S} \|x_i - \bar{x}_S\|^2. \quad (10) \end{aligned}$$

Now, using Jensen's inequality, we can write for all  $i \in S$ ,

$$\|x_{k^*} - \bar{x}_S\|^2 \leq 2 \|x_{k^*} - x_i\|^2 + 2 \|x_i - \bar{x}_S\|^2.$$

Therefore, by rearranging the terms, we have for all  $i \in S$ ,

$$\|x_{k^*} - x_i\|^2 \geq \frac{1}{2} \|x_{k^*} - \bar{x}_S\|^2 - \|x_i - \bar{x}_S\|^2.$$

Together with the fact that  $|S \cap \mathcal{N}_{k^*}| = |S| + |\mathcal{N}_{k^*}| - |S \cup \mathcal{N}_{k^*}| \geq (n - f) + (n - f) - n = n - 2f$ , the previous inequality implies that

$$\begin{aligned} \sum_{i \in \mathcal{N}_{k^*}} \|x_{k^*} - x_i\|^2 &\geq \sum_{i \in S \cap \mathcal{N}_{k^*}} \|x_{k^*} - x_i\|^2 \geq |S \cap \mathcal{N}_{k^*}| \frac{1}{2} \|x_{k^*} - \bar{x}_S\|^2 - \sum_{i \in S \cap \mathcal{N}_{k^*}} \|x_i - \bar{x}_S\|^2 \\ &\geq \frac{n - 2f}{2} \|x_{k^*} - \bar{x}_S\|^2 - \sum_{i \in S \cap \mathcal{N}_{k^*}} \|x_i - \bar{x}_S\|^2. \end{aligned}$$



By rearranging the terms, and invoking (10) we can write

$$\begin{aligned}
 \|x_{k^*} - \bar{x}_S\|^2 &\leq \frac{2}{n-2f} \left[ \sum_{i \in \mathcal{N}_{k^*}} \|x_{k^*} - x_i\|^2 + \sum_{i \in S \cap \mathcal{N}_{k^*}} \|x_i - \bar{x}_S\|^2 \right] \\
 &\leq \frac{2}{n-2f} \left[ \sum_{i \in \mathcal{N}_{k^*}} \|x_{k^*} - x_i\|^2 + \sum_{i \in S} \|x_i - \bar{x}_S\|^2 \right] \\
 &\leq \frac{2}{n-2f} \left[ 2 \sum_{i \in S} \|x_i - \bar{x}_S\|^2 + \sum_{i \in S} \|x_i - \bar{x}_S\|^2 \right] \\
 &= \frac{6}{n-2f} \sum_{i \in S} \|x_i - \bar{x}_S\|^2 = \frac{6(n-f)}{n-2f} \frac{1}{|S|} \sum_{i \in S} \|x_i - \bar{x}_S\|^2.
 \end{aligned}$$

We conclude by remarking that  $\frac{n-f}{n-2f} = 1 + \frac{f}{n-2f}$ .  $\square$

### 8.1.3 Geometric Median

The Geometric Median of  $x_1, \dots, x_n \in \mathbb{R}^d$  denoted by  $\text{GM}(x_1, \dots, x_n)$ , is defined to be a vector that minimizes the sum of the  $\ell_2$ -distances to these vectors. Specifically, we have

$$\text{GM}(x_1, \dots, x_n) \in \arg \min_{y \in \mathbb{R}^d} \sum_{i=1}^n \|y - x_i\|.$$

**Proposition 4.** *Let  $n \in \mathbb{N}^*$  and  $f < n/2$ . GM is  $(f, \kappa)$ -robust with  $\kappa = 4 \left(1 + \frac{f}{n-2f}\right)^2$ .*

*Proof.* The proof is similar to the analysis of Geometric Median in (Karimireddy et al., 2022). For completeness, we provide the full proof adapted to the definition of  $(f, \kappa)$ -robustness. Let us denote by  $x^* := \text{GM}(x_1, \dots, x_n)$ , the geometric median of the input vectors. Consider any subset  $S \subseteq [n]$  of size  $|S| = n - f$ . By the reverse triangle inequality, for any  $i \in S$ , we have

$$\|x^* - x_i\| \geq \|x^* - \bar{x}_S\| - \|x_i - \bar{x}_S\|. \quad (11)$$

Similarly, for any  $i \in [n] \setminus S$ , we obtain

$$\|x^* - x_i\| \geq \|x_i - \bar{x}_S\| - \|x^* - \bar{x}_S\|. \quad (12)$$

Summing up (11) and (12) over all input vectors we obtain

$$\sum_{i \in [n]} \|x^* - x_i\| \geq (n-2f) \|x^* - \bar{x}_S\| + \sum_{i \in [n] \setminus S} \|x_i - \bar{x}_S\| - \sum_{i \in S} \|x_i - \bar{x}_S\|.$$

Rearranging the terms, we obtain

$$\|x^* - \bar{x}_S\| \leq \frac{1}{n-2f} \left( \sum_{i \in [n]} \|x^* - x_i\| - \sum_{i \in [n] \setminus S} \|x_i - \bar{x}_S\| + \sum_{i \in S} \|x_i - \bar{x}_S\| \right)$$

Note that by the definition of the geometric median, we have

$$\sum_{i \in [n]} \|x^* - x_i\| \leq \sum_{i \in [n]} \|\bar{x}_S - x_i\|.$$

Therefore,

$$\|x^* - \bar{x}_S\| \leq \frac{1}{n-2f} \left( \sum_{i \in [n]} \|x_i - \bar{x}_S\| - \sum_{i \in [n] \setminus S} \|x_i - \bar{x}_S\| + \sum_{i \in S} \|x_i - \bar{x}_S\| \right) \leq \frac{2}{n-2f} \sum_{i \in S} \|x_i - \bar{x}_S\|.$$

Squaring both sides and using Jensen's inequality, we obtain

$$\begin{aligned} \|x^* - \bar{x}_S\|^2 &\leq \frac{4(n-f)}{(n-2f)^2} \sum_{i \in S} \|x_i - \bar{x}_S\|^2 = \frac{4(n-f)^2}{(n-2f)^2} \cdot \frac{1}{n-f} \sum_{i \in S} \|x_i - \bar{x}_S\|^2 \\ &= 4 \left(1 + \frac{f}{n-2f}\right)^2 \frac{1}{|S|} \sum_{i \in S} \|x_i - \bar{x}_S\|^2. \end{aligned}$$

This is the desired result.  $\square$

#### 8.1.4 Median

For input vectors  $x_1, \dots, x_n$ , their coordinate-wise median, denoted by  $\text{CWMed}(x_1, \dots, x_n)$ , is defined to be a vector whose  $k$ -th coordinate, for all  $k \in [d]$ , is defined to be

$$[\text{CWMed}(x_1, \dots, x_n)]_k := \text{Median}([x_1]_k, \dots, [x_n]_k). \quad (13)$$

**Proposition 5.** *Let  $n \in \mathbb{N}^*$  and  $f < n/2$ .  $\text{CWMed}$  is  $(f, \kappa)$ -robust with  $\kappa = 4 \left(1 + \frac{f}{n-2f}\right)^2$ .*

*Proof.* Given that Geometric Median coincides with Median for one-dimensional inputs, it follows from Proposition 4 that, for one-dimensional inputs, Median is  $(f, \kappa)$ -robust with  $\kappa = 4 \left(1 + \frac{f}{n-2f}\right)^2$ . Since Median is coordinate-wise, we deduce from Lemma 2 that the latter holds for any  $d$ -dimensional inputs, and conclude the proof.  $\square$

## 8.2 Lower Bounds

In this section, we establish the tightness of our analysis by proving a universal lower bound on  $\kappa$  in Proposition 6, and an aggregation-specific lower bound in Proposition 7.

**Proposition 6.** *Let  $n \in \mathbb{N}^*$ ,  $f < n$  and  $\kappa > 0$ . If  $F$  is  $(f, \kappa)$ -robust, then  $n > 2f$  and  $\kappa \geq \frac{f}{n-2f}$ .*

*Proof.* Let  $n \in \mathbb{N}^*$ ,  $f < n$  and  $\kappa > 0$ . Assume that  $F$  is  $(f, \kappa)$ -resilient averaging aggregation rule. Consider  $x_1, \dots, x_n$  such that  $x_1 = \dots = x_{n-f} = 0$ , and  $x_{n-f+1} = \dots = x_n = 1$ . Let us first consider a set  $S_0 = \{1, \dots, n-f\}$ . Since  $|S_0| = n-f$ , by definition, we have

$$|F(x_1, \dots, x_n) - \bar{x}_{S_0}|^2 \leq \kappa \frac{1}{n-f} \sum_{i \in S_0} |x_i - \bar{x}_{S_0}|^2 = 0.$$

Thus,  $F(x_1, \dots, x_n) = \bar{x}_{S_0} = 0$ .

Now, consider another set  $S_1 = \{f+1, \dots, n\}$ . Observe that we necessarily have  $n > 2f$ . Assume by contradiction that  $2f \geq n$ , i.e.,  $f+1 \geq n-f+1$ . This implies that, for every  $i \in S_1$ ,  $x_i = 1$ . Therefore,  $\bar{x}_{S_1} = 1$ . And, since  $F$  is  $(f, \kappa)$ -robust, we must have

$$|F(x_1, \dots, x_n) - \bar{x}_{S_1}|^2 \leq \kappa \frac{1}{n-f} \sum_{i \in S_1} |x_i - \bar{x}_{S_1}|^2.$$

Therefore,  $0 = \bar{x}_{S_0} = F(x_1, \dots, x_n) = \bar{x}_{S_1} = 1$ , which is a contradiction.

As a result, we must have  $n > 2f$ . This implies that  $\bar{x}_{S_1} = \frac{f}{n-f}$ . Thus,

$$|F(x_1, \dots, x_n) - \bar{x}_{S_1}|^2 = \left(\frac{f}{n-f}\right)^2. \quad (14)$$

Since  $F$  is  $(f, \kappa)$ -resilient averaging rule, we have

$$\begin{aligned} |F(x_1, \dots, x_n) - \bar{x}_{S_1}|^2 &\leq \kappa \frac{1}{n-f} \sum_{i \in S_1} |x_i - \bar{x}_{S_1}|^2 = \frac{\kappa}{n-f} \left( (n-2f) \left|0 - \frac{f}{n-f}\right|^2 + f \left|1 - \frac{f}{n-f}\right|^2 \right) \\ &= \frac{\kappa}{n-f} \left( (n-2f) \frac{f^2}{(n-f)^2} + f \frac{(n-2f)^2}{(n-f)^2} \right) \\ &= \frac{\kappa f(n-2f)}{n-f} \left( \frac{f+n-2f}{(n-f)^2} \right) = \kappa \frac{f(n-2f)}{(n-f)^2}. \end{aligned}$$

Plugging this inequality back in (14), we conclude  $\kappa \geq \frac{f}{n-2f}$ .  $\square$

**Proposition 7.** *Let  $n \in \mathbb{N}^*$ ,  $f \geq 1$  such that  $n > 2f$ , and  $\kappa \geq 0$ . For any  $F \in \{\text{GM}, \text{CWMed}, \text{Krum}\}$ , if  $F$  is  $(f, \kappa)$ -robust then  $\kappa \in \Omega(1)$ .*

*Proof.* Let  $n \in \mathbb{N}^*$ ,  $f \geq 1$  such that  $n > 2f$  and  $\kappa \geq 0$ . Let  $F \in \{\text{GM}, \text{CWMed}, \text{Krum}\}$ . We will prove that if  $F$  is  $(f, \kappa)$ -robust, then there exists an *absolute* constant  $c > 0$  such that  $\kappa \geq c$ .

Consider  $x_1, \dots, x_n \in \mathbb{R}$  such that  $x_1, \dots, x_{\lfloor \frac{n-f}{2} \rfloor} = -1$  and  $x_{\lfloor \frac{n-f}{2} \rfloor + 1}, \dots, x_{n-f} = 1$ . Moreover, we set  $x_{n-f+1}, \dots, x_n = 1$ . Observe that, because  $f \geq 1$ , there is a strict majority of the inputs taking value 1: the number of such inputs is  $(n-f - \lfloor \frac{n-f}{2} \rfloor) + f = \lceil \frac{n-f}{2} \rceil + f \geq \lfloor \frac{n-f}{2} \rfloor + f > \lfloor \frac{n-f}{2} \rfloor$ .

Besides, recall that by setting  $S = \{1, \dots, n-f\}$ , any  $(f, \kappa)$ -robust function  $F$  should verify

$$|F(x_1, \dots, x_n) - \bar{x}_S|^2 \leq \kappa \frac{1}{n-f} \sum_{i \in S} |x_i - \bar{x}_S|^2. \quad (15)$$

However, the average  $\bar{x}_S$  is equal to

$$\begin{aligned} \bar{x}_S &= \frac{1}{n-f} \sum_{i \in S} x_i = \frac{1}{n-f} \left[ \left\lfloor \frac{n-f}{2} \right\rfloor \times (-1) + \left( n-f - \left\lfloor \frac{n-f}{2} \right\rfloor \right) \times (+1) \right] \\ &= \frac{1}{n-f} \left( n-f - 2 \left\lfloor \frac{n-f}{2} \right\rfloor \right) = \begin{cases} 0 & \text{if } n-f \text{ is even,} \\ \frac{1}{n-f} & \text{if } n-f \text{ is odd.} \end{cases} \end{aligned} \quad (16)$$

Besides, the empirical average  $\frac{1}{n-f} \sum_{i \in S} |x_i - \bar{x}_S|^2$  is equal to

$$\begin{aligned} \frac{1}{n-f} \sum_{i \in S} |x_i - \bar{x}_S|^2 &= \frac{1}{n-f} \sum_{i \in S} (x_i^2 + \bar{x}_S^2 - 2x_i \bar{x}_S) = \frac{1}{n-f} \sum_{i \in S} x_i^2 - \bar{x}_S^2 \\ &= 1 - \bar{x}_S^2 = \begin{cases} 1 & \text{if } n-f \text{ is even,} \\ 1 - \left( \frac{1}{n-f} \right)^2 & \text{if } n-f \text{ is odd.} \end{cases} \end{aligned} \quad (17)$$

Now, observe that the value of  $F(x_1, \dots, x_n)$  is equal to 1 for every  $F \in \{\text{GM}, \text{CWMed}, \text{Krum}\}$ . Indeed, since  $f \geq 1$ , a strict majority of the inputs take the value 1, while the remaining inputs take the value  $-1$ . Recall also that GM is identical to CWMed in one dimension.

By plugging this in (15), using (16) and (17), and rearranging terms, we obtain for every  $F \in \{\text{GM}, \text{CWMed}, \text{Krum}\}$

$$\kappa \geq \frac{|F(x_1, \dots, x_n) - \bar{x}_S|^2}{\frac{1}{n-f} \sum_{i \in S} |x_i - \bar{x}_S|^2} = \frac{|1 - \bar{x}_S|^2}{\frac{1}{n-f} \sum_{i \in S} |x_i - \bar{x}_S|^2} = \begin{cases} 1 & \text{if } n-f \text{ is even,} \\ \frac{|1 - \frac{1}{n-f}|^2}{1 - \left( \frac{1}{n-f} \right)^2} & \text{if } n-f \text{ is odd.} \end{cases} \quad (18)$$

Note however that  $\frac{|1 - \frac{1}{n-f}|^2}{1 - \left( \frac{1}{n-f} \right)^2} = \frac{1 - \frac{2}{n-f} + \frac{1}{n^2}}{1 - \frac{1}{n^2}} = \frac{n-f-1}{n-f+1} = 1 - \frac{2}{n-f+1}$ . Since  $f \geq 1$ , then  $n \geq 2f + 1 \geq 3$ , and thus  $\frac{|1 - \frac{1}{n-f}|^2}{1 - \left( \frac{1}{n-f} \right)^2} = 1 - \frac{2}{n-f+1} \geq \frac{1}{3}$ . Therefore, in both cases of (18), we have  $\kappa \geq \frac{1}{3}$  for every  $F \in \{\text{GM}, \text{CWMed}, \text{Krum}\}$ . This concludes the proof.  $\square$

We can now see that the robustness coefficients given in Table 1 are tight in order of magnitude. Assume that  $n \geq (2 + \nu)f$  for some absolute constant  $\nu > 0$ . The coefficient of CWTM is of order  $\mathcal{O}\left(\frac{f}{n-2f}\right)$ , which is optimal following Proposition 6. The coefficients of CWMed, GM and Krum are of order  $\mathcal{O}(1)$ , which is optimal following Proposition 7.

### 8.3 Unifying Robustness Definitions

In this section, we prove that  $(f, \kappa)$ -robustness unifies existing robustness definitions in the literature (Farhadkhani et al., 2022b; Karimireddy et al., 2022). Specifically, we prove that verifying our definition implies verifying the other definitions. The reason behind this is that, in  $(f, \kappa)$ -robustness, we control the error on estimating the average with a smaller quantity compared to existing works.

#### 8.3.1 $(f, \lambda)$ -resilient Averaging

Recall the definition of  $(f, \lambda)$ -resilient averaging (Farhadkhani et al., 2022b) below.

**Definition 3 ( $(f, \lambda)$ -Resilient averaging).** For  $f < n$  and real value  $\lambda \geq 0$ , an aggregation rule  $F$  is called  $(f, \lambda)$ -resilient averaging if for any collection of  $n$  vectors  $x_1, \dots, x_n$ , and any set  $S \subseteq \{1, \dots, n\}$  of size  $n - f$ ,

$$\|F(x_1, \dots, x_n) - \bar{x}_S\| \leq \lambda \max_{i,j \in S} \|x_i - x_j\|,$$

where  $\bar{x}_S := \frac{1}{|S|} \sum_{i \in S} x_i$ , and  $|S|$  is the cardinality of  $S$ .

The following proposition shows that  $(f, \kappa)$ -robustness implies  $(f, \lambda)$ -resilient averaging.

**Proposition 8.** Let  $n \in \mathbb{N}^*$ ,  $f < n$  and  $\kappa > 0$ . If  $F$  is  $(f, \kappa)$ -robust, then  $F$  is  $(f, \lambda)$ -resilient with  $\lambda = \sqrt{\kappa/2}$ .

*Proof.* Let  $n \in \mathbb{N}^*$ ,  $f < n$  and  $\lambda, \kappa > 0$ . Let  $x_1, \dots, x_n \in \mathbb{R}^d$  and  $S \subseteq \{1, \dots, n\}$  such that  $|S| = n - f$ .

If  $F$  is  $(f, \kappa)$ -robust, then we have

$$\|F(x_1, \dots, x_n) - \bar{x}_S\|^2 \leq \kappa \frac{1}{n-f} \sum_{i \in S} \|x_i - \bar{x}_S\|^2 = \kappa \frac{1}{2(n-f)^2} \sum_{i,j \in S} \|x_i - x_j\|^2 \leq \frac{\kappa}{2} \max_{i,j \in S} \|x_i - x_j\|^2. \quad (19)$$

The equality above is due to

$$\begin{aligned} \frac{1}{2(n-f)^2} \sum_{i,j \in S} \|x_i - x_j\|^2 &= \frac{1}{2(n-f)^2} \sum_{i,j \in S} \|x_i - \bar{x}_S + \bar{x}_S - x_j\|^2 \\ &= \frac{1}{2(n-f)^2} \sum_{i,j \in S} \left( \|x_i - \bar{x}_S\|^2 + \|\bar{x}_S - x_j\|^2 - 2(x_i - \bar{x}_S)(x_j - \bar{x}_S) \right) \\ &= \frac{1}{2(n-f)^2} \left[ (n-f) \sum_{i \in S} \|x_i - \bar{x}_S\|^2 + (n-f) \sum_{j \in S} \|\bar{x}_S - x_j\|^2 \right] \\ &\quad - \frac{1}{(n-f)^2} \sum_{i,j \in S} (x_i - \bar{x}_S)(x_j - \bar{x}_S) \\ &= \frac{1}{n-f} \sum_{i \in S} \|x_i - \bar{x}_S\|^2 - \frac{1}{(n-f)^2} \sum_{i \in S} (x_i - \bar{x}_S) \underbrace{\sum_{j \in S} (x_j - \bar{x}_S)}_{=0} \\ &= \frac{1}{n-f} \sum_{i \in S} \|x_i - \bar{x}_S\|^2. \end{aligned} \quad (20)$$

Taking the square root of both sides in (19) concludes the proof.  $\square$

As a consequence of Proposition 8, one can measure the significant improvement over the analysis of aggregation rules in (Farhadkhani et al., 2022b). For example, the coefficient  $\lambda$  proved for CWMed in (Farhadkhani et al., 2022b) is  $\frac{n}{2(n-f)} \min\{2\sqrt{n-f}, \sqrt{d}\}$ , which is either growing with the dimension or  $n$ . In contrast, Proposition 8 shows that our analysis implies the coefficient  $\lambda = \sqrt{\kappa/2} = \sqrt{2} \left(1 + \frac{f}{n-2f}\right)$ , which is bounded whenever  $n \geq (2 + \nu)f$  for some  $\nu > 0$ . A similar observation can be made for CWTM.

### 8.3.2 $(\delta_{\max}, c)$ -agnostic Robust Aggregation

We recall the definition of an agnostic robust aggregator (ARAgg) (Karimireddy et al., 2022) below. Note that the so-called "good" subset in the original definition of  $(\delta_{\max}, c)$ -ARAgg (Karimireddy et al., 2022) is only required to be of size  $|S| \geq (1 - \delta)n$ , where  $\delta$  is an *upper bound* on the fraction of Byzantine workers. In our formalism,  $f$  is the actual number of Byzantine workers, and thus we directly require  $|S| = n - f$  without loss of generality.

**Definition 4** ( $(\delta_{\max}, c)$ -ARAgg). *Given inputs  $X_1, \dots, X_n$  such that a subset  $S, |S| = n - f$  with  $\frac{f}{n} \leq \delta_{\max} < 0.5$  satisfies  $\mathbb{E} [\|X_i - X_j\|^2] \leq \rho^2$  for all  $i, j \in S$ . Then, the output  $\hat{X}$  of a  $(\delta_{\max}, c)$ -ARAgg satisfies  $\mathbb{E} \left[ \left\| \hat{X} - \bar{X}_S \right\|^2 \right] \leq c \frac{f}{n} \rho^2$ .*

The following proposition shows that  $(f, \kappa)$ -robustness implies  $(\delta_{\max}, c)$ -agnostic robust aggregation. We also show that to obtain  $(\delta_{\max}, c)$ -agnostic robust aggregation with  $c = \mathcal{O}(1)$ , although there is no such theoretical requirement in (Karimireddy et al., 2022), it is sufficient to have  $\kappa \in \mathcal{O}\left(\frac{f}{n-2f}\right)$ .

**Proposition 9.** *Let  $n \in \mathbb{N}^*, 0 < f < n/2, \kappa > 0$ , and  $\frac{f}{n} \leq \delta_{\max} < \frac{1}{2}$ . If  $F$  is  $(f, \kappa)$ -robust, then  $F$  is  $(\delta_{\max}, c)$ -ARAgg with  $c = \kappa \frac{n}{2f}$ . Furthermore, if  $\kappa \in \mathcal{O}\left(\frac{f}{n-2f}\right)$  then  $c = \mathcal{O}(1)$ .*

*Proof.* Let  $n \in \mathbb{N}^*, 0 < f < n/2, \kappa > 0$ , and  $0 < \delta_{\max} < \frac{1}{2}$ . Assume that  $F$  is  $(f, \kappa)$ -robust. Consider any  $x_1, \dots, x_n \in \mathbb{R}^d$  and  $S \subseteq [n]$  such that  $|S| = n - f$ .

As  $F$  is  $(f, \kappa)$ -robust, using (20), we have

$$\|F(x_1, \dots, x_n) - \bar{x}_S\|^2 \leq \frac{\kappa}{n-f} \sum_{i \in S} \|x_i - \bar{x}_S\|^2 = \frac{\kappa}{2(n-f)^2} \sum_{i, j \in S} \|x_i - x_j\|^2. \quad (21)$$

For any random variables  $X_1, \dots, X_n$ , integrating over (21) with the joint probability measure of these variables then gives

$$\begin{aligned} \mathbb{E} \left[ \|F(X_1, \dots, X_n) - \bar{X}_S\|^2 \right] &\leq \frac{\kappa}{2} \mathbb{E} \left[ \frac{1}{(n-f)^2} \sum_{i, j \in S} \|X_i - X_j\|^2 \right] = \frac{\kappa}{2} \frac{1}{|S|^2} \sum_{i, j \in S} \mathbb{E} \left[ \|X_i - X_j\|^2 \right] \\ &\leq \frac{\kappa}{2} \max_{i, j \in S} \mathbb{E} \left[ \|X_i - X_j\|^2 \right] = c \frac{f}{n} \max_{i, j \in S} \mathbb{E} \left[ \|X_i - X_j\|^2 \right]. \end{aligned} \quad (22)$$

where  $c := \kappa \frac{n}{2f}$ . Thus, if  $X_1, \dots, X_n$  are such that there exists a subset  $S \subseteq [n], |S| = n - f$ , for which  $\mathbb{E} \left[ \|X_i - X_j\|^2 \right] \leq \rho^2$  for all  $i, j \in S$ . Then, it holds that  $\rho^2 \geq \max_{i, j \in S} \mathbb{E} \left[ \|X_i - X_j\|^2 \right]$ . This fact together with (22) allows to conclude the desired result; that is,

$$\mathbb{E} \left[ \|F(X_1, \dots, X_n) - \bar{X}_S\|^2 \right] \leq c \frac{f}{n} \rho^2.$$

Furthermore, if  $\kappa = \mathcal{O}\left(\frac{f}{n-2f}\right)$ , then we can write  $\kappa \leq 2m(1 - 2\delta_{\max})\frac{f}{n-2f}$  for some absolute constant  $m > 0$ , since  $\delta_{\max} \in (0, \frac{1}{2})$ . However, since  $\frac{f}{n} \leq \delta_{\max}$  and  $\delta \mapsto \frac{2\delta}{1-2\delta}$  is non-decreasing, we have  $\frac{f}{n-2f} = \frac{f}{n} \left(1 + \frac{2f}{n-2f}\right) \leq \frac{f}{n} \left(1 + \frac{2\delta_{\max}}{1-2\delta_{\max}}\right)$ , and thus  $\kappa \leq 2m(1 - 2\delta_{\max})\frac{f}{n-2f} \leq 2m \frac{f}{n}$ . As a result, we have  $c = \kappa \frac{n}{2f} \leq m = \mathcal{O}(1)$ . This concludes the proof.  $\square$

Our analysis improves over that of (Karimireddy et al., 2022) for CWMed and Krum. For CWMed, their rate is dimension-dependent unlike ours. For Krum, they only prove robustness assuming  $n > 4f$ , while our analysis holds for  $n > 2f$ .

## 9 PROOF OF LEMMA 1: ANALYSIS OF NNM

### 9.1 Proof Overview

The proof of Lemma 1 relies on the observation that NNM brings the inputs closer to the true average. This is formalized in Lemma 5 where we show that the empirical variance is reduced by a factor of order  $f/n$ . Note that although outputs of NNM have smaller variance than the original inputs, their average may deviate from the original average. Then, in the proof of Lemma 1, we control this bias introduced by the NNM operation, and use the reduction proved in Lemma 5 to conclude.

**Notation.** In the following, for every set  $S \subseteq [n]$ , and every vectors  $x_1, \dots, x_n \in \mathbb{R}^d$ , we denote by  $\bar{x}_S$  the average  $\frac{1}{|S|} \sum_{i \in S} x_i$ . Let  $\mu \in \mathbb{R}^d$ . We denote by  $y_\mu$  the average of the  $n - f$  nearest neighbors of  $\mu$  in  $\{x_1, \dots, x_n\}$ :

$$y_\mu := \frac{1}{|\mathcal{N}_\mu|} \sum_{i \in \mathcal{N}_\mu} x_i,$$

where  $\mathcal{N}_\mu \subseteq \{1, \dots, n\}$ ,  $|\mathcal{N}_\mu| = n - f$ , is the set of indices of the  $n - f$  nearest neighbors of  $\mu$ .

### 9.2 Proof of Supporting Lemmas

We first prove a general lemma allowing us to control the distance between the nearest neighbor average  $y_\mu$  and the true average  $\bar{x}_S$  with the dispersion of  $(x_i)_{i \in S}$  around the pivot  $\mu$ .

**Lemma 4.** *Let  $n \in \mathbb{N}^*$ ,  $f < n$ , and  $\mu \in \mathbb{R}^d$ . For any set  $S \subseteq \{1, \dots, n\}$ ,  $|S| = n - f$ , we have for any vectors  $x_1, \dots, x_n \in \mathbb{R}^d$ ,*

$$\|y_\mu - \bar{x}_S\|^2 \leq \frac{4f}{n-f} \frac{1}{|S|} \sum_{i \in S} \|x_i - \mu\|^2.$$

*Proof.* Let  $n \in \mathbb{N}^*$ ,  $f < n$ ,  $\mu \in \mathbb{R}^d$ ,  $x_1, \dots, x_n \in \mathbb{R}^d$ , and  $S \subseteq \{1, \dots, n\}$ ,  $|S| = n - f$ . Recall that, by definition, we have  $y_\mu = \frac{1}{n-f} \sum_{i \in \mathcal{N}_\mu} x_i$ , where  $\mathcal{N}_\mu \subseteq \{1, \dots, n\}$ ,  $|\mathcal{N}_\mu| = n - f$ , is the set of indices of the  $n - f$  nearest neighbors of  $\mu$ . We then have

$$\begin{aligned} \|y_\mu - \bar{x}_S\|^2 &= \left\| \frac{1}{n-f} \sum_{i \in \mathcal{N}_\mu} x_i - \frac{1}{n-f} \sum_{i \in S} x_i \right\|^2 = \frac{1}{(n-f)^2} \left\| \sum_{i \in \mathcal{N}_\mu} x_i - \sum_{i \in S} x_i \right\|^2 \\ &= \frac{1}{(n-f)^2} \left\| \sum_{i \in \mathcal{N}_\mu \setminus S} x_i - \sum_{i \in S \setminus \mathcal{N}_\mu} x_i \right\|^2. \end{aligned}$$

Observe that, since  $|S| = |\mathcal{N}_\mu| = n - f$ , we have  $|S \setminus \mathcal{N}_\mu| = |\mathcal{N}_\mu \setminus S| = |S \cup \mathcal{N}_\mu| - |S| \leq n - (n - f) \leq f$ . As a result, by applying Jensen's inequality, we have

$$\begin{aligned} \|y_\mu - \bar{x}_S\|^2 &= \frac{1}{(n-f)^2} \left\| \sum_{i \in \mathcal{N}_\mu \setminus S} x_i - \sum_{i \in S \setminus \mathcal{N}_\mu} x_i \right\|^2 = \frac{1}{(n-f)^2} \left\| \sum_{i \in \mathcal{N}_\mu \setminus S} (x_i - \mu) - \sum_{i \in S \setminus \mathcal{N}_\mu} (x_i - \mu) \right\|^2 \\ &\leq \frac{|S \setminus \mathcal{N}_\mu| + |\mathcal{N}_\mu \setminus S|}{(n-f)^2} \left[ \sum_{i \in \mathcal{N}_\mu \setminus S} \|x_i - \mu\|^2 + \sum_{i \in S \setminus \mathcal{N}_\mu} \|x_i - \mu\|^2 \right] \\ &\leq \frac{2f}{(n-f)^2} \left[ \sum_{i \in \mathcal{N}_\mu \setminus S} \|x_i - \mu\|^2 + \sum_{i \in S \setminus \mathcal{N}_\mu} \|x_i - \mu\|^2 \right]. \end{aligned}$$

On one hand, since  $\mathcal{N}_\mu$  is the set of  $n - f$  nearest neighbors to  $\mu$ , the first term can be bounded by

$$\sum_{i \in \mathcal{N}_\mu \setminus S} \|x_i - \mu\|^2 \leq \sum_{i \in \mathcal{N}_\mu} \|x_i - \mu\|^2 \leq \sum_{i \in S} \|x_i - \mu\|^2.$$

On the other hand, the second term can be bounded by

$$\sum_{i \in S \setminus \mathcal{N}_\mu} \|x_i - \mu\|^2 \leq \sum_{i \in S} \|x_i - \mu\|^2.$$

We finally conclude that

$$\|y_\mu - \bar{x}_S\|^2 \leq \frac{4f}{(n-f)^2} \sum_{i \in S} \|x_i - \mu\|^2.$$

□

The second lemma crucially shows that the sum of the bias and the variance is reduced by a factor  $\frac{8f}{n-f}$ . To do so, we specialize the first lemma by setting the pivot to be the element  $x_i, i \in S$ .

**Lemma 5.** *Let  $n \in \mathbb{N}^*$ ,  $f < n$ . For any set  $S \subseteq \{1, \dots, n\}$ ,  $|S| = n - f$ , for any vectors  $x_1, \dots, x_n \in \mathbb{R}^d$ , the vectors  $(y_1, \dots, y_n) = \text{NNM}(x_1, \dots, x_n)$  verify*

$$\underbrace{\frac{1}{|S|} \sum_{i \in S} \|y_i - \bar{y}_S\|^2}_{\text{variance}} + \underbrace{\|\bar{y}_S - \bar{x}_S\|^2}_{\text{bias}} \leq \frac{8f}{n-f} \frac{1}{|S|} \sum_{i \in S} \|x_i - \bar{x}_S\|^2.$$

*Proof.* Let  $n \in \mathbb{N}^*$ ,  $f < n$ ,  $x_1, \dots, x_n \in \mathbb{R}^d$ , and  $S \subseteq \{1, \dots, n\}$ ,  $|S| = n - f$ . We first prove the following bias-variance decomposition

$$\frac{1}{|S|} \sum_{i \in S} \|y_i - \bar{y}_S\|^2 + \|\bar{y}_S - \bar{x}_S\|^2 = \frac{1}{|S|} \sum_{i \in S} \|y_i - \bar{x}_S\|^2. \quad (23)$$

We develop the first term of the l.h.s. of the equality above and obtain

$$\begin{aligned} \frac{1}{n-f} \sum_{i \in S} \|y_i - \bar{y}_S\|^2 &= \frac{1}{n-f} \sum_{i \in S} \|y_i - \bar{x}_S + \bar{x}_S - \bar{y}_S\|^2 \\ &= \frac{1}{n-f} \sum_{i \in S} \|y_i - \bar{x}_S\|^2 + \|\bar{x}_S - \bar{y}_S\|^2 + 2 \frac{1}{n-f} \sum_{i \in S} \langle y_i - \bar{x}_S, \bar{x}_S - \bar{y}_S \rangle. \end{aligned}$$

However, we have

$$\frac{1}{n-f} \sum_{i \in S} \langle y_i - \bar{x}_S, \bar{x}_S - \bar{y}_S \rangle = \left\langle \underbrace{\frac{1}{n-f} \sum_{i \in S} y_i - \bar{x}_S}_{=\bar{y}_S}, \bar{x}_S - \bar{y}_S \right\rangle = -\|\bar{x}_S - \bar{y}_S\|^2.$$

As a result, we then have

$$\begin{aligned} \frac{1}{n-f} \sum_{i \in S} \|y_i - \bar{y}_S\|^2 &= \frac{1}{n-f} \sum_{i \in S} \|y_i - \bar{x}_S\|^2 + \|\bar{x}_S - \bar{y}_S\|^2 - 2\|\bar{x}_S - \bar{y}_S\|^2 \\ &= \frac{1}{n-f} \sum_{i \in S} \|y_i - \bar{x}_S\|^2 - \|\bar{x}_S - \bar{y}_S\|^2. \end{aligned}$$

This proves (23). Besides, for any  $i \in S$ , we know from Lemma 4, with  $\mu = x_i$ , that

$$\|y_i - \bar{x}_S\|^2 \leq \frac{4f}{n-f} \frac{1}{n-f} \sum_{j \in S} \|x_j - x_i\|^2.$$

As a consequence, using (23), we have

$$\begin{aligned}
 \frac{1}{|S|} \sum_{i \in S} \|y_i - \bar{y}_S\|^2 + \|\bar{y}_S - \bar{x}_S\|^2 &= \frac{1}{n-f} \sum_{i \in S} \|y_i - \bar{x}_S\|^2 \leq \frac{4f}{n-f} \frac{1}{(n-f)^2} \sum_{i,j \in S} \|x_j - x_i\|^2 \\
 &= \frac{4f}{n-f} \frac{1}{(n-f)^2} \sum_{i,j \in S} \|x_j - \bar{x}_S - (x_i - \bar{x}_S)\|^2 \\
 &= \frac{4f}{n-f} \frac{1}{(n-f)^2} \sum_{i,j \in S} \left( \|x_j - \bar{x}_S\|^2 + \|x_i - \bar{x}_S\|^2 - 2 \langle x_j - \bar{x}_S, x_i - \bar{x}_S \rangle \right) \\
 &= \frac{4f}{n-f} \frac{1}{(n-f)^2} \left[ \sum_{i,j \in S} \|x_j - \bar{x}_S\|^2 + \sum_{i,j \in S} \|x_i - \bar{x}_S\|^2 - 2 \sum_{i,j \in S} \langle x_j - \bar{x}_S, x_i - \bar{x}_S \rangle \right] \\
 &= \frac{4f}{n-f} \frac{1}{(n-f)^2} \left[ 2(n-f) \sum_{i \in S} \|x_i - \bar{x}_S\|^2 - 2 \sum_{i,j \in S} \langle x_j - \bar{x}_S, x_i - \bar{x}_S \rangle \right] \\
 &= \frac{4f}{n-f} \frac{1}{(n-f)^2} \left[ 2(n-f) \sum_{i \in S} \|x_i - \bar{x}_S\|^2 - 2 \sum_{i \in S} \underbrace{\left\langle \sum_{j \in S} (x_j - \bar{x}_S), x_i - \bar{x}_S \right\rangle}_{=0} \right] \\
 &= \frac{8f}{n-f} \frac{1}{n-f} \sum_{i \in S} \|x_i - \bar{x}_S\|^2.
 \end{aligned}$$

This concludes the proof.  $\square$

### 9.3 Proof of Lemma 1

We can now prove Lemma 1, mainly thanks to the bias-variance trade-off guaranteed by Lemma 5.

**Lemma 1.** *Let  $f < n/2$  and  $F: \mathbb{R}^{d \times n} \rightarrow \mathbb{R}^d$ . If  $F$  is  $(f, \kappa)$ -robust, then  $F \circ \text{NNM}$  is  $(f, \kappa')$ -robust with  $\kappa' \leq \frac{8f}{n-f}(\kappa + 1)$ .*

*Proof.* Let  $n \in \mathbb{N}^*$ ,  $f < n$ ,  $x_1, \dots, x_n \in \mathbb{R}^d$ ,  $S \subseteq \{1, \dots, n\}$ ,  $|S| = n - f$ . Assume  $F$  is  $(f, \kappa)$ -robust. First, denote  $(y_1, \dots, y_n) := F \circ \text{NNM}(x_1, \dots, x_n)$  and  $\bar{y}_S := \frac{1}{n-f} \sum_{i \in S} y_i$ . Since  $F$  is  $(f, \kappa)$ -robust, the vectors  $y_1, \dots, y_n$  satisfy

$$\|F \circ \text{NNM}(x_1, \dots, x_n) - \bar{y}_S\|^2 = \|F(y_1, \dots, y_n) - \bar{y}_S\|^2 \leq \kappa \frac{1}{|S|} \sum_{i \in S} \|y_i - \bar{y}_S\|^2. \quad (24)$$

Now, thanks to Young's inequality we have for  $c = 1/\kappa$ ,  $(a + b)^2 \leq (1 + c)a^2 + (1 + 1/c)b^2$ . We can then write using (24)

$$\begin{aligned}
 \|F \circ \text{NNM}(x_1, \dots, x_n) - \bar{x}_S\|^2 &= \|F \circ \text{NNM}(x_1, \dots, x_n) - \bar{y}_S + \bar{y}_S - \bar{x}_S\|^2 \\
 &\leq \left(1 + \frac{1}{\kappa}\right) \|F \circ \text{NNM}(x_1, \dots, x_n) - \bar{y}_S\|^2 + (1 + \kappa) \|\bar{y}_S - \bar{x}_S\|^2 \\
 &\leq \left(1 + \frac{1}{\kappa}\right) \kappa \frac{1}{|S|} \sum_{i \in S} \|y_i - \bar{y}_S\|^2 + (1 + \kappa) \|\bar{y}_S - \bar{x}_S\|^2 \\
 &= \frac{1 + \kappa}{n-f} \sum_{i \in S} \|y_i - \bar{y}_S\|^2 + (1 + \kappa) \|\bar{y}_S - \bar{x}_S\|^2. \quad (25)
 \end{aligned}$$

Recall that Lemma 5 shows that  $\frac{1}{|S|} \sum_{i \in S} \|y_i - \bar{y}_S\|^2 + \|\bar{y}_S - \bar{x}_S\|^2 \leq \frac{8f}{n-f} \frac{1}{|S|} \sum_{i \in S} \|x_i - \bar{x}_S\|^2$ . Plugging this in (25) allows to conclude

$$\|F \circ \text{NNM}(x_1, \dots, x_n) - \bar{x}_S\|^2 \leq \frac{8f}{n-f} \frac{1 + \kappa}{n-f} \sum_{i \in S} \|x_i - \bar{x}_S\|^2.$$

$\square$



## 10 ANALYSIS OF BUCKETING

In this section, we compare the robustness guarantees of the Bucketing algorithm (Karimireddy et al., 2022) with NNM. Bucketing is a pre-aggregation step that maps input vectors  $\{x_1, \dots, x_n\}$  to output vectors  $\{y_1, \dots, y_{\lceil n/s \rceil}\}$ , where  $s$  is a parameter called the bucket size. The Bucketing algorithm works as follows. First, it samples a random permutation  $\pi: [n] \rightarrow [n]$ . Then, for each  $i \in \{1, \dots, \lceil n/s \rceil\}$ , it computes the output of bucket  $i$  as<sup>7</sup>

$$y_i := \frac{1}{s} \sum_{k=s(i-1)+1}^{\min(s \cdot i, n)} x_{\pi(k)}.$$

Recall that the property that enables NNM to boost the robustness of aggregations is that the heterogeneity of the output vectors  $\frac{1}{|S|} \sum_{i \in S} \|y_i - \bar{x}_S\|^2$  is a factor  $\mathcal{O}(f/n)$  smaller than the heterogeneity of the input vectors  $\frac{1}{|S|} \sum_{i \in S} \|x_i - \bar{x}_S\|^2$  (see Lemma 5). But, unlike NNM, Bucketing only reduces the heterogeneity *in expectation over the random permutations* (see Lemma 1 in (Karimireddy et al., 2022)). Thus, there are iterations of the learning algorithm where Bucketing may not reduce the heterogeneity, and these iterations are the best opportunity for the Byzantine workers to cause the most damage to the learning.

We give in Observation 1 below a simple example showing that deterministically reducing the heterogeneity of inputs is impossible using Bucketing in general, even in the absence of malicious inputs.

**Observation 1.** *Using Bucketing, it is impossible to provide a worst-case heterogeneity reduction guarantee regardless of the value of  $s$ , even in the absence of Byzantine inputs.*

*Proof.* Let  $\pi$  be any permutation over  $[n]$ . We will construct an instance of  $n$  inputs such that an execution of Bucketing using permutation  $\pi$  does not reduce the heterogeneity. Consider the inputs  $(x_1, \dots, x_n)$  such that for all  $1 \leq i \leq \lceil n/s \rceil$  and all  $k, l \in \{s(i-1)+1, \dots, s \cdot i\}$ , it holds that  $x_{\pi(k)} = x_{\pi(l)} =: y_i$ . Thus, applying the permutation  $\pi$  on these inputs yields  $(\underbrace{y_1, \dots, y_1}_{s \text{ times}}, \dots, \underbrace{y_{\lceil n/s \rceil}, \dots, y_{\lceil n/s \rceil}}_{s \text{ times}})$  in which each vector  $y_i$  is repeated  $s$  times. Overall, the execution of Bucketing with  $\pi$  will produce  $(y_1, \dots, y_{\lceil n/s \rceil})$ , which has the same variance as the original inputs, and thus heterogeneity was not reduced.  $\square$

Moreover, given that the buckets are randomly chosen, the (expected) reduction of heterogeneity is only achieved at the expense of an increase in the fraction of Byzantine inputs:

**Observation 2.** *Bucketing increases the fraction of Byzantine workers by a factor  $s$  in the worst case.*

*Proof.* In some executions, Bucketing might assign every Byzantine input to a different bucket. In this case, since the average within each ‘‘contaminated’’ bucket is arbitrarily manipulable by a single Byzantine input, the number of Byzantine output vectors is the same as the number of Byzantine input vectors. However, the total number of output vectors is  $s$  times smaller than the total number of input vectors, thereby effectively increasing the fraction of Byzantine workers.  $\square$

This observation implies that in order to have a worst-case robustness guarantee, the bucket size  $s$  can be at most  $\lfloor \frac{n}{2f} \rfloor$ , as done in (Karimireddy et al., 2022). This instability in reducing the heterogeneity results in a poor estimation of the mean in some iterations. We observe this instability in practice as we show below on experiments on CIFAR-10.

**Experimental validation.** Recall robust D-SHB (Algorithm 3), and the corresponding notations. In Figure 2, we plot in each step  $t \in [T]$  the quantity  $\hat{\kappa}_t$  defined as follows

$$\hat{\kappa}_t^2 := \|R_t - \bar{m}_t\|^2 / \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \|m_t^{(i)} - \bar{m}_t\|^2, \quad (26)$$

where  $\bar{m}_t := \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} m_t^{(i)}$  is the average of honest momentums at step  $t$ . In words, Figure 2 shows the error in estimating the true average, scaled by the standard deviation of the honest inputs, across different steps of the learning. The quantity in Equation (26) is an empirical estimation of  $\kappa$ , a parameter of the definition of  $(f, \kappa)$ -robustness in our theory. The figure validates our insights about Bucketing and shows the superiority of NNM both in terms of stability of the error curves, as well as the quality of the average estimation. Indeed, the NNM curves are consistently below Bucketing’s.

<sup>7</sup>In the theoretical analysis, for simplicity, we assume that  $n$  is divisible by  $s$ . If this is not the case, some buckets will include less than  $s$  vectors which is less favorable for the theoretical guarantees of Bucketing but does not change the asymptotic behavior.




Figure 2: Plot of  $\hat{\kappa}_t$  (defined in (26)) on CIFAR-10 with  $f = 2$  Byzantine workers among  $n = 17$  executing the ALIE (column 1) and FOE (column 2) attacks. The heterogeneity level is  $\alpha = 1$ . The corresponding accuracies are in Figure 25.

## 11 PROOF OF THEOREM 1: Analysis of Robust D-GD

**Theorem 1.** *Let Assumption 1 hold and recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 1 with learning rate  $\gamma = \frac{1}{L}$ . If  $F$  is  $(f, \kappa)$ -robust, then for all  $T \geq 1$*

$$\left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \leq 4\kappa G^2 + \frac{4L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{T}.$$

*Proof.* Let Assumption 1 hold. Assume  $\mathcal{L}_{\mathcal{H}}$  to be  $L$ -smooth and  $F$  to be  $(f, \kappa)$ -robust. Consider Algorithm 1 with learning rate  $\gamma = \frac{1}{L}$ . The update step corresponds to  $\theta_t = \theta_{t-1} - \frac{1}{L} R_t$  for every  $t \geq 1$ .

Since  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth, we have (see (Bottou et al., 2018)) for all  $t \geq 1$ ,

$$\begin{aligned} \mathcal{L}_{\mathcal{H}}(\theta_t) - \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) &\leq \langle \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), \theta_t - \theta_{t-1} \rangle + \frac{L}{2} \|\theta_t - \theta_{t-1}\|^2 = -\gamma \langle \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), R_t \rangle + \frac{1}{2} \gamma^2 L \|R_t\|^2 \\ &= -\frac{1}{L} \langle \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), R_t \rangle + \frac{1}{2L} \|R_t\|^2. \end{aligned} \quad (27)$$

We expand the second term as follows

$$\begin{aligned} \|R_t\|^2 &= \|R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) + \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \\ &= \|R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2 \langle \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle. \end{aligned}$$

Plugging this back in (27), then simplifying, yields

$$\begin{aligned} \mathcal{L}_{\mathcal{H}}(\theta_t) - \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) &\leq -\frac{1}{L} \langle \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), R_t \rangle + \frac{1}{2L} \|R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + \frac{1}{2L} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \\ &\quad + \frac{1}{L} \langle \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle \\ &= \frac{1}{2L} \|R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 - \frac{1}{2L} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2. \end{aligned}$$

Upon rearranging terms and multiplying both sides by  $2L$  we get

$$\|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \leq \|R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2L(\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) - \mathcal{L}_{\mathcal{H}}(\theta_t)).$$

By the  $(f, \kappa)$ -robustness property of  $F$  and Assumption 1, we can bound the first term as follows

$$\begin{aligned} \|R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 &= \left\| F(g_t^{(1)}, \dots, g_t^{(n)}) - \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} g_t^{(i)} \right\|^2 \leq \frac{\kappa}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| g_t^{(i)} - \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} g_t^{(i)} \right\|^2 \\ &= \frac{\kappa}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \|\nabla \mathcal{L}_i(\theta_{t-1}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \leq \kappa G^2. \end{aligned} \quad (28)$$

As a result, we have

$$\|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \leq \kappa G^2 + 2L(\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) - \mathcal{L}_{\mathcal{H}}(\theta_t)).$$

By taking the average over  $t$  from 1 to  $T$ , and since  $\mathcal{L}_{\mathcal{H}}(\theta_T) \geq \inf_{\theta \in \mathbb{R}^d} \mathcal{L}_{\mathcal{H}}(\theta) = \mathcal{L}^*$ , we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_t)\|^2 \leq \kappa G^2 + \frac{2L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}_{\mathcal{H}}(\theta_T))}{T} \leq \kappa G^2 + \frac{2L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{T}. \quad (29)$$

**Final step.** Recall from Algorithm 1 that  $\hat{\theta} = \theta_{\tau-1}$  with  $\tau \in \arg \min_{1 \leq t \leq T} \|R_t\|$ . Thus, using Jensen's inequality, (28) and then (29), we have

$$\begin{aligned} \left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 &= \min_{1 \leq t \leq T} \|R_t\|^2 = \min_{1 \leq t \leq T} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) + (R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}))\|^2 \\ &\leq \min_{1 \leq t \leq T} \left\{ 2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2 \|R_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right\} \leq 2 \min_{1 \leq t \leq T} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2\kappa G^2 \\ &\leq 2 \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_t)\|^2 + 2\kappa G^2 \leq 2\kappa G^2 + \frac{4L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{T} + 2\kappa G^2 \\ &= 4\kappa G^2 + \frac{4L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{T}. \end{aligned}$$

This concludes the proof.  $\square$

## 12 PROOF OF COROLLARY 1: ANALYSIS OF ROBUST D-GD WITH NNM

**Corollary 1.** *Let Assumption 1 hold and recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 1 with  $\gamma = 1/L$  and aggregation  $F \circ \text{NNM}$ . If  $F$  is  $(f, \kappa)$ -robust with  $\kappa = \mathcal{O}(1)$ , then Algorithm 1 is  $(f, \varepsilon)$ -Byzantine resilient with*

$$\varepsilon = \mathcal{O}(f/nG^2 + 1/T).$$

*Proof.* Let Assumption 1 hold. Assume  $\mathcal{L}_{\mathcal{H}}$  to be  $L$ -smooth and  $F$  to be  $(f, \kappa)$ -robust. First recall that, by Lemma 1, the composition  $F \circ \text{NNM}$  is  $(f, \kappa')$ -robust with  $\kappa' = \frac{8f}{n-f}(\kappa + 1)$ .

Consider Algorithm 1 with learning rate  $\gamma = 1/L$  and aggregation rule  $F \circ \text{NNM}$ . Following Theorem 1, we have for every  $T \geq 1$ ,

$$\left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \leq 4\kappa' G^2 + \frac{4L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{T} = \frac{32f}{n-f}(\kappa + 1)G^2 + \frac{4L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{T}.$$

Recall that, as  $f < n/2$ , we have  $\frac{f}{n-f} \leq \frac{2f}{n}$ . Thus, if  $\kappa = \mathcal{O}(1)$ , we can write

$$\left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 = \mathcal{O}(f/nG^2 + 1/T).$$

This concludes the proof.  $\square$

We provide the proof of the lower bound on Byzantine resilience under heterogeneity (Proposition 1) below.

**Proposition 1.** *If a learning algorithm  $\mathcal{A}$  is  $(f, \varepsilon)$ -Byzantine resilient for every collection of smooth loss functions  $\mathcal{L}_1, \dots, \mathcal{L}_n$  satisfying Assumption 1, then  $\varepsilon = \Omega(f/nG^2)$ .*

*Proof.* The proof is similar to that of Theorem III (Karimireddy et al., 2022). Assume learning algorithm  $\mathcal{A}$  is  $(f, \varepsilon)$ -Byzantine resilient for every collection of smooth loss functions  $\mathcal{L}_1, \dots, \mathcal{L}_n$  satisfying Assumption 1. Consider the following quadratic loss functions  $\mathcal{L}_1 = \dots = \mathcal{L}_{n-f} = \theta \mapsto \frac{1}{2} \|\theta\|^2$  and  $\mathcal{L}_{n-f+1} = \dots = \mathcal{L}_n = \theta \mapsto \frac{1}{2} \|\theta - z\|^2$ , where  $z \in \mathbb{R}^d$  is such that  $\|z\|^2 = \frac{(n-f)^2}{f(n-2f)}G^2$ . Consider the two situations  $\mathcal{H}_1 = \{1, \dots, n-f\}$  and  $\mathcal{H}_2 = \{f+1, \dots, n\}$ .

We first show that the loss functions satisfy Assumption 1 in both situations. This is straightforward in situation  $\mathcal{H}_1$  since honest losses are identical. In situation  $\mathcal{H}_2$ , we have for all  $\theta \in \mathbb{R}^d$ ,

$$\nabla \mathcal{L}_{\mathcal{H}_2}(\theta) = \frac{1}{n-f} \sum_{i \in \mathcal{H}_2} \nabla \mathcal{L}_i(\theta) = \frac{n-2f}{n-f} \theta + \frac{f}{n-f}(\theta - z) = \theta - \frac{f}{n-f}z.$$

Therefore, thanks to the choice of  $z$ , we now show that Assumption 1 holds, as for all  $\theta \in \mathbb{R}^d$  we have

$$\frac{1}{|\mathcal{H}_2|} \sum_{i \in \mathcal{H}_2} \|\nabla \mathcal{L}_i(\theta) - \nabla \mathcal{L}_{\mathcal{H}_2}(\theta)\|^2 = \frac{n-2f}{n-f} \left\| \frac{f}{n-f} z \right\|^2 + \frac{f}{n-f} \left\| \frac{n-2f}{n-f} z \right\|^2 = \frac{f(n-2f)}{(n-f)^2} \|z\|^2 = G^2.$$

Now, since learning algorithm  $\mathcal{A}$  is  $(f, \varepsilon)$ -Byzantine resilient, it outputs  $\hat{\theta}$  such that  $\|\nabla \mathcal{L}_{\mathcal{H}_1}(\hat{\theta})\|^2 \leq \varepsilon$  and  $\|\nabla \mathcal{L}_{\mathcal{H}_2}(\hat{\theta})\|^2 \leq \varepsilon$ . Note that situations  $\mathcal{H}_1$  and  $\mathcal{H}_2$  are indistinguishable to algorithm  $\mathcal{A}$  because it ignores the Byzantine identities, and thus  $\hat{\theta}$  is the same in both situations. Therefore, invoking Jensen's inequality, we have

$$\begin{aligned} \varepsilon &\geq \max \left\{ \|\nabla \mathcal{L}_{\mathcal{H}_1}(\hat{\theta})\|^2, \|\nabla \mathcal{L}_{\mathcal{H}_2}(\hat{\theta})\|^2 \right\} \geq \frac{1}{2} \left( \|\nabla \mathcal{L}_{\mathcal{H}_1}(\hat{\theta})\|^2 + \|\nabla \mathcal{L}_{\mathcal{H}_2}(\hat{\theta})\|^2 \right) \\ &= \frac{1}{2} \left( \|\hat{\theta}\|^2 + \left\| \hat{\theta} - \frac{f}{n-f} z \right\|^2 \right) \geq \frac{1}{4} \left\| \frac{f}{n-f} z \right\|^2 = \frac{1}{4} \left( \frac{f}{n-f} \right)^2 \frac{(n-f)^2}{f(n-2f)} G^2 = \frac{1}{4} \frac{f}{n-2f} G^2. \end{aligned}$$

Since  $n-2f \leq n$ , we obtain  $\varepsilon \geq \frac{1}{4} \frac{f}{n} G^2$ , which concludes the proof.  $\square$

## 13 PROOF OF THEOREM 2: ANALYSIS OF ROBUST D-SHB

### 13.1 Proof Outline

Our analysis of robust D-SHB (Algorithm 3) follows (Farhadkhani et al., 2022b) and consists of four elements: (i) *Momentum drift* (Lemma 6), (ii) *Aggregation error* (Lemma 7), (iii) *Momentum deviation* (Lemma 8), and (iv) *Descent bound* (Lemma 9). We combine these elements to obtain the final convergence result stated in Theorem 2. The originality of our analysis, compared to (Farhadkhani et al., 2022b), is (i) the tighter analysis of the aggregation error thanks to  $(f, \kappa)$ -robustness, and (ii) the extension of the momentum drift analysis to the heterogeneous setting. There are other subtle differences, such as the choice of the learning rate.

**Notation.** Recall that for each step  $t$ , for each honest worker  $w_i$ ,

$$m_t^{(i)} = \beta m_{t-1}^{(i)} + (1-\beta) g_t^{(i)} \quad (30)$$

where  $m_0^{(i)} = 0$  by convention. As we analyze Algorithm 3 with aggregation  $F$ , we denote

$$R_t := F(m_t^{(1)}, \dots, m_t^{(n)}), \quad (31)$$

and

$$\theta_t = \theta_{t-1} - \gamma R_t. \quad (32)$$

We denote by  $\mathcal{P}_t$  the history from steps 1 to  $t$ . Specifically,

$$\mathcal{P}_t := \left\{ \theta_0, \dots, \theta_{t-1}; m_1^{(i)}, \dots, m_{t-1}^{(i)}; i = 1, \dots, n \right\}.$$

By convention,  $\mathcal{P}_1 = \{\theta_0\}$ . We denote by  $\mathbb{E}_t[\cdot]$  and  $\mathbb{E}[\cdot]$  the conditional expectation  $\mathbb{E}[\cdot | \mathcal{P}_t]$  and the total expectation, respectively. Thus,  $\mathbb{E}[\cdot] = \mathbb{E}_1[\dots \mathbb{E}_T[\cdot]]$ .

#### 13.1.1 Momentum Drift

Along the trajectory  $\theta_0, \dots, \theta_{t-1}$ , the honest workers' local momentums may drift away from each other. This is in part due to the heterogeneity between the training sets. This induces a dissimilarity between honest workers' local gradients that we can control thanks to Assumption 1. The drift is also due to the stochasticity of the local gradients.

We show in Lemma 6 below that the bound on the *drift* between the honest workers' momentums can be controlled by tuning the momentum coefficient  $\beta$ . In fact, the smaller  $(1-\beta)$  the smaller the bound on the drift. Recall that we denote by  $\mathcal{H} \subseteq [n]$  the set of  $n-f$  honest workers. We define

$$\bar{m}_t := \frac{1}{(n-f)} \sum_{i \in \mathcal{H}} m_t^{(i)}, \quad (33)$$

the average of honest workers' local momentums. We obtain the following bound on the momentum drift, i.e., Lemma 6, proof of which can be found in Appendix 13.4.1.

**Lemma 6.** *Suppose that assumptions 1 and 2 hold true. Consider Algorithm 3. For each  $t \in [T]$ , we obtain that*

$$\mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| m_t^{(i)} - \bar{m}_t \right\|^2 \right] \leq 3 \left( \frac{1 - \beta}{1 + \beta} \right) \left( 1 + \frac{1}{n - f} \right) \sigma^2 + 3G^2.$$

### 13.1.2 Aggregation Error

By building upon this first lemma and the  $(f, \kappa)$ -robustness property, we can obtain a bound on the error between the aggregate  $R_t$  and  $\bar{m}_t$  the average momentum of honest workers for the case. Specifically, when defining the error

$$\xi_t := R_t - \bar{m}_t, \quad (34)$$

we get the following bound on the error in Lemma 7, proof of which can be found in Appendix 13.4.2.

**Lemma 7.** *Suppose that assumptions 1 and 2 hold true. Assume  $F$  is  $(f, \kappa)$ -robust. Consider Algorithm 3 with aggregation  $F$ . For each step  $t \in [T]$ , we obtain that*

$$\mathbb{E} \left[ \|\xi_t\|^2 \right] \leq 6\kappa \frac{1 - \beta}{1 + \beta} \sigma^2 + 3\kappa G^2.$$

### 13.1.3 Momentum Deviation

Next, we study the momentum deviation; i.e., the distance between the average honest momentum  $\bar{m}_t$  and the true gradient  $\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})$  in an arbitrary step  $t$ . Specifically, we define *deviation* to be

$$\delta_t := \bar{m}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), \quad (35)$$

and obtain in Lemma 8 below an upper bound on the growth of the deviation over the learning steps  $t \in [T]$ . (Proof of Lemma 8 can be found in Appendix 13.4.3.)

**Lemma 8.** *Suppose that Assumption 2 holds. Recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 3 with  $T > 1$ . For all  $t \geq 2$  we obtain that*

$$\begin{aligned} \mathbb{E} \left[ \|\delta_t\|^2 \right] &\leq \beta^2 c \mathbb{E} \left[ \|\delta_{t-1}\|^2 \right] + 4\gamma L(1 + \gamma L)\beta^2 \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})\|^2 \right] + (1 - \beta)^2 \frac{\sigma^2}{(n - f)} \\ &\quad + 2\gamma L(1 + \gamma L)\beta^2 \mathbb{E} \left[ \|\xi_{t-1}\|^2 \right], \end{aligned}$$

where  $c := (1 + \gamma L)(1 + 4\gamma L)$ .

### 13.1.4 Descent Bound

Finally, we analyze the fourth element, i.e., the growth of cost function  $\mathcal{L}_{\mathcal{H}}(\theta)$  along the trajectory of Algorithm 3. From (32) and (31), we obtain that, for each step  $t$ ,

$$\theta_t = \theta_{t-1} - \gamma R_t = \theta_{t-1} - \gamma \bar{m}_t - \gamma (R_t - \bar{m}_t),$$

Furthermore, by (34),  $R_t - \bar{m}_t = \xi_t$ . Thus, for all  $t$ ,

$$\theta_t = \theta_{t-1} - \gamma \bar{m}_t - \gamma \xi_t. \quad (36)$$

This means that Algorithm 3 can actually be treated as distributed SGD with a momentum term that is subject to perturbation proportional to  $\xi_t$  at each step  $t$ . This perspective leads us to Lemma 9, proof of which can be found in Appendix 13.4.4.

**Lemma 9.** *Recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 3. For all  $t \in [T]$ , we obtain that*

$$\begin{aligned} \mathbb{E} \left[ 2\mathcal{L}_{\mathcal{H}}(\theta_t) - 2\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \right] &\leq -\gamma(1 - 4\gamma L) \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + 2\gamma(1 + 2\gamma L) \mathbb{E} \left[ \|\delta_t\|^2 \right] \\ &\quad + 2\gamma(1 + \gamma L) \mathbb{E} \left[ \|\xi_t\|^2 \right]. \end{aligned}$$

Putting all of the previous lemmas together, we prove Theorem 2 in Section 13.2 below. As a corollary of Theorem 2 and Lemma 1, we prove Corollary 2 in Section 13.3.

### 13.2 Proof of Theorem 2

We recall the theorem statement below for convenience. Recall that

$$\mathcal{L}^* = \inf_{\theta \in \mathbb{R}^d} \mathcal{L}_{\mathcal{H}}(\theta), \quad a_1 = 36, \quad a_2 = 6\sqrt{\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*}, \quad a_3 = 1728L, \quad a_4 = 288L, \quad \text{and} \quad a_5 = 6La_2^2. \quad (37)$$

**Theorem 2.** *Let assumption 1 and 2 hold and recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Let  $F$  be a  $(f, \kappa)$ -robust aggregation rule. Consider Algorithm 3 with momentum coefficient  $\beta = \sqrt{1 - 24\gamma L}$ , and learning rate*

$$\gamma = \min \left\{ \frac{1}{24L}, \frac{a_2}{2a_\kappa \sigma \sqrt{T}} \right\},$$

with  $a_\kappa^2 := a_3\kappa + \frac{a_4}{n-f}$ . For all  $T \geq 1$

$$\mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \right] \leq a_1 \kappa G^2 + \frac{a_2 a_\kappa \sigma}{\sqrt{T}} + \frac{a_5}{T} + \frac{a_2 a_4 \sigma}{na_\kappa T^{3/2}},$$

where the expectation is over the randomness of the algorithm.

*Proof.* Define

$$\gamma_o := \frac{1}{18L}. \quad (38)$$

Note that as specified in the theorem statement, by definition of  $\gamma$ , we have

$$\gamma \leq \frac{1}{24L} \leq \gamma_o. \quad (39)$$

Therefore,  $\beta = \sqrt{1 - 24\gamma L}$  (as defined) is a well-defined real value in  $(0, 1)$ .

To obtain the convergence result we define the Lyapunov function to be

$$V_t := \mathbb{E} \left[ 2\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) + z \|\delta_t\|^2 \right] \quad \text{and} \quad z = \frac{1}{8L}. \quad (40)$$

We consider an arbitrary  $t \in [T]$ .

**Invoking Lemma 8.** Upon substituting from Lemma 8, we obtain that

$$\begin{aligned} \mathbb{E} \left[ z \|\delta_{t+1}\|^2 - z \|\delta_t\|^2 \right] &\leq z\beta^2 c \mathbb{E} \left[ \|\delta_t\|^2 \right] + 4z\gamma L(1 + \gamma L)\beta^2 \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + z(1 - \beta)^2 \frac{\sigma^2}{n - f} \\ &\quad + 2z\gamma L(1 + \gamma L)\beta^2 \mathbb{E} \left[ \|\xi_t\|^2 \right] - z \mathbb{E} \left[ \|\delta_t\|^2 \right]. \end{aligned} \quad (41)$$

Recall that

$$c = (1 + \gamma L)(1 + 4\gamma L) = 1 + 5\gamma L + 4\gamma^2 L^2. \quad (42)$$

**Invoking Lemma 9.** Substituting from Lemma 9 we obtain that

$$\begin{aligned} \mathbb{E} \left[ 2\mathcal{L}_{\mathcal{H}}(\theta_t) - 2\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \right] &\leq -\gamma(1 - 4\gamma L) \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + 2\gamma(1 + 2\gamma L) \mathbb{E} \left[ \|\delta_t\|^2 \right] \\ &\quad + 2\gamma(1 + \gamma L) \mathbb{E} \left[ \|\xi_t\|^2 \right]. \end{aligned} \quad (43)$$

Substituting from (41) and (43) in (40) we obtain that

$$\begin{aligned} V_{t+1} - V_t &= \mathbb{E} \left[ 2\mathcal{L}_{\mathcal{H}}(\theta_t) - 2\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \right] + \mathbb{E} \left[ z \|\delta_{t+1}\|^2 - z \|\delta_t\|^2 \right] \\ &\leq -\gamma(1 - 4\gamma L) \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + 2\gamma(1 + 2\gamma L) \mathbb{E} \left[ \|\delta_t\|^2 \right] + 2\gamma(1 + \gamma L) \mathbb{E} \left[ \|\xi_t\|^2 \right] \\ &\quad + z\beta^2 c \mathbb{E} \left[ \|\delta_t\|^2 \right] + 4z\gamma L(1 + \gamma L)\beta^2 \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + z(1 - \beta)^2 \frac{\sigma^2}{n - f} \\ &\quad + 2z\gamma L(1 + \gamma L)\beta^2 \mathbb{E} \left[ \|\xi_t\|^2 \right] - z \mathbb{E} \left[ \|\delta_t\|^2 \right]. \end{aligned} \quad (44)$$

Upon re-arranging the R.H.S. in (44) we obtain that

$$V_{t+1} - V_t \leq -\gamma \left( (1 - 4\gamma L) - 4zL(1 + \gamma L)\beta^2 \right) \mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \right\|^2 \right] + z(1 - \beta)^2 \frac{\sigma^2}{n - f} \\ + (2\gamma(1 + 2\gamma L) + z\beta^2 c - z) \mathbb{E} \left[ \left\| \delta_t \right\|^2 \right] + 2\gamma(1 + \gamma L + zL(1 + \gamma L)\beta^2) \mathbb{E} \left[ \left\| \xi_t \right\|^2 \right].$$

For simplicity, we define

$$A := (1 - 4\gamma L) - 4zL(1 + \gamma L)\beta^2, \quad (45)$$

$$B := 2\gamma(1 + 2\gamma L) + z\beta^2 c - z, \quad (46)$$

and

$$C := 2\gamma(1 + \gamma L + zL(1 + \gamma L)\beta^2), \quad (47)$$

Thus,

$$V_{t+1} - V_t \leq -A\gamma \mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \right\|^2 \right] + B \mathbb{E} \left[ \left\| \delta_t \right\|^2 \right] + C \mathbb{E} \left[ \left\| \xi_t \right\|^2 \right] + z(1 - \beta)^2 \frac{\sigma^2}{n - f}. \quad (48)$$

We now analyse below the terms  $A$ ,  $B$  and  $C$ .

**Term A.** Recall from (39) that  $\gamma \leq \gamma_o = \frac{1}{18L}$ . Upon using this in (45), and the facts that  $z = \frac{1}{8L}$  and  $\beta^2 < 1$ , we obtain that

$$A \geq 1 - 4\gamma_o L - \frac{4L}{8L}(1 + \gamma_o L) \geq \frac{1}{2} - \frac{9\gamma_o L}{2} \geq \frac{1}{4}. \quad (49)$$

**Term B.** Substituting  $c$  from (42) in (46) we obtain that

$$B = 2\gamma(1 + 2\gamma L) + z\beta^2(1 + 5\gamma L + 4\gamma^2 L^2) - z \\ = -(1 - \beta^2)z + \gamma(2 + 4\gamma L + 5z\beta^2 L + 4z\beta^2 L\gamma L).$$

Using the facts that  $\beta^2 \leq 1$  and  $\gamma \leq \gamma_o \leq \frac{1}{18L}$ , and then substituting  $z = \frac{1}{8L}$  we obtain that

$$B \leq \frac{-(1 - \beta^2)}{8L} + \gamma \left( 2 + \frac{4}{18} + \frac{5}{8} + \frac{4}{18 \times 8} \right) \leq \frac{-(1 - \beta^2)}{8L} + 3\gamma \leq \frac{-(1 - \beta^2) + 24\gamma L}{8L} = 0, \quad (50)$$

where the last equality follows from the fact that  $1 - \beta^2 = 24\gamma L$ .

**Term C.** Substituting  $z = \frac{1}{8L}$  in (47), and then using the fact that  $\beta^2 < 1$ , we obtain that

$$C = 2\gamma \left( 1 + \gamma L + \frac{1}{8}(1 + \gamma L) \right) \leq \frac{9\gamma}{4}(1 + \gamma L).$$

As  $\gamma \leq \gamma_o \leq \frac{1}{18L}$ , from above we obtain that

$$C \leq \frac{9\gamma}{4} \left( 1 + \frac{1}{18} \right) \leq 3\gamma. \quad (51)$$

**Combining terms A, B, and C.** Finally, substituting from (49), (50), and (51) in (48) (and recalling that  $z = \frac{1}{8L}$ ) we obtain that

$$V_{t+1} - V_t \leq -\frac{\gamma}{4} \mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \right\|^2 \right] + 3\gamma \mathbb{E} \left[ \left\| \xi_t \right\|^2 \right] + (1 - \beta)^2 \frac{\sigma^2}{8L(n - f)}.$$

As the above is true for an arbitrary  $t \in [T]$ , by taking summation on both sides from  $t = 1$  to  $t = T$  we obtain that

$$V_{T+1} - V_1 \leq -\frac{\gamma}{4} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + 3\gamma \sum_{t=1}^T \mathbb{E} \left[ \|\xi_t\|^2 \right] + (1 - \beta)^2 \frac{\sigma^2}{8L(n-f)} T.$$

Thus,

$$\frac{\gamma}{4} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq V_1 - V_{T+1} + 3\gamma \sum_{t=1}^T \mathbb{E} \left[ \|\xi_t\|^2 \right] + (1 - \beta)^2 \frac{\sigma^2}{8L(n-f)} T. \quad (52)$$

Note that, as  $\beta > 0$ , and  $1 - \beta^2 = 24\gamma L$ , we have

$$(1 - \beta)^2 = \frac{(1 - \beta^2)^2}{(1 + \beta)^2} \leq (1 - \beta^2)^2 = 576\gamma^2 L^2.$$

Substituting from above in (52) we obtain that

$$\frac{\gamma}{4} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq V_1 - V_{T+1} + 3\gamma \sum_{t=1}^T \mathbb{E} \left[ \|\xi_t\|^2 \right] + \frac{576\gamma^2 L^2 \sigma^2}{8L(n-f)} T.$$

Multiplying both sides by  $4/\gamma$  we obtain that

$$\sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq \frac{4(V_1 - V_{T+1})}{\gamma} + 12 \sum_{t=1}^T \mathbb{E} \left[ \|\xi_t\|^2 \right] + \frac{576\gamma L \sigma^2}{2(n-f)} T. \quad (53)$$

**Invoking Lemma 7.** Next, we use Lemma 7 to derive an upper bound on  $\sum_{t=1}^T \mathbb{E} \left[ \|\xi_t\|^2 \right]$ . Since  $F$  is  $(f, \kappa)$ -robust, we have from Lemma 7 that for all  $t \in [T]$ ,

$$\mathbb{E} \left[ \|\xi_t\|^2 \right] \leq 6\kappa \frac{1 - \beta}{1 + \beta} \sigma^2 + 3\kappa G^2.$$

By summing over  $t$  from 1 to  $T$ , we obtain that

$$\sum_{t=1}^T \mathbb{E} \left[ \|\xi_t\|^2 \right] \leq 6\kappa \frac{1 - \beta}{1 + \beta} \sigma^2 T + 3\kappa G^2 T. \quad (54)$$

As  $\beta > 0$ , and the fact that  $1 - \beta^2 = 24\gamma L$ , we have

$$\frac{1 - \beta}{1 + \beta} = \frac{1 - \beta^2}{(1 + \beta)^2} \leq 1 - \beta^2 = 24\gamma L.$$

Substituting the above in (54), we obtain that

$$\sum_{t=1}^T \mathbb{E} \left[ \|\xi_t\|^2 \right] \leq (24 \times 6) \sigma^2 \kappa \gamma L T + 3\kappa G^2 T = 144\sigma^2 \kappa \gamma L T + 3\kappa G^2 T.$$

Substituting from above in (53) we obtain that

$$\sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq \frac{4(V_1 - V_{T+1})}{\gamma} + (12 \times 144) \sigma^2 \kappa \gamma L T + (12 \times 3) \kappa G^2 T + \frac{288\gamma L \sigma^2}{(n-f)}.$$

Recall that

$$a_1 = (12 \times 3) = 36, a_3 = (12 \times 144)L = 1728L, \text{ and } a_4 = 288L$$



Thus, from above we obtain that

$$\sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq \frac{4(V_1 - V_{T+1})}{\gamma} + a_3 \kappa \sigma^2 \gamma T + \frac{a_4 \sigma^2}{(n-f)} \gamma T + a_1 \kappa G^2 T.$$

Diving both sides by  $T$  we obtain that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq \frac{4(V_1 - V_{T+1})}{\gamma T} + a_3 \kappa \sigma^2 \gamma + \frac{a_4 \sigma^2}{(n-f)} \gamma + a_1 \kappa G^2. \quad (55)$$

**Analyzing  $V_t$ .** Recall that  $\mathcal{L}^* = \inf_{\theta \in \mathbb{R}^d} \mathcal{L}_{\mathcal{H}}(\theta)$ . Note that for an arbitrary  $t$ , by definition of  $V_t$  in (40),

$$V_t - 2\mathcal{L}^* = 2\mathbb{E} [\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) - \mathcal{L}^*] + z \mathbb{E} [\|\delta_t\|^2] \geq 0 + z \mathbb{E} [\|\delta_t\|^2] \geq 0.$$

Thus,

$$V_1 - V_{T+1} = V_1 - 2\mathcal{L}^* - (V_{T+1} - 2\mathcal{L}^*) \leq V_1 - 2\mathcal{L}^*. \quad (56)$$

Moreover,

$$V_1 = 2\mathcal{L}_{\mathcal{H}}(\theta_0) + z \mathbb{E} [\|\delta_1\|^2]. \quad (57)$$

By definition of  $\delta_t$  in (35), the definition of  $\bar{m}_t$  in (33), and the fact that  $m_0^{(i)} = 0$  for all  $i \in \mathcal{H}$ , we obtain that

$$\mathbb{E} [\|\delta_1\|^2] = \mathbb{E} [\|\bar{m}_1 - \nabla \mathcal{L}_{\mathcal{H}}(\theta_0)\|^2] = \mathbb{E} [\|(1-\beta)\bar{g}_1 - \nabla \mathcal{L}_{\mathcal{H}}(\theta_0)\|^2]$$

where  $\bar{g}_1$ , defined in (59), is the average of  $n-f$  honest workers' stochastic gradients in step 1. Expanding the R.H.S. above we obtain that

$$\mathbb{E} [\|\delta_1\|^2] = (1-\beta)^2 \mathbb{E} [\|\bar{g}_1 - \nabla \mathcal{L}_{\mathcal{H}}(\theta_0)\|^2] + \beta^2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_0)\|^2 - 2\beta(1-\beta) \langle \mathbb{E} [\bar{g}_1] - \nabla \mathcal{L}_{\mathcal{H}}(\theta_0), \nabla \mathcal{L}_{\mathcal{H}}(\theta_0) \rangle.$$

Recall that  $\mathbb{E} [\bar{g}_1] = \nabla \mathcal{L}_{\mathcal{H}}(\theta_0)$ , and that (due to Assumption 2)  $\mathbb{E} [\|\bar{g}_1 - \nabla \mathcal{L}_{\mathcal{H}}(\theta_0)\|^2] \leq \sigma^2/(n-f)$ . Therefore,

$$\mathbb{E} [\|\delta_1\|^2] \leq \frac{(1-\beta)^2 \sigma^2}{(n-f)} + \beta^2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_0)\|^2.$$

Recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Thus,  $\|\nabla \mathcal{L}_{\mathcal{H}}(\theta_0)\|^2 \leq 2L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)$  (see (Nesterov et al., 2018), Theorem 2.1.5). Therefore,

$$\mathbb{E} [\|\delta_1\|^2] \leq \frac{(1-\beta)^2 \sigma^2}{(n-f)} + 2\beta^2 L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*).$$

Substituting from above in (57) we obtain that

$$V_1 \leq 2\mathcal{L}_{\mathcal{H}}(\theta_0) + z \left( \frac{(1-\beta)^2 \sigma^2}{(n-f)} + 2\beta^2 L(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*) \right).$$

Recall that  $(1-\beta)^2 \leq (1-\beta^2)^2 = 576\gamma^2 L^2$ . Using this, and the facts that  $\beta^2 < 1$  and  $z = \frac{1}{8L}$ , we obtain that

$$V_1 \leq 2\mathcal{L}_{\mathcal{H}}(\theta_0) + \frac{1}{4}(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*) + \frac{72\gamma^2 L^2 \sigma^2}{L(n-f)}.$$

Recall that  $a_4 = 288L$ . Therefore,

$$V_1 \leq 2\mathcal{L}_{\mathcal{H}}(\theta_0) + \frac{1}{4}(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*) + \frac{a_4 \sigma^2}{4(n-f)} \gamma^2.$$

Substituting the above in (56) we obtain that

$$V_1 - V_{T+1} \leq 2\mathcal{L}_{\mathcal{H}}(\theta_0) - 2\mathcal{L}^* + \frac{1}{4}(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*) + \frac{a_4\sigma^2}{4(n-f)}\gamma^2 = \frac{9}{4}(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*) + \frac{a_4\sigma^2}{4(n-f)}\gamma^2.$$

Substituting from above in (55) we obtain that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq \frac{(4 \times \frac{9}{4})(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{\gamma T} + \left( \frac{a_4\sigma^2}{n-f} \right) \frac{\gamma}{T} + a_3\kappa\sigma^2\gamma + \frac{a_4\sigma^2}{(n-f)}\gamma + a_1\kappa G^2.$$

Upon re-arranging the terms on R.H.S. above we obtain that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq a_1\kappa G^2 + \frac{9(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)}{\gamma T} + \left( a_3\kappa + \frac{a_4}{n-f} \right) \sigma^2\gamma + \left( \frac{a_4\sigma^2}{n-f} \right) \frac{\gamma}{T}.$$

Recall that  $a_2^2 = 36(\mathcal{L}_{\mathcal{H}}(\theta_0) - \mathcal{L}^*)$  and  $a_\kappa^2 := a_3\kappa + \frac{a_4}{n-f}$ , we obtain that

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq a_1\kappa G^2 + \frac{a_2^2}{4\gamma T} + a_\kappa^2\sigma^2\gamma + \left( \frac{a_4\sigma^2}{n-f} \right) \frac{\gamma}{T} \quad (58)$$

**Final step.** Recall that by definition

$$\gamma = \min \left\{ \frac{1}{24L}, \frac{a_2}{2a_\kappa\sigma\sqrt{T}} \right\},$$

and thus  $\frac{1}{\gamma} = \max \left\{ 24L, \frac{2a_\kappa\sigma\sqrt{T}}{a_2} \right\} \leq 24L + \frac{2a_\kappa\sigma\sqrt{T}}{a_2}$ .

Upon substituting this value of  $\gamma$  in (58), and recalling that  $a_5 = 6La_2^2$ , we obtain that

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] &\leq a_1\kappa G^2 + \frac{a_2 a_\kappa \sigma}{\sqrt{T}} + \frac{24La_2^2}{4T} + \frac{a_2 a_4 \sigma^2}{2(n-f)a_\kappa T^{3/2}} \\ &\leq a_1\kappa G^2 + \frac{a_2 a_\kappa \sigma}{\sqrt{T}} + \frac{a_5}{T} + \frac{a_2 a_4 \sigma}{na_\kappa T^{3/2}}. \end{aligned} \quad (\text{since } n \geq 2f)$$

Finally, recall from Algorithm 3 that  $\hat{\theta}$  is chosen randomly from the set of computed parameter vectors  $(\theta_0, \dots, \theta_{T-1})$ .

Thus,  $\mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta})\|^2 \right] = \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right]$ . Substituting this above proves the theorem.  $\square$

### 13.3 Proof of Corollary 2: Analysis of Robust D-SHB with NNM

**Corollary 2.** *Let Assumption 1 hold and recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 3 with aggregation  $F \circ \text{NNM}$ , under the same setting as Theorem 2. If  $F$  is  $(f, \kappa)$ -robust with  $\kappa = \mathcal{O}(1)$ , then we have*

$$\mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \right] = \mathcal{O}(f/nG^2 + 1/\sqrt{T}),$$

where the expectation is over the randomness of the algorithm.

*Proof.* Let Assumption 1 hold. Assume  $\mathcal{L}_{\mathcal{H}}$  to be  $L$ -smooth and  $F$  to be  $(f, \kappa)$ -robust. First recall that, by Lemma 1, the composition  $F \circ \text{NNM}$  is  $(f, \kappa')$ -robust with  $\kappa' = \frac{8f}{n-f}(\kappa + 1)$ .

Consider Algorithm 1 with aggregation rule  $F \circ \text{NNM}$ , and learning rate and momentum coefficient set as in Theorem 2. Following Theorem 2, and recalling the constants defined in (37), we have for every  $T \geq 1$ ,

$$\begin{aligned} \mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \right] &\leq a_1 \kappa' G^2 + \frac{a_2 a_{\kappa'} \sigma}{\sqrt{T}} + \frac{a_5}{T} + \frac{a_2 a_4 \sigma}{n a_{\kappa'} T^{3/2}} \\ &= \frac{8a_1 f}{n-f} (\kappa + 1) G^2 + \frac{a_2 a_{\kappa'} \sigma}{\sqrt{T}} + \frac{a_5}{T} + \frac{a_2 a_4 \sigma}{n a_{\kappa'} T^{3/2}}. \end{aligned}$$

Recall that  $a_{\kappa'}^2 = a_3 \kappa' + \frac{a_4}{n-f} = \frac{8a_3 f}{n-f} (\kappa + 1) + \frac{a_4}{n-f}$ . As  $\kappa \geq 0$  and  $\kappa = \mathcal{O}(1)$  by assumption, we have  $a_{\kappa'} = \Theta(1)$ . Now, ignoring constants and the last two terms on the RHS above (since they are dominated by  $\frac{1}{\sqrt{T}}$ ), we obtain

$$\mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \right] = \mathcal{O} \left( \frac{f}{n-f} (\kappa + 1) G^2 + \frac{1}{\sqrt{T}} \right).$$

Recall that, as  $f < n/2$ , we have  $\frac{f}{n-f} \leq \frac{2f}{n}$ . Thus, if  $\kappa = \mathcal{O}(1)$ , we can write

$$\mathbb{E} \left[ \left\| \nabla \mathcal{L}_{\mathcal{H}}(\hat{\theta}) \right\|^2 \right] = \mathcal{O}(f/nG^2 + 1/\sqrt{T}).$$

This concludes the proof. □

### 13.4 Proof of Supporting Lemmas

#### 13.4.1 Proof of Lemma 6

We now recall Lemma 6 below, and present its proof.

**Lemma 6.** *Suppose that assumptions 2 and 1 hold true. Consider Algorithm 3. For each  $t \in [T]$ , we have*

$$\mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| m_t^{(i)} - \bar{m}_t \right\|^2 \right] \leq 3 \left( \frac{1-\beta}{1+\beta} \right) \left( 1 + \frac{1}{n-f} \right) \sigma^2 + 3G^2.$$

*Proof.* Recall from (33) that

$$\bar{m}_t := \frac{1}{(n-f)} \sum_{i \in \mathcal{H}} m_t^{(i)}.$$

We consider an arbitrary  $i \in \mathcal{H}$ . For simplicity we define

$$\bar{g}_t := \frac{1}{(n-f)} \sum_{j \in \mathcal{H}} g_t^{(j)}. \quad (59)$$

Now, we consider an arbitrary step  $t \in [T]$ . Expanding the sum in (30) we obtain that

$$m_t^{(i)} = (1-\beta) \sum_{k=1}^t \beta^{t-k} g_k^{(i)}.$$

Therefore, applying Jensen's inequality, we write

$$\begin{aligned} \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| m_t^{(i)} - \bar{m}_t \right\|^2 \right] &= (1-\beta)^2 \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| \sum_{k=1}^t \beta^{t-k} (g_k^{(i)} - \bar{g}_t) \right\|^2 \right] \\ &\leq 3(1-\beta)^2 \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| \sum_{k=1}^t \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})) \right\|^2 \right] \\ &\quad + 3(1-\beta)^2 \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| \sum_{k=1}^t \beta^{t-k} (\bar{g}_k - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{k-1})) \right\|^2 \right] \\ &\quad + 3(1-\beta)^2 \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| \sum_{k=1}^t \beta^{t-k} (\nabla \mathcal{L}_i(\theta_{k-1}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{k-1})) \right\|^2 \right]. \quad (60) \end{aligned}$$

We obtain below upper bounds for each of the three terms on the right-hand side of (60).

For the first term on the R.H.S in (60) we denote

$$A_t := \mathbb{E} \left[ \left\| \sum_{k=1}^t \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})) \right\|^2 \right].$$

We obtain that

$$\begin{aligned} A_t &= \mathbb{E} \left[ \left\| \sum_{k=1}^{t-1} \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})) + (g_t^{(i)} - \nabla \mathcal{L}_i(\theta_{t-1})) \right\|^2 \right] \\ &= \mathbb{E} \left[ \left\| \sum_{k=1}^{t-1} \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})) \right\|^2 \right] + \mathbb{E} \left[ \left\| g_t^{(i)} - \nabla \mathcal{L}_i(\theta_{t-1}) \right\|^2 \right] \\ &\quad + 2 \mathbb{E} \left[ \left\langle \sum_{k=1}^{t-1} \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})), g_t^{(i)} - \nabla \mathcal{L}_i(\theta_{t-1}) \right\rangle \right] \end{aligned}$$

The last term on the right-hand side is zero by the total law of expectation, and the unbiasedness of stochastic gradients (Assumption 2):

$$\begin{aligned} \mathbb{E} \left[ \left\langle \sum_{k=1}^{t-1} \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})), g_t^{(i)} - \nabla \mathcal{L}_i(\theta_{t-1}) \right\rangle \right] &= \mathbb{E} \left[ \mathbb{E}_t \left[ \left\langle \sum_{k=1}^{t-1} \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})), g_t^{(i)} - \nabla \mathcal{L}_i(\theta_{t-1}) \right\rangle \right] \right] \\ &= \mathbb{E} \left[ \left\langle \sum_{k=1}^{t-1} \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})), \underbrace{\mathbb{E}_t [g_t^{(i)} - \nabla \mathcal{L}_i(\theta_{t-1})]}_{=0} \right\rangle \right] \\ &= 0. \end{aligned}$$

By Assumption 2, we have  $\mathbb{E} \left[ \|g_t^{(i)} - \nabla \mathcal{L}_i(\theta_{t-1})\|^2 \right] \leq \sigma^2$ . Thus, we have

$$A_t \leq \mathbb{E} \left[ \left\| \sum_{k=1}^{t-1} \beta^{t-k} (g_k^{(i)} - \nabla \mathcal{L}_i(\theta_{k-1})) \right\|^2 \right] + \sigma^2 = \beta^2 A_{t-1} + \sigma^2.$$

By recursion, we obtain

$$A_t \leq \beta^{2(t-1)} A_1 + \sigma^2 \sum_{l=0}^{t-2} \beta^{2l}.$$

As  $A_1 = \mathbb{E} \left[ \|g_1^{(i)} - \nabla \mathcal{L}_i(\theta_0)\|^2 \right] \leq \sigma^2$ , from above we obtain that

$$A_t \leq \sigma^2 \sum_{l=0}^{t-1} \beta^{2l} \leq \frac{\sigma^2}{1 - \beta^2}. \quad (61)$$

For the second term on the R.H.S in (60), we denote

$$B_t := \mathbb{E} \left[ \left\| \sum_{k=1}^t \beta^{t-k} (\bar{g}_k - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{k-1})) \right\|^2 \right].$$

By Assumption 2 and the mutual independence of stochastic gradients we have

$$\mathbb{E} \left[ \|\bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] = \frac{1}{(n-f)^2} \sum_{i \in \mathcal{H}} \mathbb{E} \left[ \|g_t^{(i)} - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq \frac{\sigma^2}{n-f}.$$

Thus, similar to the bound on  $A_t$ , we obtain that

$$B_t \leq \frac{\sigma^2}{(1 - \beta^2)(n-f)}. \quad (62)$$

For the third term on the R.H.S in (60), we denote

$$C_t := \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| \sum_{k=1}^t \beta^{t-k} (\nabla \mathcal{L}_i(\theta_{k-1}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{k-1})) \right\|^2 \right].$$

By applying Jensen's inequality, followed by Assumption 1, we have

$$\begin{aligned}
 C_t &\leq \left( \sum_{k=1}^t \beta^{t-k} \right) \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \sum_{k=1}^t \beta^{t-k} \|\nabla \mathcal{L}_i(\theta_{k-1}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{k-1})\|^2 \right] \\
 &= \left( \sum_{k=1}^t \beta^{t-k} \right) \mathbb{E} \left[ \sum_{k=1}^t \beta^{t-k} \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \|\nabla \mathcal{L}_i(\theta_{k-1}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{k-1})\|^2 \right] \\
 &\leq \left( \sum_{k=1}^t \beta^{t-k} \right)^2 \cdot G^2 \leq \frac{G^2}{(1-\beta)^2}.
 \end{aligned} \tag{63}$$

Finally, substituting from (61), (62) and (63) in (60) we obtain that

$$\begin{aligned}
 \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| m_t^{(i)} - \bar{m}_t \right\|^2 \right] &\leq 3(1-\beta)^2 \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} (A_t + B_t + C_t) \\
 &\leq 3(1-\beta)^2 \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left( \frac{\sigma^2}{1-\beta^2} + \frac{\sigma^2}{(1-\beta^2)(n-f)} + \frac{G^2}{(1-\beta)^2} \right) \\
 &= 3 \frac{1-\beta}{1+\beta} \left( 1 + \frac{1}{n-f} \right) \sigma^2 + 3G^2.
 \end{aligned}$$

The above proves the lemma.  $\square$

### 13.4.2 Proof of Lemma 7

**Lemma 7.** *Suppose that assumptions 1 and 2 hold true. Assume  $F$  is  $(f, \kappa)$ -robust. Consider Algorithm 3 with aggregation  $F$ . For each step  $t \in [T]$ , we obtain that*

$$\mathbb{E} \left[ \|\xi_t\|^2 \right] \leq 6\kappa \frac{1-\beta}{1+\beta} \sigma^2 + 48\kappa G^2.$$

*Proof.* Recall from (31) and (34), respectively, that

$$m_t := F \left( m_t^{(1)}, \dots, m_t^{(n)} \right) \quad \text{and} \quad \xi_t := m_t - \bar{m}_t.$$

We consider an arbitrary step  $t$ . Since  $F$  is  $(f, \kappa)$ -robust, we obtain that

$$\|\xi_t\|^2 = \|m_t - \bar{m}_t\|^2 \leq \frac{\kappa}{n-f} \sum_{i \in \mathcal{H}} \left\| m_t^{(i)} - \bar{m}_t \right\|^2. \tag{64}$$

Upon taking total expectations on both sides we obtain that

$$\mathbb{E} \left[ \|\xi_t\|^2 \right] \leq \frac{\kappa}{n-f} \sum_{i \in \mathcal{H}} \mathbb{E} \left[ \left\| m_t^{(i)} - \bar{m}_t \right\|^2 \right]. \tag{65}$$

From Lemma 6, under Assumption 2, we have

$$\begin{aligned}
 \mathbb{E} \left[ \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} \left\| m_t^{(i)} - \bar{m}_t \right\|^2 \right] &\leq 3 \frac{1-\beta}{1+\beta} \left( 1 + \frac{1}{n-f} \right) \sigma^2 + 3G^2 \\
 &\leq 6 \frac{1-\beta}{1+\beta} \sigma^2 + 3G^2. \quad (n-f \geq 1)
 \end{aligned}$$

Substituting from above in (65) proves the lemma, i.e., we conclude that

$$\mathbb{E} \left[ \|\xi_t\|^2 \right] \leq 6\kappa \frac{1-\beta}{1+\beta} \sigma^2 + 3\kappa G^2.$$

This concludes the proof.  $\square$

### 13.4.3 Proof of Lemma 8

The proof of Lemma 8 is similar to that of Lemma 3 in (Farhadkhani et al., 2022b). We recall the lemma and the proof below for completeness.

**Lemma 8.** *Suppose that Assumption 2 holds. Recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 3 with  $T > 1$ . For all  $t \geq 2$  we obtain that*

$$\begin{aligned} \mathbb{E} \left[ \|\delta_t\|^2 \right] &\leq \beta^2 c \mathbb{E} \left[ \|\delta_{t-1}\|^2 \right] + 4\gamma L(1 + \gamma L)\beta^2 \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})\|^2 \right] + (1 - \beta)^2 \frac{\sigma^2}{(n - f)} \\ &\quad + 2\gamma L(1 + \gamma L)\beta^2 \mathbb{E} \left[ \|\xi_{t-1}\|^2 \right], \end{aligned}$$

where  $c := (1 + \gamma L)(1 + 4\gamma L)$ .

*Proof.* Recall from (35) that

$$\delta_t := \bar{m}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}).$$

Consider an arbitrary step  $t > 1$ . Substituting from (30) and (33) we obtain that

$$\delta_t = \beta \bar{m}_{t-1} + (1 - \beta) \bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}).$$

Upon adding and subtracting  $\beta \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})$  and  $\beta \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})$  on the R.H.S. above we obtain that

$$\begin{aligned} \delta_t &= \beta \bar{m}_{t-1} - \beta \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) + (1 - \beta) \bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_t) + \beta \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) + \beta \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \beta \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \\ &= \beta (\bar{m}_{t-1} - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})) + (1 - \beta) \bar{g}_t - (1 - \beta) \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) + \beta (\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})). \end{aligned}$$

As  $\bar{m}_{t-1} - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) = \delta_{t-1}$  (by (35)), from above we obtain that

$$\delta_t = \beta \delta_{t-1} + (1 - \beta) (\bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})) + \beta (\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})).$$

Therefore,

$$\begin{aligned} \|\delta_t\|^2 &= \beta^2 \|\delta_{t-1}\|^2 + (1 - \beta)^2 \|\bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + \beta^2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2\beta(1 - \beta) \langle \delta_{t-1}, \bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle \\ &\quad + 2\beta^2 \langle \delta_{t-1}, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle + 2\beta(1 - \beta) \langle \bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle. \end{aligned}$$

By taking conditional expectation  $\mathbb{E}_t[\cdot]$  on both sides, and recalling that  $\delta_{t-1}$ ,  $\theta_{t-1}$  and  $\theta_{t-2}$  are deterministic values when the history  $\mathcal{P}_t$  is given, we obtain that

$$\begin{aligned} \mathbb{E}_t \left[ \|\delta_t\|^2 \right] &= \beta^2 \|\delta_{t-1}\|^2 + (1 - \beta)^2 \mathbb{E}_t \left[ \|\bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + \beta^2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + \\ &\quad 2\beta(1 - \beta) \langle \delta_{t-1}, \mathbb{E}_t[\bar{g}_t] - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle + 2\beta^2 \langle \delta_{t-1}, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle \\ &\quad + 2\beta(1 - \beta) \langle \mathbb{E}_t[\bar{g}_t] - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle. \end{aligned}$$

Recall that  $\bar{g}_t := 1/(n-f) \sum_{j \in \mathcal{H}} g_t^{(j)}$ . Thus, we have  $\mathbb{E}_t[\bar{g}_t] = \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})$ . Using this above we obtain that

$$\begin{aligned} \mathbb{E}_t \left[ \|\delta_t\|^2 \right] &= \beta^2 \|\delta_{t-1}\|^2 + (1 - \beta)^2 \mathbb{E}_t \left[ \|\bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + \beta^2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \\ &\quad + 2\beta^2 \langle \delta_{t-1}, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle. \end{aligned}$$

Also, by Assumption 2 and the fact that  $g_t^{(j)}$ 's for  $j \in \mathcal{H}$  are independent of each other, we obtain that  $\mathbb{E}_t \left[ \|\bar{g}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] \leq \frac{\sigma^2}{(n-f)}$ . Thus,

$$\mathbb{E}_t \left[ \|\delta_t\|^2 \right] \leq \beta^2 \|\delta_{t-1}\|^2 + (1 - \beta)^2 \frac{\sigma^2}{(n-f)} + \beta^2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2\beta^2 \langle \delta_{t-1}, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle.$$

By the Cauchy-Schwartz inequality,  $\langle \delta_{t-1}, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle \leq \|\delta_{t-1}\| \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|$ . Since  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth, we have  $\|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\| \leq L \|\theta_{t-1} - \theta_{t-2}\|$ . Recall from (32) that  $\theta_t = \theta_{t-1} - \gamma m_t$ . Thus,  $\|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2}) - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\| \leq \gamma L \|m_{t-1}\|$ . Using this above we obtain that

$$\mathbb{E}_t \left[ \|\delta_t\|^2 \right] \leq \beta^2 \|\delta_{t-1}\|^2 + (1 - \beta)^2 \frac{\sigma^2}{(n - f)} + \gamma^2 \beta^2 L^2 \|m_{t-1}\|^2 + 2\gamma\beta^2 L \|\delta_{t-1}\| \|m_{t-1}\|.$$

As  $2ab \leq a^2 + b^2$ , from above we obtain that

$$\begin{aligned} \mathbb{E}_t \left[ \|\delta_t\|^2 \right] &\leq \beta^2 \|\delta_{t-1}\|^2 + (1 - \beta)^2 \frac{\sigma^2}{(n - f)} + \gamma^2 \beta^2 L^2 \|m_{t-1}\|^2 + \gamma L \beta^2 \left( \|\delta_{t-1}\|^2 + \|m_{t-1}\|^2 \right) \\ &= (1 + \gamma L) \beta^2 \|\delta_{t-1}\|^2 + (1 - \beta)^2 \frac{\sigma^2}{(n - f)} + \gamma L (1 + \gamma L) \beta^2 \|m_{t-1}\|^2. \end{aligned} \quad (66)$$

By definition of  $\xi_t$  in (34), we have  $m_{t-1} = \xi_{t-1} + \bar{m}_{t-1}$ . Thus, owing to the triangle inequality and the fact that  $2ab \leq a^2 + b^2$ , we have  $\|m_{t-1}\|^2 \leq 2\|\xi_{t-1}\|^2 + 2\|\bar{m}_{t-1}\|^2$ . Similarly, by definition of  $\delta_t$  in (35), we have  $\|\bar{m}_{t-1}\|^2 \leq 2\|\delta_{t-1}\|^2 + 2\|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})\|^2$ . Thus,  $\|m_{t-1}\|^2 \leq 2\|\xi_{t-1}\|^2 + 4\|\delta_{t-1}\|^2 + 4\|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})\|^2$ . Using this in (66) we obtain that

$$\begin{aligned} \mathbb{E}_t \left[ \|\delta_t\|^2 \right] &\leq (1 + \gamma L) \beta^2 \|\delta_{t-1}\|^2 + (1 - \beta)^2 \frac{\sigma^2}{(n - f)} \\ &\quad + 2\gamma L (1 + \gamma L) \beta^2 \left( \|\xi_{t-1}\|^2 + 2\|\delta_{t-1}\|^2 + 2\|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})\|^2 \right). \end{aligned}$$

By re-arranging the terms on the R.H.S. we get

$$\begin{aligned} \mathbb{E}_t \left[ \|\delta_t\|^2 \right] &\leq \beta^2 (1 + \gamma L) (1 + 4\gamma L) \|\delta_{t-1}\|^2 + 4\gamma L (1 + \gamma L) \beta^2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})\|^2 + (1 - \beta)^2 \frac{\sigma^2}{(n - f)} \\ &\quad + 2\gamma L (1 + \gamma L) \beta^2 \|\xi_{t-1}\|^2. \end{aligned}$$

Substituting  $c = (1 + \gamma L) (1 + 4\gamma L)$  above we obtain that

$$\mathbb{E}_t \left[ \|\delta_t\|^2 \right] \leq \beta^2 c \|\delta_{t-1}\|^2 + 4\gamma L (1 + \gamma L) \beta^2 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-2})\|^2 + (1 - \beta)^2 \frac{\sigma^2}{(n - f)} + 2\gamma L (1 + \gamma L) \beta^2 \|\xi_{t-1}\|^2.$$

Recall that  $t$  in the above is an arbitrary value in  $[T]$  greater than 1. Hence, upon taking total expectation on both sides above proves the lemma.  $\square$

#### 13.4.4 Proof of Lemma 9

The proof of Lemma 9 is similar to that of Lemma 4 in (Farhadkhani et al., 2022b). We recall the lemma and the proof below for completeness.

**Lemma 9.** *Recall that  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth. Consider Algorithm 3. For all  $t \in [T]$ , we obtain that*

$$\begin{aligned} \mathbb{E} \left[ 2\mathcal{L}_{\mathcal{H}}(\theta_t) - 2\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \right] &\leq -\gamma (1 - 4\gamma L) \mathbb{E} \left[ \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right] + 2\gamma (1 + 2\gamma L) \mathbb{E} \left[ \|\delta_t\|^2 \right] \\ &\quad + 2\gamma (1 + \gamma L) \mathbb{E} \left[ \|\xi_t\|^2 \right]. \end{aligned}$$

*Proof.* Consider an arbitrary step  $t$ . Since  $\mathcal{L}_{\mathcal{H}}$  is  $L$ -smooth, we have (see (Bottou et al., 2018))

$$\mathcal{L}_{\mathcal{H}}(\theta_t) - \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \leq \langle \theta_t - \theta_{t-1}, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle + \frac{L}{2} \|\theta_t - \theta_{t-1}\|^2.$$

Substituting from (36), i.e.,  $\theta_t = \theta_{t-1} - \gamma \bar{m}_t - \gamma \xi_t$ , we obtain that

$$\begin{aligned} \mathcal{L}_{\mathcal{H}}(\theta_t) - \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) &\leq -\gamma \langle \bar{m}_t, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle - \gamma \langle \xi_t, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle + \gamma^2 \frac{L}{2} \|\bar{m}_t + \xi_t\|^2 \\ &= -\gamma \langle \bar{m}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) + \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}), \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle - \gamma \langle \xi_t, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle + \gamma^2 \frac{L}{2} \|\bar{m}_t + \xi_t\|^2. \end{aligned}$$



By Definition (35),  $\bar{m}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) = \delta_t$ . Thus, from above we obtain that (scaling by factor of 2)

$$2\mathcal{L}_{\mathcal{H}}(\theta_t) - 2\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \leq -2\gamma \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 - 2\gamma \langle \delta_t, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle - 2\gamma \langle \xi_t, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle + \gamma^2 L \|\bar{m}_t + \xi_t\|^2. \quad (67)$$

Now, we consider the last three terms on the R.H.S. separately. Using Cauchy-Schwartz inequality, and the fact that  $2ab \leq \frac{1}{c}a^2 + cb^2$  for any  $c > 0$ , we obtain that (by substituting  $c = 2$ )

$$2 |\langle \delta_t, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle| \leq 2 \|\delta_t\| \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\| \leq \frac{2}{1} \|\delta_t\|^2 + \frac{1}{2} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2. \quad (68)$$

Similarly,

$$2 |\langle \xi_t, \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \rangle| \leq 2 \|\xi_t\| \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\| \leq \frac{2}{1} \|\xi_t\|^2 + \frac{1}{2} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2. \quad (69)$$

Finally, using triangle inequality and the fact that  $2ab \leq a^2 + b^2$  we have

$$\begin{aligned} \|\bar{m}_t + \xi_t\|^2 &\leq 2 \|\bar{m}_t\|^2 + 2 \|\xi_t\|^2 = 2 \|\bar{m}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_t) + \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2 \|\xi_t\|^2 \\ &\leq 4 \|\delta_t\|^2 + 4 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2 \|\xi_t\|^2. \quad [\text{since } \bar{m}_t - \nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1}) = \delta_t] \end{aligned} \quad (70)$$

Substituting from (68), (69) and (70) in (67) we obtain that

$$\begin{aligned} 2\mathcal{L}_{\mathcal{H}}(\theta_t) - 2\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) &\leq -2\gamma \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + \gamma \left( 2\|\delta_t\|^2 + \frac{1}{2} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right) + \gamma \left( 2\|\xi_t\|^2 + \frac{1}{2} \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 \right) \\ &\quad + \gamma^2 L \left( 4 \|\delta_t\|^2 + 4 \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2 \|\xi_t\|^2 \right). \end{aligned}$$

Upon re-arranging the terms in the R.H.S. we obtain that

$$2\mathcal{L}_{\mathcal{H}}(\theta_t) - 2\mathcal{L}_{\mathcal{H}}(\theta_{t-1}) \leq -\gamma(1 - 4\gamma L) \|\nabla \mathcal{L}_{\mathcal{H}}(\theta_{t-1})\|^2 + 2\gamma(1 + 2\gamma L) \|\delta_t\|^2 + 2\gamma(1 + \gamma L) \|\xi_t\|^2.$$

As  $t$  is arbitrarily chosen from  $[T]$ , taking expectation on both sides above proves the lemma.  $\square$

## 14 EXPERIMENTAL SETUP

### 14.1 Detailed Experimental Setup

The architecture of the models, as well as additional details on the experimental setup, are presented in Table 3. Note that CNN stands for convolutional neural network, and NLL refers to the negative log likelihood loss. In order to present the architecture of the models used, we introduce the following compact notation.

L(#outputs) represents a **fully-connected linear layer**, R stands for **ReLU activation**, S stands for **log-softmax**, C(#channels) represents a *fully-connected 2D-convolutional layer* (kernel size 5, padding 0, stride 1), M stands for **2D-maxpool** (kernel size 2), B stands for **batch-normalization**, and D represents **dropout** (with fixed probability 0.25).

<i>Dataset</i>	MNIST & Fashion-MNIST	CIFAR-10
<i>Model type</i>	CNN	CNN
<i>Model architecture</i>	C(20)-R-M-C(20)-R-M-L(500)-R-L(10)-S	(3,32×32)-C(64)-R-B-C(64)-R-B-M-D-C(128)-R-B-C(128)-R-B-M-D-L(128)-R-D-L(10)-S
<i>Loss</i>	NLL	NLL
<i>Gradient clipping</i>	2	5
<i>ℓ<sub>2</sub>-regularization</i>	10 <sup>-4</sup>	10 <sup>-2</sup>
<i>Number of steps</i>	$T = 800$	$T = 2000$
<i>Learning rate</i>	$\gamma_t = \frac{0.75}{1 + \lfloor \frac{t}{50} \rfloor}$	$\gamma_t = \begin{cases} 0.25 & t \leq 1500 \\ 0.025 & 1500 < t \leq 2000 \end{cases}$
<i>Momentum parameter</i>	$\beta = 0.9$	$\beta = 0.9, 0.99$
<i>Batch size</i>	$b = 25$	$b = 50$
<i>Total number of workers</i>	$n = 17$	$n = 17$
<i>Number of Byzantine workers</i>	$f = 4, 6, 8$	$f = 2, 3, 4$

Table 3: Model Architectures, Hyperparameters, and Distributed Settings

On all datasets, we implement NNM and Bucketing with four aggregation rules namely geometric median (GM), coordinate-wise median (CWMed), coordinate-wise trimmed mean (CWTM), and Krum. In every setting, we execute the Bucketing algorithm with bucket size  $s = \lfloor \frac{n}{2f} \rfloor$  (Karimireddy et al., 2022). We also implement the vanilla aggregation rules. Note that when  $f = 4$  (out of  $n = 17$ ), Bucketing+CWMed and Bucketing+CWTM are exactly equivalent. In fact, when  $f = 4$ , buckets are of size  $s = 2$ . Therefore, applying Bucketing results in a total of  $n' = 9$  buckets, out of which  $f' = 4$  buckets might be potentially Byzantine (i.e., contaminated by at least one Byzantine worker). Therefore, we can clearly see that executing CWMed and CWTM with  $n' = 9$  and  $f' = 4$  (i.e., post Bucketing) outputs the same vector. Accordingly, we only show the performance of the learning under Bucketing+CWMed in the plots of Appendix 15 when  $f = 4$ .

### 14.2 Dataset Preprocessing

MNIST receives an input image normalization of mean 0.1307 and standard deviation 0.3081, while Fashion-MNIST is expanded with horizontally flipped images. Furthermore, the images of CIFAR-10 are horizontally flipped, and per channel normalization is also applied with means 0.4914, 0.4822, 0.4465 and standard deviations 0.2023, 0.1994, 0.2010.

### 14.3 Byzantine Attacks

In our experiments, the Byzantine workers execute five state-of-the-art gradient attacks, namely A Little is Enough (ALIE) (Baruch et al., 2019), Fall of Empires (FOE) (Xie et al., 2019), Sign-flipping (SF) (Allen-Zhu et al., 2020), Label-flipping (LF) (Allen-Zhu et al., 2020), and Mimic (Karimireddy et al., 2022). The first three attacks rely on the same primitive explained below.

Let  $a_t$  be the attack vector in step  $t$  and  $\eta \geq 0$  a fixed real number. In every step  $t$ , the Byzantine workers send to the server the Byzantine vector  $B_t = \bar{s}_t + \eta a_t$ , where  $\bar{s}_t$  is an estimation of the true gradient (or momentum) at step  $t$ .

Experimentally, when running the gradient descent (GD) algorithm, we set  $\bar{s}_t = \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} g_t^{(i)}$ , where  $g_t^{(i)}$  is the gradient computed by honest worker  $w_i$  in step  $t$ . However, when executing the stochastic heavy ball (SHB) method, we set  $\bar{s}_t = \frac{1}{|\mathcal{H}|} \sum_{i \in \mathcal{H}} m_t^{(i)}$ , where  $m_t^{(i)}$  is the momentum vector sent by honest worker  $w_i$  in step  $t$ .

- **ALIE:** In this attack,  $a_t = \sigma_t$ , where  $\sigma_t$  is coordinate-wise standard deviation of  $\bar{s}_t$ .
- **FOE:** In this attack,  $a_t = -\bar{s}_t$ . All Byzantine workers thus send  $(1 - \eta)\bar{s}_t$  in step  $t$ .
- **SF:** In this attack,  $a_t = -\bar{s}_t$ , and  $\eta = 2$ . All Byzantine workers thus send  $B_t = a_t = -\bar{s}_t$  in step  $t$ .

In our experiments, as done in (Shejwalkar and Houmansadr, 2021), we implement optimized versions of ALIE and FOE, where the optimal  $\eta$  is determined greedily by linearly searching over a defined range of values. In every step  $t$ , we pick the value of  $\eta$  that maximises the L2-distance between the output of the aggregation rule  $R_t$  at the server and the average of honest gradients or momentums, i.e.,  $\bar{s}_t$ . In other words, we pick the value of  $\eta$  that causes the maximum damage (in terms of L2-distance from the average of honest inputs) by the Byzantine workers.

LF and Mimic on the other hand are executed as follows:

- **LF:** Every Byzantine worker computes its gradient on flipped labels. Since the labels  $l$  for MNIST, Fashion-MNIST, and CIFAR-10 are in  $\{0, 1, \dots, 9\}$ , the Byzantine workers flip the labels by computing  $l' = 9 - l$  on the batch, where  $l'$  is the flipped/modified label.
- **Mimic:** All Byzantine workers “mimic” a certain honest worker by simply sending its gradient or momentum to the server. In order to chose which honest worker to mimic during the learning, we adopt the heuristic introduced by (Karimireddy et al., 2022).

### 14.4 Data Heterogeneity




Figure 3: Distribution of data samples across workers and class labels on MNIST/CIFAR-10 when sampling from a Dirichlet distribution of parameter  $\alpha$ . *Left:*  $\alpha = 0.1$ , *Middle:*  $\alpha = 1$ , *Right:*  $\alpha = 10$ .

In order to simulate heterogeneity in our experiments, we sample from the original datasets (i.e., MNIST, Fashion-MNIST, and CIFAR-10) using a **Dirichlet** distribution of parameter  $\alpha$ , which indicates the level of heterogeneity induced in the workers’ datasets. The smaller is  $\alpha$ , the more heterogeneous is the setting (i.e., the more likely it is that workers possess samples from only one class). We consider three heterogeneity regimes in our experiments, namely *extreme* where  $\alpha = 0.1$ , *moderate* where  $\alpha = 1$ , and *low* where  $\alpha = 10$ . The corresponding distributions of the number of samples across workers and class labels (depending on  $\alpha$ ) are shown in Figure 3.

## 15 FULL EXPERIMENTAL RESULTS

Section 15.1 presents our results on the MNIST dataset. Section 15.2 contains our results on Fashion-MNIST. Section 15.3 shows the performance of our algorithm on CIFAR-10.

### 15.1 Comprehensive Results on MNIST

In this section, we present the entirety of our results on the MNIST dataset.

#### 15.1.1 Results on D-SHB

We consider three Byzantine regimes:  $f = 4$ ,  $f = 6$ , and  $f = 8$  (largest possible) out of  $n = 17$  workers in total. We also consider three heterogeneity regimes:  $\alpha = 0.1$  (extreme),  $\alpha = 1$  (moderate), and  $\alpha = 10$  (low). We compare the performance of NNM and Bucketing when executed with four aggregation rules namely Krum (Blanchard et al., 2017), GM (Small, 1990), CWMed (Yin et al., 2018), and CWTM (Yin et al., 2018). The plots are presented below and complement our (partial) results in Table 2 in Section 6.2 of the main paper. We note that in the two strongest Byzantine settings, i.e., when  $f > 4$ , Bucketing can only be applied with bucket size  $s = 1$  (Karimireddy et al., 2022). This means that when  $f = 6, 8$ , executing Bucketing boils down to running the vanilla aggregation rules.

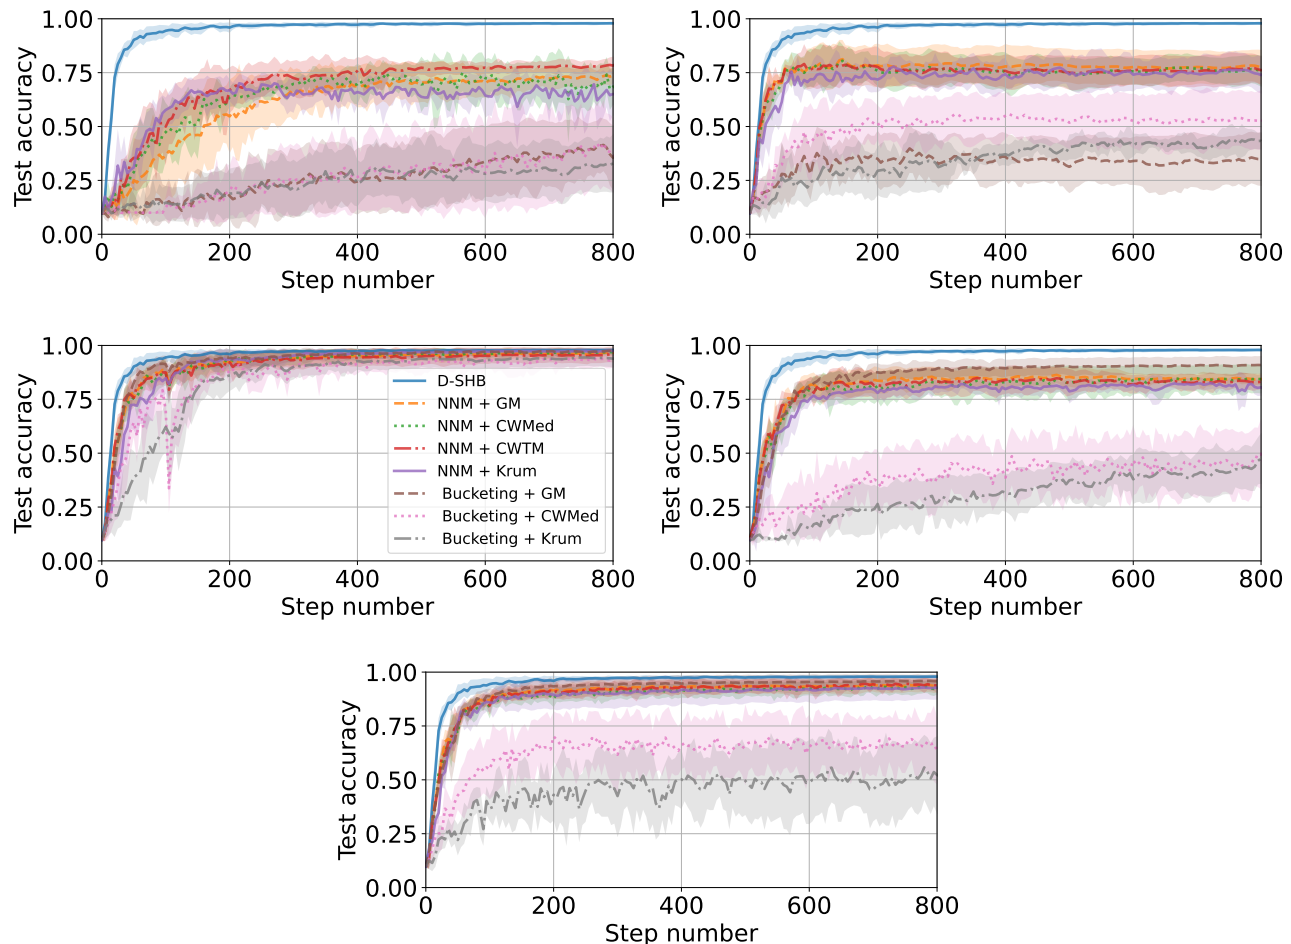


Figure 4: Experiments on MNIST using robust D-SHB with  $f = 4$  among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 0.1$ . The Byzantine workers execute FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3).

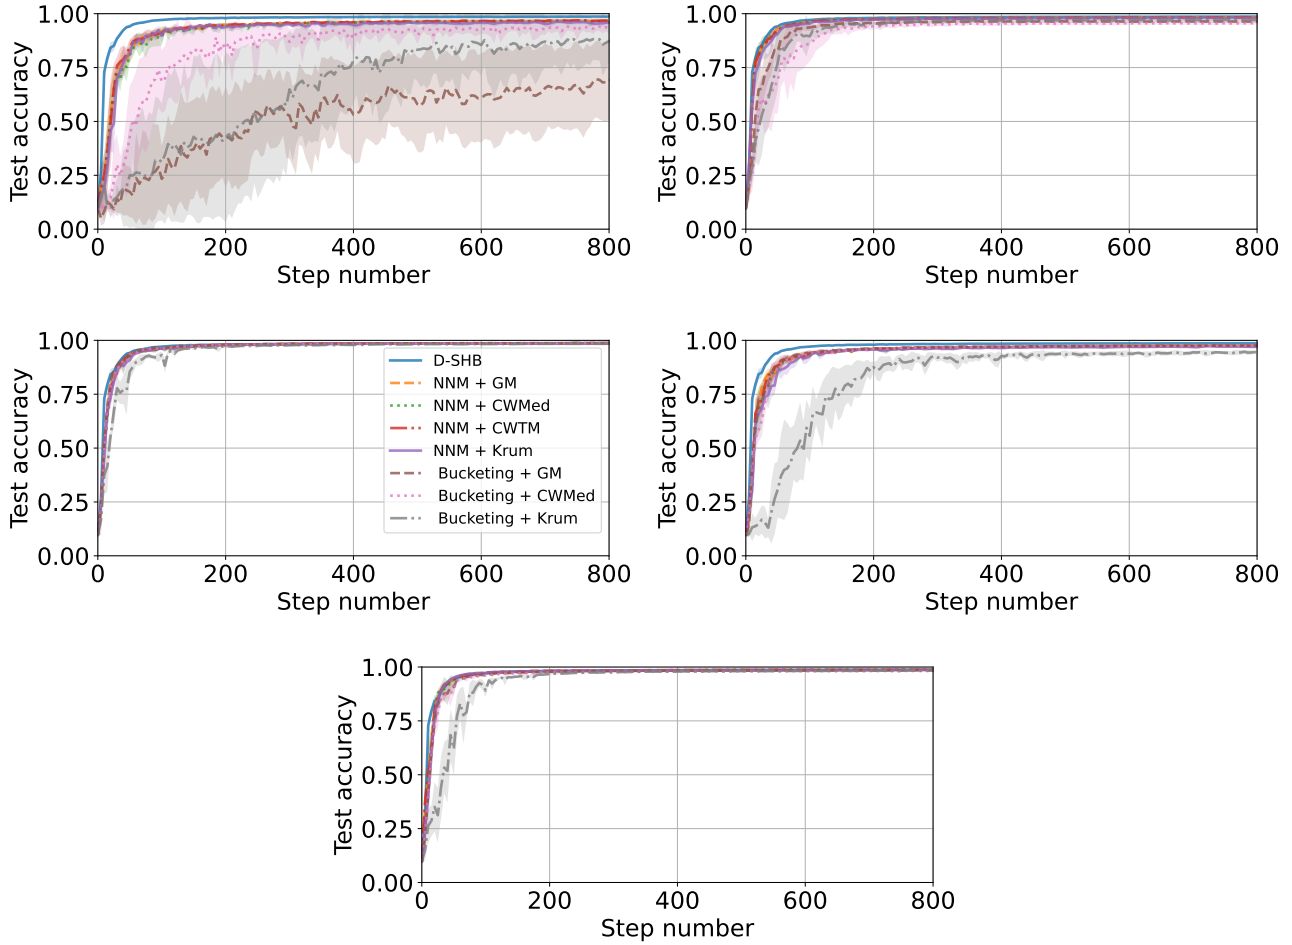


Figure 5: Experiments on MNIST using robust D-SHB with  $f = 4$  among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3).




Figure 6: Experiments on MNIST using robust D-SHB with  $f = 4$  among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3).




Figure 7: Experiments on MNIST using robust D-SHB with  $f = 6$  among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3).




Figure 8: Experiments on MNIST using robust D-SHB with  $f = 6$  among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3).



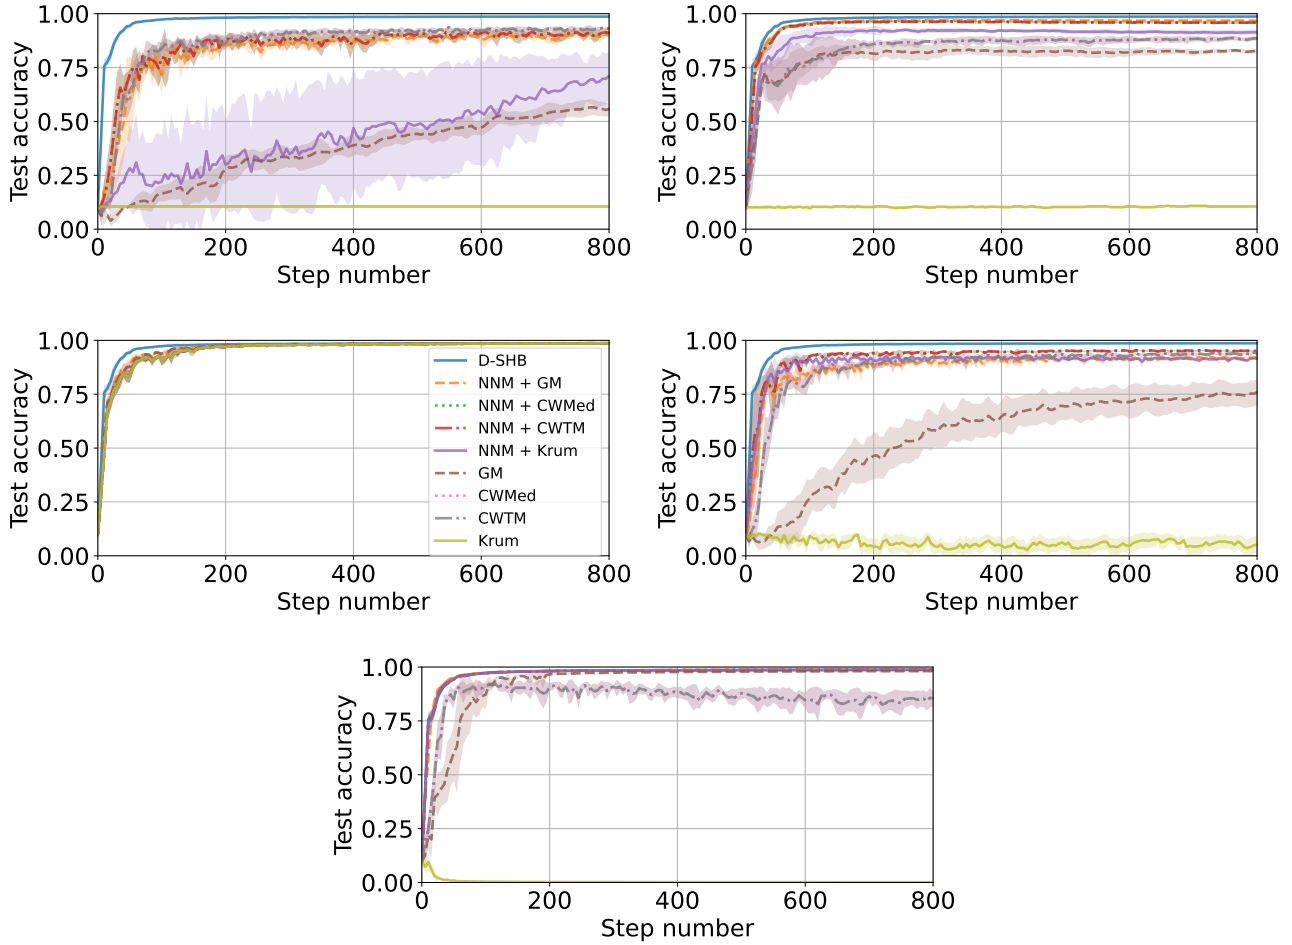


Figure 9: Experiments on MNIST using robust D-SHB with  $f = 8$  among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3).

15.1.2 Results on D-GD

We also test our algorithm with robust D-GD on MNIST in one Byzantine regime:  $f = 4$  Byzantine workers among a total of  $n = 17$  workers. We also consider three heterogeneity settings:  $\alpha = 0.1$ ,  $\alpha = 1$ , and  $\alpha = 10$ . The results are shown below, and convey the same observations made in Appendix 15.1.

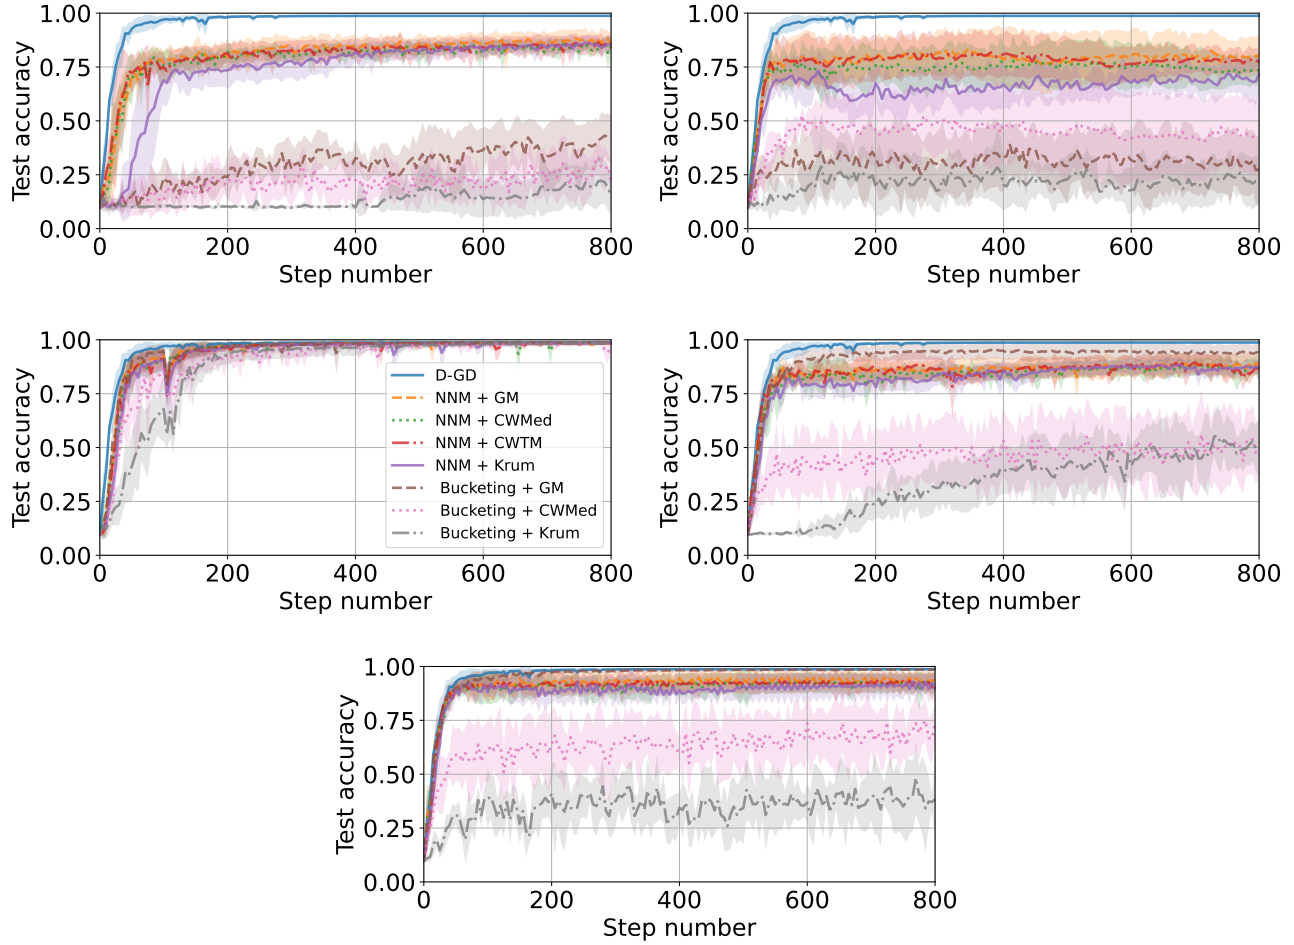


Figure 10: Experiments on MNIST using robust D-GD with  $f = 4$  among  $n = 17$  workers, with  $\alpha = 0.1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.




Figure 11: Experiments on MNIST using robust D-GD with  $f = 4$  among  $n = 17$  workers, with  $\alpha = 1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.




Figure 12: Experiments on MNIST using robust D-GD with  $f = 4$  among  $n = 17$  workers, with  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

### 15.1.3 Results on Other Aggregation Rules: MDA, MultiKrum, and MeaMed

We also execute NNM with three additional aggregation rules namely MDA (Rousseeuw, 1985; El Mhamdi et al., 2021a), MeaMed (Xie et al., 2018), and Multi-Krum (Blanchard et al., 2017). The experimental setup is the following. We consider three Byzantine regimes:  $f = 4$ ,  $f = 6$ , and  $f = 8$  (largest possible) out of  $n = 17$  workers in total. We also consider three heterogeneity regimes:  $\alpha = 0.1$  (extreme),  $\alpha = 1$  (moderate), and  $\alpha = 10$  (low). We compare the performance of NNM and Bucketing when executed with MDA, MultiKrum, and MeaMed. The results (shown below) are very similar to the ones presented in Appendix 15.1, and convey exactly the same observations made in Section 6. Indeed, NNM combined with these methods provides consistently good performance on MNIST in three regimes of data heterogeneity and under five state-of-the-art Byzantine attacks. Given that no guarantees have yet been derived for these aggregation techniques, an interesting future direction would be to prove that MDA, MultiKrum, and MeaMed are  $(f, \kappa)$ -robust with  $\kappa \in \mathcal{O}(1)$ , confirming the modularity of our solution as a reliable framework for heterogeneous Byzantine learning.




Figure 13: Experiments on MNIST using robust D-SHB with  $f = 4$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 0.1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.




Figure 14: Experiments on MNIST using robust D-SHB with  $f = 4$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

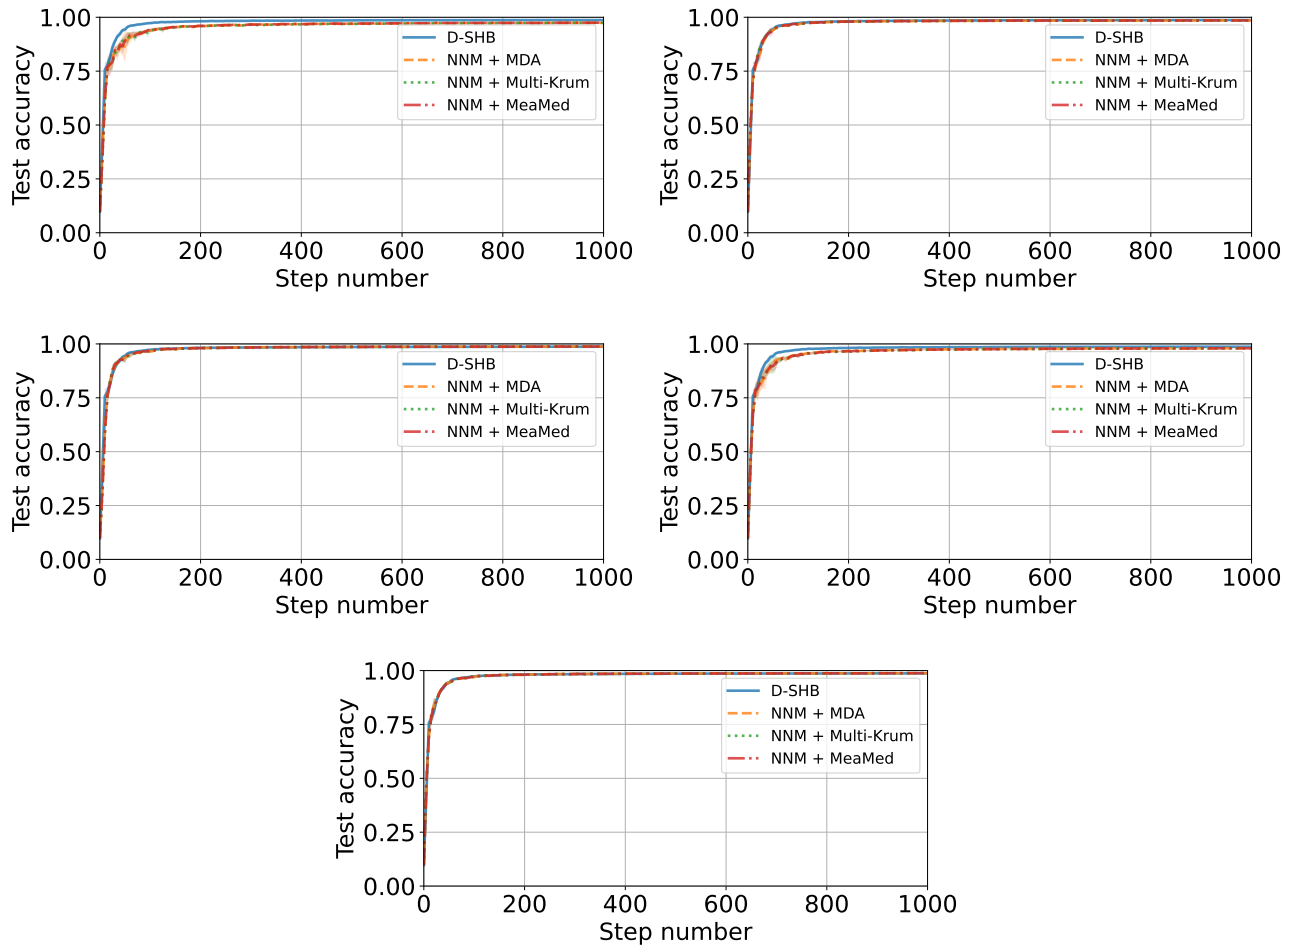


Figure 15: Experiments on MNIST using robust D-SHB with  $f = 4$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

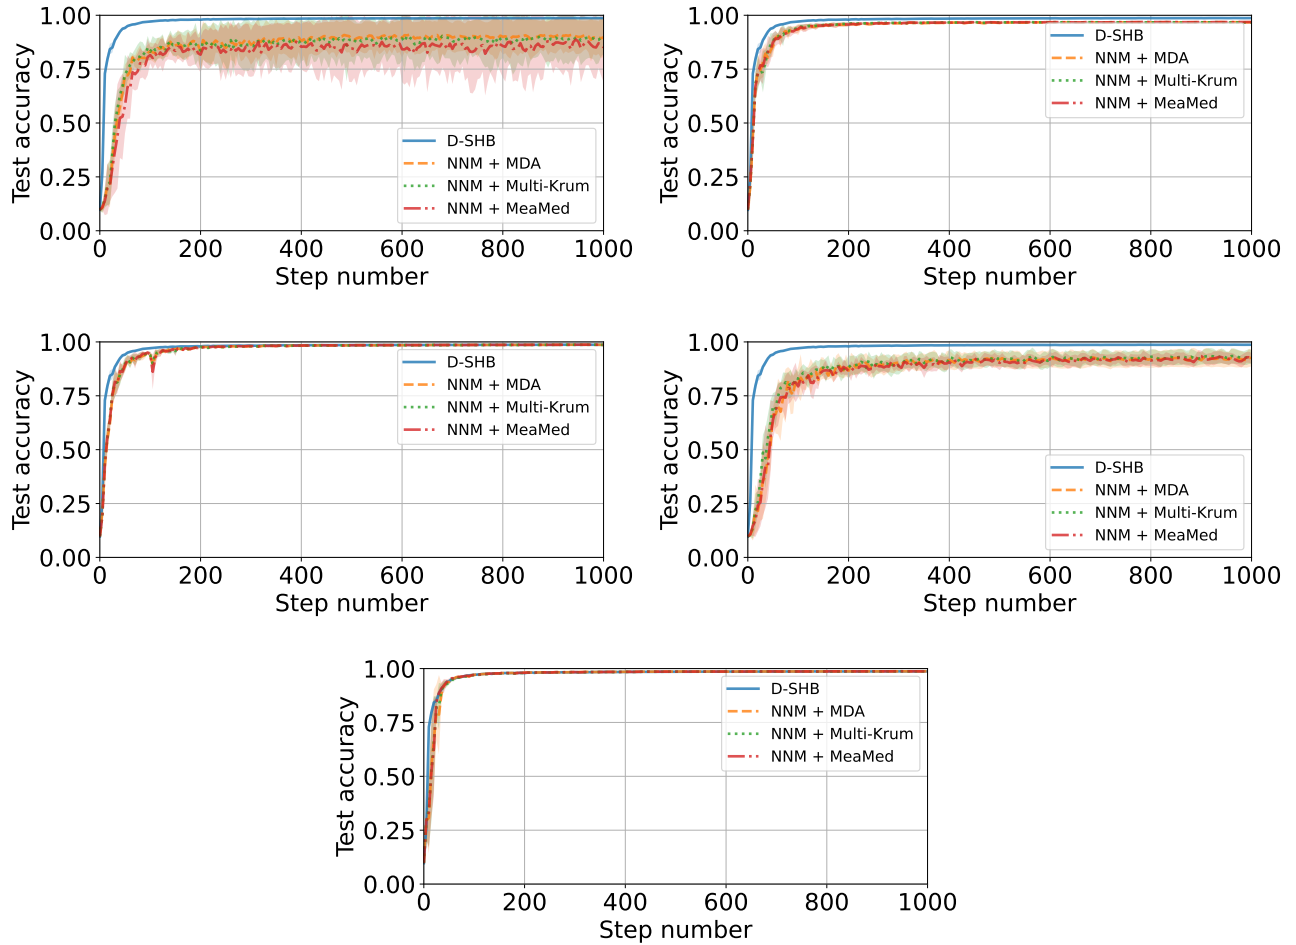


Figure 16: Experiments on MNIST using robust D-SHB with  $f = 6$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.



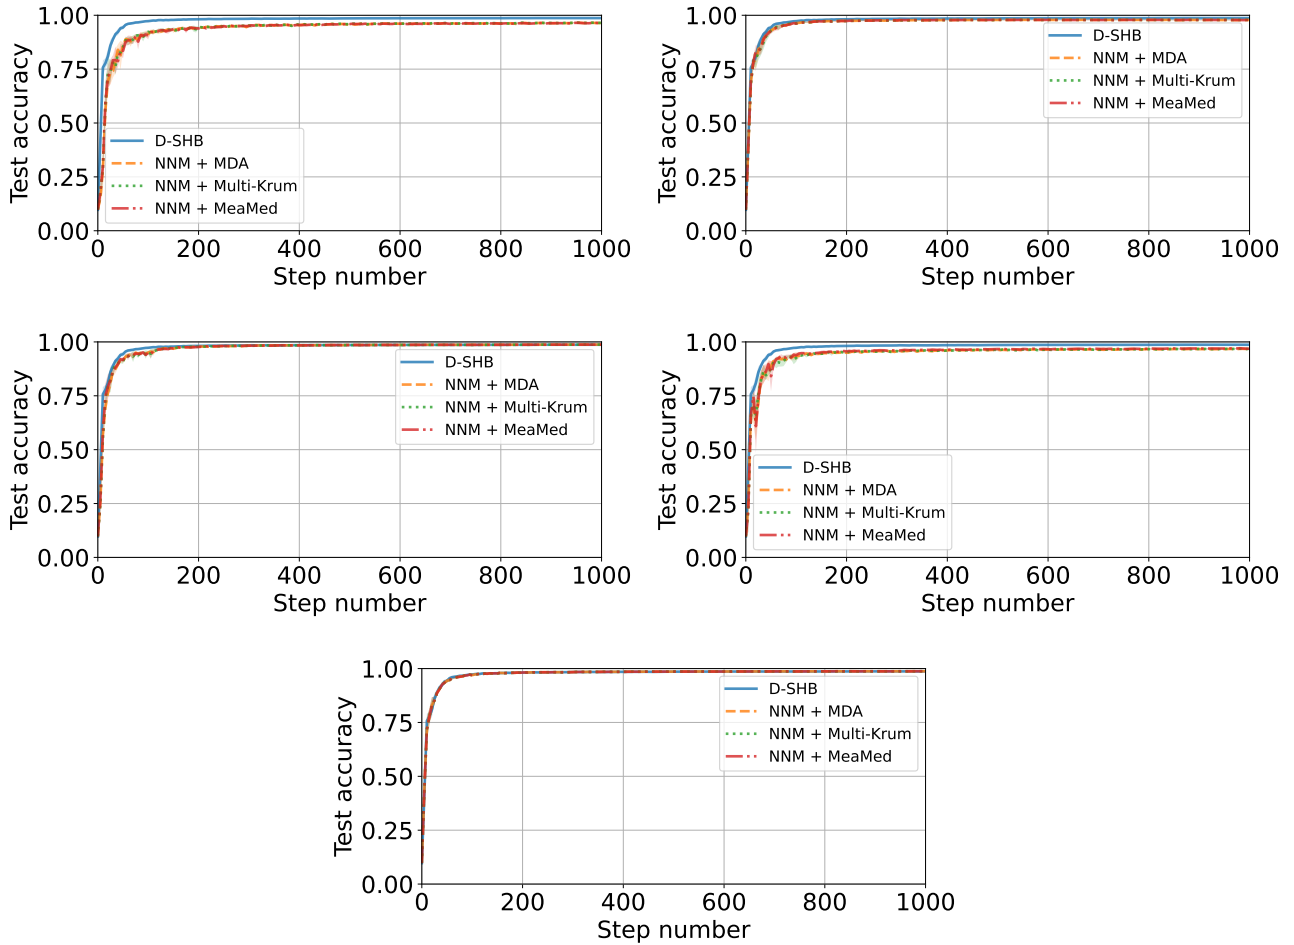


Figure 17: Experiments on MNIST using robust D-SHB with  $f = 6$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

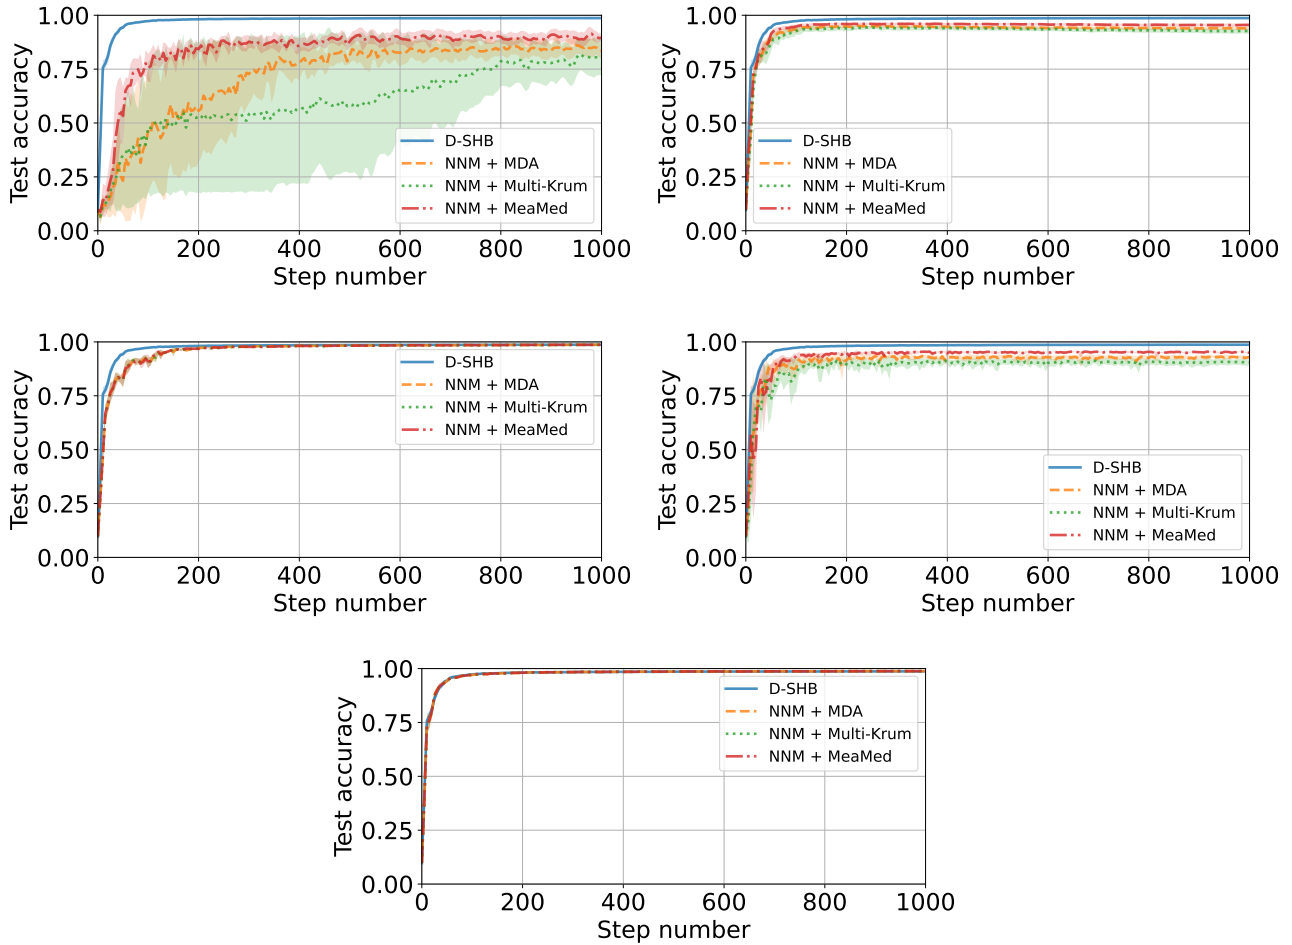


Figure 18: Experiments on MNIST using robust D-SHB with  $f = 8$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

### 15.1.4 Results without Clipping

We also run experiments on MNIST without clipping the gradients on honest workers, as suggested by the reviewers. We consider  $f = 4$  Byzantine workers among  $n = 17$  workers, and a moderate heterogeneity setting where  $\alpha = 1$ . The results are shown in Figure 19, and convey the same observations made from Figure 5 in Section 15.1. Note that we also obtain similar results when running with  $\alpha = 10$ .

Finally, we also consider the case  $\alpha = 0.1$  without clipping honest gradients. However, we do not show the results since no technique (including Bucketing and NNM) is able to withstand the considered Byzantine attacks in such an extreme heterogeneity setting. In fact, we believe that exploding gradients are even more detrimental when considering Byzantine attacks, and thus gradient clipping is important in our setting, although not needed in theory.

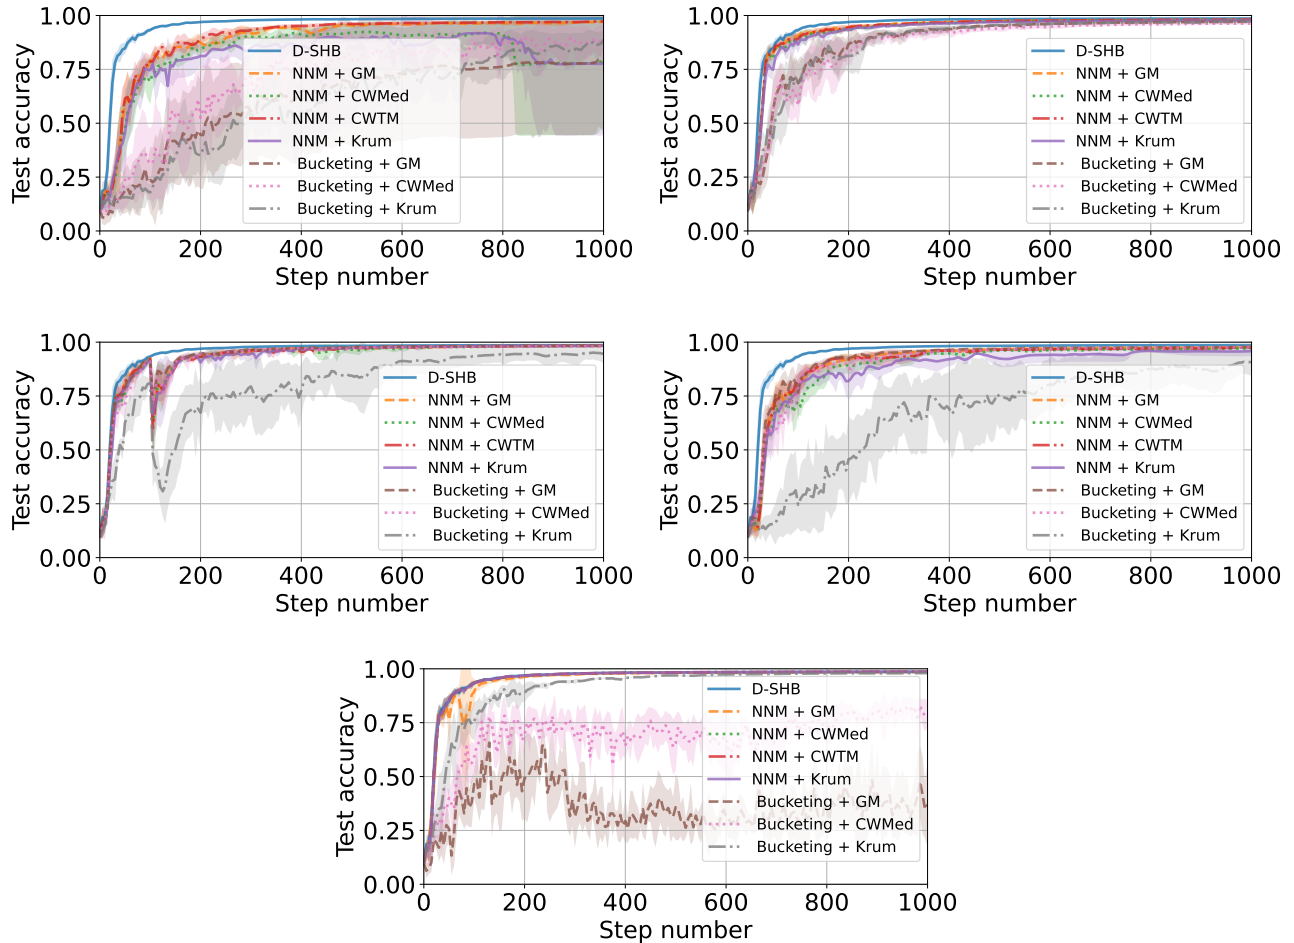


Figure 19: Experiments on MNIST without clipping and using robust D-SHB with  $f = 4$  among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3).

## 15.2 Comprehensive Results on Fashion-MNIST

In this section, we showcase our results on the Fashion-MNIST dataset. We consider two Byzantine regimes:  $f = 4$  and  $f = 6$  out of  $n = 17$  workers in total. We also consider three heterogeneity regimes:  $\alpha = 0.1$  (extreme),  $\alpha = 1$  (moderate), and  $\alpha = 10$  (low). We plots compare the performance of NNM and Bucketing when executed with four aggregation rules namely Krum, GM, CWMed, and CWTM. The plots are presented below. We note that in the strongest Byzantine setting, i.e., when  $f > 4$ , Bucketing can only be applied with bucket size  $s = 1$  (Karimireddy et al., 2022). This means that when  $f = 6$ , executing Bucketing boils down to running the vanilla aggregation rules.

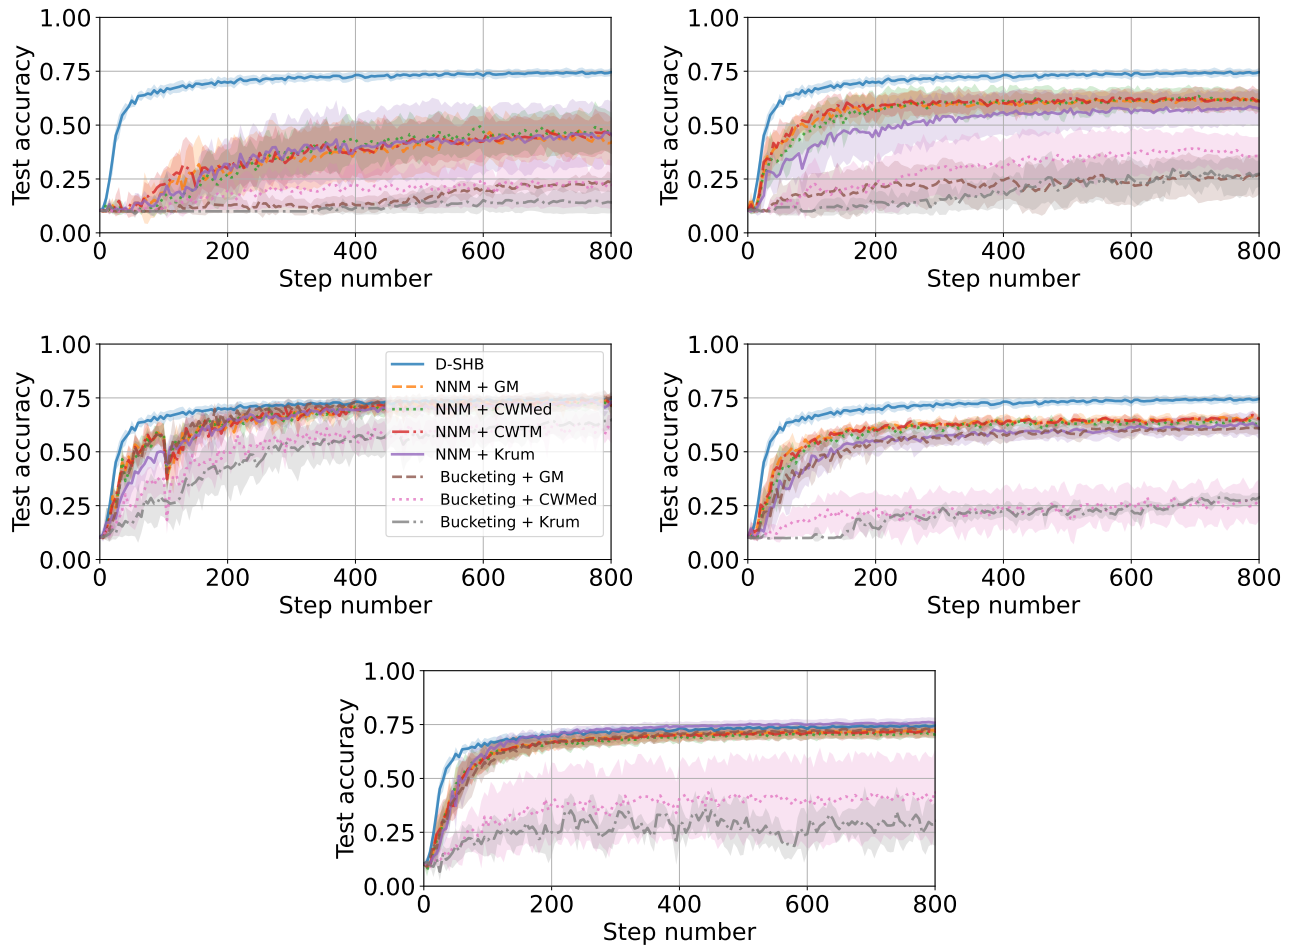


Figure 20: Experiments on Fashion-MNIST using robust D-SHB with  $f = 4$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 0.1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

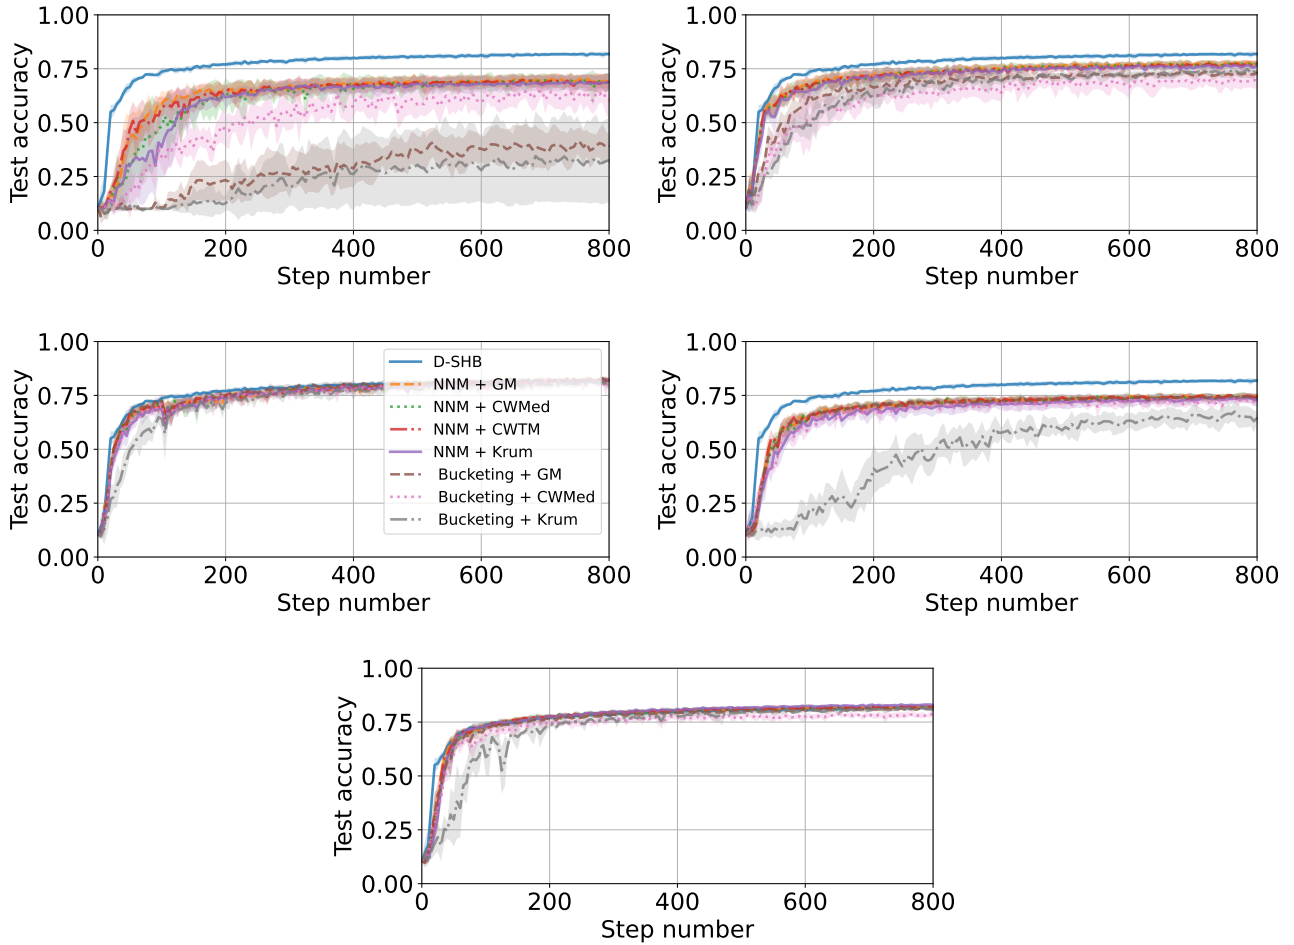


Figure 21: Experiments on Fashion-MNIST using robust D-SHB with  $f = 4$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

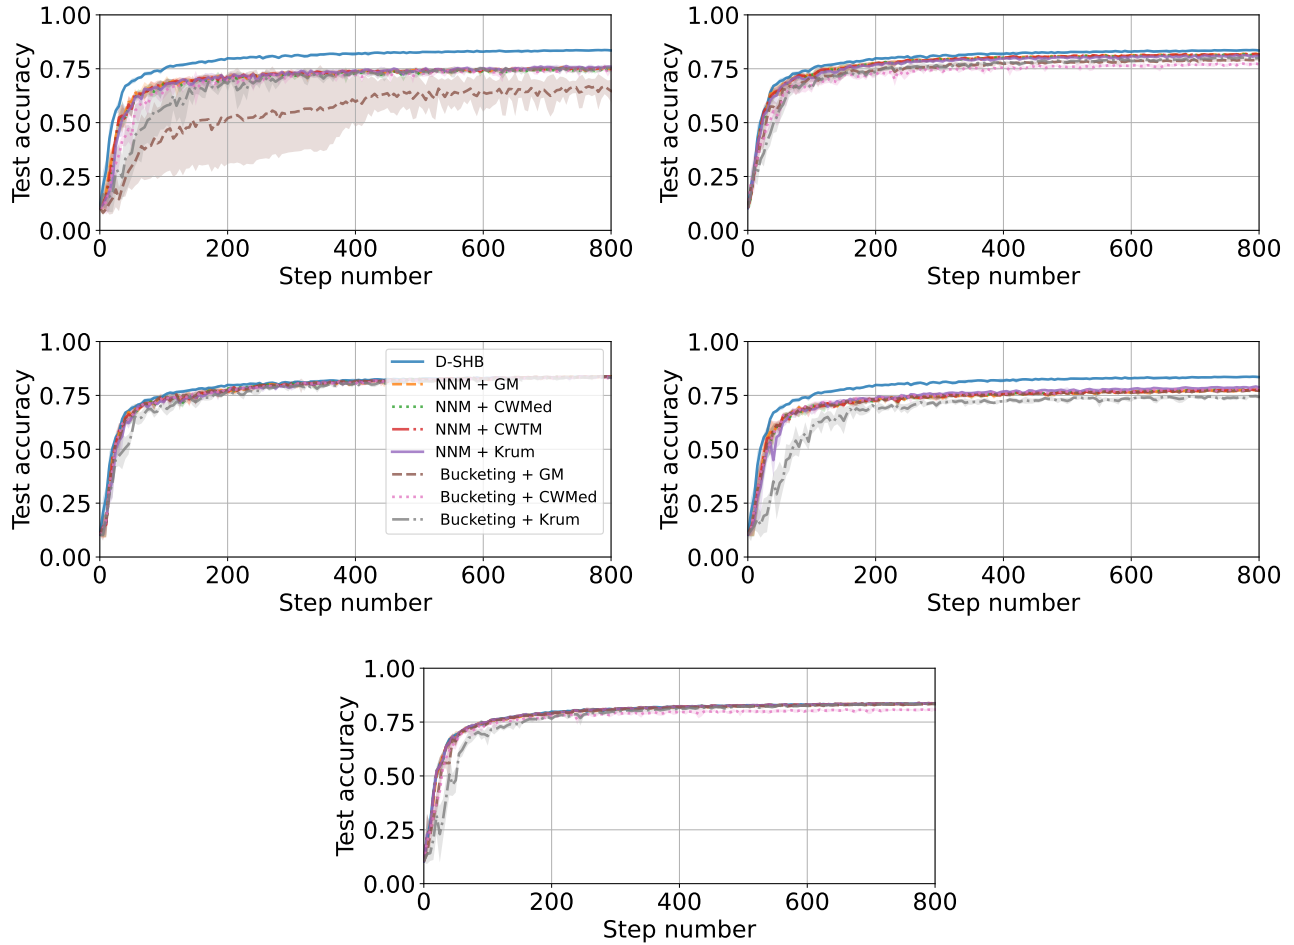


Figure 22: Experiments on Fashion-MNIST using robust D-SHB with  $f = 4$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

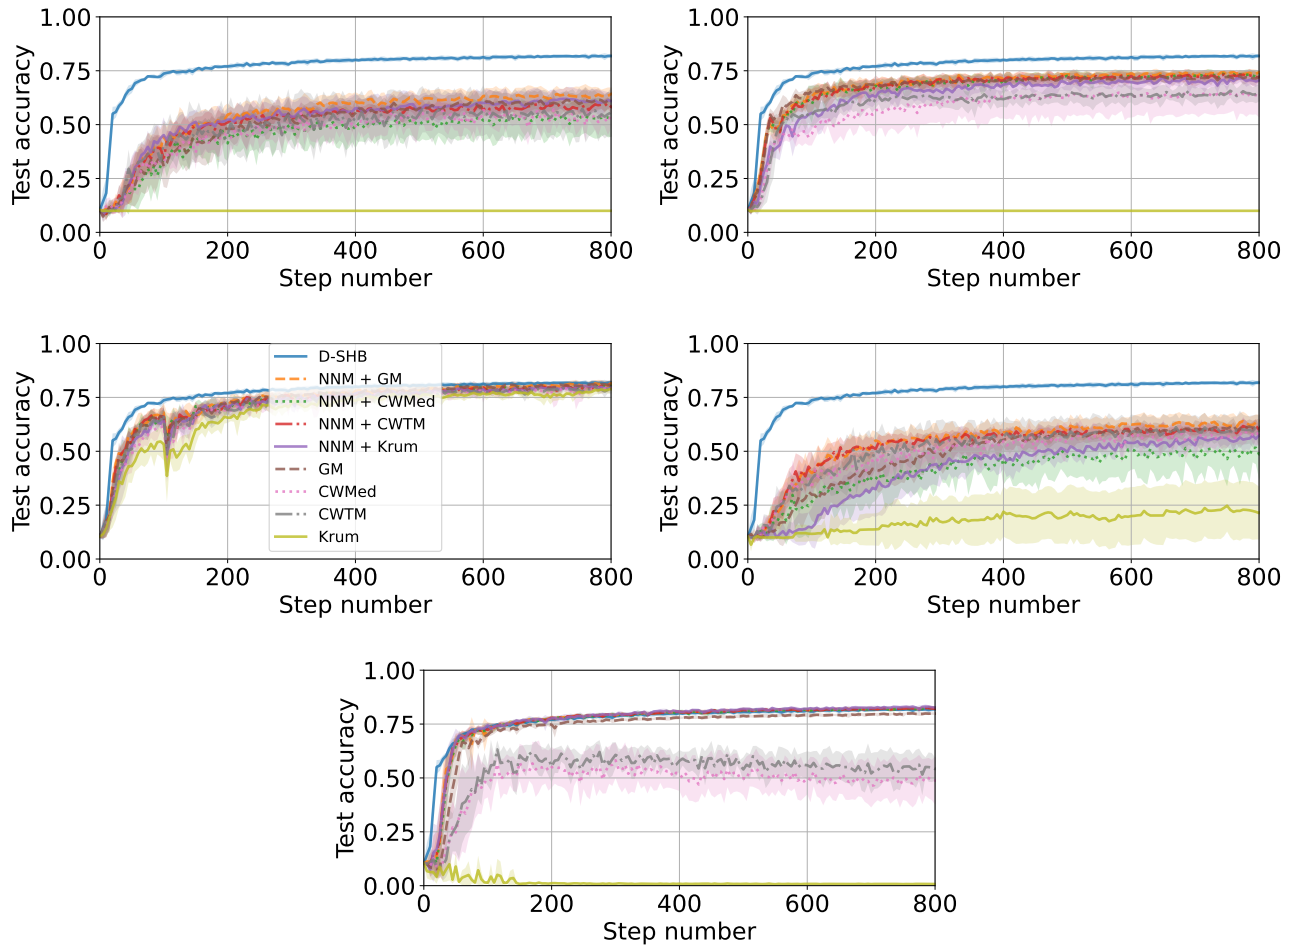


Figure 23: Experiments on Fashion-MNIST using robust D-SHB with  $f = 6$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.




Figure 24: Experiments on Fashion-MNIST using robust D-SHB with  $f = 6$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.



### 15.3 Comprehensive Results on CIFAR-10

In this section, we present the entirety of our results on CIFAR-10. We consider three Byzantine regimes:  $f = 2$ ,  $f = 3$ , and  $f = 4$  out of  $n = 17$  workers in total. We also consider two heterogeneity regimes:  $\alpha = 1$  (moderate), and  $\alpha = 10$  (low). We compare the performance of NNM and Bucketing when executed with four aggregation rules namely Krum, GM, CWMed, and CWTM. The plots are presented below and complement our (partial) results in Figure 1 in Section 6.3 of the main paper.

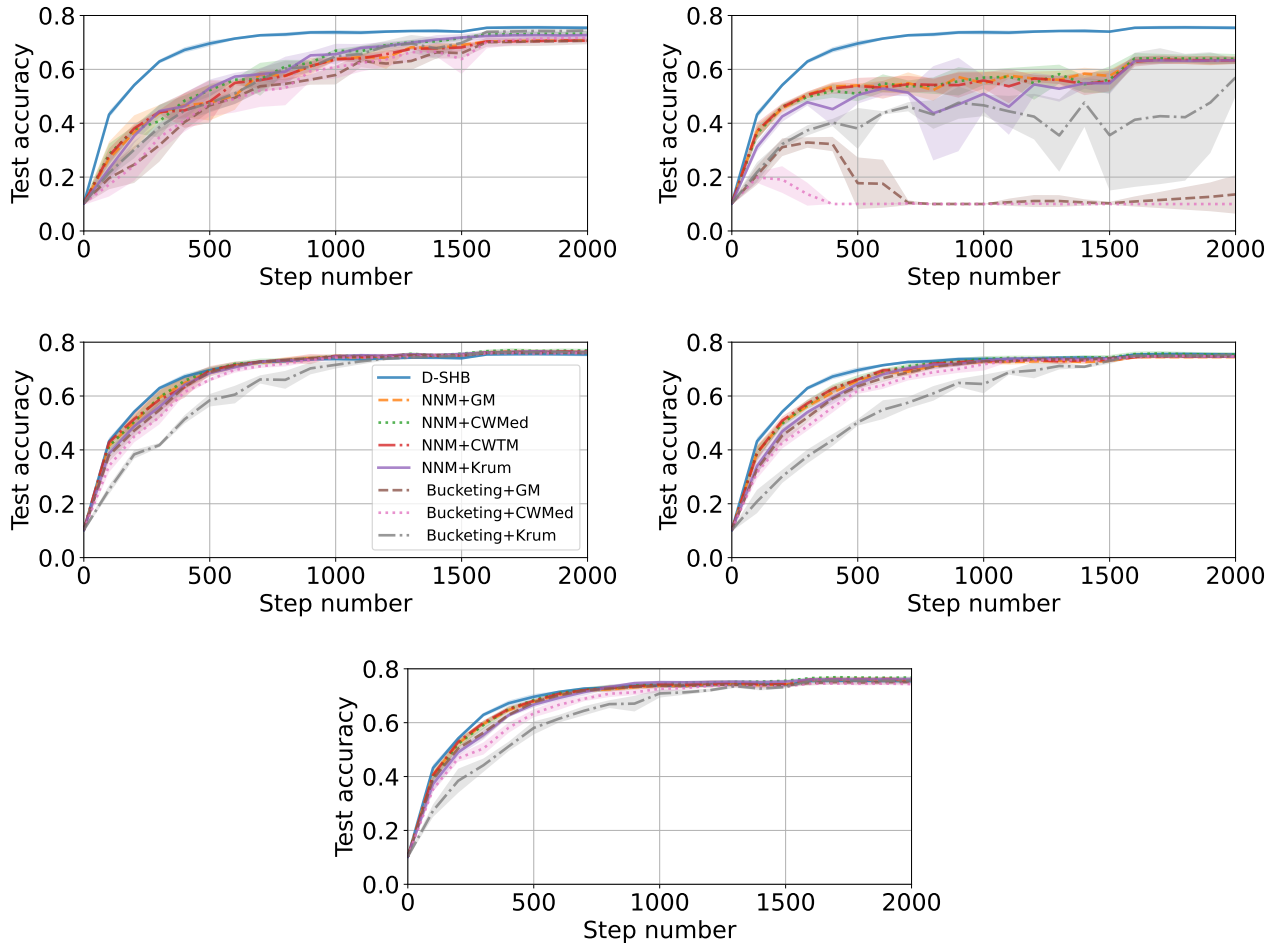


Figure 25: Experiments on CIFAR-10 using robust D-SHB with  $f = 2$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

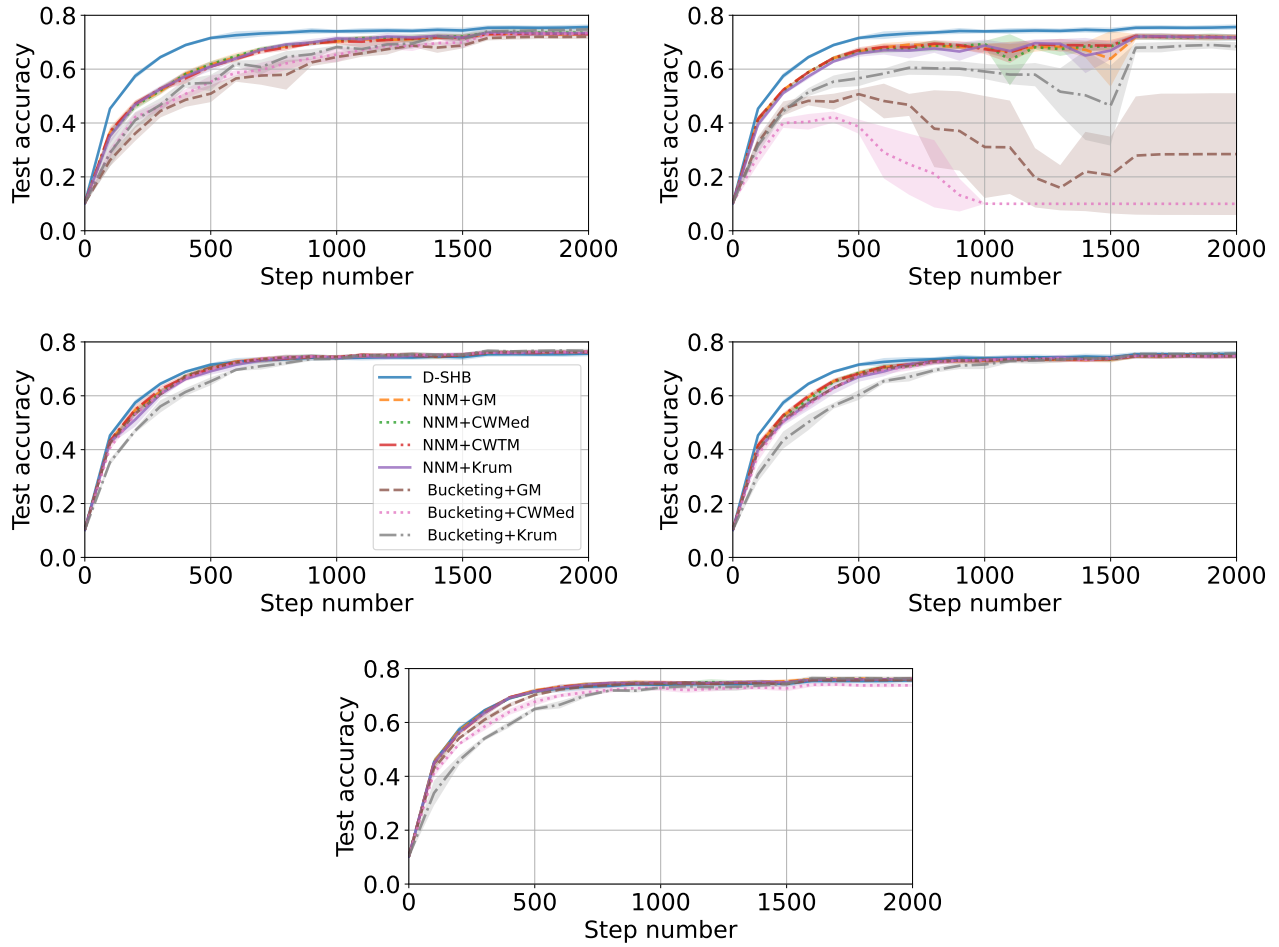


Figure 26: Experiments on CIFAR-10 using robust D-SHB with  $f = 2$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

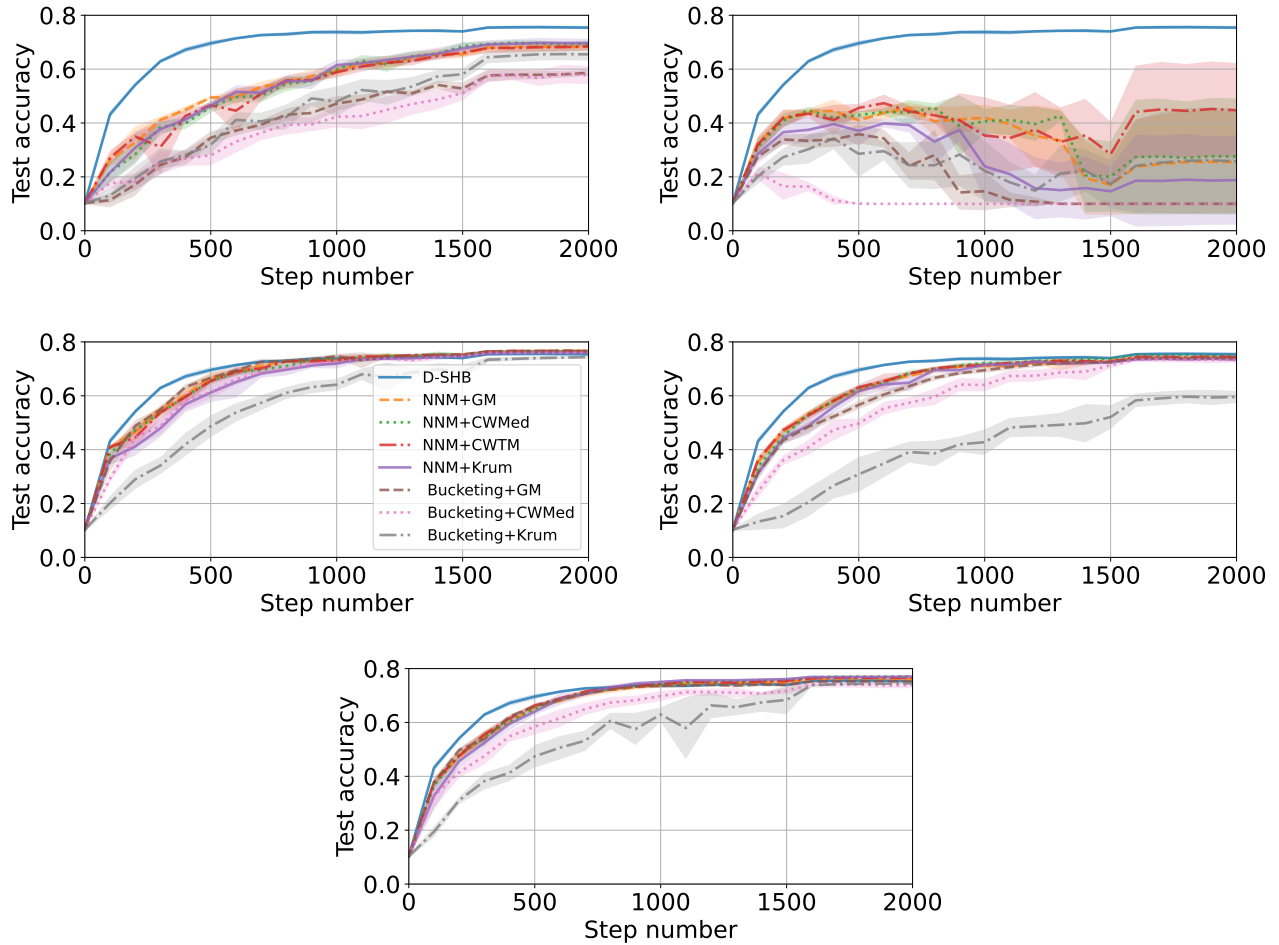


Figure 27: Experiments on CIFAR-10 using robust D-SHB with  $f = 3$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 1$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

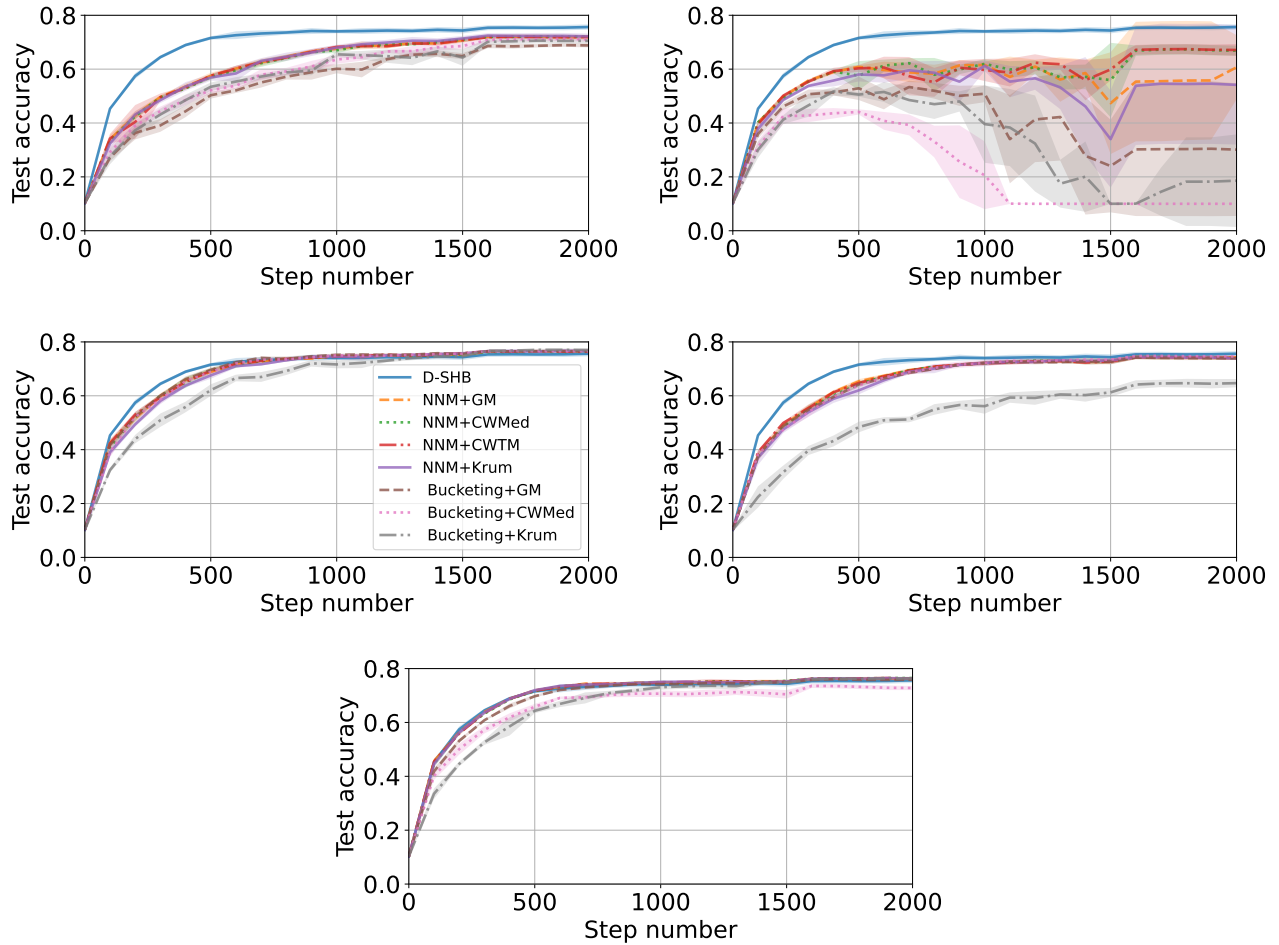


Figure 28: Experiments on CIFAR-10 using robust D-SHB with  $f = 3$  Byzantine among  $n = 17$  workers, with  $\beta = 0.9$  and  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.

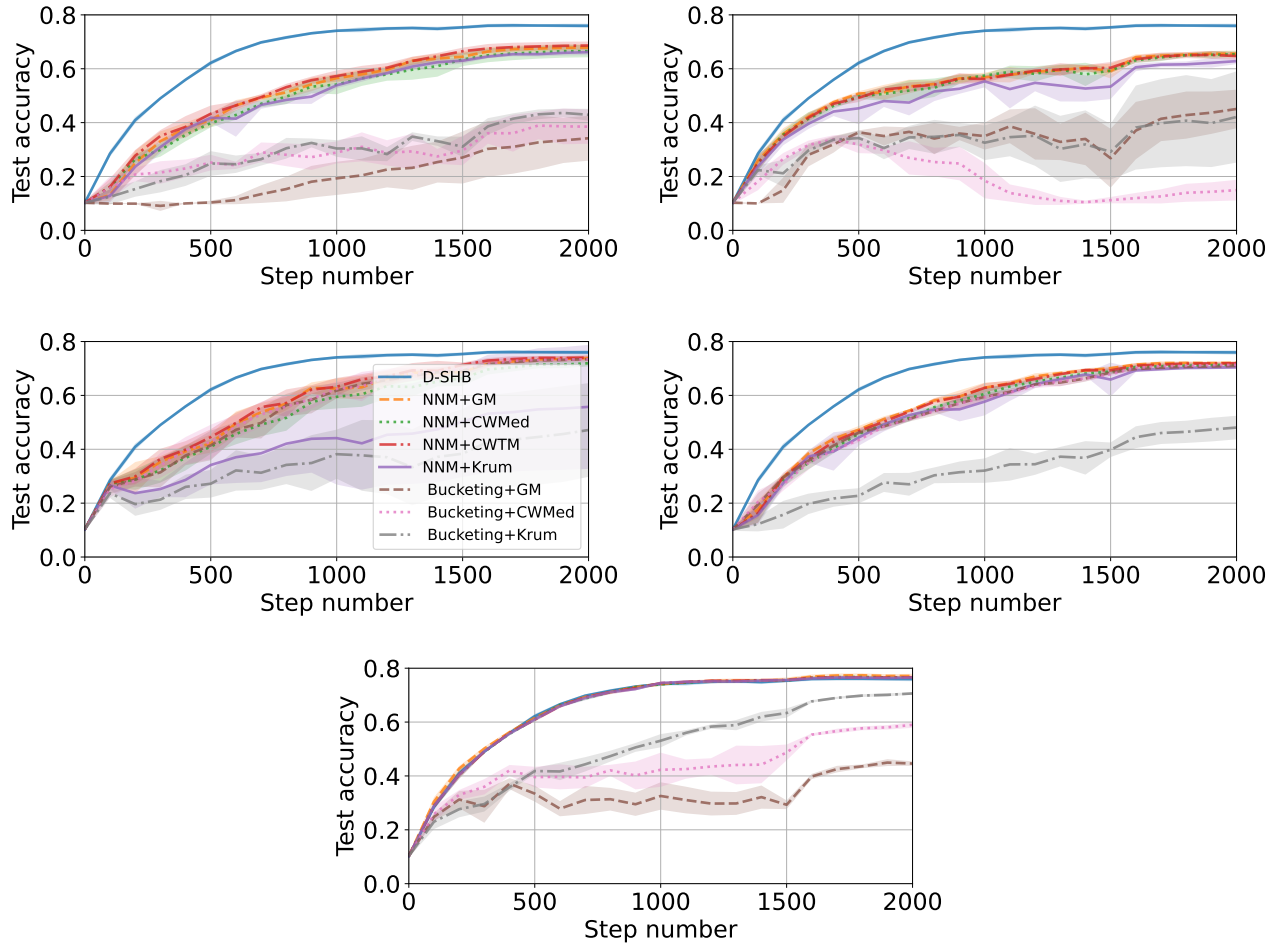


Figure 29: Experiments on CIFAR-10 using robust D-SHB with  $f = 4$  Byzantine among  $n = 17$  workers, with  $\beta = 0.99$  and  $\alpha = 10$ . The Byzantine workers execute the FOE (row 1, left), ALIE (row 1, right), Mimic (row 2, left), SF (row 2, right), and LF (row 3) attacks.