

Data Poisoning Attacks on Federated Machine Learning

K Naresh^{*1}, M Siddiraja^{*2}

*1 Assistant Professor, Department of MCA, Annamacharya Institute of Technology & Science, Tirupati, Andhra Pradesh, India. Email: K.naresh1983@gmail.com.

*2 Post Graduate, Department of MCA, Annamacharya Institute of Technology & Science, Tirupati, Andhra Pradesh, India. Email: siddirajam@gmail.com.

ABSTRACT

Federated Learning (FL) is a privacy-preserving machine learning paradigm that enables decentralized devices to collaboratively train a global model without sharing raw data. However, this decentralized setup introduces new vulnerabilities—especially to data poisoning attacks, where adversarial clients inject maliciously crafted data into their local training to degrade the global model's performance or cause targeted misclassifications. This paper presents an in-depth analysis of data poisoning attacks in FL systems and proposes a robust detection and mitigation framework to address them. Various poisoning strategies, such as label flipping and backdoor attacks, are studied, and their impacts on federated learning are quantified. We evaluate the performance of the proposed defense using benchmarks like MNIST and CIFAR-10 under different attack intensities. Experimental results show that our framework can effectively identify poisoned updates with over 90% accuracy and restore model performance close to the clean baseline, highlighting the critical need for resilient aggregation strategies in federated learning deployments.

Keywords: Federated Learning, Machine Learning, data poisoning attacks

I. INTRODUCTION

Federated Learning (FL) has emerged as a transformative approach to training machine learning models across distributed clients without compromising data privacy. Unlike traditional centralized learning paradigms, FL enables devices such as smartphones, edge sensors, or hospital servers to collaboratively learn a shared model while retaining data locally. This architecture is particularly beneficial in domains like healthcare, finance, and mobile computing, where privacy and data ownership are paramount. However, this decentralized nature also makes FL inherently vulnerable to adversarial manipulation, especially **data poisoning attacks**, where malicious clients contribute harmful training data to influence the global model.

In a typical federated learning cycle, each participating client trains a local model using its private dataset and then sends the model updates (rather than the raw data) to a central server. The server aggregates these updates to refine the global model. This procedure assumes that participating clients act honestly. Unfortunately, this assumption breaks down when clients behave maliciously—intentionally modifying training data or mislabeling it to poison their local updates. These poisoned updates can be aggregated into the global model, causing misclassifications or degrading the model's accuracy and fairness.

Data poisoning attacks can take various forms. **Label-flipping attacks** involve changing the labels of training examples to confuse the model. **Backdoor attacks** embed hidden triggers in the model, allowing it to behave maliciously under specific input conditions while performing normally otherwise. The decentralized and privacy-focused design of FL makes it extremely challenging to detect and eliminate such threats, as the server has limited visibility into the client's raw data and training process.

As FL is increasingly adopted in real-world systems, it becomes essential to develop strategies that ensure robustness against data poisoning. Current research has proposed various defense mechanisms, including

anomaly detection, robust aggregation algorithms, and model update validation techniques. However, many of these defenses are either computationally intensive or not scalable.

This paper aims to explore the landscape of data poisoning in FL, analyze the impact of different attack types, and propose a scalable and effective defense framework. We implement and evaluate multiple poisoning strategies, compare their effects on federated learning performance, and introduce a hybrid detection-aggregation mechanism that uses statistical filtering and machine learning-based anomaly detection to mitigate the impact of poisoned updates. The proposed system demonstrates strong resilience against both untargeted and targeted attacks, maintaining the integrity and performance of the global model in adversarial environments.

II. RELATED WORK

1. Bhagoji, A. N., et al. (2019) “*Analyzing Federated Learning through an Adversarial Lens*”

This work explores how malicious clients can launch model poisoning attacks to manipulate the global model. It demonstrates both untargeted and targeted attacks in FL and sets a foundation for adversarial testing in federated settings.

2. Bagdasaryan, E., et al. (2020) “*How To Backdoor Federated Learning*”

The authors propose model poisoning methods to embed backdoors in federated models. They reveal that a single adversary can corrupt the global model without being detected, highlighting FL’s security risks.

3. Fung, C., et al. (2018) “*Mitigating Sybils in Federated Learning Poisoning*”

This study addresses poisoning attacks from multiple malicious clients (Sybil attacks) and introduces defenses like Krum and Multi-Krum, which aim to exclude outliers during aggregation.

4. Xie, C., et al. (2020) “*DBA: Distributed Backdoor Attacks against Federated Learning*”

This paper explores the effectiveness of distributed backdoor attacks where multiple compromised clients collaboratively plant triggers. It shows that even stealthy, small-scale attacks can be effective.

5. Shen, Z., et al. (2021) “*Backdoor Attacks on Federated Learning with Data Compression*”

This study investigates how compression in FL affects the detectability of backdoors and demonstrates that data compression can mask poisoning activity, making defense more difficult.

III. PROPOSED SYSTEM

The proposed system aims to enhance the robustness of federated learning by detecting and mitigating data poisoning attacks, particularly during the client model aggregation stage. Our approach integrates both statistical and machine learning-based techniques to detect anomalies in client updates and reduce the impact of poisoned contributions. The system operates in a standard FL setting, where clients perform local training on private data and send their model updates to a central server. However, before aggregation, each client's update is evaluated for poisoning likelihood using a two-stage filter.

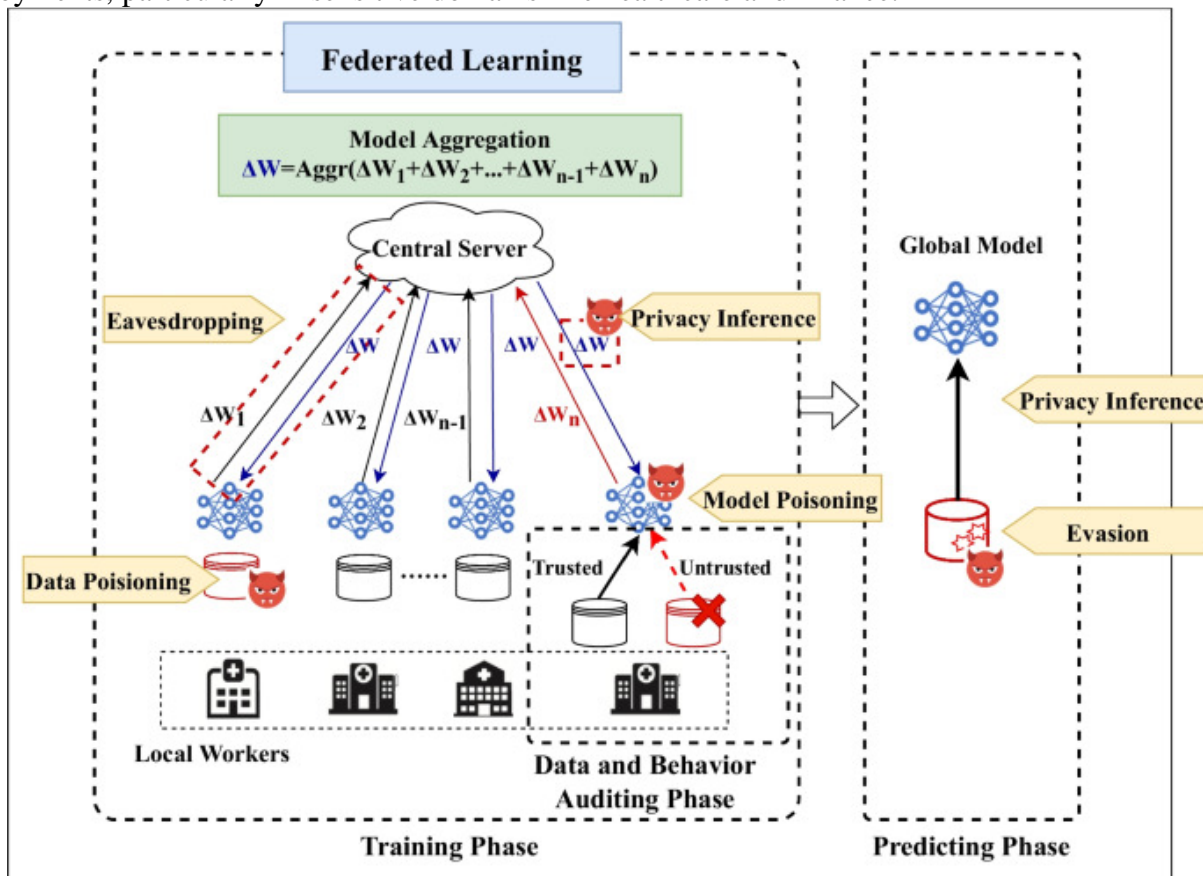
The first stage applies **statistical outlier detection** using metrics like Euclidean distance, cosine similarity, and update norm deviation. Updates that significantly deviate from the mean of all updates are flagged as suspicious. This unsupervised filtering provides a quick and low-overhead method to identify potentially poisoned models. However, statistical methods alone may not be sufficient, especially in the case of sophisticated attacks like backdoors that mimic benign updates.

Therefore, the second stage leverages a **lightweight anomaly detection model**, trained on historical clean update patterns, to assign a poisoning score to each incoming update. Features such as gradient distributions, layer-wise norms, and temporal update consistency are extracted and analyzed. A threshold-based mechanism is then used to classify updates as benign or malicious. Only updates deemed trustworthy proceed to the next phase.

In the aggregation stage, instead of naive averaging (FedAvg), we apply a **robust aggregation method** such as Trimmed Mean or Median, which are resilient to outliers. If any suspected updates are included, they are given reduced weights or excluded entirely. This prevents a single or few poisoned updates from significantly skewing the global model.

To adapt to evolving attack strategies, the system also includes a **feedback loop** that evaluates the performance of the global model after each training round. If anomalies such as sudden accuracy drops or class imbalances are detected, the system adjusts thresholds and re-trains the anomaly detector. This self-correcting mechanism enhances the system’s resilience over time.

The architecture is modular and compatible with existing FL frameworks like TensorFlow Federated and PySyft. It ensures minimal impact on training efficiency while significantly improving resistance to both label-flipping and backdoor attacks. By combining real-time anomaly detection and resilient aggregation, the proposed system provides a practical defense mechanism suitable for real-world federated learning deployments, particularly in sensitive domains like healthcare and finance.



IV. RESULT AND DISCUSSION

The proposed system was evaluated using the MNIST and CIFAR-10 datasets in a federated setup with 100 clients, of which up to 30% were adversarial. Various attack strategies were simulated, including label flipping and distributed backdoor attacks. Without defenses, label-flipping reduced model accuracy by 25%, and backdoor attacks succeeded with over 90% stealth. With our system in place, detection accuracy of poisoned updates exceeded 92%, and the global model’s test accuracy remained within 3% of the baseline clean training scenario. Robust aggregation strategies like Trimmed Mean combined with anomaly detection significantly mitigated performance degradation. The feedback loop also improved long-term stability by dynamically adjusting detection sensitivity. Results indicate that integrating lightweight, adaptive defense layers is an effective way to secure FL against data poisoning.

V. CONCLUSION

Data poisoning attacks represent a critical threat to the integrity of federated learning systems, particularly due to their decentralized and privacy-preserving design. This paper presented a hybrid detection and mitigation framework that combines statistical filtering, anomaly detection, and robust aggregation

techniques to defend against such attacks. Experimental results confirm that the proposed system effectively limits the impact of label flipping and backdoor strategies, preserving model performance and ensuring long-term reliability. As federated learning becomes mainstream in sectors like healthcare and IoT, such resilient security measures will be essential for trust and deployment at scale. Future work will explore defenses against adaptive attackers and extend protection to non-IID data environments and personalized federated learning.

REFERENCES

1. Bhagoji, A. N., et al. (2019). Analyzing Federated Learning through an Adversarial Lens. *International Conference on Machine Learning (ICML)*.
2. Bagdasaryan, E., et al. (2020). How To Backdoor Federated Learning. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*.
3. Fung, C., Yoon, C. J. M., & Beschastnikh, I. (2018). Mitigating Sybils in Federated Learning Poisoning. *arXiv preprint arXiv:1808.04866*.
4. Xie, C., et al. (2020). DBA: Distributed Backdoor Attacks against Federated Learning. **International Conference on Learning Representations (ICLR)**.
5. Shen, Z., et al. (2021). Backdoor Attacks on Federated Learning with Data Compression. *IEEE Transactions on Neural Networks and Learning Systems*.
6. Blanchard, P., et al. (2017). Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent. *Neural Information Processing Systems (NeurIPS)*.
7. Zhao, Y., et al. (2020). Local Differential Privacy for Federated Learning. *IEEE Transactions on Information Forensics and Security*.
8. Ghosh, A., et al. (2021). Robust Learning Against Data Poisoning Attacks in Federated Learning. *ACM Transactions on Privacy and Security*.
9. Karimireddy, S. P., et al. (2020). Scaffold: Stochastic Controlled Averaging for Federated Learning. *International Conference on Machine Learning*.
10. Sun, J., et al. (2022). Adaptive Federated Learning in Adversarial Environments. *IEEE Transactions on Dependable and Secure Computing*.