

Decoding FL Defenses: Systemization, Pitfalls, and Remedies

MOMIN AHMAD KHAN, University of Massachusetts Amherst, USA

VIRAT SHEJWALKAR, Google, USA

YASRA CHANDIO, University of Massachusetts Amherst, USA

AMIR HOUMANSADR, University of Massachusetts Amherst, USA

FATIMA MUHAMMAD ANWAR, University of Massachusetts Amherst, USA

While the community has designed various defenses to counter the threat of poisoning attacks in Federated Learning (FL), there are no guidelines for evaluating these defenses. These defenses are prone to subtle pitfalls in their experimental setups that lead to a false sense of security, rendering them unsuitable for practical deployment. In this paper, we systematically understand, identify, and provide a better approach to address these challenges. First, we design a comprehensive systemization of FL defenses along three dimensions: i) how client updates are processed, ii) what the server knows, and iii) at what stage the defense is applied. Next, we thoroughly survey 50 top-tier defense papers and identify the commonly used components in their evaluation setups. Based on this survey, we uncover six distinct pitfalls and study their prevalence. For example, we discover that around 30% of these works solely use the intrinsically robust MNIST dataset, and 40% employ simplistic attacks, which may inadvertently portray their defense as robust. Using three representative defenses as case studies, we perform a critical reevaluation to study the impact of the identified pitfalls and show how they lead to incorrect conclusions about robustness. We provide actionable recommendations to help researchers overcome each pitfall.

ACM Reference Format:

Momin Ahmad Khan, Virat Shejwalkar, Yasra Chandio, Amir Houmansadr, and Fatima Muhammad Anwar. 2018. Decoding FL Defenses: Systemization, Pitfalls, and Remedies. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 32 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Federated learning (FL) [59] is an emerging approach in machine learning (ML) where multiple data owners, called *clients*, collaboratively train a shared model, known as the *global model*, while keeping their individual training data private. The central *server* (service provider) iteratively aggregates *model updates* from each client, which are generated based on their local data. The server merges these updates using an *aggregation rule* (AGR) and uses them to update the global model. Following each training iteration (also known as *round*), the refined global model is distributed to the clients participating in the next round. Prominent distributed platforms such as Google's Gboard [2] for next-word prediction, Apple's Siri [1] for automatic speech recognition [72], and WeBank [93] for credit risk predictions, have adopted this FL mechanism. Its intrinsic characteristic of promoting collaboration while preserving privacy has rendered

Authors' Contact Information: Momin Ahmad Khan, makh@umass.edu, University of Massachusetts Amherst, Amherst, Massachusetts, USA; Virat Shejwalkar, Google, San Jose, USA, vshejwalkar@google.com; Yasra Chandio, ychandio@umass.edu, University of Massachusetts Amherst, Amherst, Massachusetts, USA; Amir Houmansadr, amir@cs.umass.edu, University of Massachusetts Amherst, Amherst, Massachusetts, USA; Fatima Muhammad Anwar, fanwar@umass.edu, University of Massachusetts Amherst, Amherst, Massachusetts, USA.

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 the author(s) 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.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Manuscript submitted to ACM




Fig. 1. FL defense evaluation pipeline. We display common choices for each stage in the pipeline, e.g., FedSGD or FedAvg as the FL algorithm, and highlight the associated pitfall, i.e., using a slow-converging algorithm. We also indicate interdependencies between stages, e.g., large-scale and cross-device in FL type limit the number of malicious clients in the attack.

it indispensable in critical applications, notably in medical diagnosis [29, 46, 75], activity recognition [26, 68, 86, 107], and next-character prediction [87].

FL is gaining popularity due to its privacy-preserving and collaborative nature, yet it faces vulnerabilities to *poisoning attacks* [28, 83, 84, 91], where malicious or *compromised clients* intentionally corrupt FL training and *poison* the global model. This can result in a poisoned model that performs poorly on all inputs in *untargeted poisoning* attacks or on specific inputs in *targeted poisoning* attacks. To address these threats, the community has developed various defense mechanisms. Robust AGRs such as Multi-krum [14] and Trimmed-mean [103], detect and discard malicious updates. Certified defenses like CRFL [97] and Ensemble FL [19] provide robustness certifications. Tools like FLDetector [106] proactively identify and remove malicious clients during training. Meanwhile, FedRecover [20] focuses on post-poisoning recovery after an attack, aiming to restore the global model’s performance. Ditto [52] integrates fairness and robustness by regularizing the local training objective, and Cronus [23] enhances security and privacy through knowledge distillation.

Systemizing FL Defenses: Besides these few defenses, the variety of available options (Table 4) poses a challenge for practitioners, i.e., determining the right defense for a specific use case or integrating multiple defenses for enhanced robustness becomes complex without a clear understanding of where a defense and its dependencies fit in the FL pipeline.

To address these research gaps, we conduct a comprehensive systemization (§3) of FL defenses, organizing them along three crucial dimensions: processing of client updates, server’s knowledge, and defense phase. While existing works [10, 13, 38, 40, 80, 84] have designed taxonomies primarily focused on adversarial ML, including those that guide the selection of the appropriate attacks and settings for FL, a dedicated systemization for FL defenses has been lacking.

To the best of our knowledge, we are the first to propose such a systemization. Our framework simplifies the selection and integration of defenses, clarifying when and where each defense is applied (§3.1). For example, Figure 2 shows that FLDetector operates in the pre-aggregation phase (“when”) at the server (“where”). Moreover, the systemization highlights underrepresented defense types, encouraging further exploration and innovation. For example,

our analysis identified FedRecover as the only post-training, update-modification technique using estimation, prompting exploration of not only other estimation-based but also post-training defenses. From our systemization in Table 1, we can see that most defenses are on the server side, operate during pre-aggregation, and employ metric-based processing of client updates. Importantly, our systemization is designed to be expandable to incorporate more attributes in the future.

Pitfalls in experimental setups of FL defenses: During our systemization, we find that defenses are evaluated across various experimental setups with different choices of datasets, data distributions, attacks, etc. We thoroughly survey 50 top-tier FL defense works (§4) and report frequently used choices of six experimental setup components: *data and distribution, FL type and algorithm, attack, and evaluation type*. We find questionable trends in these choices, e.g., most works use the intrinsically robust MNIST dataset and resort to simple attacks such as label flipping. This motivates the need to uncover *pitfalls* in the experimental setups of existing FL defense works and provide guidelines for future evaluations.

Existing literature provides guidelines for robustness evaluations. For instance, [6] offers insights and recommendations, focusing on pitfalls in centralized ML-based security system evaluations. Similarly, [22] identifies pitfalls in evaluating adversarial robustness and suggests mitigation guidelines. In [84], authors question trends in existing attack threat models, demonstrating FL robustness under practical limitations. However, a comprehensive exploration of FL defenses and their experimental setups is lacking.

To fill this gap in the literature, we identify six distinct pitfalls in the evaluation setups of existing works based on our survey in §4. We take inspiration from [43] in identifying pitfalls in FL defenses and building on top of their work. Not only do we explore additional pitfalls and analyze their impact in much more diverse settings, but we also first perform an extensive systemization to justify the choices we make in pitfall analysis. For example, Figure 2 and Table 1 help us to select and justify three distinct defenses for the impact analysis of pitfalls in §6.

Figure 1 shows the pitfalls at the point of their occurrence in the FL training pipeline. For instance, choosing an intrinsically robust dataset is a pitfall that occurs in the dataset selection stage. We also show the common choices for each stage of the pipeline, such as global or personalized evaluation in the last stage. Figure 1 also highlights the interdependencies between stages. For example, non-i.i.d. distribution in the distribution stage influences the choice of evaluation metrics (§6.6), while large-scale limits the threat model in the attack stage (§6.4). We thoroughly explain each pitfall and its prevalence in §4, and provide actionable recommendations to overcome them. Finally, in §6, we perform a thorough impact analysis of the identified pitfalls using three representative defenses and show how we can avoid them by following our recommendations. These pitfalls *can also be applied to attacks*, e.g., only evaluating an attack in cross-silo settings and not considering its efficacy in the cross-device setting is a pitfall. Similarly, other pitfalls can arise from attacks, as attacks and defenses are inherently interconnected—essentially two sides of the same coin working together. *However, for the purpose of this paper, we will analyze the pitfalls from the lens of defenses.* In summary, we make the following contributions:

- 1. Systemization of FL Defenses:** We perform a comprehensive systemization of FL defenses along three dimensions: processing of client updates, server’s knowledge, and the defense phase (§3).
- 2. Identifying Major Pitfalls:** We comprehensively review 50 top-tier FL defense works and identify six prevalent pitfalls. In response to each pitfall, we provide actionable recommendations to guide future research efforts (§4).
- 3. Dissecting the Impact of Pitfalls:** Guided by our systemization, we choose three representative FL defenses and use them to perform a thorough impact analysis of our identified pitfalls. We show how the pitfalls lead to a false sense of security, and by following our recommendations, the research community can overcome them (§6).

This work is not meant as finger-pointing, particularly to the defenses under evaluation. We have chosen them as representative contributions to the field, humbly employing them as testbeds to offer constructive guidelines for future research.

Table 1. A systematic overview of FL defenses, helping practitioners in 1) selecting defenses aligned with their use cases, 2) combining multiple defenses for heightened performance, and 3) designing new defense approaches by gaining insights into the FL defenses.

Dimension	Types	Attributes	Description	Defenses
Processing of Client Updates	Filtering	Update-based	Filters updates by comparing update values with each other.	TrMean [103], Krum [14], GeometricMean [73]
		Metric-based	Filters updates by comparing metrics associated with each update.	MST-AD & Density-AD [77], LF-Fighter [39], ERR & LFR [28], FLDetector [106]
		Similarity-based	Filters updates by comparing their similarity with a reference.	CONTRA [8], FLTrust [18]
	Update re-weighting	Loss-based	Applies weights based on loss for an update.	Sageflow [71], Anomaly Detection [51]
		Optimization-based	Applies weights based on an optimization problem.	SmartFL [100]
	Update modification	Scaling/Clipping	Scales or clips an update if it exceeds a certain threshold.	Norm Bounding [88], Sigeguard [102]
		Distillation	Distills an update into a low-dimensional vector.	Cronus [23], Auror [85]
		Regularization	Introduces regularization to control robustness and privacy.	Ditto [52]
		Estimation	Estimates a benign update from historical information.	FedRecover [20]
Server's Knowledge	Knowledge of data	No-knowledge	Server has no knowledge of data or its distribution at the client side.	TrMean [103], FedRecover [20], FLDetector [106], Cronus [23]
		Partial knowledge	The server uses a small auxiliary dataset.	Sageflow [71], SmartFL [100], FLTrust [18]
	Knowledge of updates	No-knowledge	Server does not have knowledge of local model updates.	Auror [85], Cronus [23]
		Full knowledge	Server has complete knowledge of local model updates.	FLCert [21], Krum [14], Anomaly Detection [51]
Defense Phase	Aggregation	Pre-aggregation	Processing of client updates is done before updates are aggregated.	MST-AD & Density-AD [77], Bulyan [60], TrMean [103], Krum [14], Sageflow [71]
		Post-aggregation	Processing of client updates is done after aggregation.	FLCert [21], ERR & LFR [28]
	Non-aggregation	Local training	The defense component is part of the local training.	FLIP [105], Ditto [52], Cronus [23]
		Post-training	The defense is applied after training.	FedRecover [20]

2 Background

2.1 Federated Learning (FL)

In FL [41, 44, 59], a service provider, called *server*, trains a *global model*, θ^g , on the private data from multiple collaborating clients, all without directly collecting their individual data. During the t^{th} FL round, the server selects n out of total N clients and shares the most recent global model (θ_g^t) with them. Then, a client k uses their local data D_k to compute an update ∇_k^t and shares it with the server. These updates serve as a client's contribution towards refining a global model. Depending on how a client computes their update, FL algorithms can be broadly divided [59] into *FedSGD* and *FedAvg*. In *FedSGD*, a client computes the update by sampling a subset b from their local data and calculating a gradient of loss $\ell(b; \theta_g^t)$ of the global model on the subset, i.e., $\nabla_k^t = \partial \ell(b; \theta_g^t) / \partial \theta_g^t$. In *FedAvg*, a client k *fine-tunes* θ_g^t on their local data using stochastic gradient descent (SGD) for a fixed number of local epochs E , resulting in an updated local model θ_k^t . The client then computes their update as the difference $\nabla_k^t = \theta_k^t - \theta_g^t$ and shares ∇_k^t with the server. Next, the server computes an aggregate of client updates using an AGR, f_{agg} (such as mean), i.e., using $\nabla_{\text{agg}}^t = f_{\text{agr}}(\nabla_{\{k \in [n]\}}^t)$. Finally, the server updates the global model of the $(t+1)^{\text{th}}$ round using SGD as $\theta_g^{t+1} \leftarrow \theta_g^t + \eta \nabla_{\text{agg}}^t$ with server's learning rate η . Due to these differences, FedAvg achieves faster convergence and attains higher accuracy than FedSGD [59].

After discussing the update process of FL models and the collaboration between servers and clients, we present the FL applications (*where it is deployed*) and setups (*how it is used*). There are two main types of deployments: **cross-device** and **cross-silo**, as explained in [41]. In *cross-device FL*, N is very large (ranging from a few thousand to billions) [78], and only a tiny fraction of them are chosen in each FL training round ($n \ll N$). These clients are typically resource-constrained devices such as mobile phones, smartwatches, and other IoT devices [2, 107]. Contrastingly, in *cross-silo FL*, N is moderate (up to 100) [49], and all clients are selected in each round ($n = N$). These clients are typically large corporations, including banks [93] and hospitals [66].

2.2 Poisoning Attacks in FL

There are various poisoning attacks in literature [9, 11, 12, 14, 28, 57, 60, 63, 83, 99]. An *untargeted* poisoning attack aims to lower the test accuracy for all test inputs indiscriminately [11, 28, 57, 60, 99]. A *targeted* poisoning attack [9, 12] lowers the accuracy on specific test inputs. For instance, in *backdoor attacks* [9] (a sub-category of targeted attacks), the goal is to misclassify only those test inputs that have an embedded *backdoor trigger*. Since these attacks only affect a subset of inputs, they are much weaker than untargeted attacks.

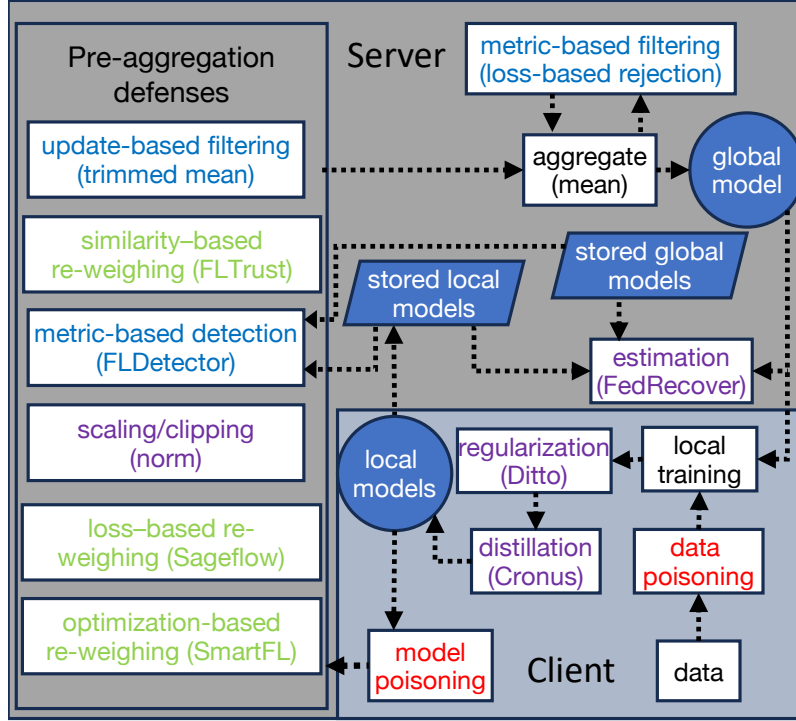


Fig. 2. Systemizing FL Defenses: Categorization of defenses based on their *defense phases*. Processing operations for client updates are color-coded: filtering (blue), re-weighting (green), and modification (purple). Defenses belonging to the same category of processing may have different phases, e.g., FedRecover performs update modification at the server, while Ditto performs update modification at the client. Dependencies are also highlighted, such as FedRecover utilizing stored local and global models for estimation.

We can also divide the attacks based on the adversary’s capabilities. In model poisoning attacks [9, 11, 12, 28, 60, 83, 99], the adversary is strong enough to access and perturb the model gradients on malicious devices before they are sent to the server in every training round. A data-poisoning adversary [63] is much weaker than the model poisoning adversary [11, 28, 83, 84] as it can only poison the datasets on malicious devices.

2.2.1 Attacks used in our Study. In our evaluation, we focus on untargeted model poisoning attacks *as they are stronger* [84].

Stat-Opt [28]: provides a generic model poisoning method and tailors it to specific AGRs such as TrMean [103], Median [103], and Krum [14]. The adversary first calculates the mean of the benign clients’ updates, ∇^b , and finds the *static* malicious direction $w = -\text{sign}(\nabla^b)$. It directs the benign average along the calculated direction and scales it with γ to obtain the final poisoned update, $-\gamma w$.

Dyn-Opt [83]: proposes a general poisoning method and tailors it to specific AGRs, similar to Stat-Opt. The key distinction lies in the *dynamic* and *data-dependent* nature of the perturbation. The attack starts by computing the mean of benign updates, ∇^b , and a data-dependent direction, w . The final poisoned update is calculated as $\nabla' = \nabla^b + \gamma w$, where the attack finds the largest γ that can bypass the AGR. They compare their attack with Stat-Opt and show that the dataset-tailored w and optimization-based scaling factor γ make their attack a much stronger one.

2.2.2 *Threat Model in Our Study.* Below, we detail the threat models for poisoning attacks used in our study.

Goal: Our adversary is *untargeted*, crafts malicious updates and sends them to the server, where, upon aggregation with other updates, it indiscriminately lowers the accuracy of the global model for all test inputs.

Knowledge: In line with most defense approaches, we assume that the adversary knows the AGR used by the server. Unless explicitly specified, the adversary has complete knowledge of the gradients of both malicious and benign clients. In some cases, we employ a partial knowledge adversary, where the adversary knows the server’s AGR but not the gradients of benign clients.

Capabilities: Our untargeted model poisoning adversary controls m out of N clients. The adversary is strong enough to manipulate model updates of the malicious clients it controls and has access to the global model parameters shared every round. We set the proportion of malicious clients at 20% (unless stated otherwise), a common benchmark in prior studies [18, 20, 28, 83], which also examined how varying this percentage impacts the severity of attacks. To ensure consistency and comparability with these works, we adhere to the same 20% setting in our implementation.

3 Systemization of Defenses Against FL Poisoning

Here, we introduce a systematization for FL defenses and use it to rationalize the selection of three representative defenses from the literature. Later, in §6, we use these chosen defenses to conduct a comprehensive impact analysis of the pitfalls outlined in §4.

3.1 Classification of Defenses

Here, we present the three key dimensions along which we classify FL defenses in literature, as shown in Table 1. In Figure 2, we group defenses according to the third dimension, i.e., *defense phase*, and highlight their dependencies.

3.1.1 *Processing of client updates.* Client model updates undergo several processing steps before they are aggregated at the server. The commonly used processing operations are:

Filtering updates to entirely or partially eliminate local updates. *Filtering* defenses fall into two main categories: those based on the values of local model updates, termed *update-based filtering* [14, 73, 103], and those relying on some metrics associated with local models, known as *metric-based filtering* [28, 39, 77, 106].

a) *Update-based filtering* [14, 73, 103] is based on the values of local model updates. It further divides into *dimension-wise* and *vector-wise* filtering. Dimension-wise filtering defenses, such as TrMean [103] and Median [103], filter out malicious values along each update dimension, while vector-wise filtering defenses, such as, Krum [14], remove entire malicious updates.

b) *Metric-based filtering* [28, 39, 77, 106] relies on some metrics associated with local models, e.g., FLDetector [106] uses a *suspicious score* as the metric to identify and remove malicious clients from the training process. We describe FLDetector in detail in §3.2.1. In *loss-based rejection* [28], the loss associated with and without incorporating an update for aggregation is calculated, and updates with higher loss are removed. Similarly, *error-based rejection* [28] removes updates by assessing the error instead of the loss.

Update re-weighting involves assigning a weight to each local update, reflecting its perceived level of maliciousness. Various re-weighting approaches exist, including:

a) *Similarity-based re-weighting* [8, 18] is shown by FLTrust [18], where the server assigns a trust score to each client based on the similarity of its updates to the server’s update, computed on a small dataset.

b) *Loss-based re-weighting* [51, 71] illustrated by Sageflow [71], involves the server assigning a weight to each update based on local model loss on a small dataset at the server.

c) *Optimization-based re-weighting* is demonstrated by SmartFL [100], which assigns weights through an optimization problem with the same number of parameters as that of clients.

Update Modification changes the update itself to safeguard the global model from the impact of malicious updates. The key techniques within this approach are:

a) *Scaling Updates* limits a local update by clipping it if it exceeds a certain threshold, for example, defenses like Norm-bounding [88, 102].

b) *Distilling update knowledge* [23, 85] is another facet of update modification that involves avoiding the transmission of the entire local model update to the server due to the *curse of dimensionality* [23], which increases the risk of higher impact from poisoning attacks. Instead, clients send distilled information to the server. In Cronus [23], clients send soft labels to the server, and the aggregate is used to update local models. This defense mitigates the risk of poisoning attacks and directly prevents whitebox inference attacks, as the server cannot access the local model parameters.

c) *Regularization* defenses perform regularization to achieve personalized [35, 36] client models. An example is Ditto [52], which modifies the local training objective by introducing a regularization term to control the tradeoff between privacy and robustness of the local model.

d) *Estimation* is exemplified by FedRecover [20], which leverages past information and estimation to recover the poisoned global model. We discuss the mechanism of FedRecover in detail in §3.2.1.

3.1.2 Server’s Knowledge. The defense is generally applied at the server, where it collects all local updates and strives to obtain the best possible global model that is least affected by malicious updates. The server’s knowledge varies across different defenses and can be described in terms of data and local model updates:

Knowledge of Data: In the *no-knowledge* setting [20, 23, 103, 106], as the name suggests, the server lacks information about the data used for training and testing. It only possesses knowledge of the collected local model updates, examples of which include TrMean [103], Krum [14], and Median [103]. Conversely, in the *partial-knowledge* setting [18, 71, 100], the server possesses a small dataset, which it deploys in various ways, such as calculating entropy associated with updates [71] or assigning a trust score to each client [18] to enhance aggregation robustness against attacks.

Knowledge of Local Model Updates: In the *full-knowledge* setting [14, 21, 51], the server has complete access to the model parameters of all clients, representing the widely used scenario. In the *partial-knowledge*, or *distilled-knowledge* setting [23, 85], the server only has access to some distilled form of the model parameters, such as the output layers [23].

3.1.3 Defense Phase. We categorize defenses based on the training *phase*, specifying when and where in the training pipeline the defense is applied.

Aggregation-based defenses are applied during the aggregation phase. These defenses can be further categorized into *pre-aggregation* [14, 60, 71, 77, 103] defenses that perform processing of client updates such as dimension-wise filtering before aggregating updates, or *post-aggregation* [19, 21, 28] defenses, e.g., Ensemble FL [19] that creates all possible aggregations of k models from N clients, then selects the most frequent predicted label as the correct one.

Non-aggregation-based defenses that are not performed at aggregation can be further categorized based on their *phase*:

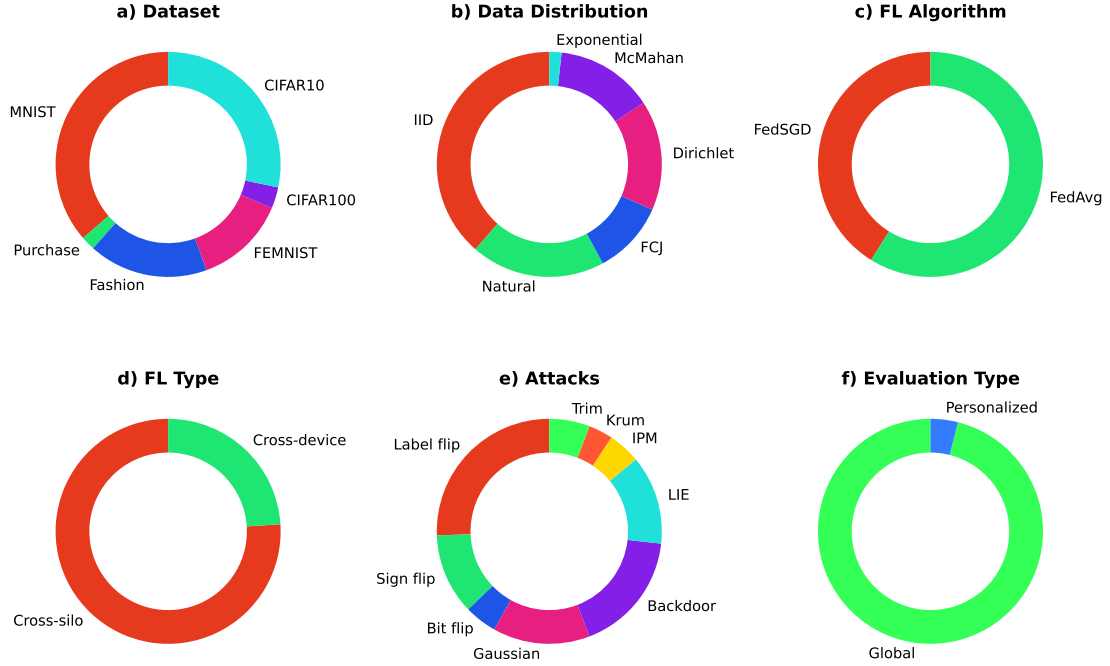


Fig. 3. Frequency of choices of the six key components of robustness evaluation setup: dataset, distribution of clients’ data, FL algorithm, FL type, attacks, and evaluation. § 6 discusses the impacts of choices on the robustness of FL poisoning defenses.

a) *During Local Training:* [23, 52, 105] The defense takes place at the client’s side during local training. For instance, Ditto [52] uses regularization in client training to control the deviation of benign local models from poisoned global models.

b) *Post-Training:* FedRecover [20], employing a recovery-based mechanism, falls into this category as it requires historical information from one training session to estimate the un-poisoned global model during the *recovery* phase.

3.2 Our Selected Case Study Defenses

We choose three defenses for brevity to analyze the impact of pitfalls in §6; TrMean, FLDetector, and FedRecover. We first provide each of their descriptions and then give a detailed justification for choosing them for our impact analysis.

3.2.1 *Descriptions of our chosen defenses.* **Trimmed Mean (TrMean)** [103] is a foundational defense used in advanced defenses [20, 83, 106]. It sorts each input dimension j of the client updates, discards the m largest and smallest values (where m indicates compromised clients), and averages the rest.

FLDetector [106] is designed to *detect* and eliminate malicious clients, ensuring a byzantine-robust FL system obtains a precise global model. FLDetector operates on the principle that malicious updates, tainted by adversaries, differ statistically from benign ones.

For discerning these updates, the server estimates a global model update for client k at round t using the L-BFGS algorithm: $(\hat{\nabla}_t^k = \nabla_{t-1}^k + \hat{H}^t \cdot (\theta_t - \theta_{t-1}))$. Here, ∇_{t-1}^k . The server retains past N global model differences $(\Delta\theta_t)$ and updates differences $(\Delta\nabla_{t-1})$ to compute the *HVP (Hessian Vector Product)* with *L-BFGS*. It then gauges a client’s *suspicious*

score by comparing actual and predicted updates through their Euclidean distance. With scores from the last N rounds, clients are clustered via Gap statistics [89] and K-means. The group with higher average scores is deemed malicious. Upon detecting a rogue client, the server ceases training, removes the offender, and restarts training to achieve enhanced accuracy.

FedRecover [20] aims to *recover* an FL global model compromised by a poisoning attack. In the *original training phase*, for each round t , FedRecover saves ∇_k^t from client k and global models θ_g^t . This data is used as *historical information* in the *recovery phase*, which consists of the following stages. In the *warmup phase*, the server requests clients' *exact updates* for the initial T_w rounds. In the *estimation phase*, the server *estimates* client updates each round, with $\hat{\nabla}_k^t$ representing client k 's estimated update at round t . This estimate is derived using the L-BFGS algorithm [67] based on the original global model, client update, and recovered global model. The model's estimated update is defined as $\hat{\nabla}_k^t = \bar{\nabla}_k^t + \hat{H}_k^t(\hat{\theta}_g^t - \bar{\theta}_g^t)$, where the latter term is the HVP(Hessian Vector Product). Every T_c rounds, to ensure the estimated global model $\hat{\theta}_g^t$ aligns closely with the accurate model, the server initiates a *periodic correction* by requesting *exact updates*. If any component of a client's estimated update surpasses the *abnormality threshold* τ , that client is prompted for an exact update. In the final *fine-tuning phase*, spanning T_f rounds, clients are asked to provide their exact updates, ∇_k^t , to refine the global model by eliminating potential estimation errors.

3.2.2 Justification for our choice of representative defenses. We justify that the chosen defenses are distinct along the dimensions in Table 1. TrMean [103] is a **pre-aggregation, update-based filtering** defense that removes malicious components of client updates dimension-wise during training. FLDetector [106] is a **pre-aggregation, metric-based filtering** defense designed to detect and remove malicious clients during training by analyzing their updates. The metric FLDetector uses is the *suspicious score*, which measures the consistency between the actual update and an estimated one. It is important to note that although TrMean and FLDetector seem similar in their filtering mechanisms, they are different and have distinct approaches. TrMean filters *components* of updates before aggregating them, while FLDetector removes entire *clients* from the training process if they deviate too much from an estimated reference update, which is calculated using historical information. FedRecover [20] is a **post-training, update modification** defense that uses *estimation* to *recover* from a previously poisoned global model. It is an advanced defense that uses TrMean in its aggregation phase and relies on a detection mechanism to remove malicious clients before it starts the recovery process. In §6.1.2 and §6.5.3, we show the performance of FedRecover when we perform Stat-Opt on TrMean and FLDetector, respectively.

All three defenses do not require any auxiliary dataset at the server, and we prefer this setting because the server might not have access to the dataset in practical, real-world scenarios. Similarly, we prefer the *full-knowledge* of client updates scenario, as most of the works [14, 18, 21, 51, 52, 71] use this setting, and it leads to a stronger adversarial setting. Although the server in all three of our chosen defenses has full knowledge of the client model updates, the defenses differ in the amount of historical information needed. Figure 2 clearly shows these different dependencies. TrMean does not require model updates from the past rounds and performs filtering using client model updates of the current round. FLDetector requires updates from the past *few* rounds and uses them to calculate the malicious score for updates in the current round. Since FedRecover is a post-training defense, it requires updates from all the rounds in the original training. Therefore, the volume of updates required is highest for FedRecover and lowest for TrMean.

4 Pitfalls in FL Defense Evaluations

After presenting the detailed systemization of defenses, it is imperative to unveil critical pitfalls in FL robustness evaluations. By scrutinizing 50 defenses (Table 4 in Appendix), we link each pitfall to specific components in the FL workflow (Figure 1). We examine each pitfall’s prevalence (Figure 3) across the 50 works, discuss their implications, and conclude with practical recommendations.

Pitfall-1: Intrinsically robust datasets. The chosen datasets are intrinsically robust and lead to incorrect conclusions about a defense’s performance.

Description: A designed defense may seem to perform well against specific attacks upon evaluation [103]. However, the evaluation dataset might be inherently robust because it is simple and lacks complexity. Therefore, in such situations, we cannot tell if an attack is mitigated due to the inherent robustness of the dataset or the effectiveness of the defense.

Prevalence and implications: While it is intuitive that a defense mechanism’s performance inevitably varies across datasets, using overly simple datasets like MNIST fails to yield meaningful insights into the true robustness of a defense mechanism, (§6.1). MNIST is a class-balanced dataset used in FL using synthetic techniques such as Dirichlet Distribution [61]. Conversely, real-world FL tasks are complex and characterized by highly class-imbalanced datasets (§6.2.1).

Despite efforts to create open-source datasets mirroring real-world scenarios [17, 25], *our survey reveals that MNIST remains predominant, constituting 30% of the works* [19, 42, 48, 55, 94, 103] (Figure 3a). CIFAR10 and FashionMNIST, though commonly used, lack true FL representation due to their class-balanced nature (§6.2.1). FEMNIST [17], a real-world dataset specifically curated for FL, is used by only 20%. Another interesting observation from Figure 3a is the *exclusive reliance on image-classification datasets*, despite the popularity of language and vision-language models in contemporary research.

Recommendations: Future Evaluations should consider using FL tasks of varying complexities, such as FEMNIST and CIFAR10, for classification and exploring other modalities, such as language, for NLP tasks. In our evaluations in §6 we use image classification datasets, FEMNIST and CIFAR10, and a language dataset, StackOverflow, in §6.4.1 for large-scale FL evaluation.

Pitfall-2: Homogeneous client data distributions. Using i.i.d. (homogeneous) distributions with low heterogeneity may create a deceptive sense of system robustness that does not reflect real-world complexities.

Description: Evaluating a defense on a particular dataset may yield perceived robustness due to inherent homogeneity (§6.2.1) in the dataset distribution rather than the efficacy of the defense technique itself.

Prevalence and implications: We find that *around 50% of the works, use i.i.d. distributed data, despite evidence that it is easier to defend against such distributions* [11, 28, 83, 104] (as seen in Figure 3b). The second most common approach involves a natural data distribution where each sample is associated with a client such as StackOverflow [3]. Other artificial distributions include *FCJ* [28], *Dirichlet (Dir)* [9, 78] and *Mcmahan* [59]. In §6.2.1, we prove that Dirichlet more closely aligns with real-world distributions, informing our subsequent analysis in §6.2 on how defense performance varies with different distributions and levels of heterogeneity.

Recommendations: Future works should prioritize the use of real-world datasets to provide a more realistic evaluation. When working with class-balanced datasets like MNIST, FashionMNIST, and CIFAR10, it is crucial to distribute them

among clients heterogeneously. This approach aims to mimic, to some extent, the diversity found in real-world distributions.

Pitfall-3: Slow-converging algorithms. Evaluations often overlook the use of state-of-the-art, fast-converging algorithms, thereby compromising robustness.

Description: The robustness evaluation of an FL defense may heavily rely on the choice of FL algorithms, such as FedAvg or FedSGD, used in its implementation. Using a superior algorithm can enhance system robustness by addressing potential weaknesses in the FL algorithm rather than focusing solely on defense improvement.

Prevalence and implications: Despite FedAvg’s recognized advantages in performance, faster convergence, and lower communication overhead compared to FedSGD [59], **approximately 40% of prior works employ the slow-converging FedSGD algorithm for evaluations** (Figure 3c). This suboptimal choice contributes to a larger window for attacks and results in a more significant accuracy drop (§6.3).

Recommendations: Future evaluations should prioritize state-of-the-art, fast-converging FL algorithms to remove any weaknesses (such as slow convergence) associated with the FL algorithm.

Pitfall-4: Limited FL settings. Considering only limited scenarios while ignoring practical limitations and factors related to scalability, such as computation, communication, cost, and storage.

Description: The performance of a defense can vary when constrained by real-world limitations, e.g., cost constraints may lead to selecting a very low percentage of malicious clients. Also, the computation, communication, cost, and storage overhead associated with scaling up the system might not be feasible in practical scenarios.

Prevalence and implications: Only 24% of the works in our survey use the cross-device setting (Figure 3d). As demonstrated by [84] on FEMNIST, CIFAR10, and Purchase datasets with the cross-device setting, using a low percentage of malicious clients due to cost constraints reduces attack performance. We show that on an even larger scale using the naturally distributed Stackoverflow dataset in the cross-device setting, the attack shows *shows no effect on the non-robust mean AGR* (§6.4.1). Additionally, defenses that rely on consistent historical information [20, 106] are incompatible with the cross-device setting (§6.4.2) because a client is not selected in every round. We thoroughly discuss the scalability issues of FedRecover and FLDetector in §6.4.2.

Recommendations: Future evaluations should include deployment conditions of a much larger scale (cross-device) and ensure that the defense provides significant utility compared to the computation and communication cost it incurs. While we do not label the cross-silo setting as impractical, we emphasize designing defenses that are compatible with the cross-device setting as well.

Pitfall-5: Naive attacks. Evaluating defenses solely against simple and naive attacks rather than incorporating strong state-of-the-art attacks makes a defense seem robust.

Description: The true robustness of a defense emerges when tested against strong and adaptive attacks, i.e., attacks tailored for a defense algorithm. Relying on attacks known to be weak and naive for evaluating a new defense will not give us a true picture of the defense’s robustness.

Prevalence and implications: Figure 3e shows the frequency of various attacks used in our survey. Despite the existence of several strong poisoning attacks in the literature [11, 28, 83, 91, 99], our analysis reveals that **about 40% of the works opt for simplistic approaches** such as random Gaussian [14], label flipping [52], sign flipping [48], bit




Fig. 4. Comparative analysis of TrMean AGR with FedSGD and FedAvg under trim attack. TrMean is more susceptible to poisoning with FedSGD due to slow convergence, which gives adversaries more time for poisoning (threat model in §2.2.2)

flipping [98]), even though prior works have shown their poor performance even under strong adversarial settings [28, 48, 74, 84]. We discuss the impact of this pitfall extensively in §6.5 by using state-of-the-art attacks, and in §6.5.2, we design our own adaptive attack against FLDetector. Our impact analysis reveals that this is one of the most crucial components of FL evaluation.

Recommendations: To correctly evaluate the robustness of a defense, one should use 1) strong adaptive attacks that are tailored to the defense algorithm and maximally reduce its performances and 2) strong state-of-the-art attacks from existing literature.

Pitfall-6: Unfair metrics. Not capturing clients’ performances separately and only reporting global accuracy metrics does not give a good measure of per-client robustness.

Description: Data heterogeneity across clients in practical FL systems results in varying performance across clients. This phenomenon is also known as *representation disparity* [37]. Global accuracy does not give us any idea of the individual performances of clients. In addition to this, real-world datasets are class imbalanced [17] as opposed to synthetic datasets such as MNIST and CIFAR10. The commonly reported overall accuracy metric lacks information on per-class performance.

Prevalence and implications: Despite heterogeneity being a well-known issue, only 4% of the surveyed works incorporate personalized evaluations (Figure 3f). Our per-client analysis of TrMean and FedRecover demonstrates that per-client performances vary a lot, highlighting the need to account for these variations in real-world FL systems (§6.6). We also highlight the importance of reporting per-class accuracies, as we show that the overall accuracy is a misleading metric (§6.6.3).

Recommendations: Future evaluations should include personalized evaluations along with global evaluations to capture the variation in clients’ performances. In addition to this per-client evaluation, future works should also report per-class performance, especially when using class-imbalanced datasets.

5 Experimental setup

In this section, we provide the details of our experimental setup.

5.1 FEMNIST [17, 24]

FEMNIST (Figure 5) is a character recognition classification task with 3,400 clients, 62 classes (52 for upper and lower case letters and 10 for digits), and 671,585 grayscale images. Each client has data of their own handwritten digits or




Fig. 5. FEMNIST histogram from leaf.cmu.edu

letters. We use 300 randomly selected clients with their original data in a cross-silo fashion, as FedRecover uses the cross-silo setting in its implementation. We use the CNN used by [20] and use the Xavier weight initialization.

Hyperparameters for re-eval: For FEMNIST, we run over 200 epochs with 300 clients. In the attack setting, 60 clients are malicious. The results in Figure 6 use $T_w = 10$, and $T_c = 10$. The FL algorithm used here is FedAVG with a local learning rate of 0.05 and a global learning rate of 1. We keep the batch size to 32. The number of local epochs is kept at 1. For Figure 12a, we consider the possibility of benign clients being misclassified as malicious or malicious clients being misclassified as benign, so we vary the false negative and false positive rates between 0.1 and 0.5.

5.2 CIFAR10 [45]

CIFAR10 is a 10-class classification task with 60,000 total RGB images, each of size 32×32 . We divide all the data among 100 clients using either Dirichlet [78] or FCJ [28] distributions, which are the two most popular synthetic strategies to generate the FL dataset. We use a Resnet20 model with the CIFAR dataset.

Hyperparameters for re-eval: We run over 100 epochs with 100 clients. In the attack setting, 20 clients are malicious. The FL algorithm used here is FedAVG, with a local learning rate of 0.01 and a global learning rate of 1. We keep the batch size to 16. The number of local epochs is kept at 2. The results in Figure 6 use $T_w = 10$, and $T_c = 5$ and the fang distribution. Contrary to the rest of the datasets used, we use $T_c = 5$ because CIFAR10 was a much more challenging learning task.

5.3 MNIST [47]

MNIST is a 10-class digit image classification dataset, which contains 70,000 grayscale images of size 28×28 . We consider 100 FL clients and divide all data using Dirichlet or FCJ distributions. We use the same CNN as the FEMNIST dataset.

Hyperparameters for re-eval: For MNIST, we run over 2000 epochs with 100 clients, a learning rate of 0.03, and a batch size of 32. In the attack setting, 20 clients are malicious. We set $T_w = 20$, and $T_c = 10$. The FL algorithm used here is FedSGD. The results reported in Figure 6 use the FCJ distribution.

5.4 Fashion-MNIST [96]

Fashion-MNIST is a 10-class image classification dataset with grayscale images of clothing of size 28×28 . It contains 70,000 total images. We consider 100 FL clients and divide all 70,000 images using Dirichlet or FCJ distributions. For CIFAR10, MNIST, and FashionMNIST, we divide each client’s data in train/test/validation splits in the ratio of 10 : 1 : 1. We combine clients’ validation data and use it for validation and hyperparameter tuning and report accuracy on test data. We use the same CNN as the FEMNIST dataset.

Hyperparameters for re-eval: We run over 2000 epochs with 100 clients, a learning rate of 3×10^{-3} ¹, and a batch size of 32. In the attack setting, 20 clients are malicious. We set $T_w = 20$, and $T_c = 10$. The FL algorithm used here is FedSGD. The results reported in Figure 6 use the FCJ distribution.

5.5 StackOverflow [7]

StackOverflow is a language-modeling dataset that is used for tag prediction and next-word prediction. It consists of 342,477 users who are used as clients, and the training data consists of 135,818,730 examples. We use an RNN with a 96-dimensional embedding and a 10000-word vocabulary. The complete network consists of an input layer followed by an embedding layer, an LSTM layer, and two dense layers. We use the cross-device setting to obtain the baseline in [78] by using the fedjax [79] framework. The FL algorithm is *FedAdam*, and the training consists of 1500 rounds with 50 clients chosen every round with one local epoch. We keep the batch size to 16, the client optimizer as SGD with a learning rate of 10^{-3} , and Adam as the server optimizer with a learning rate of 10^{-2} .

6 Impact Analysis of Pitfalls

In this section, we analyze the impact of each identified pitfall. We test representative defenses across diverse setups by following the recommendations outlined in §4 and scrutinizing their implications. This systematic exploration is intended to help researchers make informed decisions about the robustness of FL defenses.

6.1 Intrinsically Robust Datasets

First, we evaluate how the selection of datasets impacts the robustness of our three representative FL defenses: TrMean, FedRecover, and FLDetector.

6.1.1 TrMean is only robust with MNIST. To assess TrMean’s sensitivity to different datasets, we employ FedAvg and FedSGD on MNIST, FashionMNIST, CIFAR10, and FEMNIST, both without attacks and under Stat-Opt [28]. Results in Figure 4a highlight MNIST’s intrinsic robustness. Despite a high (20%) amount of malicious clients, MNIST-trained FL model accuracy drops less than 1% in FedAvg, while other datasets experience more substantial declines, peaking at 50% for CIFAR10.

The variation in performance can be attributed to task complexity. Baseline accuracies in Figure 4 reflect this complexity, with MNIST having the highest no-attack accuracy and CIFAR10 the lowest. From this set of experiments, we can conclude that *TrMean is highly robust using MNIST-based evaluations, but not with other datasets as the evaluations of TrMean using other three datasets, FashionMNIST, FEMNIST, and CIFAR10 show.*

6.1.2 FedRecover works better with simple datasets. We evaluate FedRecover on FashionMNIST, FEMNIST, and CIFAR10, in addition to MNIST, since MNIST is heavily evaluated in [20]. We test FedRecover under *recovery-from-benign* and

¹We could not achieve the same accuracy reported in [20] using their reported 3×10^{-4} learning rate, hence we use 3×10^{-3} .





Fig. 6. Performance of FedRecover on four datasets, both with and without trim attack in the original training.

Table 2. Impact of data-level perturbations on FLDetector.


dataset	perturbation	FedSGD		FedAvg	
		FPR	FNR	FPR	FNR
MNIST	Noisy-features	0	0	0	0
	Noisy-label	0	0	0	0
Fashion	Noisy-features	0	0	0	0
	Noisy-label	0	0	0	0
FEMNIST	Noisy-features	0	1	0	1
	Noisy-label	0	1	0	1

recovery-from-attack scenarios. In the former, no attack occurs during original training, while in the latter, Stat-Opt attack is applied during original training but not during recovery (we discuss attack during recovery in §6.5.3). Consistent results with [20] on MNIST and FashionMNIST validate our implementation.

Figure 6 shows that even without attacks, FedRecover does not fully recover for complex datasets like FEMNIST and CIFAR10. In recovery from an attack for MNIST (Fashion-MNIST), FedRecover achieves 91% (72.7%) recovery accuracy from a post-attack accuracy of 80% (64%), where the original training accuracy is 92% (75.2%). For FEMNIST and CIFAR10, differences between the baseline accuracy and recovery accuracies are 11% and 19%, respectively, indicating an incomplete recovery in the attack setting. This performance variation stems from dataset complexity, with estimation errors higher for complex tasks such as FEMNIST. The impact of *periodic correction* and *warmup* phase on estimation error and recovery performance is discussed in §6.3, §6.4, and §6.5.



(a) Histogram of number of samples per client.



(b) Histogram of number of classes per client.

Fig. 7. Comparison of Sample and Class Distribution in FL Datasets: Histograms illustrating (a) the number of samples and (b) the number of classes per client, generated using FCJ and Dir distributions. From left to right, the non-i.i.d. degree of generated datasets increases, reflecting the impact of higher FCJ (or Dir) parameters in generating more (or less) non-i.i.d. datasets. We note that all FCJ client datasets are almost the same size, while Dir client datasets are widely varying. Similarly, with FCJ, all clients have all the classes, while with Dir, the number of classes varies widely.

6.1.3 FLDetector’s performance varies with task complexity. We assess FLDetector’s performance across varying task complexities using MNIST, FashionMNIST, and FEMNIST. To enhance task complexity, we introduce minor perturbations in features and labels, as we already know the results of these datasets in the unperturbed setting [106]. Table 2 shows that MNIST and FashionMNIST datasets remain robust, with FLDetector achieving perfect detection, i.e., zero FNR (False Negative Rate) and FPR (False Positive Rate). However, for FEMNIST, FLDetector shows an FPR of 0 and FNR of 1 across all conditions, indicating that it fails to detect any attacks, classifying all malicious clients as benign and allowing them to remain in the training process.

Performance variations across datasets stem from task complexity, extensively discussed in §6.1.1 and §6.1.2. MNIST and FashionMNIST, less affected by the Stat-Opt attack (Figure 4), have closely clustered benign updates, making it difficult for malicious updates to be stealthy. This proximity aids FLDetector in distinguishing between malicious

and benign updates. Conversely, FEMNIST, a naturally distributed and more heterogeneous dataset than MNIST and FashionMNIST (§6.2.1), results in client updates being further apart. This increased distance enables a malicious update to blend in seamlessly, evading detection by FLDetector.

6.2 Homogeneous Data Distribution

In this section, we first show that the Dir (Dirichlet) distribution is more *real-world* than the FCJ distribution through statistical analyses and visualization. Subsequently, we show the effect of these distributions and their varying levels of heterogeneity on the robustness of TrMean and FedRecover. The original work on FLDetector [106] already analyzed different levels of heterogeneity, so we skip this here.

6.2.1 Statistical analyses of FCJ and Dir distributions. We consider a classification task with a total of C classes; we generate client datasets using Dir and FCJ for 100 clients with varying degrees of non-i.i.d. We provide analyses for CIFAR10, but it applies to other datasets. We then plot the following three statistics of the datasets:

(Stat-1): We plot the number of samples per user, which is motivated by the visualizations of real-world FL datasets in Leaf (Figure 5). Figure 7a shows the results for three degrees of non-i.i.d. For Dir we use $\alpha \in \{0.1, 0.3, 0.5\}$ and for FCJ we use bias $b \in \{0.9, 0.5, 0.1\}$. We note that Dir produces client datasets with heterogeneous sizes; the histograms are similar to real-world datasets in the Leaf repository (Figure 5). However, FCJ makes client datasets with almost equal size; note that the FCJ histograms are always concentrated around 500 (the total number of samples in the CIFAR10 dataset / total number of clients).


(Stat-2): In Figure 7b, we show the number of classes each client has when we use Dir or FCJ. We observe that for FCJ, all clients have all the classes except when the bias is 0.9. In the real-world datasets, all clients generally do not have all the classes [59]. Similar to these real-world datasets, the clients have a widely varying number of classes in Dir distribution.

(Stat-3): For each client, we compute a C -dimensional vector where the i^{th} dimension represents the number of samples from class i ; here, we use CIFAR10 with 100 clients; hence we get 100 10-dimensional vectors. Then, we plotted T-SNE projections of these 100 vectors; we plotted them for both Dir and FCJ distributed client datasets. Figures 8b and 8c show the projections for Dir and FCJ, respectively. For reference, we also show in Figure 8a how the T-SNE projections look like for (1) 100 clients with perfectly i.i.d. datasets and (2) the real-world FEMNIST dataset.


For FCJ distribution, we note that, for bias values greater than 0.2 (Figure 8b center and right), the client datasets form local clusters, i.e., within these clusters, the clients have highly i.i.d. datasets. This is expected because FCJ forms C groups of clients where i^{th} group gets the bias fraction of data from i^{th} class and $\text{bias}/(C - 1)$ fraction of data from other classes. This data is then randomly assigned to clients within the i^{th} group, which makes these clients' datasets i.i.d. On the other hand, globally, these clusters form potentially non-i.i.d. structures. However, for a bias of 0.1, we observe almost i.i.d. datasets as expected; Figure 8a-left shows a perfectly i.i.d. dataset. To summarize, *although both are globally non-i.i.d., we have shown that the FCJ distribution is locally-i.i.d., while Dir is locally non-i.i.d.* Next, we study their impact.

6.2.2 TrMean is more robust with lower heterogeneity. We demonstrate the impact of FCJ and Dir on TrMean's robustness for FashionMNIST² in Figure 9a. In the no-attack setting, FCJ (DIR) shows little difference, with accuracy going from 89% (88%) at 0.1 (0.5) bias to 84% (86%) at 0.9 (0.1) bias. However, the distinction emerges in the attack setting where


²We observe similar trends for other datasets, but for brevity, we only include FashionMNIST here.



(a) (Left) i.i.d. client datasets, (right) real-world FEMNIST datasets.



(b) FCJ distributed client datasets.



(c) Dir distributed client datasets.

Fig. 8. T-SNE projections of class frequency vectors of client datasets generated using FCJ (b) and Dir (c) distributions. From left to right, non-i.i.d. degree increases.

FCJ is more robust. For the Stat-Opt attack, FCJ (DIR) accuracy decreases from 84% (80%) at 0.1 (0.5) bias to 75% (59%) at 0.9 (0.1) bias. A similar trend is observed for the Dyn-Opt attack. The performance change is rooted in the nature of the distributions (§6.2.1). FCJ is locally i.i.d. but globally non-i.i.d., while Dir is non-i.i.d. both locally and globally.




Fig. 9. The effect of varying heterogeneity levels for FCJ and Dir distributions on the FashionMNIST dataset.

Due to Dir’s greater heterogeneity, benign updates are more dispersed, making it easier for a malicious update to hide. Consequently, TrMean struggles to detect malicious updates, leading to lower global model performance.

6.2.3 FedRecover performs better under lower heterogeneity. We show the impact of FCJ and Dir distributions on FedRecover for the FashionMNIST dataset in Figure 9b. We observe that the performance of FedRecover lowers when we increase the level of heterogeneity and that FCJ performs slightly better than Dir. Specifically, the difference between no-attack accuracies for FCJ at bias 0.3 and 0.7 is 7% and 12.6%, respectively, whereas, for Dir, that difference is slightly higher; 8.2% and 13.3% corresponding to non-i.i.d. parameters of 0.7 and 0.3 respectively³.

The reason for FedRecover’s performance reduction under high heterogeneity is that a higher heterogeneity means that client updates are far apart and drift away from the global model. Since *the update estimation in FedRecover uses knowledge of the current and past global models, the HVP calculation step (§3.2.1) incurs significant estimation errors if the local model drifts away from the global model.*

6.3 Slow-converging Algorithms

Here, we will study the impact of the choice of the two widely used FL algorithms, FedSGD and FedAvg, on the robustness of TrMean and FedRecover. We combine our analysis of algorithms with the choice of attacks for FLDetector in §6.5.2; therefore, we do not discuss it in this section.

6.3.1 Fast algorithms make TrMean more robust. In Figure 4, we show the accuracy of FedSGD and FedAvg with TrMean in benign and Stat-Opt attack scenarios. For FedSGD, we align our implementation with recent defenses [20, 28]. In contrast, FedAvg is optimized for greater accuracy with only 5% to 20% of FedSGD’s communication (the number of rounds).

³For FCJ, a higher value of the i.i.d. parameter means higher heterogeneity, but it is the opposite case for Dir.

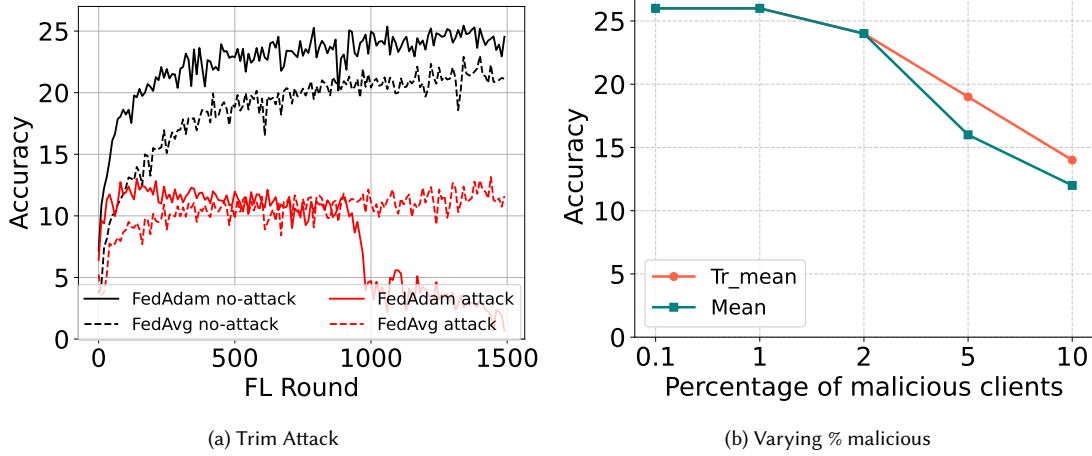


Fig. 10. Stackoverflow in practical settings.

Under both benign and adversarial conditions, FedAvg greatly surpasses FedSGD in performance, convergence, and communication for all datasets, as shown in Figure 4. We can conclude here that in adversarial situations, *FedAvg proves more resilient to untargeted poisoning due to its rapid convergence*, giving adversaries minimal time for poisoning.

6.3.2 Fast algorithms lower estimation errors in FedRecover. We examine FedRecover’s performance under different algorithms, FedSGD and FedAvg. The original study [20] applies the Stat-Opt attack to MNIST with FedSGD over 2000 rounds, reducing accuracy from 96% to 81%. By employing FedAvg with carefully chosen hyperparameters (Appendix 5), as shown in Figure 4a, we achieve 98% accuracy in 50 rounds, with the Stat-Opt attack lowering it to only 96%. This allows for perfect recovery, as depicted in Figure 11b, where both MNIST and FashionMNIST (no-attack accuracy: 87%) show perfect recovery under the fast-converging FedAvg. This perfect recovery holds even with variations in the periodic correction periodicity T_c . FedRecover excels in this scenario because *FedAvg’s rapid convergence (§2.1) provides a high starting accuracy during FedRecover’s warmup phase*. Throughout periodic correction and abnormality fixing phases, *client updates computed over multiple local epochs (in contrast to FedSGD’s single local epoch) lead to lower estimation errors, ensuring perfect recovery*.

6.4 Limited FL Settings

To show the effect of *scale* on FL defenses, we first assess the impact of a scale-constrained threat model on TrMean. We then show why FedRecover and FLDetector are incompatible with cross-device settings and discuss the practicality of storage, computation, and communication for FedRecover (or, any mechanism requiring stored historical information).

6.4.1 Mean AGR is robust on a large scale. In this section, we showcase the robustness of the mean AGR, typically not robust in cross-silo settings [103], within a large-scale, cross-device environment. *Given the growing popularity of language models [15, 90], although underrepresented in our defense survey (Figure 3a), we leverage the StackOverflow [7] dataset for large-scale evaluation* (setup details in Appendix A). The no-attack baselines in Figure 10a are obtained following the settings in [78]. We can see that with the standard amount of 20% malicious clients, the Stat-Opt attack significantly impacts StackOverflow in cross-device.

The StackOverflow FL setting has about 300,000 clients and 20% amounts to 60,000 malicious clients. Accessing and modifying such a vast number of devices is considered impractical, factoring in operational and financial costs [84]. Figure 10b shows a decline in attack performance as the percentage of malicious clients decreases. Notably, **with less than 5% malicious clients, the Mean AGR remains unaffected** by the Stat-Opt attack. *It is important to highlight the difference between our findings and [84] where the attack impact with Mean AGR is > 0% even for 0.01% of malicious clients for FEMNIST, CIFAR10, and Purchase, while in our Stack Overflow case, the attack impact is 0% for < 2% malicious clients.* With our results, **we strongly emphasize using datasets especially designed for FL and evaluating defenses under constraints imposed by scaling up the system in addition to small-scale experimentation.** We do not dismiss the significance of small-scale experiments. Cross-silo FL is indeed widely used. Instead, we emphasize the critical need for evaluating FL defenses within scaled-up settings.

6.4.2 Resource overheads for FedRecover. Here, we first explain the reasons behind our conviction that FedRecover and FLDetector are incompatible in the cross-device setting. Consequently, we do not evaluate these two defenses under the cross-device setting. Nevertheless, in this section, we comment on some of the *practical* aspects of FedRecover, such as computation, communication, and storage costs, to assess the feasibility of scaling up such systems.

Compatibility of FedRecover and FLDetector with the cross-device FL: FedRecover’s reliance on historical information from clients’ past updates is hindered in cross-device, where clients participate in a few rounds, limiting available historical updates. Similarly, FLDetector faces challenges in the cross-device setting due to a lack of consistent historical information. Therefore, *we find FedRecover and FLDetector incompatible with the cross-device setting.*

Communication-accuracy tradeoff for FedRecover: In Figure 11a, we show the tradeoff between the recovery accuracy and communication of FedRecover. The recovery accuracy increases as it relies more on exact updates (locally computed updates instead of estimated ones) from the clients, where the *minimum number of exact updates* is $T_w + \frac{T_{total} - T_w - T_f}{T_c} + T_f$. Interestingly, using the same amount of exact updates (effectively the same amount of communication as the server would communicate with clients for the same number of rounds) gives us more accuracy by training from scratch, i.e., not using FedRecover but restarting training with benign clients only. For instance, with only 20% exact updates and $T_w = 20$, FedRecover achieves $\approx 76\%$ accuracy while training from scratch with an equivalent 40 rounds achieves $\approx 80\%$.

In Figure 11b, we can see that with a small percentage of exact updates, FedRecover completely recovers for simpler datasets like MNIST and FashionMNIST with FedAvg and is unaffected by the variation in periodicity T_c . With the same amount of exact updates, we can achieve the same results without FedRecover as well. This is because a fast algorithm gives less time for the adversary to attack and a higher starting point for FedRecover to recover, leading to lower estimation errors (§6.3.2).

Based on these observations, we conclude that with fast converging algorithms and proper settings, we might not need to use FedRecover, especially when it comes with an additional computational and storage cost. In our experiments, for the slow baseline that uses FedSGD over 1000 epochs for MNIST, we require $\approx 200GB$ of storage for saving the client model updates every round. By extension, this applies to any defense that requires knowledge of past updates. This cost would significantly increase with the number of clients and by using larger models for more complex tasks.

6.5 Naive Attacks

This component is critical in our research, given the prevalence of simple attacks highlighted in §4. We demonstrate TrMean’s vulnerability to powerful model poisoning attacks such as Stat-Opt and Dyn-Opt. We also test our adaptive




Fig. 11. Communication-accuracy tradeoff for FedRecover. Despite the presence of four lines, their overlap is discernible as we can achieve baseline accuracy with exact updates.




Fig. 12. Performance of FedRecover, represented by the difference between no-attack and post-recovery accuracies, under non-zero FPRs and FNRs caused by adaptive attacks.

attack on FLDetector, revealing high rates of *imperfect client detection* and showcasing its impact on FedRecover, which relies on FLDetector to identify malicious clients before recovery.

6.5.1 Choice of attacks greatly impact $TrMean$'s performance. We conduct Stat-Opt [28] and Dyn-Opt [83] attacks on FashionMNIST, varying the heterogeneity for both FCJ and Dir distributions. The results are depicted in Figure 9a. Generally, Dyn-Opt is stronger than Stat-Opt, particularly noticeable at higher heterogeneity levels. For instance, at Dir bias level 0.1, the accuracy drops to 51% for Dyn-Opt compared to 59% for Stat-Opt. Dyn-Opt's strength lies in finding optimal perturbations tailored for the dataset at every FL round, making it more potent compared to the static nature of perturbation in Stat-Opt, as discussed in §2.2.1.

6.5.2 Overcoming FLDetector with our adaptive attack. We introduce a novel attack to assess FLDetector’s resilience against *adaptive attacks*. Malicious clients craft updates using our attack and evade detection by FLDetector, resulting in non-zero FNRs and FPRs.

Our attack adds a carefully crafted perturbation vector to client model updates, so they are close enough to the estimated model updates, thereby bypassing FLDetector’s detection. Here we describe our attack formulation in detail.

Attack Formulation: FLDetector computes the estimated update (§3.2.1) for client k as:

$$\hat{\nabla}_t^k = \nabla_{t-1}^k + \hat{H}^t \cdot (\theta_t - \theta_{t-1}) \quad (1)$$

In this attack, we introduce a perturbation vector, \mathcal{P} , that modifies the updates sent from malicious clients. A malicious client computes its update so that its final update is computed as the sum of its previous update, the *HVP* or *Hessian Vector Product* ($\hat{H}^t \cdot (\theta_t - \theta_{t-1})$), and the perturbation vector. This can be written as:

$$\hat{\nabla}_t^k = \nabla_{t-1}^k + \hat{H}^t \cdot (\theta_t - \theta_{t-1}) + \mathcal{P} \quad (2)$$

The server estimates an update by adding the *HVP* to the last round’s exact update and compares it with the actual update in that round (§3.2.1). The malicious update, therefore, deviates from the estimated update by \mathcal{P} . We proceed to detail the calculation of this perturbation vector \mathcal{P} . To calculate \mathcal{P} , we first calculate a *good distance range*, \mathcal{R} , that is a safe perturbation distance for the perturbation vector by taking the norm between the old and new client updates. The good distance range, \mathcal{R}_t^k , for client k , at round t is given by:

$$\mathcal{R}_t^k = \|\hat{\nabla}_t^k - \nabla_t^k\| \quad (3)$$

Here, $\hat{\nabla}_t^k$ is the estimated update for client k at round t , and ∇_t^k is the actual update for client k at round t . The deviation, \mathcal{D} , for the perturbation vector is calculated by following deviation strategies in [83]. It can either be *unit vector*, *sign*, or *std*. Finally, \mathcal{P} is computed by taking the average of all the *good distance ranges*, \mathcal{R} , and directing it in the direction of the deviation, \mathcal{D} :

$$\mathcal{P} = \frac{\mathcal{D}}{\|\mathcal{D}\|} \cdot \frac{1}{N} \sum_{k=1}^N \mathcal{R}_t^k \quad (4)$$

The impact of our adaptive attack on the performance of FLDetector is shown in Table 3 and across all cases (different combinations of dataset, FL algorithm, and percentage of compromised clients), except one, we find that the FNR is non-zero, which means that malicious clients have not been detected, and they continue to be part of the training process. Since the attack is designed to craft malicious updates that are statistically close to the benign ones, it leads to a non-zero FPR as well, which means that many benign clients are falsely detected as malicious and are removed from the training process. Since *we achieve a 0 FPR and a 1 FNR for FEMNIST, this means that the attack never gets detected, no benign or malicious client is removed, and all malicious clients seem benign to FLDetector on the server*. We discuss the impact of non-zero FNR and FPR on further training and recovery in §6.5.3.

6.5.3 Imperfect detection leads to lower recovery performance. Our adaptive attack results in imperfect client detection (§6.5.2), allowing escaped clients into the recovery process. The escaped clients are denoted by the FNR, while benign clients, incorrectly classified as malicious and subsequently removed from training, are represented by the FPR.

Figure 12 shows that non-zero FNRs and FPRs challenge FedRecover’s ability to reach the no-attack accuracy of 82%. For instance, at $FNR = FPR = 0.3$ for $T_w = 10$, the recovery accuracy is only 67%. This is because malicious

Table 3. Impact of our adaptive attack on FLDetector. Here $%m$ represents the percentage of malicious clients.

% m	Baseline	MNIST		Fashion		CIFAR10		FEMNIST	
		FPR	FNR	FPR	FNR	FPR	FNR	FPR	FNR
5	FedSGD	0.4	1	0.52	0.4	0.21	0	0	1
	FedAvg	0.02	1	0	1	0.13	1	0	1
10	FedSGD	0.39	1	0.56	0.6	0.37	0	0	1
	FedAvg	0.02	1	0.02	1	0.23	0.67	0	1
15	FedSGD	0.48	0.8	0.54	0.3	0.35	0.35	0	1
	FedAvg	0.02	1	0.05	1	0.18	0.67	0	1
20	FedSGD	0.45	1	0.62	0.3	0.35	0.33	0	1
	FedAvg	0.01	1	0.06	1	0.17	0.67	0	1

updates in the recovery process cause higher estimation errors and deviate the recovery model from the benign one. A marginal improvement is observed with more warmup rounds, achieving a post-recovery accuracy of 70% for $T_w = 20$, as increased warmup rounds provide a higher starting point for the model, resulting in lower overall estimation errors.

6.6 Unfair Metrics

A key feature of FL is heterogeneity (§6.2.1), which makes testing and reporting individual client’s performance essential. Unfortunately, almost all of the surveyed works only report the global model accuracy (§4). To show that this is unfair, we show how personalized performance differs from that of the global model using TrMean and FedRecover under benign and adversarial settings.

6.6.1 Different clients have different levels of robustness under TrMean. We train on FEMNIST for 200 rounds, achieving a global model accuracy of 82%. Figure 13a displays the accuracy trends, where *global w/o attack* represents the global model accuracy, and *per-client w/o attack* shows individual client test accuracies. *Per-client w/o attack* clusters around the global accuracy, indicating the model learns from the combined data. Due to heterogeneity, we see a lot of variation in the performance of individual clients. In the attack scenario (*global w/ attack* at 66% and *per-client w/ attack*), a similar trend persists, but most clients now fall below the global attack accuracy. *Based on our observations due to heterogeneity across clients, we strongly advocate using per-client metrics in future evaluations.*

6.6.2 FedRecover’s performance greatly varies on a per-client basis. Similar to the personalized evaluations for TrMean in §6.6.1, we apply personalized evaluations to FedRecover for the FEMNIST dataset. After performing FedRecover on the global model with 66% accuracy (§6.6.1), we achieve a post-recovery accuracy of 74% (*post-recovery* in Figure 13b). This is illustrated on a per-client basis with the lines *fnr0.1fpr0.1*, *fnr0.3fpr0.3*, *fnr0.5fpr0.5* in Figure 13b. The post-recovery accuracies for various FNRs and FPRs indicate that most client models fall below the global post-recovery accuracy. At an FNR and FPR of 0.5, all clients lie below both the post-recovery accuracy and the post-attack accuracy, signifying a failure in achieving recovery. This aligns with our findings in §6.5.3 that *the presence of malicious clients during recovery leads to higher estimation errors, as mirrored in our per-client evaluations. This consistency underscores the necessity for personalized evaluations, particularly in non-i.i.d. datasets.*

6.6.3 Impact of class-imbalance on per-class accuracies. We examine the impact of class imbalance in CIFAR10 on the performance of TrMean under the no-attack and Stat-Opt attack scenarios. To set up our baseline, we train an Alexnet




Fig. 13. Personalized evaluations, i.e., per-client accuracy for FEMNIST with TrMean and FedRecover. Note that the accuracy does not increase monotonically with the client number, rather we plot it in an ascending order here since order does not matter when we want to show variation in per-client accuracy. This should not be perceived as a relationship between accuracy and client ID/number. We are simply showing the range of individual client accuracies and have plotted them in a monotonic fashion for visual comparison.




Fig. 14. Overall and per-class accuracies for balanced CIFAR10 before and after Stat-Opt attack.

model with 100 clients over 100 epochs. We use the standard CIFAR10 dataset that has 50000 training samples and 10000 test samples. Figure 14 shows that we achieve an overall accuracy of 70.82% without attack and 42.66% with attack in the balanced setting. It also shows the respective per-class accuracies. We define the *overall accuracy* as the total number of correct samples out of the total number of samples in the test dataset and the *mean per-class accuracy* as the mean of all the per-class accuracies in the test dataset. In this case, since the test dataset is perfectly balanced, i.e., each class has the same number of test samples, the overall accuracy and the mean accuracy are the same.

Next, we create an imbalance in the CIFAR10 dataset by removing 90% of the samples of all classes except class zero. This produces 5000 samples of class zero and 500 samples each for other classes in the training set. Figure 15 shows that in the imbalanced setting, the overall accuracy in no-attack is 68%, which is close to its balanced accuracy of 70.82%. However, this accuracy is biased towards class 0, which has the highest per-class accuracy of 86.2%(75.8% in the balanced scenario) since it has ten times more samples than the rest. The mean per-class accuracy is 52% as it ignores the class imbalance. This indicates that *the mean per-class accuracy is a better metric than overall accuracy in this case, as it shows that the model loses utility with reduced samples.*

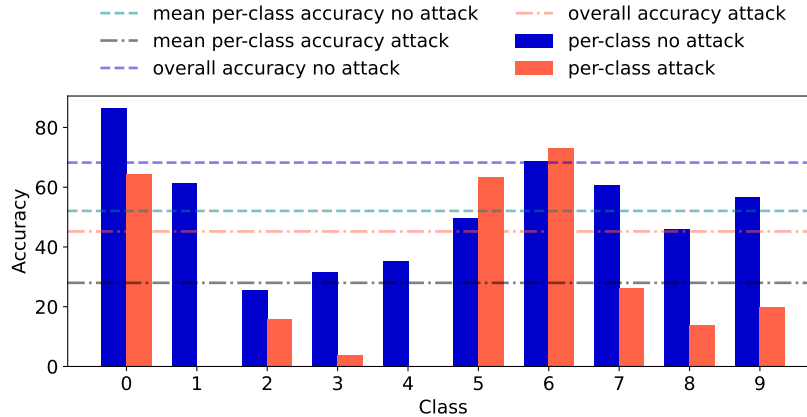


Fig. 15. Overall, per-class, and mean per-class accuracies for imbalanced CIFAR10 before and after Stat-Opt attack. Class 0 has the most samples.

The imbalanced Stat-Opt scenario yields an overall accuracy of 48%, surpassing the balanced counterpart at 42.66%. Figure 15 shows that the overall accuracy is significantly influenced by class 0, dropping from 86.2% to 64.3% under attack, compared to a drop from 75.8% to 17.3% in the balanced scenario. The mean per-class accuracy for the imbalanced attack is 27%. This highlights that *dominant classes are less affected by attacks*, influencing the overall accuracy to reflect their performance. *While overall accuracy is crucial, we stress the importance of reporting per-class and mean per-class accuracies, especially in real-world datasets with class imbalances, as neglecting this aspect can lead to misjudgments about a system’s robustness.*

7 Related Work

Previous systemizations: In prior research, various taxonomies in adversarial ML have been developed, extending beyond FL [10, 13, 32, 38, 40, 80] and their main focus is on the attacks. [80] presents a taxonomy specifically for FL defenses, categorizing them based on their occurrence at the client, server, or communication channel. Our contribution extends beyond this by providing a more comprehensive and multidimensional systemization of the existing defenses in the literature. Additionally, we leverage our systemization to identify representative defenses for in-depth pitfall analyses.

Shejwalkar et al. [84] have performed the systemization of FL attacks, highlighting misconceptions about the robustness of FL systems that may arise from overlooking practical considerations about the threat model in deployment scenarios. Their conclusion emphasizes the high robustness of FL with simple and cost-effective defenses in practical threat models. Our work complements this perspective by focusing on defenses. We systematically categorize and re-evaluate representative defenses, uncovering common pitfalls, and providing actionable recommendations to address each.

Pitfalls and guidelines: Arp et al. [6] have extensively studied the pitfalls associated with ML in computer security, identifying and providing recommendations for several challenges in this domain. It is crucial to emphasize the distinctions between our work and theirs. While we may identify similar pitfalls, such as Inappropriate Baseline (§6.3), Inappropriate Performance Measures (§6.6), and Lab-only Evaluation (§6.4), the underlying reasons for these pitfalls and

their impact are specific to FL. For example, a pitfall, “inappropriate performance measures,” addresses inappropriate performance measures like overall accuracy due to its limited capture of information about false positives and false negatives. In our context (§6.6), we discuss inappropriate performance measures, particularly in relation to fairness, and advocate for the use of personalized metrics, especially when dealing with non-i.i.d. data.

In the space of adversarial ML, Carlini et al. [22] identified numerous pitfalls associated with evaluating defenses against adversarial learning. While there are some commonalities with our work, such as considerations regarding the choice of attacks and testing against adaptive attacks, the primary distinction lies in the fact that our evaluation is specifically conducted in FL. Certain evaluation components are exclusive to FL, such as dealing with client heterogeneity and personalized evaluations.

Improvements over [43]: We take inspiration from [43], which initiated the identification of FL pitfalls and used FedRecover as a case study to showcase the impact of some of these pitfalls. Before building upon their work, we set out to perform a systemization of FL defenses to understand the FL defenses space. This helps us to select representative defenses for pitfall analysis. We have expanded [43]’s analysis by exploring additional pitfalls and studying their impact. Additionally, we broadened the evaluation spectrum by incorporating additional representative defenses like FLDetector and TrMean. Our study extends to a large-scale evaluation of FL with the language modality using StackOverflow [7].

8 Limitations and Future Work

Scope of our evaluation: Performing a thorough exploration and evaluation of all the defenses in Table 4, along with additional ones, goes beyond the scope of a single paper. Performing a thorough exploration of one defense from each dimension, type, and attribute is also beyond the scope of one single paper as it requires running the full set of experiments for every defense and then comparing them.

Expandability of our systemization: The defenses chosen for evaluation serve as representatives of the broader defense literature; however, it is acknowledged that alternative evaluations may differ based on different sets of defenses. Also, a newly designed defense based on a novel technique might not fit exactly along the attributes in our systemization. Therefore, our systemization remains adaptable (as detailed in §1), allowing for the incorporation of additional type, attributes, and even dimensions in the future.

Incorporating other modalities: As highlighted in §1, our aim is to expand the scope of our evaluation across diverse dimensions, encompassing data distributions, data modalities, and the nature of the ML task. While we have addressed image classification and NLP, other modalities, such as multimodal time-series tasks [107], and the emerging paradigm of *vision-language models* [76], remain unexplored. The evaluation of these modalities, along with more complex vision-language models, under traditional threat models could prove intriguing. Such exploration might lead to the development of new attacks and defenses, potentially uncovering novel pitfalls in the process.

9 Conclusion

Our study contributes a crucial systemization of FL defenses, offering valuable insights for researchers and practitioners in the selection, combination, and design of defenses. Additionally, we address an often-overlooked aspect in the FL poisoning defense literature—the *experimental setups employed* to assess defense efficacy. After reviewing a variety of defense works, we identify prevalent questionable experimental trends. Through case studies featuring well-known defenses such as TrMean, FLDetector, and FedRecover, we illustrate how the choice of experimental setups can significantly impact robustness claims.

References

- [1] [n. d.]. How Apple personalizes Siri without hoovering up your data — technologyreview.com. <https://www.technologyreview.com/2019/12/11/131629/apple-ai-personalizes-siri-federated-learning/>.
- [2] 2017. Federated Learning: Collaborative Machine Learning without Centralized Training Data. <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>.
- [3] 2019. The Stack Overflow Data. <https://www.kaggle.com/datasets/stackoverflow/stackoverflow>.
- [4] Zeyuan Allen-Zhu, Faeze Ebrahimi, Jerry Li, and Dan Alistarh. 2020. Byzantine-resilient non-convex stochastic gradient descent. *arXiv preprint arXiv:2012.14368* (2020).
- [5] Sebastien Andreina, Giorgia Azzurra Marson, Helen Möllering, and Ghassan Karame. 2021. Baffle: Backdoor detection via feedback-based federated learning. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 852–863.
- [6] Daniel Arp, Erwin Quiring, Feargus Pendlebury, Alexander Warnecke, Fabio Pierazzi, Christian Wressnegger, Lorenzo Cavallaro, and Konrad Rieck. 2022. Dos and don'ts of machine learning in computer security. In *31st USENIX Security Symposium (USENIX Security 22)*. 3971–3988.
- [7] The TensorFlow Federated Authors. 2019. TensorFlow Federated Stack Overflow dataset.
- [8] Sana Awan, Bo Luo, and Fengjun Li. 2021. Contra: Defending against poisoning attacks in federated learning. In *Computer Security—ESORICS 2021: 26th European Symposium on Research in Computer Security, Darmstadt, Germany, October 4–8, 2021, Proceedings, Part I 26*. Springer, 455–475.
- [9] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. 2020. How to backdoor federated learning. In *AISTATS*.
- [10] Marco Barreno, Blaine Nelson, and Anthony D. Joseph. 2010. The security of machine learning. *Machine Learning* (2010).
- [11] Moran Baruch, Baruch Gilad, and Yoav Goldberg. 2019. A Little Is Enough: Circumventing Defenses For Distributed Learning. In *NeurIPS*.
- [12] Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. 2019. Analyzing federated learning through an adversarial lens. In *ICML*.
- [13] Battista Biggio and Fabio Roli. 2018. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognition* (2018).
- [14] Peva Blanchard, Rachid Guerraoui, Julien Stainer, et al. 2017. Machine learning with adversaries: Byzantine tolerant gradient descent. In *NeurIPS*.
- [15] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [16] Lukas Burkhalter, Hidde Lycklama, Alexander Viand, Nicolas Küchler, and Anwar Hithnawi. 2021. Rofl: Attestable robustness for secure federated learning. *arXiv preprint arXiv:2107.03311* (2021).
- [17] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. 2018. Leaf: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097* (2018).
- [18] Xiaoyu Cao, Minghong Fang, Jia Liu, and Neil Zhenqiang Gong. 2021. FLTrust: Byzantine-robust Federated Learning via Trust Bootstrapping. In *NDSS*.
- [19] Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. 2021. Provably Secure Federated Learning against Malicious Clients. In *AAAI*.
- [20] X. Cao, J. Jia, Z. Zhang, and N. Gong. 2023. FedRecover: Recovering from Poisoning Attacks in Federated Learning using Historical Information. In *2023 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 326–343. <https://doi.org/10.1109/SP46215.2023.00019>
- [21] Xiaoyu Cao, Zaixi Zhang, Jinyuan Jia, and Neil Zhenqiang Gong. 2022. FLCert: Provably Secure Federated Learning Against Poisoning Attacks. *IEEE Transactions on Information Forensics and Security* 17 (2022), 3691–3705.
- [22] Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian Goodfellow, Aleksander Madry, and Alexey Kurakin. 2019. On evaluating adversarial robustness. *arXiv preprint arXiv:1902.06705* (2019).
- [23] Hongyan Chang, Virat Shejwalkar, Reza Shokri, and Amir Houmansadr. 2019. Cronus: Robust and Heterogeneous Collaborative Learning with Black-Box Knowledge Transfer. *arXiv:1912.11279* (2019).
- [24] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. 2017. EMNIST: Extending MNIST to handwritten letters. In *IJCNN*.
- [25] Jean Ogier Du Terrail, Samy-Safwan Ayed, Edwige Cyffers, Felix Grimberg, Chaoyang He, Regis Loeb, Paul Mangold, Tanguy Marchand, Othmane Marfoq, Erum Mushtaq, et al. 2022. FLamby: Datasets and Benchmarks for Cross-Silo Federated Learning in Realistic Healthcare Settings. In *NeurIPS, Datasets and Benchmarks Track*.
- [26] Sannara Ek, François Portet, Philippe Lalanda, and German Vega. 2020. Evaluation of federated learning aggregation algorithms: application to human activity recognition. In *Adjunct proceedings of the 2020 ACM international joint conference on pervasive and ubiquitous computing and proceedings of the 2020 ACM international symposium on wearable computers*. 638–643.
- [27] El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Louis Alexandre Rouault. 2021. Distributed momentum for byzantine-resilient stochastic gradient descent. In *9th International Conference on Learning Representations (ICLR)*.
- [28] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. 2020. Local Model Poisoning Attacks to Byzantine-Robust Federated Learning. In *USENIX*.
- [29] Ines Feki, Sourour Ammar, Yousri Kessentini, and Khan Muhammad. 2021. Federated learning for COVID-19 screening from Chest X-ray images. *Applied Soft Computing* 106 (2021), 107330.
- [30] Shuhao Fu, Chulin Xie, Bo Li, and Qifeng Chen. 2019. Attack-resistant federated learning with residual-based reweighting. *arXiv:1912.11464* (2019).
- [31] Clement Fung, Chris JM Yoon, and Ivan Beschastnikh. 2020. The limitations of federated learning in sybil settings. In *RAID*.

- [32] Micah Goldblum, Dimitris Tsipras, Chulin Xie, et al. 2020. Dataset Security for Machine Learning: Data Poisoning, Backdoor Attacks, and Defenses. *arXiv:2012.10544* (2020).
- [33] Eduard Gorbunov, Samuel Horváth, Peter Richtárik, and Gauthier Gidel. 2022. Variance Reduction is an Antidote to Byzantines: Better Rates, Weaker Assumptions and Communication Compression as a Cherry on the Top. *arXiv preprint arXiv:2206.00529* (2022).
- [34] Hanxi Guo, Hao Wang, Tao Song, Yang Hua, Zhangcheng Lv, Xiulang Jin, Zhengui Xue, Ruhui Ma, and Haibing Guan. 2021. Siren: Byzantine-robust federated learning via proactive alarming. In *Proceedings of the ACM Symposium on Cloud Computing*. 47–60.
- [35] Filip Hanzely, Slavomir Hanzely, Samuel Horváth, and Peter Richtárik. 2020. Lower bounds and optimal algorithms for personalized federated learning. *Advances in Neural Information Processing Systems* 33 (2020), 2304–2315.
- [36] Filip Hanzely and Peter Richtárik. 2020. Federated learning of a mixture of global and local models. *arXiv preprint arXiv:2002.05516* (2020).
- [37] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. 2018. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*. PMLR, 1929–1938.
- [38] Ling Huang, Anthony D Joseph, Blaine Nelson, Benjamin IP Rubinstein, and J Doug Tygar. 2011. Adversarial machine learning. In *AISeC*.
- [39] Najeeb Moharram Jebreel, Josep Domingo-Ferrer, David Sánchez, and Alberto Blanco-Justicia. 2022. Defending against the label-flipping attack in federated learning. *arXiv preprint arXiv:2207.01982* (2022).
- [40] Malhar S Jere, Tyler Farnan, and Farinaz Koushanfar. 2020. A taxonomy of attacks on federated learning. *IEEE Security & Privacy* (2020).
- [41] Peter Kairouz, H Brendan McMahan, Brendan Avent, et al. 2019. Advances and open problems in federated learning. *arXiv:1912.04977* (2019).
- [42] Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. 2020. Byzantine-robust learning on heterogeneous datasets via bucketing. *arXiv preprint arXiv:2006.09365* (2020).
- [43] Momin Ahmad Khan, Virat Shejwalkar, Amir Houmansadr, and Fatima M Anwar. 2023. On the Pitfalls of Security Evaluation of Robust Federated Learning. In *2023 IEEE Security and Privacy Workshops (SPW)*. IEEE, 57–68.
- [44] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. 2016. Federated learning: Strategies for improving communication efficiency. *NIPS Workshop on Private Multi-Party ML* (2016).
- [45] Alex Krizhevsky. 2009. *Learning multiple layers of features from tiny images*. Technical Report. University of Toronto.
- [46] Hanchao Ku, Willy Susilo, Yudi Zhang, Wenfen Liu, and Mingwu Zhang. 2022. Privacy-preserving federated learning in medical diagnosis with homomorphic re-encryption. *Computer Standards & Interfaces* 80 (2022), 103583.
- [47] Yann LeCun and Corinna Cortes. 1998. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/> (1998).
- [48] Liping Li, Wei Xu, Tianyi Chen, Georgios B Giannakis, and Qing Ling. 2019. RSA: Byzantine-robust stochastic aggregation methods for distributed learning from heterogeneous datasets. In *AAAI*.
- [49] Qinyin Li, Bingsheng He, and Dawn Song. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10713–10722.
- [50] Suyi Li, Yong Cheng, Yang Liu, Wei Wang, and Tianjian Chen. 2019. Abnormal client behavior detection in federated learning. *arXiv preprint arXiv:1910.09933* (2019).
- [51] Suyi Li, Yong Cheng, Wei Wang, Yang Liu, and Tianjian Chen. 2020. Learning to detect malicious clients for robust federated learning. *arXiv preprint arXiv:2002.00211* (2020).
- [52] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. 2021. Ditto: Fair and robust federated learning through personalization. In *ICML*.
- [53] Zhuohang Li, Luyang Liu, Jiaxin Zhang, and Jian Liu. 2021. Byzantine-robust federated learning through spatial-temporal analysis of local model updates. In *2021 IEEE 27th International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, 372–379.
- [54] Feng Lin, Weiyu Li, and Qing Ling. 2021. Stochastic alternating direction method of multipliers for byzantine-robust distributed learning. *arXiv preprint arXiv:2106.06891* (2021).
- [55] Yi Liu, Ruihui Zhao, Jiawen Kang, Abdulsalam Yassine, Dusit Niyato, and Jialiang Peng. 2021. Towards communication-efficient and attack-resistant federated edge learning for industrial internet of things. *ACM Transactions on Internet Technology (TOIT)* 22, 3 (2021), 1–22.
- [56] Xindi Ma, Qi Jiang, Mohammad Shojafar, Mamoun Alazab, Sachin Kumar, and Saru Kumari. 2022. DisBeZant: secure and robust federated learning against byzantine attack in IoT-enabled MTS. *IEEE Transactions on Intelligent Transportation Systems* (2022).
- [57] Saeed Mahloujifar, Mohammad Mahmoodi, and Ameer Mohammed. 2019. Universal multi-party poisoning attacks. In *ICML*.
- [58] Ranwa Al Mallah, David Lopez, Godwin Badu Marfo, and Bilal Farooq. 2021. Untargeted poisoning attack detection in federated learning via behavior attestation. *arXiv preprint arXiv:2101.10904* (2021).
- [59] H Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agueru y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*.
- [60] El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. 2018. The Hidden Vulnerability of Distributed Learning in Byzantium. In *ICML*.
- [61] Thomas Minka. 2000. Estimating a Dirichlet distribution.
- [62] Hamid Mozaffari, Virat Shejwalkar, and Amir Houmansadr. 2021. Frl: Federated rank learning. *arXiv preprint arXiv:2110.04350* (2021).
- [63] Luis Muñoz-González, Battista Biggio, Ambra Demontis, Andrea Paudice, Vasin Wonggrassamee, Emil C Lupu, and Fabio Roli. 2017. Towards poisoning of deep learning algorithms with back-gradient optimization. In *AISeC*.
- [64] Lokesh Nagalapatti and Ramasuri Narayanam. 2021. Game of gradients: Mitigating irrelevant clients in federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 9046–9054.

- [65] Mohammad Naseri, Jamie Hayes, and Emiliano De Cristofaro. 2020. Local and central differential privacy for robustness and privacy in federated learning. *arXiv preprint arXiv:2009.03561* (2020).
- [66] Dinh C Nguyen, Quoc-Viet Pham, Pubudu N Pathirana, Ming Ding, Aruna Seneviratne, Zihuai Lin, Octavia Dobre, and Won-Joo Hwang. 2022. Federated learning for smart healthcare: A survey. *ACM Computing Surveys (CSUR)* 55, 3 (2022), 1–37.
- [67] Jorge Nocedal. 1980. Updating quasi-Newton matrices with limited storage. *Mathematics of computation* 35, 151 (1980), 773–782.
- [68] Xiaomin Ouyang, Zhiyuan Xie, Jiayu Zhou, Jianwei Huang, and Guoliang Xing. 2021. Clusterfl: a similarity-aware federated learning system for human activity recognition. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*. 54–66.
- [69] Mustafa Safa Ozdayi, Murat Kantarcioglu, and Yulia R Gel. 2021. Defending against backdoors in federated learning with robust learning rate. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 9268–9276.
- [70] Xudong Pan, Mi Zhang, Duocai Wu, Qifan Xiao, Shouling Ji, and Min Yang. 2020. Justinian’s gaavernor: Robust distributed learning with gradient aggregation agent. In *Proceedings of the 29th USENIX Conference on Security Symposium*. 1641–1658.
- [71] Jungwuk Park, Dong-Jun Han, Minseok Choi, and Jaekyun Moon. 2021. Sageflow: Robust federated learning against both stragglers and adversaries. *Advances in neural information processing systems* 34 (2021), 840–851.
- [72] Matthias Paulik, Matt Seigel, Henry Mason, et al. 2021. Federated Evaluation and Tuning for On-Device Personalization: System Design & Applications. *arXiv:2102.08503* (2021).
- [73] Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. 2019. Robust aggregation for federated learning. *arXiv:1912.13445* (2019).
- [74] Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. 2020. Learning from History for Byzantine Robust Optimization. *arXiv e-prints* (2020), arXiv–2012.
- [75] Adnan Qayyum, Kashif Ahmad, Muhammad Ahtazaz Ahsan, Ala Al-Fuqaha, and Junaid Qadir. 2022. Collaborative federated learning for healthcare: Multi-modal covid-19 diagnosis at the edge. *IEEE Open Journal of the Computer Society* 3 (2022), 172–184.
- [76] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [77] Priyesh Ranjan, Ashish Gupta, Federico Coro, and Sajal K Das. 2022. Securing Federated Learning against Overwhelming Collusive Attackers. In *GLOBECOM 2022-2022 IEEE Global Communications Conference*. IEEE, 1448–1453.
- [78] Sashank J Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and Hugh Brendan McMahan. 2020. Adaptive Federated Optimization. In *ICLR*.
- [79] Jae Hun Ro, Ananda Theertha Suresh, and Ke Wu. 2021. FedJAX: Federated learning simulation with JAX. *arXiv preprint arXiv:2108.02117* (2021).
- [80] Nuria Rodríguez-Barroso, Daniel Jiménez-López, M Victoria Luzón, Francisco Herrera, and Eugenio Martínez-Cámara. 2023. Survey on federated learning threats: Concepts, taxonomy on attacks and defences, experimental study and challenges. *Information Fusion* 90 (2023), 148–173.
- [81] Amrita Roy Chowdhury, Chuan Guo, Somesh Jha, and Laurens van der Maaten. 2022. EIFFeL: Ensuring Integrity for Federated Learning. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*. 2535–2549.
- [82] Felix Sattler, Klaus-Robert Müller, Thomas Wiegand, and Wojciech Samek. 2020. On the byzantine robustness of clustered federated learning. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 8861–8865.
- [83] Virat Shejwalkar and Amir Houmansadr. 2021. Manipulating the Byzantine: Optimizing Model Poisoning Attacks and Defenses for Federated Learning. In *NDSS*.
- [84] V. Shejwalkar, A. Houmansadr, P. Kairouz, and D. Ramage. 2022. Back to the Drawing Board: A Critical Evaluation of Poisoning Attacks on Production Federated Learning. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 1117–1134. <https://doi.org/10.1109/SP46214.2022.00065>
- [85] Shiqi Shen, Shruti Tople, and Prateek Saxena. 2016. AUROR: Defending against poisoning attacks in collaborative deep learning systems. *2016 Annual Computer Security Applications Conference* (2016).
- [86] Konstantin Sozinov, Vladimir Vlassov, and Sarunas Girdzijauskas. 2018. Human activity recognition using federated learning. In *2018 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Ubiquitous Computing & Communications, Big Data & Cloud Computing, Social Computing & Networking, Sustainable Computing & Communications (ISPA/IUCC/BDCLOUD/SocialCom/SustainCom)*. IEEE, 1103–1111.
- [87] Jingwei Sun, Ang Li, Lin Duan, Samiul Alam, Xuliang Deng, Xin Guo, Haiming Wang, Maria Gorlatova, Mi Zhang, Hai Li, et al. 2022. FedSEA: A Semi-Asynchronous Federated Learning Framework for Extremely Heterogeneous Devices. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*. 106–119.
- [88] Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. 2019. Can you really backdoor federated learning? *NeurIPS FL Workshop* (2019).
- [89] Robert Tibshirani, Guenther Walther, and Trevor Hastie. 2001. Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63, 2 (2001), 411–423.
- [90] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [91] Hongyi Wang, Kartik Sreenivasan, Shashank Rajput, Harit Vishwakarma, Saurabh Agarwal, Jy-yong Sohn, Kangwook Lee, and Dimitris Papailiopoulos. 2020. Attack of the tails: Yes, you really can backdoor federated learning. In *NeurIPS*.

- [92] Ning Wang, Yang Xiao, Yimin Chen, Yang Hu, Wenjing Lou, and Y Thomas Hou. 2022. Flare: defending federated learning against model poisoning attacks via latent space representations. In *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*. 946–958.
- [93] webankcredit 2019. Utilization of FATE in Risk Management of Credit in Small and Micro Enterprises. <https://www.fedai.org/cases/utilization-of-fate-in-risk-management-of-credit-in-small-and-micro-enterprises/>.
- [94] Chen Wu, Xian Yang, Sencun Zhu, and Prasenjit Mitra. 2020. Mitigating backdoor attacks in federated learning. *arXiv:2011.01767* (2020).
- [95] Zhaoxian Wu, Qing Ling, Tianyi Chen, and Georgios B Giannakis. 2020. Federated variance-reduced stochastic gradient descent with robustness to byzantine attacks. *IEEE Transactions on Signal Processing* 68 (2020), 4583–4596.
- [96] 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).
- [97] Chulin Xie, Minghao Chen, Pin-Yu Chen, and Bo Li. 2021. CRFL: Certifiably Robust Federated Learning against Backdoor Attacks. In *ICML*.
- [98] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. 2018. Generalized byzantine-tolerant sgd. *arXiv:1802.10116* (2018).
- [99] Cong Xie, Sanmi Koyejo, and Indranil Gupta. 2019. Fall of empires: Breaking Byzantine-tolerant SGD by inner product manipulation. *arXiv:1903.03936* (2019).
- [100] Yueqi Xie, Weizhong Zhang, Renjie Pi, Fangzhao Wu, Qifeng Chen, Xing Xie, and Sunghun Kim. 2022. Robust federated learning against both data heterogeneity and poisoning attack via aggregation optimization. *arXiv preprint* (2022).
- [101] Chang Xu, Yu Jia, Liehuang Zhu, Chuan Zhang, Guoxie Jin, and Kashif Sharif. 2022. TDFL: Truth Discovery Based Byzantine Robust Federated Learning. *IEEE Transactions on Parallel and Distributed Systems* 33, 12 (2022), 4835–4848.
- [102] Jian Xu, Shao-Lun Huang, Linqi Song, and Tian Lan. 2021. Signguard: Byzantine-robust federated learning through collaborative malicious gradient filtering. *arXiv preprint arXiv:2109.05872* (2021).
- [103] Dong Yin, Yudong Chen, Kannan Ramchandran, and Peter Bartlett. 2018. Byzantine-robust distributed learning: Towards optimal statistical rates. In *ICML*.
- [104] Syed Zawad, Ahsan Ali, Pin-Yu Chen, Ali Anwar, Yi Zhou, Nathalie Baracaldo, Yuan Tian, and Feng Yan. 2021. Curse or redemption? how data heterogeneity affects the robustness of federated learning. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 10807–10814.
- [105] Kaiyuan Zhang, Guan hong Tao, Qiuling Xu, Siyuan Cheng, Shengwei An, Yingqi Liu, Shiwei Feng, Guangyu Shen, Pin-Yu Chen, Shiqing Ma, et al. 2022. FLIP: A Provable Defense Framework for Backdoor Mitigation in Federated Learning. *arXiv preprint arXiv:2210.12873* (2022).
- [106] Zaixi Zhang, Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. 2022. FLDetector: Defending federated learning against model poisoning attacks via detecting malicious clients. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2545–2555.
- [107] Yuchen Zhao, Payam Barnaghi, and Hamed Haddadi. 2022. Multimodal federated learning on iot data. In *2022 IEEE/ACM Seventh International Conference on Internet-of-Things Design and Implementation (IoTDI)*. IEEE, 43–54.

A Appendix

A.1 Our methodology to classify 50 defenses

Table 4 shows our survey of the 50 defense works and our classification methodology along the six dimensions: Datasets, Attacks, Data Distribution, FL-Algorithm, FL Type, and Evaluation Type.

A.2 Full names for datasets and attacks

Table 6. Abbreviations and full-forms of attacks in Table 4.

LF	Label Flip
IPM	Inner Product Manipulation
SF	Sign Flip
BF	Bit Flip
LIE	Little is Enough
AN	Additive Noise
BD	Backdoor [9]
RGA	Random Gaussian Attack

Table 4. Classification of 50 defense works across 6 dimensions of evaluation setup.

Work	Datasets	Attacks	Data Distribution	FL Algorithm	FL Type	Evaluation
FLDetector [106]	FA,FE,C10	Stat-Opt	Fang, Natural	FedSGD	CS	Global
FedRecover [20]	M,FA,P,H	Stat-Opt	Fang	FedSGD	CS	Global
Machine Learning with Adversaries [14]	M, spambase	RGA	IID	FedSGD	CS	Global
FLTrust [18]	M,CHM,C10,H	Krum, Stat-Opt, LF	Fang	FedAvg	CS	Global
Byzantine-Robust Distributed Learning [103]	M	RGA	IID	FedSGD	CS	Global
Provably Secure Federated Learning against [19]	M	Not applicable	Fang	FedAvg		Global
Learning to Detect Malicious Clients for [51]	M, FE, S140	SF, AN, BD	Natural, McMahan	FedAvg	CD	Global
Robust Federated Learning [100]	M,FE,C10/100,N20	LF, LIE, Fang	Dirichlet, Natural	FedAvg	CD	Global
The Hidden Vulnerability of Distributed [60]	M, C10	Specific attack	IID	FedSGD	CS/CD	Global
Sageflow [71]	M, FA, C10	SF, LF	McMahan	FedAvg	CS	Global
Mitigating Irrelevant Clients in FL [64]	M	LF	McMahan	FedAvg	CS	Global
Cronus [23]	M, C10, P, Svhn	{LF, LIE}	IID	FedAvg	CS	Global
Can You Really Backdoor Federated Learning? [88]	FE	BD	Natural	FedAvg	CS/CD	Global
Generalized Byzantine-tolerant SGD [98]	M, C10	BF, LF, LIE	IID	FedSGD	CS	Global
The Limitations of Federated Learning [31]	M, VGG, KDD, A	LF, BD	Each class to a client	Both	CS	Global
Auror [85]	M	Targeted-LF	IID	FedSGD	CS	Global
Robust Aggregation for Federated Learning [73]	FE,S140,S	Specific attacks, RGA	Natural	FedAvg	CS/CD	Global
CRFL [97]	M, FE	BD	IID	FedAvg	CS	Global
FLIP [105]	M, FA, C10	BD	Dirichlet	FedAvg	CS	Global
RoFL [16]	FE, C10	BD	Natural, Dirichlet	FedAvg	CD	Global
Securing FL against Overwhelming [77]	M, FA	LF, BD	Dirichlet	FedAvg	CS	Global
Defending against the Label-flipping Attack [39]	M, C10	LF	IID, Dirichlet	FedAvg	CS	Global
FRL [62]	M, FE, C10	Fang, Dyn-Opt	Dirichlet, Natural	FedAvg	CD	Global
CONTRA [8]	M, C10, Loan	LF, BD	Dirichlet	FedAvg	CS	Global
EIFFel [81]	M, FA, FE,C10	LIE, RGA, SF, Dyn-Opt	IID, Natural	FedAvg	CS/CD	Global
Local and central DP for robustness [65]	E, C10, s140, Reddit	BD	McMahan	FedAvg	CS/CD	Global
Signguard [102]	M, FA, C10, AGnews	LIE, RGA, SF, Dyn-Opt	IID	FedSGD	CS	Global
DisBezant [56]	{M, FA, C10}	RGA	{Fang}	FedAvg		Global
Learning from History for Byzantine [74]	M, C10	BF, LF, LIE	Exponential	FedSGD	CS	Global
Byzantine-robust learning on heterogeneous [42]	M	BF, LF, LIE, IPM, Mimic	{McMahan}	FedSGD	CS	Global
Byzantine-Resilient Non-Convex Stochastic [4]	C10, C100	SF, LF, LIE, Delayed-grad	IID	FedSGD	CS	Global
Byzantine-robust Federated Learning [53]	M, FA, C10,Spambase	LF, IPM, LIE, Uniform, arbitrary	McMahan	FedAvg	CS/CD	Global
Stochastic alternating direction method of [54]	M, Coverttype	RGA, SF, LF	IID	FedSGD	CS	Global
Variance reduction is an antidote to byzantines [33]	LIBSVM	LF, BF, LIE, IPM	{IID}	FedSGD	CS	Global
On the byzantine robustness of clustered FL [82]	M, FA, C10	RGA, LF, Uniform noise	IID	FedSGD	CS	Global
RSA [48]	M	SF	IID	FedSGD	CS	Global
Federated variance-reduced [95]	ijcnn1, covtype	RGA, SF, Zero-grad	IID	FedSGD	CS	Global
Abnormal client behavior detection in [50]	FE	SF, RGA, Grad ascent	Natural	FedAvg	CD	Global
Distributed Momentum for Byzantine [27]	M, FA, C10/100	LIE, IPM	IID	FedSGD	CS	Global
Attack-resistant FL with residual-based [30]	M, C10, Amazon, Loan	LF, BD	Dirichlet, Natural	FedAvg	CS	Global
Towards communication-efficient [55]	M	LF	IID	FedSGD	CS	Global
Justinian's GAAvernor [70]	M, C10, Yelp, Health	RGA	IID	FedSGD	CS	Global
Untargeted poisoning attack detection [58]	M, C10, MTL Trajet	BD	IID	FedSGD		Global
TDFl [101]	M, FA, C10	LF, RGA, Krum, Stat-Opt, BD	McMahan	FedAvg	CD	Global
Siren [34]	FA, C10	SF, LF, bhagoji	Fang	FedAvg	CS	Global
FLARE [92]	FA, C10, Kather	Krum, Stat-Opt	IID	FedAvg	CS	Global
Analyzing Federated Learning Through [12]	FA, UCI Census	Specific attack	IID	FedAvg		Global
BaFFLe [5]	C10, FE	BD	Dirichlet	FedAvg		Global
Defending against backdoors in FL [69]	FA, FE	BD	Both	FedAvg		Global
Ditto [52]	FA, FE, CelebA	LF, RGA, BD	Natural, McMahan	FedAvg	CS	Personalized

Table 5. Abbreviations and full-forms of datasets in Table 4.

M	MNIST
FA	FashionMNIST
FE	FEMNIST
C10	CIFAR10
C100	CIFAR100
P	Purchase
H	HAR
S140	Sentiment140
SVHN	Street-view House Numbers
VGG	VGGFace
KDD	KDDCup
N20	News 20
A	Amazon
S	Shakespeare