

Practical Differentially Private and Byzantine-resilient Federated Learning

Zihang Xiang
KAUST
zihang.xiang@kaust.edu.sa

Wanyu Lin
The Hong Kong Polytechnic University
wanyulin@comp.polyu.edu.hk

Tianhao Wang
University of Virginia
tianhao@virginia.edu

Di Wang
KAUST
di.wang@kaust.edu.sa

ABSTRACT

Privacy and Byzantine resilience are two indispensable requirements for a federated learning (FL) system. Although there have been extensive studies on privacy and Byzantine security in their own track, solutions that consider both remain sparse. This is due to difficulties in reconciling privacy-preserving and Byzantine-resilient algorithms.

In this work, we propose a solution to such a two-fold issue. We use our version of differentially private stochastic gradient descent (DP-SGD) algorithm to preserve privacy and then apply our Byzantine-resilient algorithms. We note that while existing works follow this general approach, an in-depth analysis on the interplay between DP and Byzantine resilience has been ignored, leading to unsatisfactory performance. Specifically, for the random noise introduced by DP, previous works strive to reduce its impact on the Byzantine aggregation. In contrast, we leverage the random noise to construct an aggregation that effectively rejects many existing Byzantine attacks.

We provide both theoretical proof and empirical experiments to show our protocol is effective: retaining high accuracy while preserving the DP guarantee and Byzantine resilience. Compared with the previous work, our protocol 1) achieves significantly higher accuracy even in a high privacy regime; 2) works well even when up to 90% distributive workers are Byzantine.

KEYWORDS

Federated learning, Byzantine Security, Differential privacy, Distributive computation

ACM Reference Format:

Zihang Xiang, Tianhao Wang, Wanyu Lin, and Di Wang. 2018. Practical Differentially Private and Byzantine-resilient Federated Learning. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 33 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Federated Learning (FL), a learning framework for preserving the privacy of distributed data [39], has thrived during the past few years. To comply with the privacy regulations such as General Data Protection Regulation (GDPR) [26], variants of FL frameworks have been widely studied, and recently adopted in industry, such as Apple's "FE&T" [49], Google's Gboard [25], and Alibaba's FederatedScope [71]. In an FL system, there are several local workers, each holding a dataset for local training, and a server aggregating gradient vectors from workers for global model updates.

However, current FL frameworks that seemingly can protect privacy (because the original data never leaves the local workers) are in fact vulnerable to various privacy attacks, such as membership inference attacks [56] (tries to infer whether some data samples are used in training) and model inversion attacks [78] ("reverse-engineer" sensitive data samples through gradients). These vulnerabilities drive the community to design methods that can further preserve the privacy of data held by workers. Among the privacy-enhancing techniques [23, 24, 53, 61], differential privacy (DP) [21] is a rigorous mathematical scheme that allows for rich statistical and machine learning analysis and is becoming the *de facto* notion for data privacy. Many methods have been proposed to tackle the problems of integrating DP into machine learning/deep learning from different perspectives [1, 12, 35, 47, 48, 57, 62, 63, 65]. More recently, DP has been adopted in the FL setting [2, 27, 60, 67, 74].

Besides privacy risks, FL systems are also vulnerable to adversarial manipulations from Byzantine workers, which could be fake workers injected by an attacker or genuine workers compromised by an attacker. Specifically, in a Byzantine attack, the adversary intends to sabotage the collective efforts by sending false information, such as contrived Byzantine gradients [6, 70]. To mitigate this issue, recent work proposes Byzantine-resilient machine learning approaches, such as diagnosing and rejecting gradients with abnormal features [9, 11, 15, 50, 52].

Tremendous progress on privacy and Byzantine resilience have been seen in their own track. However, all of them are not applicable to the more practical scenario where a privacy attacker is also Byzantine (a double-role attacker). Being aware of that, some recent work started to focus on such an issue yet provided unsatisfactory answers. Some of them fail to ensure both DP and Byzantine resilience simultaneously [31], while some other work tries to explore optimal parameter setups but still end up with a much-limited solution [30]. We also notice that some work [77] tries to combine existing variants in both tracks to side-step the

seeming incompatibility of DP and Byzantine security, however, their resistance is retained only when the privacy level is low and the portion of Byzantine clients is small.

Contributions: We observe that previous solutions fail to give a satisfactory answer for a common reason: neither the DP algorithm nor the Byzantine defending method is designed against both risks simultaneously. Our contribution is how we start from a co-design to form a DP and Byzantine-resilient solution, proving the synergy of combined DP and Byzantine resilience.

1) Co-design: Since random noise introduced by DP impairs the effectiveness of existing Byzantine-resilient aggregation rules, previous works tend to limit the impact of randomness by increasing the data batch size [30, 31]. In contrast, we leverage random noise to aid Byzantine aggregation: we use small batch size and accordingly construct our first-stage aggregation which effectively rejects many existing attacks.

Moreover, previous works continue to use the standard DP-SGD [1] to bound gradient sensitivity by clipping, which involves manually tuning the clipping parameter. In contrast, we ensure bounded sensitivity by normalizing, and it enables our second-stage aggregation, which provides a final sound filtering.

2) Cherry on top: We are the first to find out that bounding gradient sensitivity by normalizing is more suitable for DP learning although normalizing itself is not new¹. Specifically, we analyze its theoretical implication and also leverage it to construct a learning protocol that saves quadratic efforts² in hyper-parameter-tuning for DP learning, where a smaller amount of queries on gradient computation is more favored.

In the final evaluation, we conduct experiments to first show our contribution to DP solution and Byzantine aggregation in their own track. Then, for the core aim to preserve both privacy and Byzantine security, our experimental results show that in addition to having the DP guarantee, our protocol also remains robust against strong attackers when there are up to 90% distributive workers are Byzantine. We have released our code in supplementary material³.

2 BACKGROUND AND PRELIMINARIES

2.1 Federated Learning

In a typical setting of machine learning, we have a training dataset $\{x_1, x_2, \dots, x_m\}$ where each x_i contains a feature vector and a label, and we also have a loss function f . We aim to find the best model parameters w from a parameter space Θ which minimizes the following function through stochastic optimization:

$$\min_{w \in \Theta} F(w) = \mathbb{E}_{x \sim \mathcal{P}} [f(x; w)], \quad (1)$$

In FL, suppose there are n workers and the i -th worker has local and private data D_i , then training in FL happens in a distributed manner. Specifically, in the t -th iteration, we have:

1) Model broadcasting: The server broadcasts the current model parameters w^{t-1} to all workers.

¹Some concurrent work [10, 19, 20, 72] on DP learning use an operation similar to normalizing with different considerations.

²Instead of tuning learning rate η and clip threshold C for different ϵ , our approach only needs to turn η for any instance of ϵ .

³

<https://github.com/zihangxiang/-Practical-Differentially-Private-and-Byzantine-resilient-Federated.git>.

2) Local gradient computation: After receiving the model sent by the server, each worker will use his/her private data and the model w^{t-1} to compute his/her gradient vector $g_i^t = \nabla f(D_i; w^{t-1}) := \frac{1}{|D_i|} \sum_{x \in D_i} \nabla f(x; w^{t-1})$. Note that workers can also compute their gradients and update their model N times locally, and report the difference between the model they get locally and the last model they receive from the server. In our framework, we take $N = 1$ and this is due to the constraints of DP-SGD protocol which will be discussed later. Extending to the cases where $N > 1$ will be left for future study.

3) Gradient aggregation and model update: The server will perform an aggregation step (denoted by function **Aggregation**) on the gradient vectors reported by workers and use the result $g^t = \text{Aggregation}(g_1^t, g_2^t, \dots, g_n^t)$ to update the model by $w^t = w^{t-1} - \eta g^t$, where η is the learning rate. Note that there are variants of aggregation strategies, e.g., $g^t = \sum_i \frac{|D_i|}{\sum_j |D_j|} g_i^t$ [45].

2.2 Differential Privacy for Deep Learning

DEFINITION 1 (DIFFERENTIAL PRIVACY [21]). Given a data universe X , we say that two datasets $D, D' \subseteq X$ are neighbors if they differ by only one data sample, which is denoted as $D \sim D'$. A randomized algorithm \mathcal{A} is (ϵ, δ) -differentially private if for all neighboring datasets D, D' and for all events S in the output space of \mathcal{A} , we have $\Pr(\mathcal{A}(D) \in S) \leq e^\epsilon \Pr(\mathcal{A}(D') \in S) + \delta$.

An (ϵ, δ) -DP mechanism typically adds calibrated noise to the output of a query. In this paper we mainly use the Gaussian mechanism to guarantee (ϵ, δ) -DP:

DEFINITION 2 (GAUSSIAN MECHANISM). Given any function $q : \mathcal{X}^n \rightarrow \mathbb{R}^d$, the Gaussian mechanism is defined as $q(D) + \xi$ where $\xi \sim \mathcal{N}(0, \frac{2\Delta_2^2(q) \log(1.25/\delta)}{\epsilon^2} I_d)$, where $\Delta_2(q)$ is the ℓ_2 -sensitivity of the function q , i.e., $\Delta_2(q) = \sup_{D \sim D'} \|q(D) - q(D')\|_2$. Gaussian mechanism satisfies (ϵ, δ) -DP when $\epsilon \leq 1$.

Another notable property of DP is that DP is closed under post-processing, i.e., if we post-process the output of an (ϵ, δ) -DP algorithm, then the whole procedure will still be (ϵ, δ) -DP.

DP in deep learning: Differentially Private SGD (DP-SGD) is a widely used method in machine learning to ensure DP [1, 7, 58]. It modifies the SGD-based methods by adding Gaussian noise to perturb the (stochastic) gradient in each iteration of the training, i.e., in the centralized setting, during the t -th iteration DP-SGD will compute a noisy gradient as follows:

$$g^t = \frac{1}{|B|} \left(\sum_{x_i \in B} \hat{g}_i^t + \mathcal{N}(0, \sigma^2 C^2 I) \right), \quad (2)$$

where B is a subsampled data batch used to compute the gradient, σ is the noise multiplier, \hat{g}_i^t is the gradient vector computed by feeding one data sample to w^{t-1} which is the current model before the t -th iteration, and g^t is the (noisy) gradient used to update the model. The main reason here we use \hat{g}_i^t instead of the original gradient vector is that we wish to make the term $\sum \hat{g}_i^t$ have bounded ℓ_2 -sensitivity so that we can use the Gaussian mechanism to ensure DP. The most commonly used approach to get a \hat{g}_i^t is clipping the gradient: $\hat{g}_i^t = \nabla f(x_i; w^{t-1}) \min\{1, \frac{C}{\|\nabla f(x_i; w^{t-1})\|_2}\}$ i.e., each gradient vector is clipped by C (scale those whose ℓ_2 -norm is greater

than C to be C exactly and leave the rest untouched). Since the ℓ_2 -sensitivity of $\sum \hat{g}_i^t$ is bounded by C , after the clipping, we can add Gaussian noise to ensure DP. To prove the (ϵ, δ) -DP property of DP-SGD, there has been a line of research [4, 29, 46, 66, 75, 79]. We use *TensorFlow Privacy* [59] to search for noise multiplier given ϵ and δ . Other than the DP-SGD framework, we note that there exists others work [16] tackling the privacy issue in FL by extending the PATE framework [47]. However, their work does not apply to the case where Byzantine workers exist.

2.3 Byzantine Attacks in FL

The FL protocol is vulnerable to Byzantine attacks, as each worker can report a malicious gradient vector to deteriorate the model performance [6, 70] and even bias the model in a specific way [13, 42]. Several Byzantine-robust approaches are proposed to tackle different attacks [9, 11, 15, 32, 50, 52, 73]. And there are also new advanced Byzantine attacks that try to bypass such defenses [5, 6, 14, 15, 36, 54, 70].

A Taxonomy: To understand the features of those Byzantine attacks, we summarise three dimensions to capture their properties.

1) *Objective:* There are generally 2 types of objectives: a) Denial-of-service attack (also called *untargeted* or convergence prevention) [6] that tries to destroy the training process and makes the model unusable. b) Backdoor attack [42] that tries to poison the training data to make the model predict intended results on inputs with specific triggers while still behaving normally on other inputs.

2) *Capability:* Existing attacks assume different attackers' power. Some work assumes that the Byzantine attacker is *omniscient*, i.e., the attacker knows what the honest workers send to the server [22, 70], while others assume the attacker does not know the honest workers' data [6]. In both cases, the attacker knows the aggregation rules.

3) *Specificity of targeted defenses:* a) Some attacks are defense-specific, i.e., and they are tailored for specific defense methods [6, 70]. However, it is unclear whether such attacks still remain effective against other or new defense protocols. b) Some other attacks are universal: these attacks are either defense-agnostic [11, 77] or have a meta-method [22] that can be instantiated to attack almost all existing defense strategies upon knowing the defense rules.

Instantiating Existing Attacks: We briefly introduce some existing Byzantine attacks which have been considered in previous work. They can be categorized by the above 3 dimensions: 1) *Gaussian attack* [52, 77] uploads pure Gaussian noise trying to hurt utility. 2) *Label-flipping attack* [11, 22] first poisons the local dataset by flipping the original label I to $H - 1 - I$ (H is the total number of classes, $I = 0, 1, \dots, H - 1$ is the label) and then follows the FL protocol. Note that we can also adopt other ways to perform the label flipping (such as randomly flipping to a different label). In fact, the way to flip the label does not matter as long as it tries to reduce the overall accuracy. 3) *Optimized Local Model Poisoning attack* [22], the state-of-the-art Byzantine attack method that can be accordingly instantiated for a specific Byzantine defense method in an adversarial way given the Byzantine defense protocol first.

For ease of discussion, we use "Byzantine attacker" to refer to a master attacker which can inject several fake workers into the

system and control all of them. Hence, in this sense, attackers who can send malicious uploads to the server can possibly collude.

3 PROBLEMS AND EXISTING SOLUTIONS

3.1 Problem Setting

Attacker: FL is indeed exposed to threats from two kinds of attackers: *privacy attacker* and *Byzantine attacker*. Note that we will not discuss specific *privacy attacks* and only focus on *Byzantine attacks* for the following reasons.

Protecting privacy with DP: From an information-theoretical view, DP guarantees privacy in the worst case by limiting the maximum amount of information that any privacy attacker can extract even with side information and unlimited computational resource [21]. It has also been shown that by tuning ϵ small enough [51], adopting DP effectively rejects strong privacy attackers [56]. Following the privacy settings as in the previous work [30, 31, 77], we focus on the item-level privacy for each worker's dataset; in this case, the gradient needs to be privatized (by adding random noise) before being uploaded to the server.

Focusing on Byzantine-resilience: We are not interested in tuning ϵ to test the algorithm's strength to defend against privacy attacks. Since our algorithms are guaranteed to be (ϵ, δ) -DP theoretically, *we only need to focus on Byzantine attacks*.

In other words, we treat DP as one of the basic properties that an FL system should possess. In fact, leveraging our tailored DP protocol to defend against Byzantine attacks is one of our novelties.

Byzantine attacker: We specify the Byzantine attacker according to the taxonomy mentioned above:

- *Objective:* Our Byzantine attacker is trying to perform Denial-of-Service (DoS) attacks.
- *Capability:* Our Byzantine attacker is omniscient; it knows all the data held by honest workers, information sent by honest workers, and the aggregation rules. We assume such capability in our framework to show that our protocol still works even when facing such a strong adversary.
- *Specificity of targeted defense:* We consider a stronger version of universal attacks. Our protocol will be made public and the attacker is allowed to instantiate his attack on our protocol.

In other words, we are interested in defending the stronger *untargeted* attacks and we leave backdoor attacks as future work.

Defender: Privacy is guaranteed through DP, and each user applies DP to protect its local data. To defend against Byzantine attacks, the server needs to design new aggregation rules such that false gradients from Byzantine workers are excluded. Here we assume the server possesses a *small amount* of labeled data samples which are kept secret from attackers. Let \mathbb{X} be the data space and \mathbb{Y} be the label space in a classification task. We do not require the server to have direct access to the local-hold data, instead, we assume the data space \mathbb{X}_{aux} from which the auxiliary data is sampled is the same as that of local-hold data. Notably, this additional assumption is reasonable in real applications as getting such a tiny amount of data is relatively cheap, and there is work on DP learning [8, 28, 33, 40, 47, 48, 64, 80] and Byzantine-resilient learning [11, 52] making this assumption. In our experiments, we simulate obtaining such data by randomly drawing $2C$ sampling from a validation set

Methods	Privacy	> 50%-Resilience
Krum [9]	✗	✗
Coordinate-wise Median [73]	✗	✗
Trimmed Mean [73]	✗	✗
Bulyan [32]	✗	✗
Zhu et al. [76]	✗	✗
FLTrust [11]	✗	✓
Rachid et al. [30]	✓	✗
Xu et al. [43]	✓	✗
Heng et al. [77]	✓	✗
Our work	✓	✓

Table 1: Comparison with previous work. For privacy, ✓ means the method is guaranteed to be DP while the ✗ means the converse; for > 50%-Resilience, ✓ means the method remains resilient when the number of Byzantine attackers exceeds half of the total while ✗ means the method is no longer effective under such majority Byzantine attack.

where C is equal to the number of classes of that dataset (e.g., for the MNIST dataset, $C = 10$, thus 20 auxiliary data samples will suffice). It is also helpful to consider whether the server-own data and local-hold data follow the same label distribution \mathbb{Y} (with a slight abuse of notation), accordingly, we also conduct experiments for *i.i.d.* case (distributions on \mathbb{Y} are the same for both) and *non-i.i.d.* case (distributions on \mathbb{Y} are different).

We also assume the server knows the truth that at least γn workers are honest among all n workers. It is notable that in the paper we do not need to place any restriction on γ , while previous work [22, 73] needs to assume that $\gamma > 0.5$.

In conclusion, each local worker adopts DP to protect privacy, hence, for one worker, the privacy attacker can be anyone (including the server) except itself. The Byzantine attacker (disguised as some local workers) contrives its upload and tries to destroy the training to make the model has low utility. As a Byzantine defender, the server is honest-but-curious, *i.e.*, it wants to have a model with good utility, thus it follows the protocol. However, it may try to infer sensitive information from the uploads sent by local workers.

3.2 Existing Solutions

In Table 1, we summarize the previous methods on privacy and Byzantine security issues in FL. As we can see from the table, our method can defend against more than 50% Byzantine workers while also achieving DP. In the following, we will provide more details and discuss the limitations of previous methods.

1) The following lines of work only focus on defending against Byzantine attacks: We first recall some existing solutions to Byzantine resilience. There is a line of work focusing on designing robust aggregation rules for corrupted gradients [9, 11, 15, 32, 32, 50, 52, 73] including Krum [9], RFA [50], coordinate-wise median [73], and Trimmed Mean [73]. We summarise the detail of these four methods in supp. material. for interested readers. In general, the first two methods involve computing pair-wise distance between vectors while the latter two concentrate on robust aggregation on each coordinate of vectors. Due to their intrinsic limitations, all of these methods are only applicable when the majority of workers (> 50%) are honest. Recently, Zhu et al. [76] propose improvements to existing Byzantine-resilient methods to provide a certain level of resilience.

Recently, there is work showing that it is possible to leverage the knowledge of clean gradient computed from non-private auxiliary data [11, 52] to help the aggregation. Common behavior in such methods is that the server weights each uploaded gradient according to their similarity compared to the gradient computed by server-own auxiliary data.

All of the above methods have no DP guarantee because they are designed to defend against Byzantine attacks ignoring privacy attacks. Hence, local datasets' privacy is at risk.

2) The following lines of work try to apply Byzantine defending methods on top of DP output: Some recent work investigates the problem of maintaining both DP and Byzantine resilience in federated learning [30, 31]. Specifically, they study methods of directly combining DP-SGD with some existing aggregation methods such as Krum, *i.e.*, by applying the aggregation on the noisy gradient.

Difficulties in reconciling DP and Byzantine-resilient protocols: Previous work shows that for these methods, to become Byzantine-resilient when DP noise is injected, the fraction of Byzantine workers must decrease with \sqrt{d} , where d is the size of the model if the batch size is not large enough [30, 31]. This means such methods achieve good performance only when the number of Byzantine workers is small. Their experiments also verify that this type of method is unsatisfactory by showing that the testing accuracy deteriorates significantly even for a small model learned on a simple dataset [31]. Another line of work is based on robust stochastic model aggregation on the local workers' gradients. In these methods, the gradients of each worker are compressed into signs (1 for non-negative and -1 for negative) with DP [43, 77], however, all of them remain effective only under < 50% Byzantine attack.

4 OUR APPROACH

Observations and lessons learned : Existing solutions apply off-the-shelf Byzantine methods on top of noisy gradient to explore optimal parameter setups. They fail to reach satisfactory performance because neither the Byzantine defending protocol nor the DP protocol is designed for the scenario where privacy and Byzantine resilience are both needed.

4.1 Method Overview

We first re-design the DP protocol, there are two notable properties we enforce in our DP protocol: 1) small training batch size for each worker; 2) use normalizing instead of clipping to bound per-example gradient norm. *We are not considering privacy and Byzantine resilience in a separate manner.* We design our first-stage aggregation based on the first property and design a second-stage aggregation based on the second. As will be seen later, the first property together with the first-stage aggregation trivially yet effectively rejects some existing attacks. The second property together with the second-stage aggregation effectively rejects more advanced attacks which bypasses our first-stage aggregation. As a cherry on top, in our DP protocol, such two properties themselves enable efficient hyper-parameter tuning.

4.2 Modifying DP Protocol

Our DP protocol is summarized in Algorithm 1. The two notable properties compared with vanilla DP-SGD [1] are: 1) different from

Algorithm 1 Private and Secure Learning

Input: initial model w^0 , number of iteration T , learning rate η , datasets held by n workers $\{D_i | i = 1, 2, \dots, n\}$, gradient momentum β , noise multiplier σ , batch size b_c , loss function $f(\cdot)$

- 1: Each worker i initializes a size- b_c momentum list $\phi_i^0 = [0, \dots, 0]$
- 2: **for** $t = 1, 2, \dots, T$ **do**
- 3: Server broadcasts model w^{t-1} to all workers
- 4: **for** $i = 1, 2, \dots, n$ **do in parallel**
- 5: Sample a size- b_c mini-batch d_i
- 6: **for** $j = 1, 2, \dots, b_c$ **do in parallel**
- 7: $g_j \leftarrow \nabla f(x_j \in d_i; w^{t-1})$
- 8: $\phi_i^t[j] \leftarrow (1 - \beta)g_j + \beta\phi_i^{t-1}[j]$
- 9: **end for**
- 10: $g_i^t \leftarrow \frac{1}{b_c} \left(\sum_{j \in [b_c]} \frac{\phi_i^t[j]}{\|\phi_i^t[j]\|_2} + \mathcal{N}(0, \sigma^2 I) \right)$
- 11: Upload g_i^t to server, then $\phi_i^t[j] \leftarrow g_i^t$
- 12: **end for**
- 13: $G_s^t \leftarrow \text{FilterGradient}(\{g_i^t | i = 1, 2, \dots, n\}, w^{t-1})$
- 14: $w^t \leftarrow w^{t-1} - \eta \frac{1}{n} \sum_{g \in G_s^t} g$
- 15: **end for**

Output: learned target model w^T

existing works that adopt big batch size ($10^2 - 10^6$) [3, 20], we adopt small batch size b_c (typically 8 or 16). Note that small batch size is essential for our first-stage aggregation (see Section 4.3 for details); 2) the second, is to replace the clipping operation in vanilla DP-SGD by normalization, vanilla DP-SGD method clips the gradients by multiplying the gradient vector g by $Factor = \min\{1, \frac{C}{\|g\|_2}\}$ (C is called the clipping threshold). We modify the multiplication factor to $Factor = \frac{1}{\|g\|_2}$, which normalizes the gradients to be of unit length. Also note that inspired by [18], to have a better convergence behavior, the gradients are processed with momentum.

We will see in the following why normalizing enables efficient hyper-parameter tuning. In later sections about defending against Byzantine attacks, we will also see that using small batch size is essential in our first-stage aggregation, and normalizing also plays an important role in the second-stage aggregation.

Normalization helps hyper-parameter tuning: We now introduce Theorem 1 which supports our hyper-parameter tuning strategy. Based on our DP protocol, consider a simpler case where there is no Byzantine attacker and we only have one honest worker. The model update (without momentum) in the t -th iteration has the following form:

$$w^t = w^{t-1} - \frac{\eta}{|B^t|} \left(\sum_{g^t \in B^t} \frac{g^t}{\|g^t\|} + z \right), \quad (3)$$

where B^t is the current local batch of per-example gradient (we fix the batch size to be $|B^t| = b_c$), $z \sim \mathcal{N}(0, \sigma^2 I)$ is the DP noise and $g^t = \nabla f(x; w^{t-1})$. Since g^t is derived from only a batch of samples, there is a sampling error. Denoting the sampling error by ξ^t , we can rewrite $\nabla f(x; w^{t-1})$ as $\nabla F(w^{t-1}) + \xi^t$, where $\nabla F(w^{t-1})$ is the gradient of the stochastic function. With some assumptions⁴ that:

⁴In deep neural networks we always have $F(w) > 0$, and for the L-Lipschitz continuous and bounded variance assumptions, they have been commonly used in the previous work for convergence analysis [10, 18, 72, 74].

1) The stochastic function F is bounded below with $F(w) > 0$; 2) F has L -Lipschitz continuous gradient (defined in Assumption 1); 3) The random vector $g^t = \nabla F(w^{t-1}) + \xi^t$ has bounded variance, i.e., $\mathbb{E} \|\xi^t\|^2 \leq v^2$ with some v . We have the following result:

THEOREM 1. (Convergence Behavior) Given a learning rate η , and the model is updated according to Equation 3 (gradient is normalized), we have

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla F(w^t)\| \leq \underbrace{\frac{3F(w^0)}{T\eta} + \frac{3L\eta}{2} \left(1 + \frac{\sigma^2 d}{b_c^2}\right)}_M + 8v.$$

PROOF. See supp. material. \square

By the above result, we can see that it is sufficient to minimize the term M , whose expression provides valuable guidance on choosing T and η that need to be set before training. If the magnitude of the noise and the batch size b_c satisfy $\frac{\sigma^2 d}{b_c^2} \gg 1$, then by setting the learning rate as

$$\eta = \frac{1}{\sigma} \sqrt{\frac{2F(w^0)b_c^2}{TLd}}, \quad (4)$$

we have $M \approx \sqrt{\frac{36F(w^0)Ld\sigma^2}{2b_c^2T}}$. Note that since we always have $\sigma = \Omega\left(\frac{q\sqrt{T\log(1/\delta)}}{\epsilon}\right)$ where $q = \frac{b_c}{|D|}$ is the sampling rate ($|D|$ is the size of data) [1]. Thus, we have $M = \Omega\left(\frac{1}{\epsilon|D|} \sqrt{F(w^0)Ld\log(1/\delta)}\right)$. This implies that: 1) The lower bound of M is getting worse when ϵ becomes smaller; 2) We get this optimal bound by relating T and η via Equation (4). If we fix one, we can potentially get the other one analytically instead of going through inefficient hyper-parameter tuning. In practice, we fix T first and decide the learning rate η . With T fixed, Equation (4) suggests that the optimal learning rate should be set inversely proportional to the DP noise multiplier σ , and this leads to our efficient hyper-parameter tuning strategy which outperforms existing methods. Note that the previous analysis was built on the assumption that $\frac{\sigma^2 d}{b_c^2} \gg 1$. Thus, to satisfy the assumption, we can either use a bigger model (increase d) or adopt a smaller batch size. Hence, using a small batch size is preferred for our method and differs from existing work as mentioned before.

Hence, our DP approach saves quadratic efforts and is truly beneficial for DP learning. This is beyond only considering the running-time complexity.

4.3 First-stage Aggregation

Design strategy: Inspecting existing work on Byzantine resilience, the uploads (d -dimension vectors) by Byzantine attackers are arbitrary in \mathbb{R}^d . Hence, a single faulty inclusion on a malicious upload can totally destroy model updates. As a strategy, enforcing some constraints on the subspace where any upload should lie will be beneficial to defend against attacks. From a high-level perspective, our refactored DP protocol together with first-stage aggregation does the job of “constrain”; our second-stage aggregation does the job of “complementary aggregation”.

Specifically, our choice of small batch size leads to the phenomenon that DP noise dominates for each upload, which leads to

some expected statistical properties. Another phenomenon is that although for each worker DP noise is dominating, we can still achieve good utility overall. The reason is that the server takes the average of all aggregated uploads; such an operation reduces the DP noise variance and averages gradients to its non-zero expectation. Therefore, we can still get good utility as long as the number of honest workers is sufficiently large. As will be shown in the experiments section, 10-20 honest workers will suffice and such a number is smaller than that in previous work of FL systems [11, 22].

Forming first-stage aggregation: As we can see from Algorithm 1, an honest worker will upload $g = \tilde{g} + z$ to the server where \tilde{g} is the sum of some normalized terms, and z is the DP noise. If DP noise is dominating ($\|z\| \gg \|\tilde{g}\|$), we can approximately treat vector g as each coordinate of g is sampled from $\mathcal{N}(0, \sigma^2)$. We then have the following conclusion.

Norm test: Note that $\frac{\|g\|^2}{\sigma^2}$ follows the chi-squared distribution with degree d . As d is very large, by Central Limit Theorem, we can safely approximate the distribution of $\|g\|^2$ as Gaussian distribution: $\mathcal{N}(\sigma^2 d, 2\sigma^4 d)$. Hence, $\|g\|^2$ falls in the interval $[\sigma^2 d - 3\sigma^2 \sqrt{2d}, \sigma^2 d + 3\sigma^2 \sqrt{2d}]$ almost surely.⁵

Similarly, we can also conclude other statistical results for higher-order moments leveraging the property of Gaussian distribution. However, the real situation only allows us to use a limited number of these statistics. To further enhance the efficiency and the soundness of such checking, we will leverage non-parametric test methods which test the hypothesis that given samples follow a reference distribution. We leverage Kolmogorov–Smirnov test (KS test) [38] as described below.

KS test: Treat each coordinate of g as a sample and the null hypothesis is that these samples are sampled from the same distribution $\mathcal{N}(0, \sigma^2)$. Suppose that we are currently testing on upload g with d coordinates (d -dimension vector) and we denote the i -th coordinate of g as $g[i]$. KS test will 1) compute the empirical Cumulative Distribution Function as:

$$C_d(x) = \frac{1}{d} \sum_{i=1}^d \mathbb{1}_{g[i] < x},$$

where $\mathbb{1}_{g[i] < x}$ is the indicator function that takes value on 1 if $g[i] < x$ and 0 otherwise; 2) compute the KS statistics $D_{KS} = \sup_x |C_d(x) - \Phi_\sigma(x)|$ where $\Phi_\sigma(x)$ is the CDF of $\mathcal{N}(0, \sigma^2)$; 3) compute the P-value by D_{KS} from Kolmogorov D-statistic table [44] and there are many off-the-shelf libraries that can compute it. If the P-value is smaller than 0.05,⁶ we reject the null hypothesis (the server then treats g as one malicious upload that is not sampled from $\mathcal{N}(0, \sigma^2)$ and rejects it).

KS test confines Byzantine subspace: After we set the significance level for the P-value, we are essentially placing an upper bound on D_{KS} (if D_{KS} is too large, the corresponding P-value will be small enough to make g to be rejected). And from the definition of D_{KS} , we know that such an upper bound applies to $|C_d(x) - \Phi_\sigma(x)|$ for any $x \in \mathbb{R}$. This can also be interpreted

⁵Such interval is narrow, because $\frac{\sigma^2 \sqrt{2d}}{\sigma^2 d} \ll 1$ when d is very large. And by the 68-95-99.7 rule [68], we set such an interval to span three s.t.d. around the center so that a benign gradient falls into this interval with 99.7% probability approximately.

⁶We use the widely adopted significance level.

Algorithm 2 FirstAGG(g)

Input: g , the upload to be tested

```

1: if  $\|g\| < \sqrt{\sigma^2 d - 3\sigma^2 \sqrt{2d}}$  or  $\|g\| > \sqrt{\sigma^2 d + 3\sigma^2 \sqrt{2d}}$  then  $g \leftarrow 0$ 
2: end if
3: if  $KS(g) < 0.05$  then  $g \leftarrow 0$ 
4: end if

```

Output: g

as that the curve of $C_d(x)$ will fall into a band bounded by an upper envelop $E_u(x) = \min(1, \Phi_\sigma(x) + D_{KS})$ and a lower envelop $E_l(x) = \max(0, \Phi_\sigma(x) - D_{KS})$. If we set the significance level strictly enough, D_{KS} will be small enough such that the band will be narrow, requiring that $C_d(x)$ almost aligns with $\Phi_\sigma(x)$.

Formally, since $C_d(x)$ is a step function (with d steps) that is monotonously increasing with d steps ($\frac{i}{d} \rightarrow \frac{i+1}{d}$ for $i = 0, 1, \dots, d-1$), we have the following theorem.

THEOREM 2. (Byzantine Resilience) *If we sort all coordinates of an upload vector into a sequence by increasing order, to pass the KS test, the k -th (we count from 1 to d) element must fall into the interval:*

$$\left[E_u^{-1}\left(\frac{k}{d}\right), E_l^{-1}\left(\frac{k-1}{d}\right) \right] \quad (5)$$

PROOF. Let the k -th element be x_k , according to the definition of $C_d(x)$, then we must have $C_d(x_k) = \frac{k-1}{d}$ and $C_d(x_k + s) = \frac{k}{d}$ where s is some small real number. To pass the KS test, $\frac{k-1}{d}$ must be above $E_l(x_k)$ and $\frac{k}{d}$ must be below $E_u(x_k)$, consequently, x_k falls into the above interval. \square

In essence, our first-stage aggregation enforces that the attacker's upload must lie in the subspace of \mathbb{R}^d as described by Equation 5. This is different from existing work on Byzantine security, where the attacker's upload can be arbitrary in \mathbb{R}^d [6, 11, 22, 52, 70].

Ensuring $\|z\| \gg \|\tilde{g}\|$: Our first-stage aggregation only keeps those uploads that follow our DP protocol (approximately with the form $g = \tilde{g} + z$) and this validity builds on the assumption that $\|z\| \gg \|\tilde{g}\|$. The good news is that we can always control $\frac{\|z\|}{\|\tilde{g}\|}$. Recall that $\|z\| \approx \sigma\sqrt{d}$ and \tilde{g} is just the sum of b_c norm-bounded vectors, hence, before the training, we can compute $\frac{\|z\|}{\|\tilde{g}\|}$. To increase $\frac{\|z\|}{\|\tilde{g}\|}$, we can either 1) use a bigger model (increase d); or 2) adopt a smaller batch size for local workers. Thus, as highlighted in Section 4.2, using a small batch size is one of our technical details that differs from vanilla DP-SGD.

4.4 Second-stage Aggregation

In this part, we present our second-stage aggregation which does the job of “complementary aggregation”. In total, our first-stage and second-stage aggregation constitute our final protocol shown in Algorithm 3.

As a complement: According to the resilience analysis for our first-stage aggregation, any acceptable upload is confined to lie in a special subspace described by Equation 5. To deceive our first-stage aggregation, the Byzantine attacker can also enforce its upload has the same form as $g = g' + z$ with $\|z\| \gg \|g'\|$ where z is the DP noise and g' is malicious component. Now the question is:

Is there an effective way for the server to differentiate benign uploads from Byzantine uploads based on the different nature of \tilde{g} and g' ?

The answer is yes if the server can get some estimate on the true gradient. To have such a capability, we assume that the server has access to some auxiliary data that can be used to compute the gradient during the training. Our empirical finding shows that *two samples per class* are enough for our second-stage aggregation to be effective. The intuition is that benign \tilde{g} should update the model towards roughly the same direction as the true gradient ∇F while the malicious one does not. Quantitatively speaking, with high confidence, $\mathbb{E} \langle \nabla F, \tilde{g} \rangle > \mathbb{E} \langle \nabla F, g' \rangle$. And the server can use the gradient of non-private data to approximate ∇F .

Theoretical motivation: For simplicity, based on our DP protocol, at a certain iteration, considering one honest worker's upload is⁷:

$$g = \frac{\nabla f(x; w)}{\|\nabla f(x; w)\|} + z = \frac{\nabla F(w) + \xi}{\|\nabla F(w) + \xi\|} + z \quad (6)$$

where $z \sim \mathcal{N}(0, \sigma^2 I)$ is the DP noise and $\nabla f(x; w)$ can be written as $\nabla F(w) + \xi$ where ξ is random noise due to data sampling. We compute the expectation of the inner product between g and the true gradient $\nabla F(w)$, which has the following inequality (the proof is given by A.5.1 in supp. material.).

$$\begin{aligned} \mathbb{E} \langle \nabla F(w), g \rangle &= \mathbb{E} \left\langle \nabla F(w), \frac{\nabla F(w) + \xi}{\|\nabla F(w) + \xi\|} \right\rangle \\ &\geq \frac{\mathbb{E} \|\nabla F(w)\|}{3} - \frac{8\mathbb{E} \|\xi\|}{3}, \end{aligned} \quad (7)$$

where the expectation is taken over the randomness of the data sampling and DP noise. Note that Equation (6) is the special case where the honest worker only uses one data sample to compute the gradient and in the general case where many data samples are used, Equation (7) still holds (expectation is linear with respect to sum operation). Again, we do not have such bound if using clipping.

We then consider $\mathbb{E} \langle \nabla F, g' \rangle$ for the attacks we consider: 1) for Gaussian attack, $\mathbb{E} \langle \nabla F(w), g \rangle = 0$; 2) for Label-flipping attack, we hypothesize that $\mathbb{E} \langle \nabla F(w), g \rangle \leq 0$ as such Byzantine gradient is to destroy our learning; 3) for Optimized Local Model Poisoning attack, we have $\mathbb{E} \langle \nabla F(w), g \rangle < 0$ as this is the goal of such attack. In total, for the three attacks, we have $\mathbb{E} \langle \nabla F(w), g \rangle \leq 0$.

As $\mathbb{E} \|\xi\|^2$ is bounded, we can be confident that at least at the early phase of training, $\mathbb{E} \|\nabla F(w)\|$ is large enough to satisfy $\frac{\mathbb{E} \|\nabla F(w)\|}{3} - \frac{8\mathbb{E} \|\xi\|}{3} > 0$. Our empirical result give positive evidence on the correctness of $\frac{\mathbb{E} \|\nabla F(w)\|}{3} - \frac{8\mathbb{E} \|\xi\|}{3} > 0$. Thus, we can use this to filter Byzantine uploads and this is the foundation for our second-stage aggregation.

4.5 Final Byzantine-resilient Protocol

Combining all stages: According to the above statements, we design our second-stage aggregation which is shown in line 4-14 in Algorithm 3. In line 4, the server gets an estimation on ∇F by computing the gradient using some non-private data; Since the server has a prior belief that at least $\lceil \gamma n \rceil$ workers are honest, in line 9, server gets the average on top $\lceil \gamma n \rceil$ inner product scores

⁷We assume the batch size is 1. The iteration number is omitted for ease of notation as it is clear from the context.

Algorithm 3 FilterGradient($\{g_i^t | i = 1, 2, \dots, n\}, w_{k-1}$)

Input: gradients from each worker i at the t -th iteration g_i^t , model w^{t-1} , server-hold dataset D_p , the loss function $f(\cdot)$, server-maintained score list \mathbb{S} , server's belief of honest worker ratio γ

```

1: for  $i = 1, 2, \dots, n$  do in parallel
2:    $g_i^t \leftarrow \text{FirstAGG}(g_i^t)$ 
3: end for
4: Server computes  $g_s^t \leftarrow \nabla f(D_p; w^{t-1})$ 
5: Server initialize  $\mathbb{S}_{tmp} = [0, 0, \dots, 0]$ 
6: for  $i = 1, 2, \dots, n$  do in parallel
7:    $\mathbb{S}_{tmp}[i] = \langle g_i^t, g_s^t \rangle$   $\triangleright$  Motivated by our analysis in Section 4.4
8: end for
9:  $\hat{\mu} \leftarrow$  average of top  $\lceil \gamma n \rceil$  scores in  $\mathbb{S}_{tmp}$ 
10: for  $i = 1, 2, \dots, n$  do in parallel
11:    $\mathbb{S}_{tmp}[i] \leftarrow 0$  if  $\mathbb{S}_{tmp}[i] < \hat{\mu}$ 
12:    $\mathbb{S}[i] \leftarrow \mathbb{S}[i] + \mathbb{S}_{tmp}[i]$ 
13: end for
14: Select those upload inside  $\{g_i^t | i = 1, 2, \dots, n\}$  which correspond to top  $\lceil \gamma n \rceil$  scores inside  $\mathbb{S}$  to form set  $G_s^t$ 
Output:  $G_s^t$ 

```

among all scores computed by current upload of each worker (line 6-8). This average is used as the threshold to suppress scores lower than it to zero in line 11. By processing all scores by using the threshold, we can suppress all scores corresponding to Byzantine uploads and preserve the benign ones; the processed scores are accumulated in line 12 to be used to differentiate benign uploads from Byzantine ones which is described in line 14. Then the selected vectors are returned for the model update.

Novelties: Our first-stage aggregation is the first aggregation rule leveraging the aforementioned DP properties. Leveraging auxiliary data to aid Byzantine aggregation is not new [11, 52], nonetheless, our second-stage aggregation differs from all existing approaches in 1) theoretical support: we have a solid theoretical explanation while previous work stands on heuristics; 2) differentiation metric: we use inner product while previous work use cosine similarity (cosine similarity never leads to the lower bound in Equation 7); 3) the way to integrate any upload into model update: by a unifying language, in our protocol, the weight assigned to any upload is binary (1 or 0), while existing work use real-valued weights according to computed similarities, we find that when DP is enforced (noise is added to the upload), assigning real-valued weights to any upload results in further biasing gradient which leads to rubbish model update.

4.6 Byzantine Attacks to Our Protocol

Recall that our attacker is the stronger version. In consistency with such consideration and to test the limit of our protocol's resilience, we stand in the perspective of an attacker and form possible attacks based on our already-released protocol.

Attacker's response: First, the attacker has to pass our first-stage aggregation to possibly have a malicious impact. Hence, he only has 2 possible guidelines in general:

Guideline 1: The attacker must first generate a d -element ordered sequence according to Theorem 2. Then, the attacker will form any permuted version of such sequence to be malicious. If the attacker is content with any permutation, this would be *Gaussian*

attack [52, 77] as mentioned before. If the attacker aims to find any particular order, he will fail because it incurs $O(d!)$ computation complexity.

Guideline 2: Just like the honest workers' upload, the attacker can make his malicious upload has the form (or can be decomposed to such form) as $g = g' + z$ with $\|z\| \gg \|g'\|$ to pass the KS test.

For completeness, we will not only test on *Gaussian attack* but also test on other attacks which comply with Guideline 2. Following previous work, we will include *Label-flipping attack* [11, 22] and *Optimized Local Model Poisoning attack* [22]. The former has been described in previous sections and we will explain how to form the latter attack in the following.

The attacker forms Optimized Local Model Poisoning attack by the following meta procedure:

- Infer the aggregated result g_r by applying the aggregation rule on all benign uploads;
- Based on the result and the aggregation rule, he forms his Byzantine upload which passes the aggregation and makes the final aggregated result have the inverse direction compared to g_r .

Accordingly, the goal of the adversary is to pass the first-stage aggregation to be possibly malicious further. We formally summarize such strategy as the following optimization problem:

$$\begin{aligned} \min_{\{g_{M_i}\}} \quad & S_c \left(\sum g_{M_i} + \sum g_{B_j}, \sum g_{B_j} \right) \\ \text{s.t.} \quad & \|\text{FirstAGG}(g_{M_i})\| > 0, \end{aligned} \quad (8)$$

where the constraint means that the Byzantine upload can pass our first-stage aggregation, $\{g_{M_i}\}$ are all Byzantine uploads by the Byzantine attacker, and $\{g_{B_j}\}$ are all benign uploads by honest workers. The function $S_c(A, B) = \frac{AB}{\|A\|\|B\|}$ calculates the cosine similarity between two vectors. Suppose all benign uploads are $g_{B_1}, g_{B_2}, \dots, g_{B_m}$ by B_m honest workers and all Byzantine uploads are $g_{M_1}, g_{M_2}, \dots, g_{M_n}$ by M_n Byzantine workers. According to Equation (8), the attacker aims to reach the following goal:

$$\sum g_{M_i} = -(1 + \lambda) \sum g_{B_j} \quad (9)$$

where $\lambda > 0$ is a positive number. This leads to the term $\sum g_{M_i} + \sum g_{B_j} = -\lambda \sum g_{B_j}$ results in the inverse direction compared to $\sum g_{B_j}$. By setting:

$$g_{M_1} = g_{M_2} = \dots = -\frac{(1 + \lambda)}{M_n} \sum g_{B_j} \quad (10)$$

this goal is reached. And setting $\lambda = \frac{M_n}{\sqrt{B_m}} - 1$ will let Byzantine uploads pass our first-stage aggregation (one can check that all malicious and benign upload behaves the same when applying our first-stage aggregation on them). Note that to be able to perform such an attack, we need $M_n > \sqrt{B_m}$ (because $\lambda > 0$), that is, such a strong attack only exists when the number of Byzantine workers is sufficiently large.

Note that we do not simulate the Optimized Local Model Poisoning attack compromising our second-stage aggregation because the attacker must know the serve-hold auxiliary dataset, and this means that the attacker must fully control the server which is unrealistic. Also note that, for the goal in Equation (8), we set it to be

the inverse to the sum of all benign uploads. This is because such a goal leads to an efficient solution that can be tolerated by the attacker. In fact, the attacker can choose its goal freely as long as the constraint in Equation (8) is satisfied. However, other goals may not lead to an efficient solution. For instance, if the attacker chooses it to be orthogonal, the attacker is faced with the hard problem as discussed in Guideline 1.

Discussion on the adaptive attack: There exists another attack that copies benign uploads by honest workers for some iterations and suddenly turns to be malicious after that. We call this attack as *adaptive attack*. The way the attacker is malicious can be any instantiation of the previous three attacks we mentioned before. We will also include this attack in our experiment for completeness. Note that although Optimized Local Model Poisoning attack seems to be more advanced than Gaussian attack and Label-flipping attack, it is unclear which attack is most successful on our protocol before the experiment.

Discussion on excluded attacks: Optimized Local Model Poisoning attack performs well on attacking various existing Byzantine defense methods [22]. Another similar recent work [55] adopts the attacking intuition (the meta procedure mentioned above) of the Optimized Local Model Poisoning attack in other cases where the attacker's power is more limited (the attacker is weaker). Hence, here we only adopt the Optimized Local Model Poisoning attack in our experiments.

To the best of our knowledge, many other attacks can be trivially defended by our protocol, such as the attacks that have been considered in the existing work: "A little" attack [6] and "Inner" attack [70]. "A little" attack involves estimating the coordinate-wise mean and the s.t.d. of benign uploads to form its attack. However, our learning protocol enforces that the DP noise is dominating, hence knowing benign uploads gains the attacker no useful information when forming "A little" attack. Most importantly, naively applying such an attack will end up being rejected by first-stage aggregation. This shows the power of our protocol.

4.7 Discussions

First-stage aggregation provides critical robustness: Only using our second-stage aggregation to aggregate all worker's uploads is not enough, because, due to randomness, it is not guaranteed that Byzantine upload will never be selected for model update, and selected Byzantine upload could destroy our model in just one iteration as it is arbitrary. In contrast, there exists no such concern when we apply our first-stage aggregation, according to previous resilience analysis for our first-stage aggregation, it enforces any upload (including malicious ones) g which passes the filtering to have the form $g = \hat{g} + z$ with $\|z\| \gg \|\hat{g}\|$ where \hat{g} is strictly norm-bounded and z is the DP noise. For all malicious uploads, strictly norm-bounded \hat{g} means their detrimental impact is bounded.

DP-Byzantine-robustness interaction: we do not consider DP and Byzantine-robustness in isolation. Instead, our whole protocol is formed by leveraging each other's properties.

As mentioned previously in our design strategy, we use our first-stage rule to "constrain" the way that any upload should behave by re-designing our DP protocol so that any Byzantine upload violating it will be immediately rejected. Hence, other than only

protecting privacy, this refactoring on DP also provides the first-stage Byzantine-robustness. To deal with those Byzantine uploads that pass our first-stage aggregation, we further design our second-stage rule to do the “complementary aggregation” by leveraging the properties of our refactored DP protocol. In total, our privacy protocol and the robust aggregation rule are aware of each other, leading to a solution that is both privacy-preserving and Byzantine-resilient.

5 THEORETICAL GUARANTEES

We provide theoretical guarantees on privacy, utility, and Byzantine robustness of our protocol in this section. For convenience, we assume that the dataset of each worker has the same size which is denoted as $|D|$, and the size of non-private data held by the server is $|D_p|$. We also denote $w^* = \arg \min_{w \in \Theta} F(w)$.

Privacy guarantee: We have the following privacy guarantee.

THEOREM 3. (Privacy Guarantee) *There exist constants c_1 and c_2 such that given the sampling rate $q = \frac{b_c}{|D|}$ and the number of iteration steps T . For each worker, Algorithm 1 is (ϵ, δ) -DP for any $\delta > 0$ and $\epsilon < c_1 q^2 T$ if $\sigma \geq c_2 \frac{q \sqrt{T \ln \frac{1}{\delta}}}{\epsilon}$.*

PROOF. See A.5.2 in supp. material. \square

Utility and Byzantine robustness: Theorem 4 shows the utility and robustness of Byzantine resilience of our algorithm. Before formally introducing Theorem 4, we present some assumptions, which are commonly used in the previous work on optimization and Byzantine-robust learning [11, 15].

ASSUMPTION 1. *The expected loss function $F(w)$ is μ -strongly convex and differentiable over the space Θ with L -Lipschitz continuous gradient. Formally, we have the following for any $w, \hat{w} \in \Theta$:*

$$F(\hat{w}) \geq F(w) + \langle \nabla F(w), \hat{w} - w \rangle + \frac{\mu}{2} \|\hat{w} - w\|^2$$

$$\|\nabla F(w) - \nabla F(\hat{w})\| \leq L \|w - \hat{w}\|.$$

Moreover, the empirical loss function $f(D, w) := \frac{1}{|D|} \sum_{x \in D} f(x; w)$ is L_1 -Lipschitz continuous with high probability. Formally, for any $\zeta \in (0, 1)$, there exists an L_1 such that:

$$\Pr \left(\sup_{w, \hat{w} \in \Theta: w \neq \hat{w}} \frac{\|\nabla f(D, w) - \nabla f(D, \hat{w})\|}{\|w - \hat{w}\|} \leq L_1 \right) \geq 1 - \frac{\zeta}{3}$$

ASSUMPTION 2. *The gradient of the empirical loss function $\nabla f(D, w^*)$ at the optimal global model w^* is bounded. Moreover, the gradient difference $h(D, w) = \nabla f(D, w) - \nabla f(D, w^*)$ for any $w \in \Theta$ is bounded. Specifically, there exist positive constants σ_1 and γ_1 such that for any unit vector v , $\langle \nabla f(D, w^*), v \rangle$ is sub-exponential with σ_1 and γ_1 ; and there exist positive constants σ_2 and γ_2 such that for any $w \in \Theta$ with $w \neq w^*$ and any unit vector v , $\langle h(D, w) - \mathbb{E}[h(D, w)], v \rangle / \|w - w^*\|$ is sub-exponential with σ_2 and γ_2 . Formally, for all $|\tau| \leq 1/\gamma_1$, $|\tau| \leq 1/\gamma_2$, we have:*

$$\sup_{v \in B} \mathbb{E} [\exp(\tau \langle \nabla f(D, w^*), v \rangle)] \leq e^{\sigma_1^2 \tau^2 / 2}$$

$$\sup_{w \in \Theta, v \in B} \mathbb{E} \left[\exp \left(\frac{\tau \langle h(D, w) - \mathbb{E}[h(D, w)], v \rangle}{\|w - w^*\|} \right) \right] \leq e^{\sigma_2^2 \tau^2 / 2}$$

where B is the unit sphere $B = \{v : \|v\| = 1\}$

Note that the strongly convex and Lipschitz continuous conditions in Assumption 1 are widely adopted in the convergence analysis of optimization algorithms, and these conditions indicate the largest eigenvalue of the Hessian matrix of the loss function is between μ and L . Assumption 2 indicates that the gradient $\nabla f(D, w^*)$ is quite close to its expectation $\mathbb{E}[\nabla f(D, w^*)] = 0$, and the difference $h(D, w)$ concentrates to its expectation with high probability.

THEOREM 4. *For an arbitrary number of Byzantine workers, the difference between the global model learned by Algorithm 1 and the optimal global model w^* under no attacks is bounded. Specifically, if the parameter space $\Theta \subseteq B(0, r\sqrt{d})$, i.e., it is contained in a ball with radius $r\sqrt{d}$ and $\nabla F(w^*) = 0$. Set σ as in Theorem 3, $T = O\left(\frac{1}{\rho} \ln\left(\sqrt{n}|D|\sqrt{|D_0|}\right)\right)$ and $\eta_{t-1} \leq \|g_s^{t-1}\|_2 \eta_0$ with fixed $\eta_0 \leq \frac{\mu}{2L^2}$ in the t -th iteration in Algorithm 1, then if n , $|D_p|$ and $|D|$ are sufficiently large and η_0 is sufficiently small such that*

$$\sqrt{n} \geq \tilde{\Omega} \left(\frac{\sqrt{d \ln \frac{1}{\delta}}}{\epsilon |D|} \cdot \max \left\{ \sqrt{\ln \frac{1}{\xi}}, \frac{\ln \frac{1}{\xi}}{r \rho \sqrt{|D_p|}} \right\} \right) \quad (11)$$

and $\frac{\eta_0 \sigma_1}{\sqrt{|D_p|}} \leq O\left(\frac{r\sqrt{d}}{\rho}\right)$ with $0 < \rho < 1$. Then, with probability at least $1 - \xi$ with $\xi \in (0, 1)$, we have:

$$\|w_T - w^*\| \leq \tilde{O} \left(\frac{1}{\rho^2} \frac{d \ln \frac{1}{\xi} \sqrt{\ln \frac{1}{\delta}} \sigma_1}{|D| \sqrt{n} b_s \epsilon} + \frac{1}{\rho} \frac{\sigma_1 \sqrt{d \ln \frac{1}{\xi}}}{\sqrt{|D_p|}} \right), \quad (12)$$

where the Big- \tilde{O} and Big- $\tilde{\Omega}$ notations omit other logarithmic terms. Here $\rho = 1 - \sqrt{1 - \frac{u^2}{4L^2} - 32\eta_0\Delta_2 - 3\eta_0L}$ with $\Delta_2 = \sigma_2 \sqrt{\frac{2}{|D_p|}} \sqrt{K_1 + K_2}$ with $K_1 = d \log \frac{\max\{L, L_1\}}{\sigma_2}$ and $K_2 = \frac{d}{2} \log \frac{|D_p|}{d} + \log \frac{6\sigma_2 r \sqrt{|D_p|}}{\gamma_2 \sigma_1 \xi}$.

PROOF. In A.5.3 in supp. material. \square

Theorem 4 is on the robustness. Briefly speaking, as can be seen from Equation (12), if $|D|$, n and $|D_p|$ are large enough (nonetheless, our experiment shows that only a small number of non-private data will suffice), with some iteration number and stepsize, even there is an arbitrary number of Byzantine workers, the final model we get will be close to the optimal model (measured by the l_2 distance) with high probability.

6 EXPERIMENTAL RESULTS

6.1 Datasets and System Settings

Datasets and models: We conduct experiments on MNIST [41], Colorectal [37], Fashion [69], and USPS [34]. Details of all these benchmark datasets with various properties are summarised in supp. material. Details of neural network setup are also in supp. material. Each experiment is repeated with different random seeds $\{1, 2, 3\}$ and we report the min. max. and mean. All of our experiments are conducted under the same base learning rate $\eta_b = 0.2$ (which will be explained later), batch size $b_c = 16$ and momentum $\beta = 0.1$. We set the number of epochs $T = \lceil 10|D|/b_c \rceil$ for Colorectal and USPS, $T = \lceil 8|D|/b_c \rceil$ for MNIST and Fashion.

Data sample distribution: We consider both *i.i.d.* and *non-i.i.d.* settings. To be specific, *i.i.d.* is the case where each worker's local

dataset follows the same distribution as the whole data population while *non-i.i.d.* is the case where each worker's local dataset's distribution is arbitrary [45]. We simulate both settings following previous work [11, 22, 30, 77], and details are presented in supp. material.

For generating server-own auxiliary data, we only randomly sample 2 data samples per class from the validation dataset. As mentioned earlier, obtaining such a tiny amount of data is easy. Note that generating such auxiliary data is totally agnostic to the distribution of the whole data population while such data still enables our protocol's effectiveness (as will be confirmed by our experiments). Once the auxiliary data is generated, it is vacuous to compare the distribution of such auxiliary data to distributions of any other datasets, because the size of our auxiliary data is micro.

Byzantine setup: We fix the number of honest workers (20 for MNIST and Fashion, 10 for Colorectal and USPS), and vary the number of Byzantine workers (0%, 20%, 40%, 60%, 90% of total).

Privacy settings: We do experiments on different privacy settings $\epsilon = \{2^{-3}, 2^{-2}, 2^{-1}, 2^0, 2^1\}$ while fixed $\delta = 1/|D_i|^{1.1}$, where $|D_i|$ is the size of the local dataset possessed by worker i .

Reference Accuracy: The *Reference Accuracy* is the testing accuracy of FL under the scenario where no Byzantine threat exists and FL only adopts DP (not adopting any Byzantine defense method). Compare any private and Byzantine-resilient protocol's performance to the *Reference Accuracy*, many useful conclusions can be drawn:

- **Side-effect:** Apply a protocol under the scenario where there are no Byzantine threats, by comparing it with *Reference Accuracy*, we know how much “side-effect” caused by that protocol. The ideal case is that we expect the “medicine” causes no additional harm to the “patient” with no “illness”.
- **Efficacy:** For the scenario where there is a certain number of attackers, by comparing with *Reference Accuracy*, we know how effectively a protocol defends the attack. The ideal case is that the “medicine” eradicates the “illness” (under such case, the performance should be the same as *Reference Accuracy*).

6.2 Claims and Experimental Evidence

All of the attacks we consider have been tested. Based on the observation that our protocol remains resilient across all attacks and due to space limitation, we arbitrarily only present results for *Label-flipping attack* under *i.i.d.* in the main body. All additional results for other attacks we consider under both *i.i.d.* and *non-i.i.d.* settings is in supp. material.

A quick overview: We provide 7 claims with their corresponding evidence. By comparing with previous work, claim 1-2 show our contribution to DP learning and Byzantine resilience in their own track. Most importantly, recall our core aim is to ensure privacy and Byzantine resilience simultaneously, we use claims 3-7 to show its effectiveness.

CLAIM 1: Normalizing is better than vanilla DP-SGD (which uses clipping) at 1) gaining Byzantine resilience; 2) efficiently tuning hyper-parameter for DP deep learning.

Evidence: A thought experiment will suffice. Recall that clipping is essentially normalizing gradient vectors with ℓ_2 -norm greater

than C to be C exactly and leaving those vectors with ℓ_2 -norm smaller than C untouched. If we are guaranteed that all gradient vectors' ℓ_2 -norm is greater than C , then normalizing and clipping only differ in the learning rate scale and are essentially equivalent to each other, *i.e.*, clipping with $C = 2, \eta = 0.1$ is the same as normalizing (to be unit ℓ_2 -norm) with $\eta = 0.2$. This means that if we have that guarantee, clipping will also enjoy the lower bound in Equation 7 for gaining Byzantine resilience and will also enjoy our analysis for hyper-parameter tuning in Theorem 1.

However, it is *unfeasible* to get a *prior bound* on the gradient vector's norm for arbitrary deep-learning neural networks. Meanwhile, it is unclear whether clipping could lead to similar theoretical results which serve our purpose. Adopting normalizing circumvents such issues.

CLAIM 2: Our protocol outperforms existing solutions.

Evidence: We compare our protocol to previous work with the same aim (ensure privacy and Byzantine security simultaneously). We will show that our tailored Byzantine aggregation with DP outperforms previous solutions whose methodology is to naively apply off-the-shelf Byzantine aggregation with DP. And our result shows the contribution of our Byzantine aggregation rule.

For a fair comparison, we provide the results for the scenario where our privacy level is similar and our attacker is the same compared with existing solutions.

Method	Byz./ Privacy	“A little” attack [6]	“Inner” attack [70]
[30]	40%, $\epsilon = 3.46$.61	.75
	20%, $\epsilon = 7.58$.78	.79
Ours	60%, $\epsilon = 2.00$.79 \pm .010	.80 \pm .010
	40%, $\epsilon = 2.00$.80 \pm .005	.80 \pm .005

Table 2: Testing accuracy comparison with existing work [30] on Fashion.

Method	Byz./ Privacy	Gaussian attack
[77]	10%, $\epsilon = .21$.20
	10%, $\epsilon = .40$.43
Ours	60%, $\epsilon = .125$.86 \pm .010
	40%, $\epsilon = .125$.86 \pm .010

Table 3: Testing accuracy comparison with existing work [77] on MNIST.

Comparison with [30]: We compare our results with [30] in Table 2. We can see from Table 2 that [30] only reaches 61% accuracy under 40% “a little” Byzantine attack [6] in the privacy setting ($\epsilon = 3.46, \delta = 1.2 \times 10^{-4}$). We also notice that under the same privacy setting but a different Byzantine attack, [30] achieves 75% testing accuracy, and [30] makes comments that “a little” attack is stronger against their defense.

Applying our Byzantine defense method under the same attacks, we get around 80% testing accuracy when there are 60%, 40% Byzantine workers in the privacy setting ($\epsilon = 2, \delta = 1.4 \times 10^{-4}$). Thus, we can gain much more utility compared to [30] even when the majority worker are Byzantine and with even better privacy guarantee (we are ensuring ($\epsilon = 2, \delta = 1.4 \times 10^{-4}$)-DP instead of ($\epsilon = 3.46, \delta = 1.2 \times 10^{-4}$)-DP). The utility we gain is also better

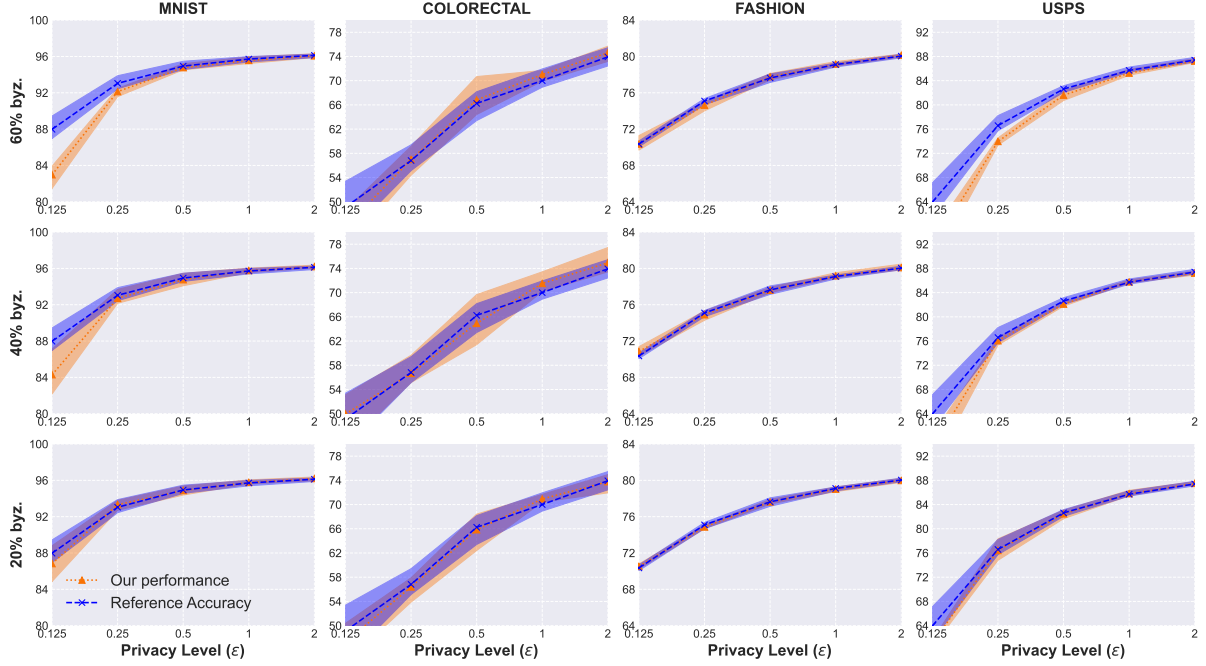


Figure 1: Byzantine-resilient performance (testing accuracy) under Label-flipping attack. The experiment is conducted under 3 different attacking levels (20%, 40%, 60% of the total workers are Byzantine).

than [30] under its weakest attack with a much weaker privacy guarantee: ($\epsilon = 7.58, \delta = 1.2 \times 10^{-4}$)-DP.

Comparison with [77]: As can be seen from Table 3, the method in [77] reaches 43% testing accuracy on MNIST when there are only 10% Byzantine workers under the privacy setting ($\epsilon = 0.4, \delta = 0$). As a comparison, our learning protocol provides 86% testing accuracy when there are 60% Byzantine workers under privacy setting ($\epsilon = 0.125, \delta = 1.4 \times 10^{-4}$). We gain much more utility even when the majority of workers are Byzantine, which is impossible for [77] to accomplish due to their intrinsic limitation of aggregation methods.

	MNIST			COLOR.			FASHION			USPS		
	$\epsilon = \frac{1}{8}$	$\frac{1}{2}$	2	$\frac{1}{8}$	$\frac{1}{2}$	2	$\frac{1}{8}$	$\frac{1}{2}$	2	$\frac{1}{8}$	$\frac{1}{2}$	2
RA	.88	.95	.96	.49	.66	.74	.69	.77	.80	.64	.82	.87
zero	.85	.94	.96	.44	.67	.74	.69	.77	.80	.58	.81	.87

Table 4: Experimental result on the test for the “side-effect” our protocol brings. RA stands for *Reference Accuracy* and *zero* stands for the scenario where all original 60% Byzantine workers turn to behave honestly (hence we have *zero* attackers) while our protocol is still applied. The performance (testing accuracy) results from taking the average of three runs with different seeds.

CLAIM 3: Our protocol brings no “side-effect” even when there is no Byzantine attack.

Evidence: we design the following experiment to test whether our protocol brings any “side-effect”. Our results are shown in Table 4.

We can see that other than at the extreme privacy level ($\epsilon = 1/8 = .125$), our protocol’s performance is almost identical to the *Reference Accuracy*, hence incurring no “side-effect”. We indeed observe a noticeable accuracy drop when at $\epsilon = 1/8 = .125$, this is because in such extreme case, noise becomes so overwhelming that the training itself is not stable.

CLAIM 4: Our protocol eradicates Byzantine attacks if not facing extreme privacy requirements.

Evidence: Figure 1 shows the performance of our method. The testing accuracy almost always aligns with the *Reference Accuracy*. Such a phenomenon can be observed not only across different privacy levels but also across different datasets.

The most discrepant results are observed for USPS and MNIST datasets when there are 60% Byzantine workers in the high privacy regime with $\epsilon = 0.125$. This is because, at this extreme privacy level, significant noise is added to the gradient. We are getting less confident in differentiating benign uploads from Byzantine ones when we are at such a high privacy level, fortunately, our first-stage aggregation guarantees that malicious upload has limited detrimental impact even if selected.

We also observe that results for Colorectal present a larger variance than the rest datasets. This is because the dataset size is much smaller (only 5,000 samples in total) than the rest and it is limited by the intrinsic limitation that DP learning requires large-scale data.

CLAIM 5: Our protocol remains robust against majority attack.

Evidence: Results are shown in Figure 2. We can observe that even when 90% workers are Byzantine (we have also simulated more stringent cases where 95%, 99% workers are Byzantine, results can be found in *supp. material*), similar results can be observed




Figure 2: Byzantine-resilient performance (testing accuracy) when 90% workers are Label-flipping Byzantine attackers.

compared with the cases where there are 60%, 40%, 20% Byzantine workers. We observe a noticeable accuracy drop for certain datasets when $\epsilon = 0.125$ and $\epsilon = 0.25$ due to overwhelming random noise which guarantees high privacy. For $\epsilon \geq 0.5$, we still gain privacy and Byzantine resilience without hurting too much performance.

CLAIM 6: As a cherry on top, our protocol enables efficient hyper-parameter tuning by saving quadratic efforts.

Evidence: For a typical DP deep learning task, vanilla DP-SGD’s running task spans on the 3-dimensional tuple (η, C, ϵ) . In contrast, adopting normalizing together with our tuning strategy only needs to tune η for one arbitrary ϵ . That is, we only need to tune the learning rate η_b for one privacy level ϵ with the corresponding noise multiplier σ_b , then we can use the learning rate $\eta = \frac{\eta_b \sigma_b}{\sigma}$ for any other privacy level with noise multiplier σ . We call η_b and σ_b at the privacy level we are tuning as “base learning rate” and “base noise multiplier”.

To evaluate the effectiveness of such a strategy, it suffices to confirm that if we find the optimal base learning rate for one privacy level, we also find the optimal learning rate for other privacy levels by setting the learning rate according to such a strategy. In this sense, we first choose the base case of $\sigma_b = 0.79$ (corresponding to $\epsilon = 2$). Then, for each privacy level, we tune the learning rate with respect to different base learning rates (the actual learning rate is computed according to the above strategy). In our experiment, we vary the base learning rate among $\{0.02, 0.04, 0.08, 0.2, 0.4, 0.8, 1\}$, so the actual learning rate will be $\{\frac{0.02\sigma_b}{\sigma}, \frac{0.04\sigma_b}{\sigma}, \frac{0.08\sigma_b}{\sigma}, \frac{0.2\sigma_b}{\sigma}, \frac{0.4\sigma_b}{\sigma}, \frac{0.8\sigma_b}{\sigma}, \frac{\sigma_b}{\sigma}\}$ for a specific privacy level with noise multiplier σ .

Results on MNIST are shown in Figure 3, and we can see that for all the privacy levels we considered, the optimal point is the same ($\frac{0.2\sigma_b}{\sigma}$ for MNIST), and a similar phenomenon can also be observed on the other three datasets.

TTBB	MNIST		COLOR.		FASHION		USPS	
	$\epsilon = 2$	$\epsilon = .125$	$\epsilon = 2$	$\epsilon = .125$	$\epsilon = 2$	$\epsilon = .125$	$\epsilon = 2$	$\epsilon = .125$
0	.96	.82	.74	.45	.80	.68	.86	.60
.2	.96	.82	.74	.41	.80	.68	.86	.60
.4	.96	.81	.73	.45	.80	.68	.86	.57
.6	.96	.81	.73	.44	.80	.69	.86	.57
.8	.96	.82	.73	.43	.80	.69	.86	.60

Table 5: Under Label-flipping attack with different TTBB.

CLAIM 7: Our protocol remains resilient against adaptive attack.

Evidence: Our robust and private learning framework is also resilient to *adaptive attack*. We evaluate that by letting 60% Byzantine workers be honest via copying the uploads of some random honest

workers from the beginning of training and turning to Byzantine at different iterations to see if they can possibly have a significant impact. Results are shown in Table 5. The first column represents the Time To Be Byzantine (TTBB), i.e., if the total iteration is T , 0.2 TTBB means that Byzantine workers behave honestly within the first $0.2T$ iterations and then start to send Byzantine uploads thereafter.

We can see that no matter when the Byzantine workers start to be Byzantine, they all have a negligible impact on the testing accuracy except for the case with extreme privacy requirements. We notice that there are some mild performance fluctuations when $\epsilon = 0.125$ for Colorectal and USPS, again, due to the large variance of DP noises.

γ	MNIST		COLOR.		FASHION		USPS	
	$\epsilon = \frac{1}{8}$	$\epsilon = 2$	$\epsilon = \frac{1}{8}$	$\epsilon = 2$	$\epsilon = \frac{1}{8}$	$\epsilon = 2$	$\epsilon = \frac{1}{8}$	$\epsilon = 2$
20%	.86	.95	.48	.73	.66	.78	.64	.85
35%	.87	.96	.47	.74	.69	.79	.63	.86
50% (exact)	.88	.96	.49	.74	.69	.80	.64	.87
65%	.85	.96	.45	.73	.70	.79	.56	.87
80%	.83	.95	.34	.74	.69	.79	.54	.85

Table 6: γ is treated as a prior belief in this experiment and we study the effect when there is a mismatch between such belief and the truth. We fix the setting that 50% workers are honest and vary γ . For the case where the belief is exactly the truth ($\gamma = 50\%$), we denote it as “exact”. All results are obtained by taking the average of three runs with different random seeds.

6.3 More Experimental Results

Convergence behavior: The convergence curve is presented in Figure 4. As can be seen in Figure 4, the training converges in the first several epochs. The convergence behavior of our protocol aligns well with “Reference Accuracy” even when we have 60% Byzantine workers. Similar to previous results in our CLAIM 4, We observe a larger variance for Colorectal than that of the rest datasets. As expected, this is due to its significantly small dataset size and the nature of training with DP.

Ablation study on γ : Recall that previously we assumed the server knows that at least γn workers are honest, what if γ is only a (prior) belief rather than the truth, and moreover, what if there is a mismatch between such belief and the truth? We further conduct an ablation study on γ if it is only a belief. We can see from Table 6 that, in the case where 50% workers are honest, as long as the server is conservative ($\gamma \leq 50\%$), we can still retain robustness. In

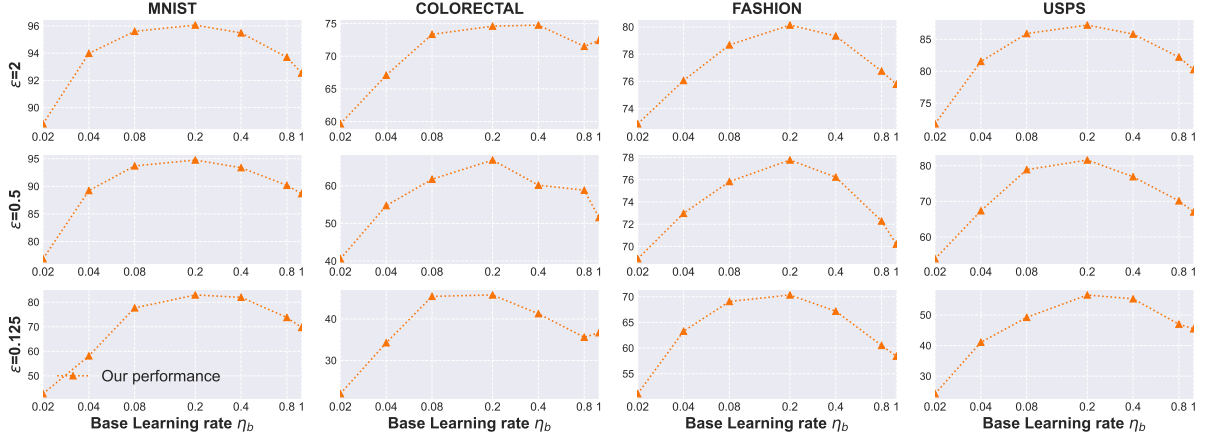


Figure 3: Our hyper-parameter tuning results when facing 60% Label-flipping attackers.

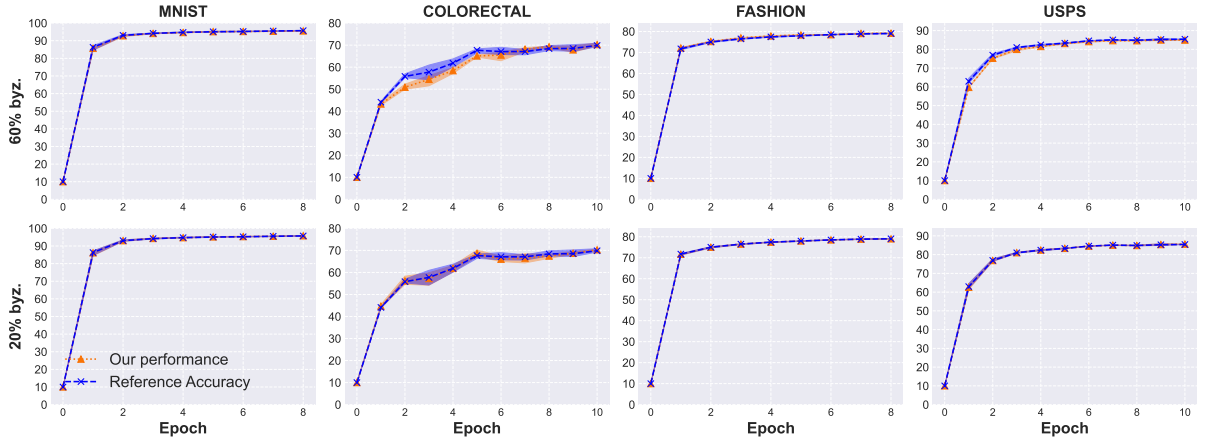


Figure 4: Byzantine-resilient convergence curves (testing accuracy) under Label-flipping attack (considering 20%, 60% of the total workers are Byzantine, fixing $\epsilon = 1$).

contrast, we observe a notable utility drop for Colorectal and USPS under privacy level $\epsilon = .125$ when the server radically believes that 80% workers are honest, this is because in our protocol, being radical (γ is greater than the true honest portion) means the server tends to aggregate malicious uploads. Hence, the more radical, the worse the utility is expected to be. Based on such observation, the learned lesson is that we can always have robustness if we are not facing extreme privacy requirements and a conservative γ is set.

7 CONCLUSION

In this paper, with the aim to ensure both DP and Byzantine resilience for FL systems, we developed a learning protocol resulting from a co-design principle. We refactor the DP-SGD algorithm and tailor the Byzantine aggregation process towards each other to form an integrated protocol. For our DP-SGD variant, the small batch size property enables our first-stage Byzantine aggregation which trivially rejects many existing Byzantine attacks; the normalization technique enables our second-stage aggregation which

provides a final sound filtering. As a cherry on top, normalizing also enables our efficient hyper-parameter tuning strategy which saves quadratic efforts. We also provide theoretical explanations behind the efficacy of our protocol.

In the experiment part, we first provide evidence to support our contribution claim to both DP learning and Byzantine security tracks in separation, we then provide evidence of the effectiveness of our protocol tackling the two-fold issue, *i.e.*, an FL system needs to be privacy-preserving and Byzantine-resilient simultaneously. We have shown that our protocol does not incur “side-effects” to a system with no Byzantine attacker, and we have also seen that our protocol remains Byzantine-resilient even when there are up to 90% distributive workers being Byzantine.

REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. 308–318.

- [2] Naman Agarwal, Peter Kairouz, and Ziyu Liu. 2021. The skellam mechanism for differentially private federated learning. *Advances in Neural Information Processing Systems* 34 (2021).
- [3] Rohan Anil, Badih Ghazi, Vineet Gupta, Ravi Kumar, and Pasin Manurangsi. 2021. Large-scale differentially private bert. *arXiv preprint arXiv:2108.01624* (2021).
- [4] Shahab Asodeh, Jiachun Liao, Flavio P Calmon, Oliver Kosut, and Lalitha Sankar. 2021. Three variants of differential privacy: Lossless conversion and applications. *IEEE Journal on Selected Areas in Information Theory* 2, 1 (2021), 208–222.
- [5] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. 2020. How to backdoor federated learning. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2938–2948.
- [6] Gilad Baruch, Moran Baruch, and Yoav Goldberg. 2019. A little is enough: Circumventing defenses for distributed learning. *Advances in Neural Information Processing Systems* 32 (2019).
- [7] Raef Bassily, Adam Smith, and Abhradeep Thakurta. 2014. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. IEEE, 464–473.
- [8] Raef Bassily, Om Thakkar, and Abhradeep Guha Thakurta. 2018. Model-agnostic private learning. *Advances in Neural Information Processing Systems* 31 (2018).
- [9] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. 2017. Machine learning with adversaries: Byzantine tolerant gradient descent. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 118–128.
- [10] Zhiqi Bu, Yu-Xiang Wang, Sheng Zha, and George Karypis. 2022. Automatic Clipping: Differentially Private Deep Learning Made Easier and Stronger. *arXiv preprint arXiv:2206.07136* (2022).
- [11] Xiaoyu Cao, Minghong Fang, Jia Liu, and Neil Zhenqiang Gong. 2020. Fltrust: Byzantine-robust federated learning via trust bootstrapping. *arXiv preprint arXiv:2012.13995* (2020).
- [12] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. 2011. Differentially private empirical risk minimization. *Journal of Machine Learning Research* 12, 3 (2011).
- [13] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526* (2017).
- [14] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526* (2017).
- [15] Yudong Chen, Lili Su, and Jiaming Xu. 2017. Distributed statistical machine learning in adversarial settings: Byzantine gradient descent. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1, 2 (2017), 1–25.
- [16] Christopher A Choquette-Choo, Natalie Dullerud, Adam Dziedzic, Yunxiang Zhang, Somesh Jha, Nicolas Papernot, and Xiao Wang. 2021. Capc learning: Confidential and private collaborative learning. *arXiv preprint arXiv:2102.05188* (2021).
- [17] Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. 2018. *Deep Learning for Classical Japanese Literature*. arXiv:cs.CV/1812.01718 [cs.CV]
- [18] Ashok Cutkosky and Harsh Mehta. 2020. Momentum improves normalized sgd. In *International Conference on Machine Learning*. PMLR, 2260–2268.
- [19] Rudrajit Das, Abolfazl Hashemi, Sujay Sanghavi, and Inderjit S Dhillon. 2021. Privacy-Preserving Federated Learning via Normalized (instead of Clipped) Updates. *arXiv preprint arXiv:2106.07094* (2021).
- [20] Soham De, Leonard Berrada, Jamie Hayes, Samuel L Smith, and Borja Balle. 2022. Unlocking high-accuracy differentially private image classification through scale. *arXiv preprint arXiv:2204.13650* (2022).
- [21] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.
- [22] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Gong. 2020. Local Model Poisoning Attacks to {Byzantine-Robust} Federated Learning. In *29th USENIX Security Symposium (USENIX Security 20)*. 1605–1622.
- [23] Fangcheng Fu, Yingxia Shao, Lele Yu, Jiawei Jiang, Huanran Xue, Yangyu Tao, and Bin Cui. 2021. VF²Boost: Very Fast Vertical Federated Gradient Boosting for Cross-Enterprise Learning. In *SIGMOD '21: International Conference on Management of Data, Virtual Event, China, June 20–25, 2021*, Guoliang Li, Zhanhui Li, Stratos Idreos, and Divesh Srivastava (Eds.). ACM, 563–576. <https://doi.org/10.1145/3448016.3457241>
- [24] Fangcheng Fu, Huanran Xue, Yong Cheng, Yangyu Tao, and Bin Cui. 2022. BlindFL: Vertical Federated Machine Learning without Peeking into Your Data. In *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12–17, 2022*, Zachary Ives, Angela Bonifati, and Amr El Abbadi (Eds.). ACM, 1316–1330. <https://doi.org/10.1145/3514221.3526127>
- [25] gboard [n.d.]. Federated Learning: Collaborative Machine Learning without Centralized Training Data. <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>.
- [26] gdpr [n.d.]. General Data Protection Regulation (GDPR). <https://gdpr.eu/what-is-gdpr/>.
- [27] Robin C Geyer, Tassilo Klein, and Moin Nabi. 2017. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557* (2017).
- [28] Aditya Golatkar, Alessandro Achille, Yu-Xiang Wang, Aaron Roth, Michael Kearns, and Stefano Soatto. 2022. Mixed Differential Privacy in Computer Vision. *arXiv preprint arXiv:2203.11481* (2022).
- [29] Sivakanth Gopi, Yin Tat Lee, and Lukas Wutschitz. 2021. Numerical Composition of Differential Privacy. *arXiv preprint arXiv:2106.02848* (2021).
- [30] Rachid Guerraoui, Nirupam Gupta, Rafael Pinot, Sebastien Rouault, and John Stephan. 2021. Combining Differential Privacy and Byzantine Resilience in Distributed SGD. *arXiv preprint arXiv:2110.03991* (2021).
- [31] Rachid Guerraoui, Nirupam Gupta, Rafael Pinot, Sebastien Rouault, and John Stephan. 2021. Differential Privacy and Byzantine Resilience in SGD: Do They Add Up?. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*. 391–401.
- [32] Rachid Guerraoui, Sébastien Rouault, et al. 2018. The hidden vulnerability of distributed learning in byzantium. In *International Conference on Machine Learning*. PMLR, 3521–3530.
- [33] Jihun Hamm, Yingjun Cao, and Mikhail Belkin. 2016. Learning privately from multiparty data. In *International Conference on Machine Learning*. PMLR, 555–563.
- [34] Jonathan J. Hull. 1994. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence* 16, 5 (1994), 550–554.
- [35] Roger Iyengar, Joseph P Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. 2019. Towards practical differentially private convex optimization. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 299–316.
- [36] Matthew Jagielski, Alina Oprea, Battista Biggio, Chang Liu, Cristina Nita-Rotaru, and Bo Li. 2018. Manipulating machine learning: Poisoning attacks and countermeasures for regression learning. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 19–35.
- [37] Jakob Nikolas Kather, Cleo-Aron Weis, Francesco Bianconi, Susanne M Melchers, Lothar R Schadt, Timo Gaiser, Alexander Marx, and Frank Gerrit Zöllner. 2016. Multi-class texture analysis in colorectal cancer histology. *Scientific reports* 6, 1 (2016), 1–11.
- [38] Andrey Kolmogorov. 1933. Sulla determinazione empirica di una legge di distribuzione. *Inst. Ital. Attuari, Giorn.* 4 (1933), 83–91.
- [39] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. 2016. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492* (2016).
- [40] Alexey Kurakin, Steve Chien, Shuang Song, Roxana Geambasu, Andreas Terzis, and Abhradeep Thakurta. 2022. Toward training at imagenet scale with differential privacy. *arXiv preprint arXiv:2201.12328* (2022).
- [41] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [42] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. 2018. Trojaning Attack on Neural Networks. In *25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18–21, 2018*. The Internet Society. http://wp.internetsociety.org/ndss/wp-content/uploads/sites/25/2018/02/ndss2018_03A-5_Liu_paper.pdf
- [43] Xu Ma, Xiaoqian Sun, Yuduo Wu, Zheli Liu, Xiaofeng Chen, and Changyu Dong. 2022. Differentially Private Byzantine-robust Federated Learning. *IEEE Transactions on Parallel and Distributed Systems* (2022).
- [44] George Marsaglia, Wai Wan Tsang, and Jingbo Wang. 2003. Evaluating Kolmogorov's distribution. *Journal of statistical software* 8 (2003), 1–4.
- [45] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*. PMLR, 1273–1282.
- [46] Ilya Mironov, Kunal Talwar, and Li Zhang. 2019. R²-enyi Differential Privacy of the Sampled Gaussian Mechanism. *arXiv preprint arXiv:1908.10530* (2019).
- [47] Nicolas Papernot, Martin Abadi, Úlfar Erlingsson, Ian Goodfellow, and Kunal Talwar. 2016. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755* (2016).
- [48] Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson. 2018. Scalable private learning with pate. *arXiv preprint arXiv:1802.08908* (2018).
- [49] Matthias Paulik, Matt Seigel, Henry Mason, Dominic Telaar, Joris Kluivers, Rogier van Dalen, Chi Wai Lau, Luke Carlson, Filip Granqvist, Chris Vandeveld, et al. 2021. Federated evaluation and tuning for on-device personalization: System design & applications. *arXiv preprint arXiv:2102.08503* (2021).
- [50] Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. 2019. Robust aggregation for federated learning. *arXiv preprint arXiv:1912.13445* (2019).
- [51] Md Atiqur Rahman, Tanzila Rahman, Robert Laganière, Noman Mohammed, and Yang Wang. 2018. Membership Inference Attack against Differentially Private Deep Learning Model. *Trans. Data Priv.* 11, 1 (2018), 61–79.

- [52] Jayanth Regatti, Hao Chen, and Abhishek Gupta. 2020. Bygars: Byzantine sgd with arbitrary number of attackers. *arXiv preprint arXiv:2006.13421* (2020).
- [53] Xuebin Ren, Liang Shi, Weiren Yu, Shusen Yang, Cong Zhao, and Zongben Xu. 2022. LDP-IDS: Local Differential Privacy for Infinite Data Streams. In *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*, Zachary Ives, Angela Bonifati, and Amr El Abbadi (Eds.). ACM, 1064–1077. <https://doi.org/10.1145/3514221.3526190>
- [54] Benjamin IP Rubinstein, Blaine Nelson, Ling Huang, Anthony D Joseph, Shing-hon Lau, Satish Rao, Nina Taft, and J Doug Tygar. 2009. Antidote: understanding and defending against poisoning of anomaly detectors. In *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement*. 1–14.
- [55] Virat Shejwalkar and Amir Houmansadr. 2021. Manipulating the byzantine: Optimizing model poisoning attacks and defenses for federated learning. In *NDSS*.
- [56] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 3–18.
- [57] Adam Smith, Abhradeep Thakurta, and Jalaj Upadhyay. 2017. Is interaction necessary for distributed private learning?. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 58–77.
- [58] Shuang Song, Kamalika Chaudhuri, and Anand D Sarwate. 2013. Stochastic gradient descent with differentially private updates. In *2013 IEEE Global Conference on Signal and Information Processing*. IEEE, 245–248.
- [59] tfp [n.d.]. TensorFlow Privacy Git repository. <https://github.com/tensorflow/privacy>.
- [60] Stacey Truex, Ling Liu, Ka-Ho Chow, Mehmet Emre Gursoy, and Wenqi Wei. 2020. LDP-Fed: Federated learning with local differential privacy. In *Proceedings of the Third ACM International Workshop on Edge Systems, Analytics and Networking*. 61–66.
- [61] Chenghong Wang, Johes Bater, Kartik Nayak, and Ashwin Machanavajjhala. 2022. IncShrink: Architecting Efficient Outsourced Databases using Incremental MPC and Differential Privacy. In *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*, Zachary Ives, Angela Bonifati, and Amr El Abbadi (Eds.). ACM, 818–832. <https://doi.org/10.1145/3514221.3526151>
- [62] Di Wang, Changyou Chen, and Jinhui Xu. 2019. Differentially private empirical risk minimization with non-convex loss functions. In *International Conference on Machine Learning*. PMLR, 6526–6535.
- [63] Di Wang, Marco Gaboardi, Adam Smith, and Jinhui Xu. 2020. Empirical risk minimization in the non-interactive local model of differential privacy. *Journal of machine learning research* 21, 200 (2020).
- [64] Di Wang, Lijie Hu, Huan Yu Zhang, Marco Gaboardi, and Jinhui Xu. 2019. Estimating smooth GLM in non-interactive local differential privacy model with public unlabeled data. *arXiv preprint arXiv:1910.00482* (2019).
- [65] Di Wang, Minwei Ye, and Jinhui Xu. 2017. Differentially private empirical risk minimization revisited: Faster and more general. *Advances in Neural Information Processing Systems* 30 (2017).
- [66] Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. 2019. Subsampled rényi differential privacy and analytical moments accountant. In *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 1226–1235.
- [67] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. 2020. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security* 15 (2020), 3454–3469.
- [68] Wikipedia contributors. 2022. 68–95–99.7 rule – Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=68%E2%80%9393.95%E2%80%9399.7_rule&oldid=1097113055 [Online; accessed 24-July-2022].
- [69] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017).
- [70] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. 2020. Fall of empires: Breaking Byzantine-tolerant SGD by inner product manipulation. In *Uncertainty in Artificial Intelligence*. PMLR, 261–270.
- [71] Yuexiang Xie, Zhen Wang, Daoyuan Chen, Dawei Gao, Liuyi Yao, Weirui Kuang, Yaliang Li, Bolin Ding, and Jingren Zhou. 2022. FederatedScope: A Comprehensive and Flexible Federated Learning Platform via Message Passing. *arXiv preprint arXiv:2204.05011* (2022).
- [72] Xiaodong Yang, Huishuai Zhang, Wei Chen, and Tie-Yan Liu. 2022. Normalized/Clipped SGD with Perturbation for Differentially Private Non-Convex Optimization. *arXiv e-prints* (2022), arXiv–2206.
- [73] Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett. 2018. Byzantine-robust distributed learning: Towards optimal statistical rates. In *International Conference on Machine Learning*. PMLR, 5650–5659.
- [74] Xinwei Zhang, Xiangyi Chen, Mingyi Hong, Zhiwei Steven Wu, and Jinfeng Yi. 2021. Understanding Clipping for Federated Learning: Convergence and Client-Level Differential Privacy. *arXiv preprint arXiv:2106.13673* (2021).
- [75] Qinqing Zheng, Jinshuo Dong, Qi Long, and Weijie Su. 2020. Sharp Composition Bounds for Gaussian Differential Privacy via Edgeworth Expansion. In

International Conference on Machine Learning. PMLR, 11420–11435.

- [76] Banghua Zhu, Lun Wang, Qi Pang, Shuai Wang, Jiantao Jiao, Dawn Song, and Michael I Jordan. 2022. Byzantine-Robust Federated Learning with Optimal Statistical Rates and Privacy Guarantees. *arXiv preprint arXiv:2205.11765* (2022).
- [77] Heng Zhu and Qing Ling. 2022. Bridging Differential Privacy and Byzantine-Robustness via Model Aggregation. *arXiv preprint arXiv:2205.00107* (2022).
- [78] Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep leakage from gradients. *Advances in Neural Information Processing Systems* 32 (2019).
- [79] Yuqing Zhu and Yu-Xiang Wang. 2019. Poission subsampled rényi differential privacy. In *International Conference on Machine Learning*. PMLR, 7634–7642.
- [80] Yuqing Zhu, Xiang Yu, Manmohan Chandraker, and Yu-Xiang Wang. 2020. Private-knn: Practical differential privacy for computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11854–11862.

A SUPPLEMENTARY MATERIAL

A.1 Implementation Detail

Hardware and software setup: The experiments are run on a machine with CentOS Linux release 7.9.2009 (Core). Hardware includes Intel(R) Xeon(R) Gold 6142 CPU @ 2.60GHz and Tesla V100-SXM2-32GB. The code is implemented with PyTorch 1.11.0 and CUDA 10.2. It takes roughly 600 GPU hours to generate all experiment results.

Model setup: We use CNN network for MNIST and Colorectal and Fully-connected layer for Fashion and USPS, the model size d is 21802, 33736, 25450, 25450 for MNIST, Colorectal, Fashion, and USPS, respectively. Detail setups of network architectures for each dataset are summarised as follows:

1) For MNIST, the network setup is shown in Table 7; For Fashion and USPS, the network setup is shown in Table 8; For Colorectal, We make the CNN for Colorectal have a residual connection, due to its complexity, we do not describe it here, the detail can be found in our code.

Layer	Setup
CNN	InputChannel=1, OutputChannel=16, KernelSize=5
ELU	-
GroupNorm	NumGroups=4, NumChannels=16
CNN	InputChannel=16, OutputChannel=16, KernelSize=5
ELU	-
GroupNorm	NumGroups=4, NumChannels=16
CNN	InputChannel=16, OutputChannel=16, KernelSize=5
ELU	-
GroupNorm	NumGroups=4, NumChannels=16
AdaptiveAvgPool	(4,4)
Linear	(256, 32)
ELU	-
Linear	(32, 10)

Table 7: Network setup for MNIST.

Layer	Setup
flatten	-
Linear	(784, 32)
ELU	-
Linear	(32, 10)

Table 8: Network setup for Fashion and USPS.

A.2 List of Notations

We summarise some useful notations we adopt throughout this paper in Table 9

Notations	Meaning
d	model size
σ	DP noise multiplier
η	learning rate
T	number of total iterations
w^{t-1}	model sent from the server to all workers at the beginning of t -th iteration
g_i^t	workers's uploading to server from i -th worker at t -th iteration
b_c	local batch size for each worker
D_p	data possessed by the server
f	loss function
G_s^t	set of selected gradient vectors by the server at round t

Table 9: Some notations used throughout the paper.

A.3 Byzantine-resilient Aggregation Methods

Giving (gradient) vectors $\{g_i\}_{i=1}^n$, we summarise Krum [9], RFA [50], coordinate-wise median [73] and Trimmed Mean [73] as follows:

- **Krum:** It returns a vector g_s , which is the closest vector to the mean after excluding $\gamma n + 2$ furthest away points (γ is some constant). Mathematically, let $\mathbb{S} \subseteq \{1, 2, \dots, n\}$ be an index set with size at least $(n - \gamma n - 2)$. Then,

$$\text{Krum}(g_1, \dots, g_n) = \arg \min_{g_s} \min_{\mathbb{S}} \sum_{j \in \mathbb{S}} \|g_s - g_j\|^2$$

- **Robust Federated Averaging (RFA):** It just returns the geometric median of all vectors:

$$\text{RFA}(g_1, \dots, g_n) = \arg \min_g \sum_{i=1}^n \|g - g_i\|_2$$

- **Coordinate-wise median:** Let $v[j]$ refer to the j -th coordinate of a vector v , then the output is just the median of each coordinate of these vectors:

$$\text{CM}(g_1, \dots, g_n)[j] = \text{median}(g_1[j], \dots, g_n[j])$$

- **Trimmed Mean:** For the j -th coordinate, it computes the average after excluding γn largest and smallest values, *i.e.*,

$$\text{TM}(g_1, \dots, g_n)[j] = \frac{1}{n - 2\gamma n} \sum_{k=\gamma n}^{n-\gamma n} g_{\Pi_j(k)}[j],$$

where $g_{\Pi_j(k)}[j]$ represents the k -th largest value among all $\{g_i[j]\}_{i=1}^n$.

A.4 Dataset Details and Distribution Setups

Datasets: 1) MNIST [41] is a hand-written digit dataset. It contains 60,000 training examples and 10,000 testing examples, equally distributed among the 10 classes. 2) Colorectal [37], which has 8 classes and 5,000 histological RGB images of size 150×150 . The image size is significantly greater than MNIST (which is 28×28). We randomly take 4600 examples for training and take the rest

Algorithm 4 GetNonIID(D, n)

Input: D , the dataset to be distributed to all workers; n , number of workers

- 1: Partition D (by class) into G_1, G_2, \dots, G_H
- 2: $D_i \leftarrow [], T_i \leftarrow []$, for $i = 1, 2, \dots, n$
- 3: **for** $k = 1, 2, \dots, H$ **do**
- 4: Generate a list V of uniform R.V.s and normalize V
- 5: Partition G_k according to V
- 6: Append parts of G_k to T_i for $i = 1, 2, \dots, n$
- 7: **end for**
- 8: Concatenate all T_i into L
- 9: $s \leftarrow \left\lceil \frac{|L|}{n} \right\rceil$
- 10: **for** $i = 1, 2, \dots, n$ **do**
- 11: $D_i \leftarrow L[(i-1) \cdot s : i \cdot s]$
- 12: **end for**

Output: D_1, D_2, \dots, D_n

400 examples for testing. 3) Fashion [69] is an image classification dataset with 10 classes and it has 60,000 training images and 10,000 testing images. 4) USPS [34]. This dataset has 10 classes and 7,291 samples for training and 2,007 for testing.

Generating non-i.i.d. data distribution: The method for generating non-i.i.d. data distributions is summarised in Algorithm 4, and the method for i.i.d. case can be trivially obtained according to line 4. The intuition for why our method can generate non-i.i.d. data distribution is that, according to line 4 in Algorithm 4, the probability that the non-i.i.d. case happens to be i.i.d. is negligible. In practice, applying Algorithm 4 gives us reasonable result which is non-i.i.d. distributed. We provide the non-i.i.d. distribution simulation result for dataset MNIST as evidence. We can see from Figure 5 that, the data distribution is non-i.i.d.: for a certain class, its sampling probability for each worker is different. Take class 1 for example, its sampling ratio for worker 0 (first sub-figure of first row) is around 0.2, whereas for worker 1 (second sub-figure of first row), it is 0 and for worker 3 (fourth sub-figure of first row), it is almost 0.3. In contrast, for the i.i.d. case, the probability for all classes and for all workers should be 0.1.

A.5 Proofs

A.5.1 Proof of Theorem 1.

PROOF. Our parameter w is a d dimensional vector, considering that our parameter updating rule in the t -th iteration is:

$$w^t = w^{t-1} - \frac{\eta^t}{|B^t|} \left(\sum_{g^t \in B^t} \frac{g^t}{\|g^t\|} + z \right) \quad (13)$$

where B^t is the current batch of per-example gradient (we fix the batch size to be $|B^t| = b_c$), $z \sim \mathcal{N}(0, \sigma^2 I)$ is the DP noise and $g^t = \nabla F(w^{t-1}) + \xi^t$ is the returned result by our stochastic gradient oracle. Considering that we assume F is L -smooth, then we have:

$$\begin{aligned} F(w^t) - F(w^{t-1}) &\leq \langle \nabla F(w^{t-1}), w^t - w^{t-1} \rangle \\ &\quad + \frac{L}{2} \|w^t - w^{t-1}\|^2 \end{aligned}$$



Figure 5: Simulation result for non-i.i.d. distribution for dataset MNIST to 20 workers. 0-9 in the horizontal axis stands for class labels, the vertical axis is the ratio of the number of a certain class to the size of the datasets each worker holds. There are 20 subplots, standing for 20 different workers.

which translates to:

$$F(w^t) - F(w^{t-1}) \leq -\frac{\eta^t}{b_c} \left\langle \nabla F(w^{t-1}), \sum_{g^t \in B^t} \frac{g^t}{\|g^t\|} + z \right\rangle + \frac{L(\eta^t)^2}{2} \left\| \frac{1}{b_c} \sum_{g^t \in B^t} \frac{g^t}{\|g^t\|} + \frac{z}{b_c} \right\| \quad (14)$$

note that each g^t inside B^t is i.i.d. (because each g^t is calculated by one data example and each data example is sampled from dataset independently with replacement), conditioned on w^{t-1} and take expectation on both sides with respect to the randomness of data sampling and DP noise:

$$\begin{aligned} \mathbb{E}(F(w^t)) - F(w^{t-1}) &\leq -\frac{\eta^t}{b_c} \sum_{g^t \in B^t} \mathbb{E} \left\langle \nabla F(w^{t-1}), \frac{g^t}{\|g^t\|} \right\rangle \\ &\quad + \frac{L(\eta^t)^2}{2} \mathbb{E} \left\| \frac{1}{b_c} \sum_{g^t \in B^t} \frac{g^t}{\|g^t\|} + \frac{z}{b_c} \right\| \\ &\leq -\frac{\eta^t}{b_c} \sum_{g^t \in B^t} \mathbb{E} \left\langle \nabla F(w^{t-1}), \frac{g^t}{\|g^t\|} \right\rangle \\ &\quad + \frac{L(\eta^t)^2}{2} \left(1 + \frac{\sigma^2 d}{b_c^2} \right) \\ &= -\eta^t \mathbb{E} \left\langle \nabla F(w^{t-1}), \frac{g^t}{\|g^t\|} \right\rangle \\ &\quad + \frac{L(\eta^t)^2}{2} \left(1 + \frac{\sigma^2 d}{b_c^2} \right) \end{aligned} \quad (15)$$

we now deal with the term $\left\langle \nabla F(\mathbf{w}^{t-1}), \frac{\mathbf{g}^t}{\|\mathbf{g}^t\|} \right\rangle$ first, if $\|\nabla F(\mathbf{w}^{t-1})\| > 2\|\xi^t\|$, we then have:

$$\begin{aligned}
 \left\langle \nabla F(\mathbf{w}^{t-1}), \frac{\mathbf{g}^t}{\|\mathbf{g}^t\|} \right\rangle &= \frac{\|\nabla F(\mathbf{w}^{t-1})\|^2 + \langle \nabla F(\mathbf{w}^{t-1}), \xi^t \rangle}{\|\nabla F(\mathbf{w}^{t-1}) + \xi^t\|} \\
 &= \frac{\|\nabla F(\mathbf{w}^{t-1})\|^2}{\|\nabla F(\mathbf{w}^{t-1}) + \xi^t\|} \\
 &\quad + \frac{\|\nabla F(\mathbf{w}^{t-1})\| \|\xi^t\| \cos(\theta)}{\|\nabla F(\mathbf{w}^{t-1}) + \xi^t\|} \\
 &\geq \frac{\|\nabla F(\mathbf{w}^{t-1})\|^2 - \frac{1}{2} \|\nabla F(\mathbf{w}^{t-1})\|^2}{\|\nabla F(\mathbf{w}^{t-1}) + \xi^t\|} \\
 &= \frac{\|\nabla F(\mathbf{w}^{t-1})\|^2}{2\|\nabla F(\mathbf{w}^{t-1}) + \xi^t\|} \\
 &\geq \frac{\|\nabla F(\mathbf{w}^{t-1})\|^2}{2(\|\nabla F(\mathbf{w}^{t-1})\| + \|\xi^t\|)} \\
 &\geq \frac{\|\nabla F(\mathbf{w}^{t-1})\|}{3} \\
 &\geq \frac{\|\nabla F(\mathbf{w}^{t-1})\|}{3} - \frac{8\|\xi^t\|}{3}
 \end{aligned}$$

if $\|\nabla F(\mathbf{w}^{t-1})\| \leq 2\|\xi^t\|$, we then have:

$$\begin{aligned}
 \left\langle \nabla F(\mathbf{w}^{t-1}), \frac{\mathbf{g}^t}{\|\mathbf{g}^t\|} \right\rangle &\geq -\|\nabla F(\mathbf{w}^{t-1})\| \\
 &= \frac{\|\nabla F(\mathbf{w}^{t-1})\|}{3} - \frac{4\|\nabla F(\mathbf{w}^{t-1})\|}{3} \\
 &\geq \frac{\|\nabla F(\mathbf{w}^{t-1})\|}{3} - \frac{8\|\xi^t\|}{3}
 \end{aligned}$$

Hence we have:

$$\left\langle \nabla F(\mathbf{w}^{t-1}), \frac{\mathbf{g}^t}{\|\mathbf{g}^t\|} \right\rangle \geq \frac{\|\nabla F(\mathbf{w}^{t-1})\|}{3} - \frac{8\|\xi^t\|}{3} \quad (16)$$

applying Equation 16 to Equation 15 yields:

$$\begin{aligned}
 \mathbb{E}(F(\mathbf{w}^t)) - F(\mathbf{w}^{t-1}) &\leq -\eta^t \frac{\mathbb{E}\|\nabla F(\mathbf{w}^{t-1})\|}{3} + \eta^t \frac{8\mathbb{E}\|\xi^t\|}{3} \\
 &\quad + \frac{L(\eta^t)^2}{2} \left(1 + \frac{\sigma^2 d}{b_c^2}\right) \quad (17)
 \end{aligned}$$

assume our loss function $F(\mathbf{w}^t) > 0$, and the result $\mathbf{g}^t = \nabla F(\mathbf{w}^{t-1}) + \xi^t$ which is returned by our stochastic gradient oracle has bounded variance, if we set a constant learning rate $\eta^t = \eta$, take iterated expectation, summing up both side from iteration 1 to T , divide both sides by $T\eta$, rearrange terms, we have:

$$\begin{aligned}
 \frac{1}{T} \sum_{t=1}^T \mathbb{E} \|\nabla F(\mathbf{w}^t)\| &\leq \frac{3F(\mathbf{w}^0)}{T\eta} + \frac{3L\eta}{2} \left(1 + \frac{\sigma^2 d}{b_c^2}\right) \\
 &\quad + \frac{8}{T} \sum_{t=1}^T \mathbb{E} \|\xi^t\| \quad (18) \\
 &\leq \frac{3F(\mathbf{w}^0)}{T\eta} + \frac{3L\eta}{2} \left(1 + \frac{\sigma^2 d}{b_c^2}\right) \\
 &\quad + 8\nu
 \end{aligned}$$

□

A.5.2 Proof of Theorem 3. The following is proof for Theorem 3.

PROOF. Algorithm 1 deals with each worker's privacy, we will show that it is (ϵ, δ) -DP for each worker-hold dataset. Note that the l_2 -sensitivity of the term $\sum_{j \in [b_c]} \frac{\phi_i[j]}{\|\phi_i[j]\|_2}$ is 2, thus by the same proof

in Theorem 1 of [1]. We can see that if $\sigma \geq c_2 \frac{q\sqrt{T \ln \frac{1}{\delta}}}{\epsilon}$, Algorithm 1 guarantee (ϵ, δ) -DP for each worker-hold dataset. □

A.5.3 Proof of Theorem 4. The following is proof for Theorem 4. Before the proof, we first recall some technical lemmas.

LEMMA 1 ([62]). For the Gaussian random variable $x \sim \mathcal{N}(0, \sigma^2 I_d)$, with probability at least $1 - \zeta$ for any $\zeta \in (0, 1)$ we have

$$\|\sigma\|_2 \leq \sqrt{2d \log \frac{1}{\zeta}} \sigma$$

LEMMA 2 ([11]). Under Assumptions 1 and 2 and assume the parameter space $\Theta \subseteq B(0, r\sqrt{d})$. Then for any $\zeta \in (0, 1)$, if $\Delta_1 \leq \frac{\sigma_1^2}{\gamma_1}$ and $\Delta_2 \leq \frac{\sigma_2^2}{\gamma_2}$, then we have the following for all $t \in [T]$:

$$\Pr\{\|g_s^{t-1} - \nabla F(\mathbf{w}^{k-1})\| \leq 8\Delta_2 \|\mathbf{w}^{k-1} - \mathbf{w}^*\| + 4\Delta_1\} \geq 1 - \zeta,$$

where $\Delta_2 = \sigma_2 \sqrt{\frac{2}{|D_p|}} \sqrt{K_1 + K_2}$ with $K_1 = d \log \frac{\max\{L, L_1\}}{\sigma_2}$ and $K_2 = \frac{d}{2} \log \frac{|D_p|}{d} + \log \frac{6\sigma^2 r \sqrt{|D_p|}}{\gamma_2 \sigma_1 \zeta}$, and $\Delta_1 = \sqrt{2} \sigma_1 \sqrt{\frac{d \log 6 + \log(3/\delta)}{|D_p|}}$.

PROOF. For a fixed iteration t , we first assume the following is true (later we will show it will hold for all $t \in [T]$):

$$\|\mathbf{w}^{t-1} - \mathbf{w}^*\| \leq \frac{r}{2} \sqrt{d}.$$

Since we have $\eta_{t-1} = \|g_s^{t-1}\|_2 \eta_0$, then

$$\begin{aligned}
\|w^{t-1} - w^*\| &= \left\| w^{t-1} - \eta_{t-1} \left(\frac{1}{n} \sum_{g \in G_s^t} g \right) - w^* \right\| \\
&= \|w^{t-1} - \eta_0 \left(\frac{\|g_s^{t-1}\|}{n} \sum_{g \in G_s^t} g \right) \\
&\quad + \eta_0 \nabla F(w^{t-1}) - \eta_0 \nabla F(w^{t-1}) - w^*\| \\
&\leq \underbrace{\|w^{t-1} - \eta_0 \nabla F(w^{t-1}) - w^*\|}_A \\
&\quad + \underbrace{\eta_0 \left\| \nabla F(w^{t-1}) - \frac{\|g_s^{t-1}\|}{n} \sum_{g \in G_s^t} g \right\|}_B.
\end{aligned} \tag{19}$$

We first consider term A and we consider the following lemma.

LEMMA 3. [11], Assume Assumption 1 holds. If we set the learning rate satisfies $\alpha \leq \mu/(2L^2)$, then we have the following in any $t \geq 1$:

$$\|w^{t-1} - w^* - \alpha \nabla F(w^{t-1})\| \leq \sqrt{1 - \frac{\mu^2}{4L^2}} \|w^{t-1} - w^*\|$$

PROOF. Since $\nabla F(w^*) = 0$, we have the following:

$$\begin{aligned}
&\|w^{t-1} - w^* - \alpha \nabla F(w^{t-1})\|^2 \\
&= \|w^{t-1} - w^* - \alpha (\nabla F(w^{t-1}) - \nabla F(w^*))\|^2 \\
&= \|w^{t-1} - w^*\|^2 + \alpha^2 \|\nabla F(w^{t-1}) - \nabla F(w^*)\|^2 \\
&\quad - 2\alpha \langle w^{t-1} - w^*, \nabla F(w^{t-1}) - \nabla F(w^*) \rangle
\end{aligned} \tag{20}$$

By Assumption 1, we have:

$$\|\nabla F(w^{t-1}) - \nabla F(w^*)\| \leq L \|w^{t-1} - w^*\| \tag{21}$$

$$\begin{aligned}
&F(w^*) + \langle \nabla F(w^*), w^{t-1} - w^* \rangle \\
&\leq F(w^{t-1}) - \frac{\mu}{2} \|w^{t-1} - w^*\|^2
\end{aligned} \tag{22}$$

$$F(w^{t-1}) + \langle \nabla F(w^{t-1}), w^* - w^{t-1} \rangle \leq F(w^*) \tag{23}$$

Summing up inequalities 22 and 23, we have:

$$\begin{aligned}
&\langle w^* - w^{t-1}, \nabla F(w^{t-1}) - \nabla F(w^*) \rangle \\
&\leq -\frac{\mu}{2} \|w^{t-1} - w^*\|^2.
\end{aligned} \tag{24}$$

Substituting inequalities 21 and 24 into 20, we have:

$$\begin{aligned}
&\|w^{t-1} - w^* - \alpha \nabla F(w^{t-1})\|^2 \\
&\leq \left(1 + \alpha^2 L^2 - \alpha \mu\right) \|w^{t-1} - w^*\|^2.
\end{aligned} \tag{25}$$

By choosing $\alpha \leq \frac{\mu}{2L^2}$, we have:

$$\|w^{t-1} - w^* - \alpha \nabla F(w^{t-1})\|^2 \leq \left(1 - \frac{\mu^2}{4L^2}\right) \|w^{t-1} - w^*\|^2, \tag{26}$$

which concludes proof for lemma 3. \square

For term B, note that,

$$\begin{aligned}
&\left\| \nabla F(w^{t-1}) - \frac{\|g_s^{t-1}\|}{n} \sum_{g \in G_s^t} g \right\|_2 \\
&\leq \left\| \|g_s^{t-1}\| \frac{1}{n} \sum_{g \in G_s^t} g - g_s^{t-1} \right\| + \|g_s^{t-1} - \nabla F(w^{t-1})\| \\
&\leq \underbrace{\|g_s^{t-1}\| \left\| \frac{1}{n} \sum_{g \in G_s^t} g \right\|}_C + \|g_s^{t-1}\| + \|g_s^{t-1} - \nabla F(w^{t-1})\|.
\end{aligned} \tag{27}$$

For term C, we can bound it as:

$$\begin{aligned}
C &= \|g_s^{t-1}\| \left\| \frac{1}{n} \sum_{g \in G_s^t} g \right\| \\
&= \|g_s^{t-1}\| \left\| \frac{1}{n} \sum_{g \in G_s^t} \tilde{g} + z \right\| \\
&\leq \underbrace{\|g_s^{t-1}\| \left\| \frac{1}{n} \sum_{g \in G_s^t} \tilde{g} \right\|}_D + \underbrace{\|g_s^{t-1}\| \|z\|}_E,
\end{aligned} \tag{28}$$

here $\tilde{g} = \frac{1}{b_c} \sum_{j \in [b_c]} \frac{\phi_i[j]}{\|\phi_i[j]\|_2}$ is just g without noise (refer to line 11) in Algorithm 1 and $z \sim \mathcal{N}(0, \frac{|G_s^t|}{n^2 b_c^2} \sigma^2)$ which is the gaussian noise ensuring DP.

For term D, $D \leq \|g_s^{t-1}\|$ since $\|\tilde{g}\| \leq 1$ and $|G_s^t| \leq n$; For term E, by Lemma 1 we have with probability at least $1 - \xi$ that,

$$\|z\| \leq \frac{\sqrt{2d|G_s^t| \ln \frac{1}{\xi}}}{nb_c} \sigma \leq \frac{\sqrt{2d \ln \frac{1}{\xi}}}{\sqrt{nb_c}} \sigma.$$

This leads to that, w.p $1 - \xi$,

$$C \leq \|g_s^{t-1}\| \left(1 + \frac{\sqrt{2d \ln \frac{1}{\xi}}}{\sqrt{nb_c}} \right) \tag{29}$$

Applying this inequality to 27 give us,

$$\begin{aligned}
B &\leq 2 \|g_s^{t-1}\| + \|g_s^{t-1} - \nabla F(w^{t-1})\| + \frac{\|g_s^{t-1}\| \sqrt{2d \ln \frac{1}{\xi}} \sigma}{\sqrt{nb_c}} \\
&\leq 3 \|g_s^{t-1} - \nabla F(w^{t-1})\| + 2 \|\nabla F(w^{t-1})\| \\
&\quad + \frac{\|g_s^{t-1} - \nabla F(w^{t-1})\| \sqrt{2d \ln \frac{1}{\xi}} \sigma}{\sqrt{nb_c}} \\
&\quad + \frac{\|\nabla F(w^{t-1})\| \sqrt{2d \ln \frac{1}{\xi}} \sigma}{\sqrt{nb_c}}
\end{aligned} \tag{30}$$

For $\|g_s^{t-1} - \nabla F(w^{k-1})\|$, by Lemma 2 we have, w.p $1 - \xi$,

$$\|g_s^{t-1} - \nabla F(w^{k-1})\| \leq 8\Delta_2 \|w^{t-1} - w\| + 4\Delta_1. \quad (31)$$

For $\|\nabla F(w^{t-1})\|$, we have,

$$\begin{aligned} \|\nabla F(w^{t-1})\| &\leq \|\nabla F(w^{t-1}) - \nabla F(w^*)\| + \|\nabla F(w^*)\| \\ &= \|\nabla F(w^{t-1}) - \nabla F(w^*)\| \\ &\leq L \|w^{t-1} - w^*\| \end{aligned} \quad (32)$$

Applying (31) and (32) to (30), then combine it with Lemma 3 to (19), we have,

$$\begin{aligned} \|w^t - w^*\| &\leq \sqrt{1 - \frac{u^2}{4L^2}} \|w^{t-1} - w^*\| \\ &+ (24\eta_0\Delta_2 + 2\eta_0L) \|w_{t-1} - w^*\| \\ &+ (8\Delta_2\eta_0 + \eta_0L) \frac{\sqrt{2d \ln \frac{1}{\xi}} \sigma}{\sqrt{nb_c}} \|w_{k-1} - w^*\| \\ &+ \frac{4\eta_0 \sqrt{2d \ln \frac{1}{\xi}} \sigma}{\sqrt{nb_c}} \Delta_1 + 12\eta_0\Delta_1. \end{aligned} \quad (33)$$

Note that for the DP noise multiplier σ (refer to Theorem 3), if we take $\sigma = c_2 \frac{b_c \sqrt{T \ln \frac{1}{\delta}}}{|D|\epsilon}$, then we have $\frac{\sqrt{2 \ln \frac{1}{\xi}} \sigma}{\sqrt{nb_c}} = O\left(\frac{\sqrt{T \ln \frac{1}{\delta} \ln \frac{1}{\xi}}}{|D|\sqrt{n\epsilon}}\right)$. Recall our assumption in Theorem 4 that $|D|\sqrt{n} \geq \Omega\left(\frac{\sqrt{T \ln \frac{1}{\delta} \ln \frac{1}{\xi}} \sqrt{d}}{\epsilon}\right)$.

Thus, we have $\frac{\sqrt{2T \ln \frac{1}{\xi}} d \sigma}{\sqrt{nb_c} \epsilon} \leq 1$. By taking σ as specified in Theorem 3, we can update (33) as

$$\begin{aligned} \|w^t - w^*\| &\leq \left(\sqrt{1 - \frac{u^2}{4L^2}} + 32\eta_0\Delta_2 + 3\eta_0L \right) \|w^{t-1} - w^*\| \\ &+ \frac{4c_1 \sqrt{\ln \frac{1}{\xi}} \ln \frac{1}{\delta} T \eta_0}{\sqrt{n}|D|\epsilon} \Delta_1 \sqrt{d} + 12\eta_0\Delta_1 \\ &= (1 - \rho) \|w^{t-1} - w^*\| + O\left(\frac{d \ln \frac{1}{\xi} \sqrt{T \ln \frac{1}{\delta}} \sigma_1}{|D|\sqrt{|D_p|n\epsilon}}\right) \\ &+ O(\eta_0\sigma_1 \sqrt{\frac{d}{|D_p|}}), \end{aligned}$$

where $\rho = 1 - \sqrt{1 - \frac{u^2}{4L^2}} - 32\eta_0\Delta_2 - 3\eta_0L$.

Since $\|w^{t-1} - w^*\| \subset B(0, r\sqrt{d})$, if $|D|\sqrt{|D_p|n}$ is sufficiently large such that $|D|\sqrt{|D_p|n} \geq \Omega\left(\frac{\sqrt{d \ln \frac{1}{\xi}} \sqrt{\ln \frac{1}{\delta}} T}{\epsilon \rho}\right)$ and $\eta_0\sigma_1/\sqrt{|D_p|} \leq O\left(\frac{r\sqrt{d}}{\rho}\right)$, then $\|w^k - w^*\| \in B(0, r\sqrt{d})$. Thus, take the union w.r.t all iterations, with probability at least $1 - \xi T$, we have,

$$\begin{aligned} \|w^T - w^*\| &\leq (1 - \rho)^T \|w^0 - w^*\| \\ &+ O\left(\frac{\eta_0 d \ln \frac{1}{\xi} \sqrt{T \ln \frac{1}{\delta}} \sigma_1}{|D|\sqrt{b_s n \rho \epsilon}} + \frac{\eta_0 \Delta_1}{\rho}\right) \end{aligned}$$

By taking $T = O\left(\ln_{1-\rho} \frac{1}{\sqrt{n}|D|\sqrt{|D_p|}}\right) = O\left(\frac{1}{\rho} \ln(\sqrt{n}|D|\sqrt{|D_p|})\right)$, we have,

$$\|w^T - w^*\| \leq \tilde{O}\left(\frac{1}{\rho^2} \frac{d \ln \frac{1}{\xi} \sqrt{\ln \frac{1}{\delta}} \sigma_1}{|D|\sqrt{nb_s \epsilon}} + \frac{1}{\rho} \frac{\sigma_1 \sqrt{d \ln \frac{1}{\xi}}}{\sqrt{|D_p|}}\right).$$

□

γ	MNIST		COLOR.		FASHION		USPS	
	$\epsilon =$		$\epsilon =$		$\epsilon =$		$\epsilon =$	
	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2
20%	.87	.95	.48	.73	.69	.78	.64	.85
35%	.88	.96	.47	.74	.69	.79	.63	.86
exact (50%)	.88	.96	.49	.74	.69	.80	.64	.87
65%	.85	.96	.48	.73	.70	.79	.64	.87
80%	.84	.95	.48	.74	.69	.79	.64	.85

Table 10: Experimental result (test accuracy) under Gaussian attack for different γ in i.i.d. setting.

γ	MNIST		COLOR.		FASHION		USPS	
	$\epsilon =$		$\epsilon =$		$\epsilon =$		$\epsilon =$	
	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2
20%	.86	.95	.48	.73	.68	.78	.64	.85
35%	.87	.96	.47	.74	.69	.79	.63	.86
exact (50%)	.88	.96	.49	.74	.69	.80	.64	.87
65%	.84	.96	.43	.73	.68	.79	.52	.87
80%	.80	.95	.30	.74	.60	.79	.47	.85

Table 11: Experimental result (test accuracy) under Optimized Local Model Poisoning attack for different γ in i.i.d. setting.

γ	MNIST		COLOR.		FASHION		USPS	
	$\epsilon =$		$\epsilon =$		$\epsilon =$		$\epsilon =$	
	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2
20%	.86	.94	.48	.73	.66	.78	.64	.85
35%	.86	.95	.47	.74	.69	.79	.64	.86
exact (50%)	.88	.95	.49	.74	.69	.80	.64	.87
65%	.84	.95	.44	.73	.70	.79	.56	.87
80%	.83	.94	.31	.74	.69	.79	.52	.85

Table 12: Experimental result (test accuracy) under Label-flipping attack for different γ in non-i.i.d. setting.

A.6 More Experimental Results

Ablation study on γ : Here, for experiments under other attacks, we present the additional results for the ablation study on γ which is treated as a belief instead of the truth. We present results under Gaussian attack in Table 10, Optimized Local Model Poisoning attack in Table 11 under i.i.d. setting. We present results under Label-flipping attack in Table 12, Gaussian attack in Table 13, Optimized Local Model Poisoning attack in Table 14 under the non-i.i.d. setting.

γ	MNIST		COLOR.		FASHION		USPS	
	$\epsilon =$		$\epsilon =$		$\epsilon =$		$\epsilon =$	
	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2
20%	.86	.95	.47	.73	.69	.78	.64	.85
35%	.87	.96	.47	.74	.69	.79	.63	.86
exact (50%)	.88	.96	.49	.74	.69	.80	.64	.87
65%	.85	.96	.49	.73	.69	.79	.64	.87
80%	.83	.95	.47	.74	.69	.79	.61	.85

Table 13: Experimental result (test accuracy) under Gaussian attack for different γ in non-i.i.d. setting.

γ	MNIST		COLOR.		FASHION		USPS	
	$\epsilon =$		$\epsilon =$		$\epsilon =$		$\epsilon =$	
	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2	$\frac{1}{8}$	2
20%	.85	.95	.45	.73	.68	.78	.64	.85
35%	.87	.96	.46	.74	.69	.79	.64	.86
50% (exact)	.87	.96	.49	.74	.69	.80	.64	.87
65%	.84	.96	.40	.73	.64	.79	.50	.87
80%	.80	.95	.32	.74	.58	.79	.43	.85

Table 14: Experimental result (test accuracy) under Optimized Local Model Poisoning attack for different γ in non-i.i.d. setting.

ϵ	MNIST	COLOR.	FASHION	USPS
Non-DP	.98 \pm .003	.80 \pm .010	.88 \pm .005	.92 \pm .004
2	.96 \pm .006	.74 \pm .013	.80 \pm .007	.87 \pm .011
1	.95 \pm .005	.70 \pm .017	.79 \pm .005	.86 \pm .012
.5	.95 \pm .011	.66 \pm .026	.78 \pm .010	.82 \pm .016
.25	.93 \pm .017	.56 \pm .023	.75 \pm .008	.76 \pm .023
.125	.88 \pm .021	.50 \pm .041	.70 \pm .007	.64 \pm .041

Table 15: Experimental result (test accuracy) on the “side-effect” that DP introduces in the i.i.d. setting.

ϵ	MNIST	COLOR.	FASHION	USPS
Non-DP	.98 \pm .004	.80 \pm .013	.88 \pm .010	.91 \pm .003
2	.96 \pm .009	.74 \pm .023	.80 \pm .005	.87 \pm .013
1	.95 \pm .009	.70 \pm .024	.79 \pm .010	.86 \pm .010
.5	.95 \pm .011	.66 \pm .021	.77 \pm .012	.82 \pm .019
.25	.93 \pm .015	.55 \pm .027	.75 \pm .017	.76 \pm .028
.125	.88 \pm .024	.50 \pm .045	.70 \pm .012	.63 \pm .041

Table 16: Experimental result (test accuracy) on the “side-effect” that DP introduces in the non-i.i.d. setting.

One notable setting is that we fix the portion of honest workers to be half of all workers. And γ is treated to be the server’s prior belief that γn workers are honest where n is the total number of workers. “50% (exact)” means that the server’s belief is exactly the same as the true situation, *i.e.*, half of all workers are honest. All results are obtained by taking the average of three runs with different seeds.

We can see similar results as that in the main body, *i.e.*, when the belief is too radical and the privacy requirement is extreme, the utility suffers, except for the Gaussian attack where we almost still have robustness, this is because Gaussian attack purely injects noise to the aggregation, the utility is not significantly affected.

Abalation study on DP: To see the “side-effect” that DP introduces, we summarise our experimental results in Table 15 for the i.i.d. case and Table 16 for the non-i.i.d. case, respectively. For a fair

(a) Under Gaussian attack

byz.	MNIST	Colorectal	Fashion	USPS
40%	.09 \pm .00	.15 \pm .05	.10 \pm .00	.10 \pm .00
20%	.12 \pm .00	.15 \pm .05	.13 \pm .00	.20 \pm .00

(b) Under Label-flipping attack

byz.	MNIST	Colorectal	Fashion	USPS
40%	.01 \pm .00	.07 \pm .05	.02 \pm .00	.04 \pm .01
20%	.07 \pm .04	.09 \pm .04	.06 \pm .03	.08 \pm .03

(c) Under Optimized Local Model Poisoning attack

byz.	MNIST	Colorectal	Fashion	USPS
40%	.09 \pm .00	.12 \pm .05	.10 \pm .00	.17 \pm .00
20%	.09 \pm .00	.12 \pm .05	.10 \pm .00	.17 \pm .00

Table 17: Testing accuracy under different Byzantine attacks at privacy level $\epsilon = 2$. In this experiment, the server-own auxiliary data is sampled from KMNIST dataset whose data space \mathbb{X} is different from that of all datasets we use for training.

comparison, all experiments are done with the same hyperparameter setup. We use “Non-DP” to stand for the case where no DP is applied. We can see from those results that by introducing DP, the utility drops compared with the non-DP case. This is due to the noise introduced by the DP-SGD protocol. The more private (more noise is added), the worse utility we will get.

Note that our Byzantine robustness is built on our refactored DP protocol, so for the non-DP case, we are only able to run the experiment without Byzantine robustness and accordingly, without Byzantine attack (or we do not have any utility otherwise). And for that specific scenario, there is existing work identifying and improving the utility of DP learning compared with non-DP. We refer the reader to them [4, 29, 46, 66, 75, 79] as they are more involved.

When the data space of auxiliary data is different from it of the local-hold data: Recall that previously we assumed the sever-own auxiliary data follows the same distribution as that of the local-hold data \mathbb{X} . Here we conduct experiments for the case that server-own data is sampled from a different data space \mathbb{X}' . The experiments are carried out as follows, for the auxiliary data, we sample it from the KMNIST [17] dataset and for the rest of the experimental settings, we keep them identical to those mentioned in our experiment section.

Results are presented in Table 17. We can see that overall, the training does not lead to useful utility if we sample the server-own auxiliary data from a different data space. We can see that when under the Gaussian attack, the training performs better than random guessing if 20% workers are Byzantine, this is because a Gaussian attack purely injects noise and such noise is not too detrimental. In contrast, when under the Label-flipping attack, the model performs even worse than random guessing, this is because the Label-flipping attack tries to let the model predict wrongly for each label.

Based on such observation, we claim that it is necessary to assume the serve-own auxiliary data is sampled from the same data

space as that of the local-hold data. However, the above negative results are possibly due to that KMNIST “looks” too different from the local-hold data. If we can find some dataset that “looks” more similar, positive results may be obtained. We leave such work as a future study.

Experimental results for more stringent cases (95% and 99% workers are Byzantine): We previously present experimental results for the case where 90% workers are Byzantine in the main context. Here we give results for more stringent cases to test our protocol’s robustness. Note that simulating stringent cases requires simulating more workers which incurs more computational burden. For instance, simulating the 99% case will be 10+ times heavier than that of the 90% case.

Results are presented in Figure 6, 7, 8, 9, 10, 11 (*i.i.d.* setting for 95% and 99%), 12, 13, 14, 15, 16, 17 (*non-i.i.d.* setting for 95% and 99%). Note that we do not have results for the Colorectal dataset, this is because we encountered CUDA memory overflow when running on that dataset. As we can see from these results, our protocol is still robust when $\epsilon = 2$. And for stronger privacy requirements, we can see the utility drops dramatically. It is unsurprising that as the number of Byzantine workers increases and the DP noise increases, the training will tend to aggregate malicious uploads, hence hurting the utility. We can also see that we still have good robustness even when the attack is Optimized Local Model Poisoning, this shows

that when the number of malicious workers is large, our protocol can better defend the Optimized Local Model Poisoning attack compared with the other two attacks.

Experimental results for i.i.d. setting: We present the experimental results on general Byzantine resilience, resilience facing 90% attackers, efficient hyper-parameter tuning in Figure 18, 19, 20 (Gaussian attack), 21, 22, 23 (Optimized Local Model Poisoning attack). As we have observed and analyzed before, similar results can also be observed when our learning protocol is under other attacks.

Experimental results for non-i.i.d. setting: We present all experiments in non-i.i.d. settings as counterparts to previous results. Results are summarise in Figure 24, 25, 26 (Label-flipping attack), 27, 28, 29 (Gaussian attack), 30, 31, 32 (Optimized Local Model Poisoning attack). As we have observed in the i.i.d. case, similar results can also be observed when our learning protocol is under various attacks in the non-i.i.d. setting.

Resilience against adaptive attacker: We present the remaining results in Figure 33, 34, 35 (*i.i.d.* setting), 36, 37, 38 (*non-i.i.d.* setting). All experiments are conducted when there are 60% Byzantine attackers. As we have observed before, our learning protocol is also resilient to adaptive attackers under various attack instantiations.




Figure 6: Byzantine-resilient performance (testing accuracy) in i.i.d. setting when 95% workers are Label-flipping Byzantine attackers.




Figure 7: Byzantine-resilient performance (testing accuracy) in i.i.d. setting when 95% workers are Gaussian Byzantine attackers.




Figure 8: Byzantine-resilient performance (testing accuracy) in i.i.d. setting when 95% workers are Optimized Local Model Poisoning Byzantine attackers.




Figure 9: Byzantine-resilient performance (testing accuracy) in i.i.d. setting when 99% workers are Label-flipping Byzantine attackers.




Figure 10: Byzantine-resilient performance (testing accuracy) in i.i.d. setting when 99% workers are Gaussian Byzantine attackers.




Figure 11: Byzantine-resilient performance (testing accuracy) in i.i.d. setting when 99% workers are Optimized Local Model Poisoning Byzantine attackers.




Figure 12: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 95% workers are Label-flipping Byzantine attackers.




Figure 13: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 95% workers are Gaussian Byzantine attackers.




Figure 14: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 95% workers are Optimized Local Model Poisoning Byzantine attackers.




Figure 15: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 99% workers are Label-flipping Byzantine attackers.




Figure 16: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 99% workers are Gaussian Byzantine attackers.




Figure 17: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 99% workers are Optimized Local Model Poisoning Byzantine attackers.

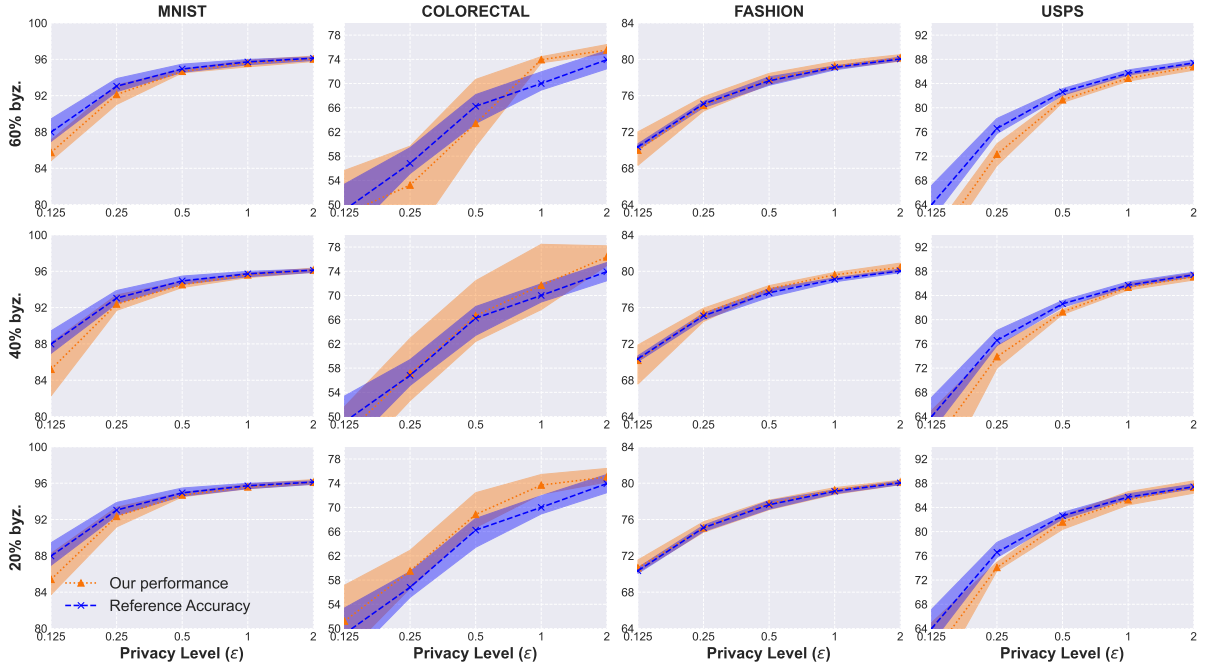


Figure 18: Byzantine-resilient performance (testing accuracy) under Gaussian attack in i.i.d. setting. The experiment is conducted under 3 different attacking levels (20%, 40%, 60% of total workers are Byzantine).




Figure 19: Byzantine-resilient performance (testing accuracy) in i.i.d. setting when 90% workers are Gaussian Byzantine attackers.

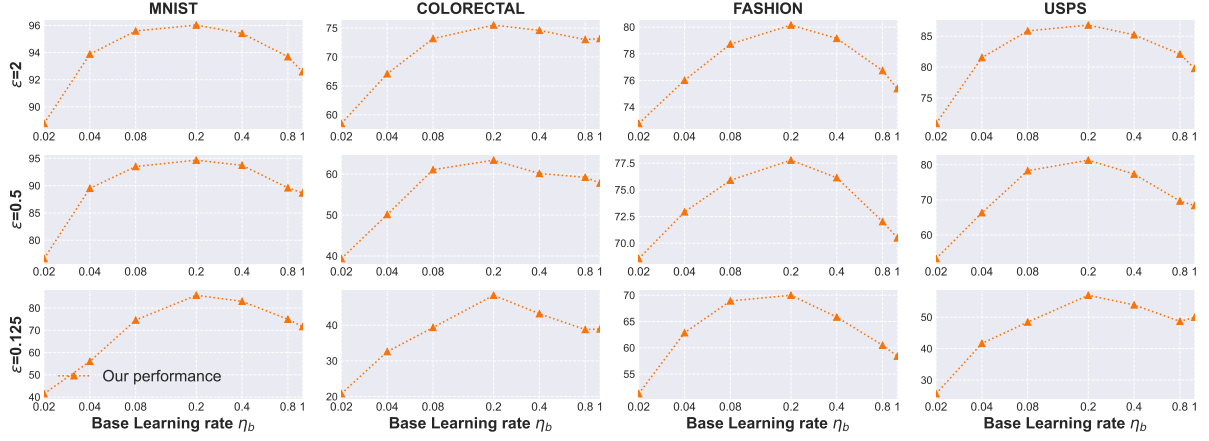


Figure 20: Our hyper-parameter tuning result under i.i.d. setting when facing 60% Gaussian attackers.




Figure 21: Byzantine-resilient performance (testing accuracy) under Optimized Local Model Poisoning attacks in i.i.d. setting. The experiment is conducted where there are 60% attackers.




Figure 22: Byzantine-resilient performance (testing accuracy) in i.i.d. setting when 90% workers are Optimized Local Model Poisoning Byzantine attackers.

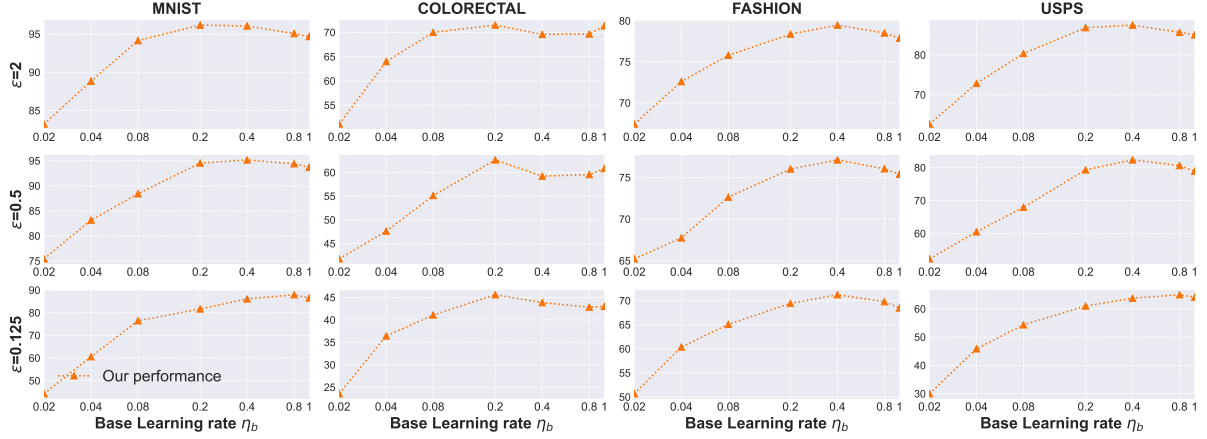


Figure 23: Our hyper-parameter tuning result under i.i.d. setting when facing 60% Optimized Local Model Poisoning attackers.

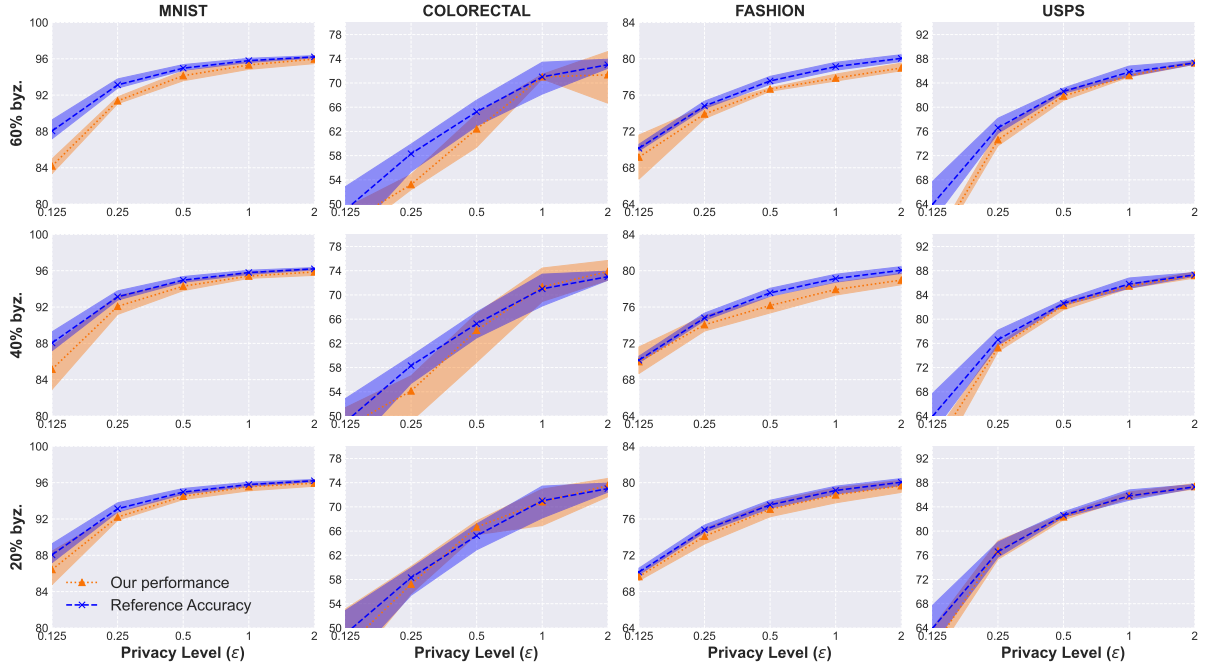


Figure 24: Byzantine-resilient performance (testing accuracy) under Label-flipping attack in non-i.i.d. setting. The experiment is conducted under 3 different attacking levels (20%, 40%, 60% of total workers are Byzantine).




Figure 25: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 90% workers are Label-flipping Byzantine attackers.

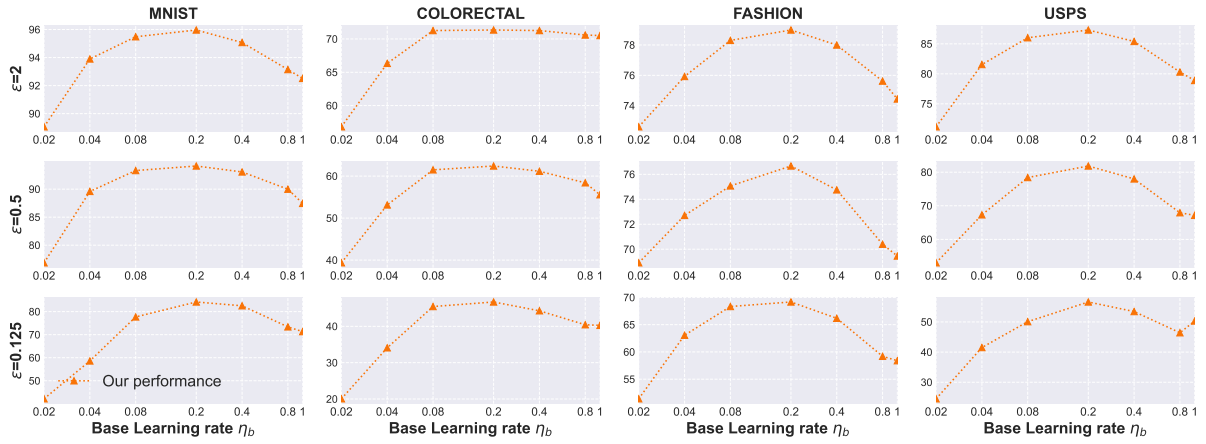


Figure 26: Our hyper-parameter tuning result under non-i.i.d. setting when facing 60% Label-flipping attackers.

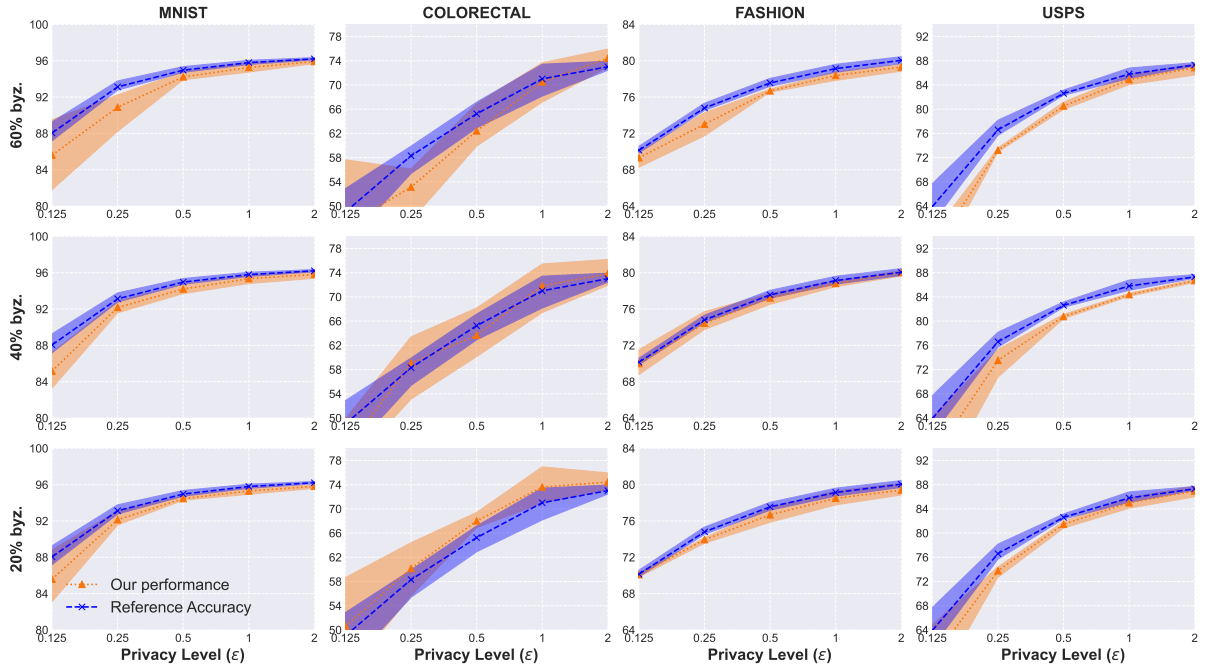


Figure 27: Byzantine-resilient performance (testing accuracy) under Gaussian attack in non-i.i.d. setting. The experiment is conducted under 3 different attacking levels (20%, 40%, 60% of total workers are Byzantine).




Figure 28: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 90% workers are Gaussian Byzantine attackers.

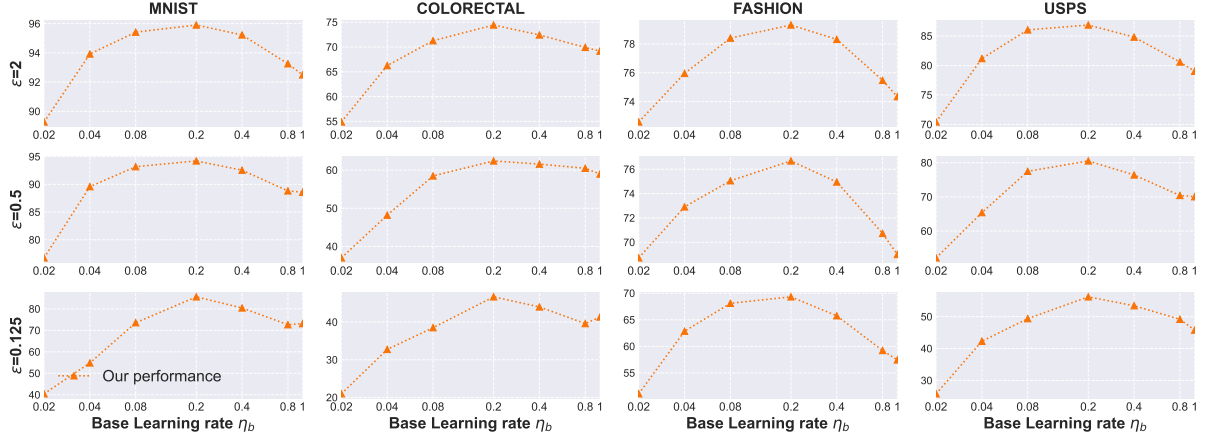


Figure 29: Our hyper-parameter tuning result under non-i.i.d. setting when facing 60% Gaussian attackers.




Figure 30: Byzantine-resilient performance (testing accuracy) under Optimized Local Model Poisoning attack in non-i.i.d. setting. The experiment is conducted where there are 60% attackers.




Figure 31: Byzantine-resilient performance (testing accuracy) in non-i.i.d. setting when 90% workers are Optimized Local Model Poisoning Byzantine attackers.

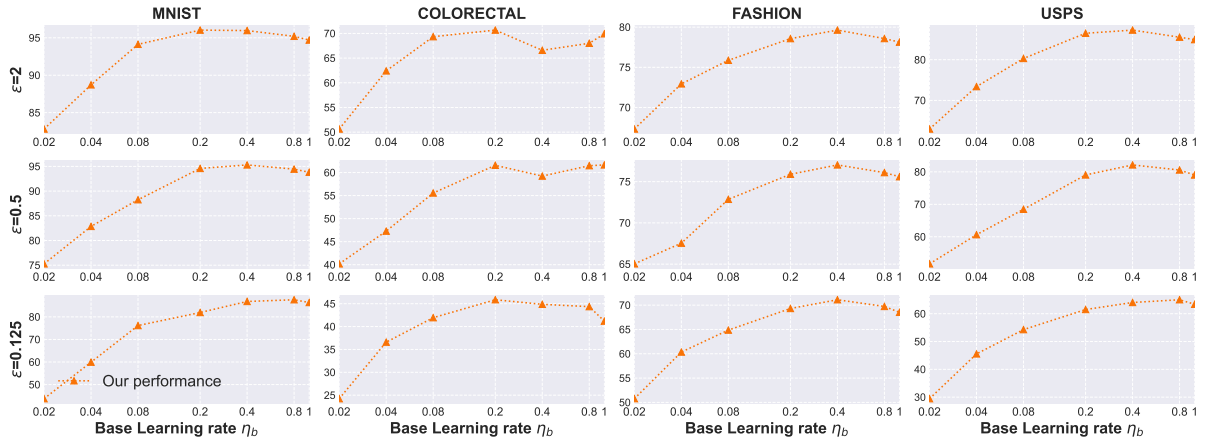


Figure 32: Our hyper-parameter tuning result under non-i.i.d. setting when facing 60% Optimized Local Model Poisoning attackers.

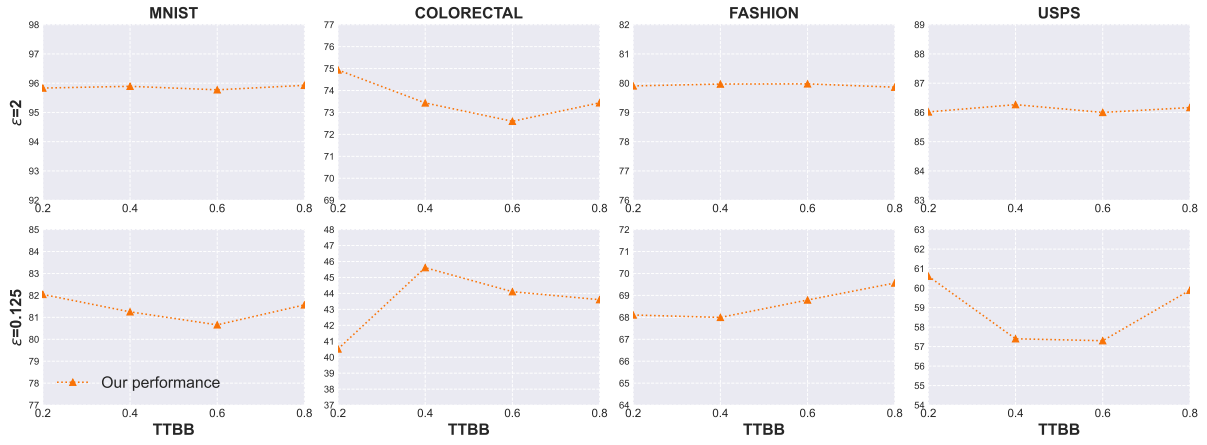


Figure 33: Experiment under Label-flipping attack in i.i.d. setting with different TTBB.

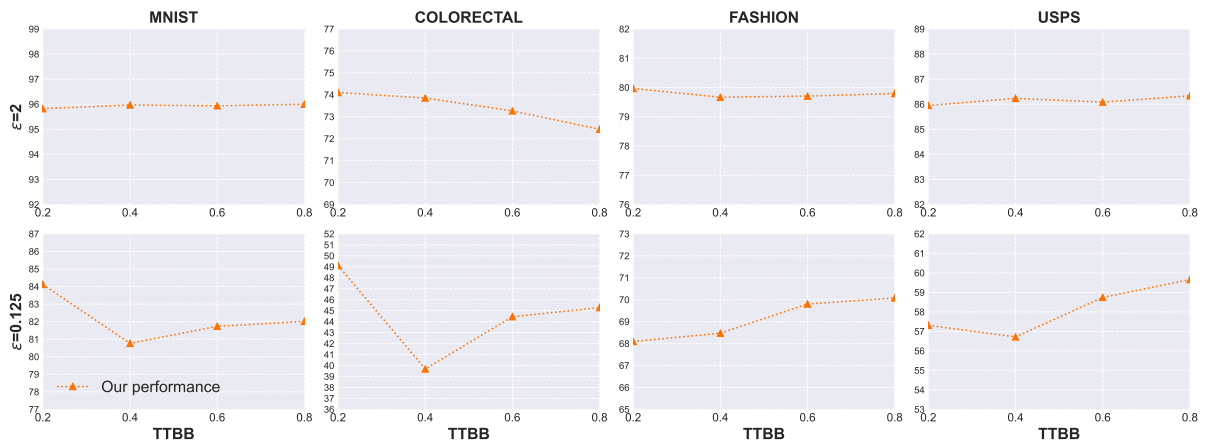


Figure 34: Experiment under Gaussian attack in i.i.d. setting with different TTBB.

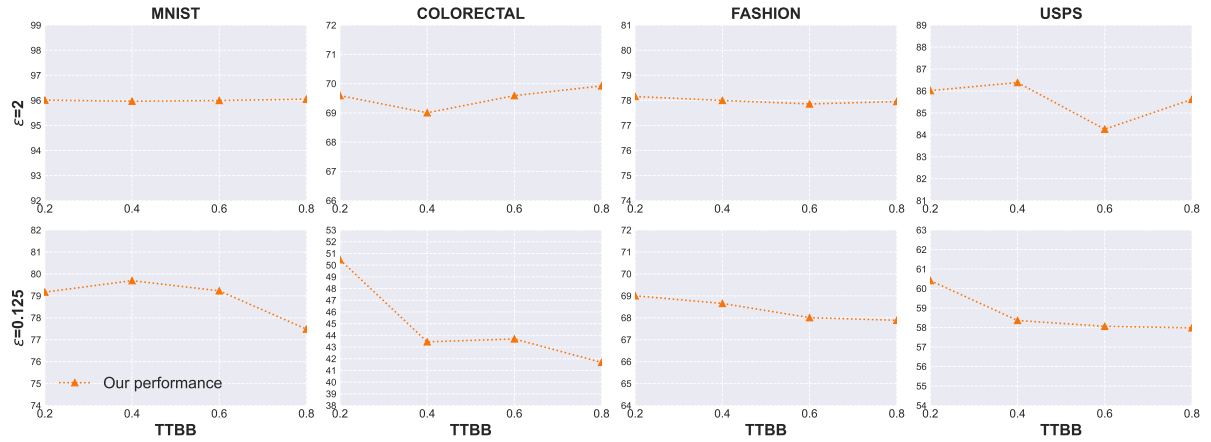


Figure 35: Experiment under Optimized Local Model Poisoning attack in i.i.d. setting with different TTBB.

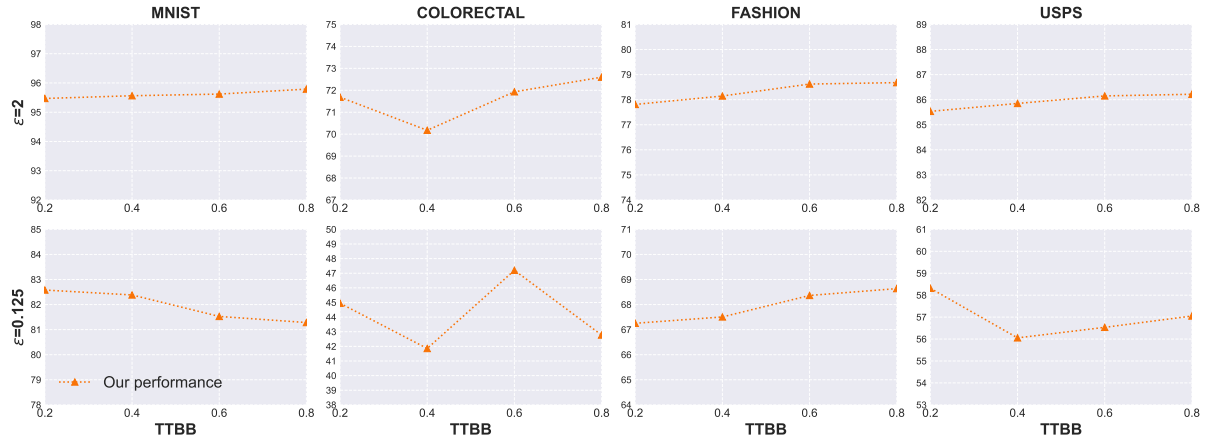


Figure 36: Experiment under Label-flipping attack in non-i.i.d. setting with different TTBB.

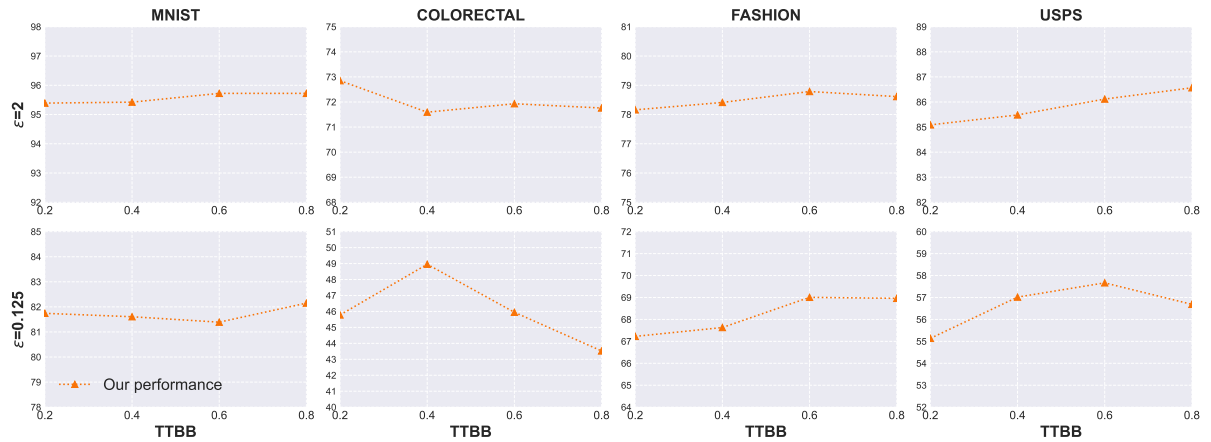


Figure 37: Experiment under Gaussian attack in non-i.i.d. setting with different TTBB.

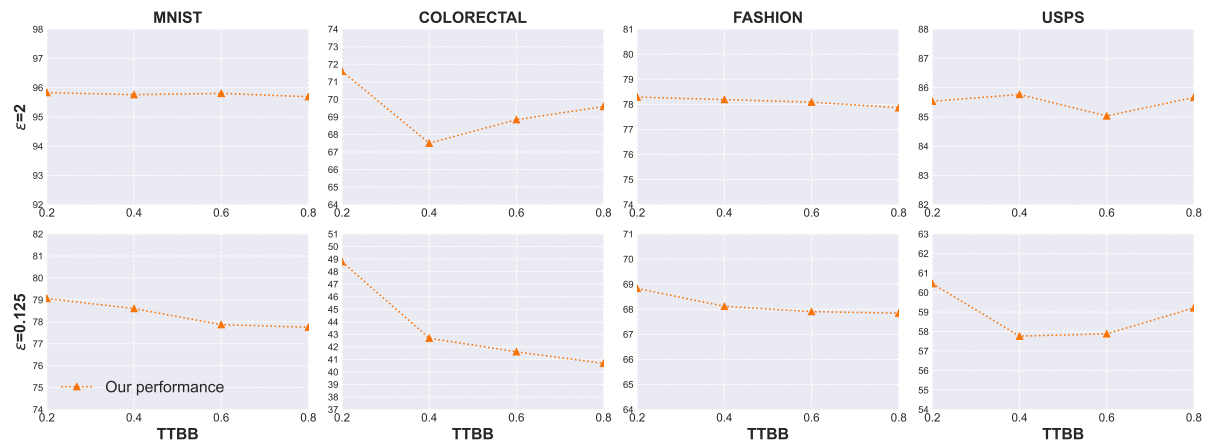


Figure 38: Experiment under Optimized Local Model Poisoning attack in non-i.i.d. setting with different TTBB.