

STDLens: Model Hijacking-resilient Federated Learning for Object Detection

Ka-Ho Chow, Ling Liu, Wenqi Wei, Fatih Ilhan, Yanzhao Wu
Georgia Institute of Technology
Atlanta, GA, USA

khchow@gatech.edu, ling.liu@cc.gatech.edu, {wenqiwei, filhan, yanzhaowu}@gatech.edu

Abstract

Federated Learning (FL) has been gaining popularity as a collaborative learning framework to train deep learning-based object detection models over a distributed population of clients. Despite its advantages, FL is vulnerable to model hijacking. The attacker can control how the object detection system should misbehave by implanting Trojaned gradients using only a small number of compromised clients in the collaborative learning process. This paper introduces STDLens, a principled approach to safeguarding FL against such attacks. We first investigate existing mitigation mechanisms and analyze their failures caused by the inherent errors in spatial clustering analysis on gradients. Based on the insights, we introduce a three-tier forensic framework to identify and expel Trojaned gradients and reclaim the performance over the course of FL. We consider three types of adaptive attacks and demonstrate the robustness of STDLens against advanced adversaries. Extensive experiments show that STDLens can protect FL against different model hijacking attacks and outperform existing methods in identifying and removing Trojaned gradients with significantly higher precision and much lower false-positive rates. The source code is available at <https://github.com/git-disl/STDLens>.

1. Introduction

Federated Learning (FL) for object detection has attracted numerous applications [12], especially in healthcare, with strict privacy regulations protecting patient medical records [25, 27]. Instead of using a centralized data curator, FL can train a global object detection model with a distributed population of clients. Each client only shares its fixed-size gradient (updated model) parameters (e.g., 246.9 MB for YOLOv3 [15]) with the FL server for aggregation while keeping its private raw data local (e.g., terabytes of videos) [13]. Such a paradigm lowers the bar for knowledge sharing and eases the recruitment of contributors, but it becomes vulnerable to model hijacking [9, 21].

Model hijacking aims at interfering with the training pro-





Scenario	Object Detection Results	AP _{person}	AP _{car}
Benign		52.43	74.12
Class-Poison Victim: Person		38.22 (↓27%)	71.79 (↓3%)
BBox-Poison Victim: Person		18.06 (↓66%)	72.82 (↓2%)
Objn-Poison Victim: Person		37.61 (↓28%)	71.59 (↓3%)

Table 1. FL-trained detectors can be hijacked by perception poisoning to misdetect objects of designated classes (e.g., person) in three different ways (2nd to 4th rows) [3], while objects of non-victim classes (e.g., car) have negligible performance degradation.

cess of a machine learning model and causes it to misbehave at the deployment stage [17]. The FL-trained global model can be indirectly hijacked by a small number of compromised clients who share Trojaned gradients to gradually confuse the learning trajectory of FL [10]. Such gradients can be generated in black-box through data poisoning. The attacker only has to poison the local training data owned by a compromised client, such as changing the ground-truth label of certain objects by dirty-label poisoning [21] or implanting a backdoor to trigger the malfunctioning of a hijacked model [1, 2, 8, 20]. The FL software will take the poisoned local training data and unintentionally generate malicious gradients for sharing. With the growing number of real-world applications, recent work has begun to understand the hijacking of FL-based object detection systems [3]. As shown in Table 1, the adversary can designate how the hijacked model should misdetect objects of certain classes (e.g., person) while ensuring other objects can still be precisely detected (e.g., car). This fine-grained control opens the door for an adversary to design a stealthy attack configuration to hurt the FL system yet remains hardly noticeable by the owner [5], leading to catastrophic consequences such as car crashes.

This paper introduces STDLens, a three-tier defense methodology against model hijacking attacks, with the following original contributions. First, we analyze existing defenses based on spatial signature analysis and show their ineffectiveness in protecting FL. Second, we introduce a pipeline with per-client spatio-temporal signature analysis to identify Trojaned gradients, track their contributors, revoke their subscriptions, and reclaim the detection performance. Third, we present a density-based confidence inspection mechanism to manage the spatio-temporal uncertainty. It avoids purging benign clients contributing useful learning signals for the FL system. Finally, we extend perception poisoning with three types of adaptiveness for analyzing STDLens in countering advanced adversaries. Extensive experiments show that STDLens outperforms competitive defense methods by a large margin in defense precision with a low false-positive rate, reducing the accuracy drop under attack from 34.47% to 0.24%, maintaining the object detection accuracy on par with the performance under the benign scenario.

2. Background and Related Work

2.1. Perception Poisoning for Model Hijacking

Threat Model. Perception poisoning adopts the following assumptions: (i) the communication between a client and the trustworthy FL server is encrypted; (ii) the benign clients faithfully follow the FL workflow; and (iii) there is only a small $m\%$ of compromised clients out of a total of N clients. An adversary on a malicious client can only manipulate its local training data, and hence, model hijacking can be done without tampering with the FL software.

Poisoning Formulation. Let F be the global model being trained in FL, f_i be the local model of client \mathcal{P}_i , and (\mathbf{x}, \mathbf{y}) be the input sample and its ground truth in the training set \mathcal{D}_i of client \mathcal{P}_i . The process of perception poisoning can be formulated as a function ψ that maps (\mathbf{x}, \mathbf{y}) to $(\mathbf{x}, \mathbf{y}')$ such that the local model f_i will generate Trojaned gradients which maximize the chance of F being fooled to make erroneous predictions with high confidence, formally:

$$\psi : \psi(\mathbf{x}, \mathbf{y}) = (\mathbf{x}, \mathbf{y}') \text{ s.t. } f_i(\mathbf{x}) = \mathbf{y}', \mathbf{y}' \neq \mathbf{y}, \max[F(\mathbf{x}) \bowtie \mathbf{y}']$$

where \mathbf{y} is a list of triplets, one for each object $\mathbf{o} = (\text{CLASS}, \text{BBOX}, \text{OBJN})$ in the training image \mathbf{x} , CLASS is the class label, BBOX is the bounding box, and OBJN is the existence of object \mathbf{o} . $f_i(\mathbf{x})$ and $F(\mathbf{x})$ denote the detected objects by the local and global models, respectively. The operator \bowtie indicates the degree of matching between two sets of objects [4]. Based on the three perceptual components of an object, one can formulate three poisoning strategies to designate malicious behaviors on the hijacked model [3]: (i) Class-Poison changes the class label of objects of a source class to a maliciously chosen target class

without altering their bounding box locations and object existence. (ii) BBox-Poison attacks objects of certain source classes by randomly changing the size and the location of their bounding boxes without altering the class label and the object existence. (iii) Objn-Poison attacks objects of certain classes by changing their existence to non-existence. As shown in Table 1, the hijacked model under Class-Poison (2nd row) mislabels the person as a potted plant in the first example and as a bird in the second example with a different attack configuration. Under BBox-Poison (3rd row), the person object is detected with a correct class label but a misaligned bounding box. Under Objn-Poison (4th row), the person vanishes in both cases. Only the designated classes suffer from performance degradation (e.g., “person” in the 4th column), while the detection of objects of other irrelevant classes is not affected (e.g., “car” in the 5th column).

2.2. Existing Defenses and Their Limitations

While no defenses have been proposed to counter perception poisoning, several solutions are being introduced for image classifiers, which are either developed for FL [21, 23, 28] or can be extended [18, 22]. They can be broadly classified into two categories. The first category aims at repairing the gradients (model updates). Contributions from all participants will be used for aggregation, but they will be processed by parameter clipping and noise injection [11, 20, 24] to wash out possibly malicious elements. However, it is challenging to determine the correct amount of noise to be injected [14]. The Trojaned gradients can still be malicious if the noise is too small, while a larger noise can remove useful learning signals from benign clients and affect the convergence of FL. The second category attempts to inspect gradients shared by participants to identify and eliminate malicious clients. MNTD [26] trains a meta-classifier to differentiate Trojaned and non-Trojaned gradients. It is a supervised defense and assumes a known distribution of Trojans to create training data. Unsupervised defenses make fewer assumptions and are more attractive in practice. AUROR [18] and its followers [21] conduct spatial clustering on gradient contributions and label the contributors to the smaller cluster to be malicious. Spectral signature [22] conducts outlier removal via SVD by purging gradient contributions deviating most significantly from a singular vector. While they neither rely on finding the right amount of noise for injection nor prior knowledge on possible attacks to be defended, our analysis in Section 3 finds that anomaly detection using solely spatial analysis on the latent projection space can fail due to the randomness in per-round client selection and the clustering uncertainty.

3. STDLens: Defense Methodology

Perception poisoning is a real threat. Adversaries can configure the attack to hijack the model without being no-

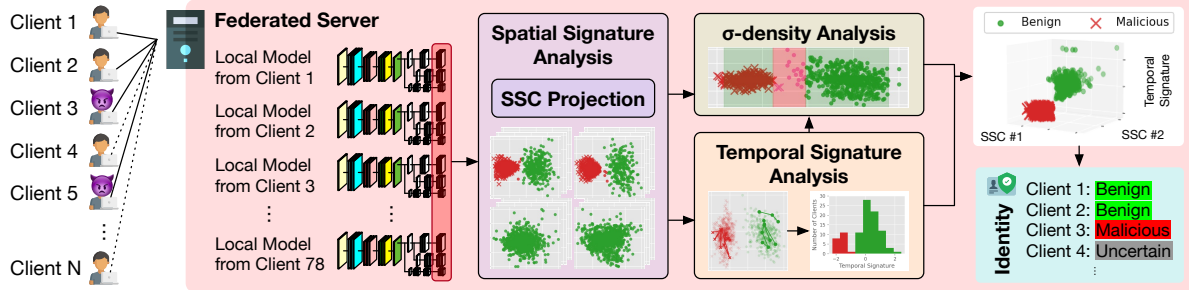


Figure 1. The overview of STDLens, a three-tier framework against model hijacking through perception poisoning.

ticed. We present a forensic framework called STDLens, to safeguard FL systems. It utilizes robust and attack-agnostic statistical properties with a three-tier forensic analysis. It does not assume prior knowledge on which poisoning attack is launched (i.e., Class-Poison, BBox-Poison, Objn-Poison, or their combinations). Figure 1 gives an overview. STDLens is operated at a system-supplied schedule, configurable at initiation (e.g., every ten FL rounds). Contributions from the participating clients during the window of rounds are accumulated for a safety inspection by the spatial signature analysis, followed by a per-client temporal signature analysis to identify Trojanged gradients and a σ -density analysis to manage the spatio-temporal uncertainty.

3.1. Spatial Signature Analysis

Under perception poisoning, it is observed that with a small $m\%$ of malicious clients, the eigenvalues of the covariance of the gradient updates for the source class consist of two sub-populations: the Trojanged gradients from malicious participants and the benign gradients from honest participants. We refer to such phenomenon as m -separable robust statistics of perception poisoning attacks.

Definition 3.1. m -Separable Robust Statistics. Given a small m with $0 < m < \frac{1}{2}$, let H and P denote the two distributions: Honest and Poisoned, respectively, with finite covariance. Let the mixing loss function be $\mathbb{G} = (1 - m)H + mP$, v be the top eigenvalue of the covariance of \mathbb{G} and $\mu_{\mathbb{G}}$ is the mean of the mixed distribution \mathbb{G} . The two distributions H and P are separable if there exists some τ such that $\Pr_{X \sim H}[|\langle X - \mu_{\mathbb{G}}, v \rangle| > \tau] < m$ and $\Pr_{X \sim P}[|\langle X - \mu_{\mathbb{G}}, v \rangle| < \tau] < m$.

Theorem 3.1. Given a small m with $0 < m < \frac{1}{2}$, let H and P denote Honest and Poisoned distributions, with mean μ_H, μ_P , and finite covariance $\Sigma_H, \Sigma_P \preceq \phi^2 \mathbb{I}$. Let the mixing loss function be $\mathbb{G} = (1 - m)H + mP$ and $\Delta = \mu_H - \mu_P$. Then, if $\|\Delta\|_2^2 \geq \frac{6\phi^2}{m}$, P and H satisfy m -separable robust statistics.

The spatial signature analysis relies on the above m -separable robust statistics, where m is represented as a decimal for brevity. We provide the proof in the appendix. The-



(a) Round 102: 7 benign clients and 3 malicious clients (b) Round 106: 4 benign clients and 6 malicious clients

Figure 2. The clustering phenomenon due to benign and malicious client participation in only one round of FL.

orem 3.1 implies that when the percentage m of compromised clients is small, the distribution of benign gradients from honest clients is separable from the distribution of poisoned gradients, especially when the distance between the means of the two distributions can be bounded by the m -normalized covariance of the two distributions [22]. Given the gradient updates are high-dimensional, we focus on the output layer to analyze the statistical properties of gradient updates from all participants because the output layer gradients are the most prominent due to the vanishing gradient in backpropagation. We project the output layer gradients onto a 2D space corresponding to the two spatial signature components (SSCs) with the largest eigenvalues for spatial clustering. Figure 2 shows the spatial signatures of the gradients collected at two different FL rounds under Class-Poison. We observe a clear clustering phenomenon (green dots and red crosses), validating the intuition about the formation of two gradient distributions. A key question is whether one can simply conclude that the smaller cluster contains the Trojanged gradients from malicious clients and thus can be removed in the upcoming FL rounds.

Inherent Limitations. Spatial signatures can offer valuable statistical indicators, especially when poisoned and benign gradient distributions can be cleanly separated. However, existing defense methods simply rely on analyzing such spatial patterns and conclude that the smaller cluster contains the poisonous gradient updates and expel those clients with contributions in the smaller cluster [18, 21]. We argue that relying solely on spatial signature analysis may result in high false positives and hurt the convergence of

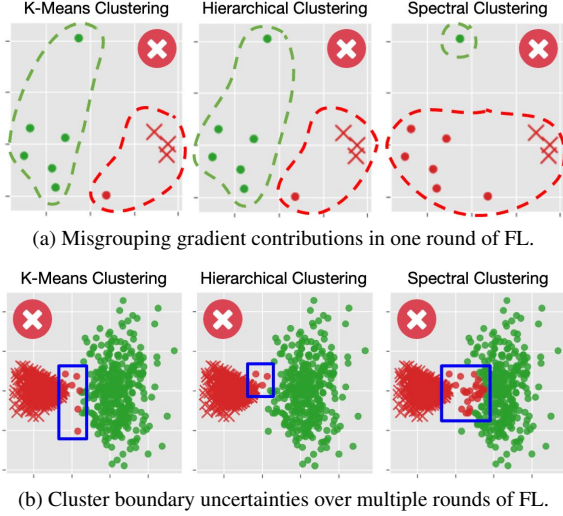
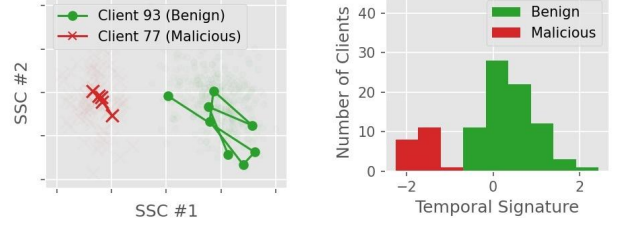


Figure 3. Three clustering algorithms on gradient contributions in (a) one round and (b) multiple rounds.

FL. We identify two reasons why such defenses can fail. First, the randomness in FL client selection can lead to an incorrect revocation, as shown in Figure 2b, where only four out of ten participants are benign. This indicates that relying on spatial signature alone is error-prone. Second, the uncertainties of clustering contribute to another dimension of complication. Figure 3a shows the spatial clustering of gradients contributed to an FL round with three Trojaned contributions (crosses) and seven benign contributions (dots). The clusters returned by three commonly used clustering methods: K-Means, hierarchical clustering, and spectral clustering are all erroneous. Rejecting the smaller cluster in all three cases will result in unacceptable errors because each contains non-poisoned contributions from honest clients. Such a problem persists even if a window of multiple rounds is used instead of a single round (Figure 3b). We observe that those gradients located between two clusters (outlined by the blue rectangle) can be ambiguous, and three different clustering methods produce different boundary cases. This further indicates that a robust defense requires high-fidelity forensic analysis with different ways to refine the inspection results and manage uncertainties.

3.2. Temporal Signature Analysis

STDLenS addresses the inherent limitations of spatial clustering by additional forensic analysis with uncertainty assessment. Instead of concluding the clients contributed to the smaller cluster as malicious and revoking their subscription to FL, STDLenS inspects each client in a given cluster over a temporal window, coined as the per-client temporal signature analysis. We narrow down the temporal analysis to those classes that are potentially under poisoning attacks based on the spatial signature analysis: the gradient updates



(a) Temporal trajectories of example benign and malicious clients (b) Unique statistical characteristics of temporal signatures

Figure 4. A benign client (green) tends to produce temporally dissimilar gradients over the course of FL while a malicious client (red) contributes similar gradients to mislead the FL process.

on a benign class will have one dense cluster, whereas the source class under poisoning attacks will exhibit at least two distinct distributions of gradient updates.

Let \mathcal{G}_i be a temporal sequence of spatial signatures contributed by the same client \mathcal{P}_i throughout a temporal window of FL rounds. Figure 4a shows the spatial signature trajectories of two clients (one benign and one malicious). We can observe that the gradient contributions from a malicious client (i.e., the red crosses) are close to each other, but those from a benign client (i.e., the green dots) scatter around. To quantify this behavior, we define the following ω -based temporal signature, denoted by Υ_ω , for a client \mathcal{P}_i :

$$\Upsilon_\omega(\mathcal{G}_i) = \min_c \left(\sum_{j=\omega+1}^{|\mathcal{G}_i^c|} \sum_{k=1}^{\omega} \frac{\text{DISSIM}(\mathcal{G}_i^c[j], \mathcal{G}_i^c[j-k])}{\omega |\mathcal{G}_i^c| - \omega^2} \right) \quad (1)$$

where \mathcal{G}_i^c denotes the sub-sequence of gradients belonging to the same cluster c from the spatial signature analysis, $\mathcal{G}_i^c[j]$ refers to its j -th element, and $\text{DISSIM}(\cdot, \cdot)$ is a dissimilarity measure, which can utilize any L_p norm-based similarity, and L_1 distance is used as the default given its computation efficiency. The parameter ω refers to the window size. The larger the window size ω is, the longer the temporal dependency is leveraged for computing temporal signatures. Figure 4b shows the distribution of temporal signatures with $\omega = 1$ (i.e., the temporal signature is computed every two consecutive rounds). This empirically verifies that malicious clients (red bars) tend to have a significantly lower temporal signature Υ_ω than benign clients (green bars). We utilize the per-client temporal signatures to refine the 2D spatial signature of a gradient update contributed by a client. This forms a 3D spatio-temporal signature space (see the top right in Figure 1) for STDLenS to re-examine each cluster from the spatial signature analysis, aiming to determine which cluster of gradient updates is from malicious clients. We first compute the mean temporal signature over all contributing clients of each cluster and mark the one with a lower temporal dissimilarity score as suspicious. Then, we perform an uncertainty assessment before concluding that those clients contributed to the suspicious cluster as malicious.

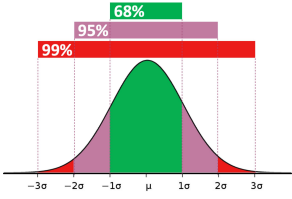


Figure 5. The empirical rule for confidence estimation.

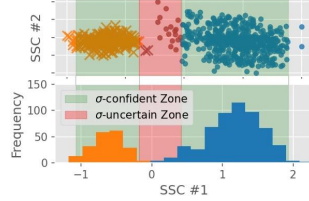


Figure 6. σ -density analysis on spatial signatures.

3.3. Density-based Uncertainty Inspection

The results from the spatial clustering analysis can be unreliable, as shown in Figure 3. Similarly, the temporal signature analysis may also incur uncertainty if the average temporal dissimilarity score for each of the two clusters is too close. We introduce the third component dedicated to uncertainty correction. Concretely, it identifies uncertain spatial and temporal zones such that a client is expelled only if we have high confidence that its gradient updates are poisoned. Our uncertainty assessment is based on the empirical rule in normal distributions that 68%, 95%, and 99% of data points lie within 1, 2, and 3 standard deviations (σ 's) from the mean (Figure 5). The confidence requirement is domain-specific. With a given confidence threshold for uncertainty control, we can define the corresponding σ -confident and σ -uncertain zones. In particular, we first identify the SSC with the largest eigenvalue. It captures the inter-cluster variance and is visualized as the x -axis by default. Figure 6 shows that the spatial signatures (top) on this component exhibit a mixture of two normal distributions (bottom). This allows STDLens to employ the empirical rule by computing the mean and the standard deviation of each cluster (orange or blue in Figure 6) and translating the confidence requirement (e.g., 99%) into the corresponding intervals (e.g., $3\text{-}\sigma$ from the mean of the respective cluster). STDLens considers the zone between two σ -confident regions to be uncertain, and any gradient updates from participating clients that fall into this zone will be put on a watchlist. The same client entering the σ -uncertain zone multiple times will be removed from FL. The same uncertainty management operation can also be done on the temporal signatures to minimize the uncertainty-induced false positives. For clients in the cluster with a smaller dissimilarity score and outside the σ -uncertain zone, STDLens will mark them as malicious and revoke their participation in the remaining rounds.

4. Experimental Evaluation

We conduct experiments on two object detection benchmark datasets: (i) PASCAL VOC [7] and (ii) INRIA Person [6]. By default, we have $N = 100$ clients and a federated server with $T = 200$ rounds for both datasets. At each round, $k = 10\%$ of the clients are randomly selected as

Perception Poison	m	Benign AP_{src} (%)	Hijacked AP_{src} (%)	
			No Defense	STDLens
Class-Poison (VOC)	5%	52.43	45.35	52.41
	10%	52.43	39.69	52.32
	20%	52.43	38.22	52.22
BBox-Poison (VOC)	5%	52.43	35.76	52.32
	10%	52.43	23.49	52.15
	20%	52.43	18.06	52.19
Objn-Poison (VOC)	5%	52.43	45.02	51.95
	10%	52.43	41.61	52.33
	20%	52.43	37.61	52.29
BBox-Poison (INRIA)	5%	63.49	45.03	61.64
	10%	63.49	32.51	61.25
	20%	63.49	28.73	61.71
Objn-Poison (INRIA)	5%	63.49	61.05	61.98
	10%	63.49	53.91	62.01
	20%	63.49	50.18	61.58

Table 2. The AP_{src} of FL-trained detectors under the benign scenario (3rd column) and three perception poisoning attacks (4th and 5th columns) with three settings of malicious client percentage m .

participants with FedAvg [12] as the aggregation algorithm. Without loss of generality, we focus on Faster R-CNN [16] with VGG16 [19] because perception poisoning is model-agnostic and black-box in nature. For each dataset, we randomly split the training set into N equal-sized partitions and assigned each to one client as its local dataset. We analyze STDLens against all three types of perception poisoning on VOC and examine BBox-Poison and Objn-Poison on INRIA as it only supports objects of the “person” class, which is the default victim source class. The target class in Class-Poison is “potted plant.” For BBox-Poison, we shrink the bounding box by a factor of 0.10. By default, $m = 20\%$ of the clients in FL are malicious. STDLens is operated with a forensic schedule of ten rounds per window, and the confidence required for uncertainty control is set to be 99%.

4.1. Effectiveness of STDLens

Trojan-agnostic Defense. We first analyze how effective STDLens is in mitigating different model hijacking attacks through perception poisoning. Table 2 reports the results measured in average precision [7] of the victim source class (AP_{src}) under three different settings of $m\%$. Given that those attacks can hijack the models to cause a drastic drop in AP_{src} (3rd and 4th columns), STDLens offers strong defensibility against all of them (5th column). Regardless of which attack is launched and how many malicious clients (reflected by m) participate in the attack process, STDLens can consistently restore the AP_{src} , on par with the one obtained in the benign scenario (i.e., 52.43% for VOC and 63.49% for INRIA). For instance, BBox-Poison, which brings down the AP_{src} to 18.06% on VOC, is mitigated with an AP_{src} of 52.19% under the same attack setting. This validates that STDLens is Trojan-agnostic, despite the divergence of different attack effects. Due to similar observa-

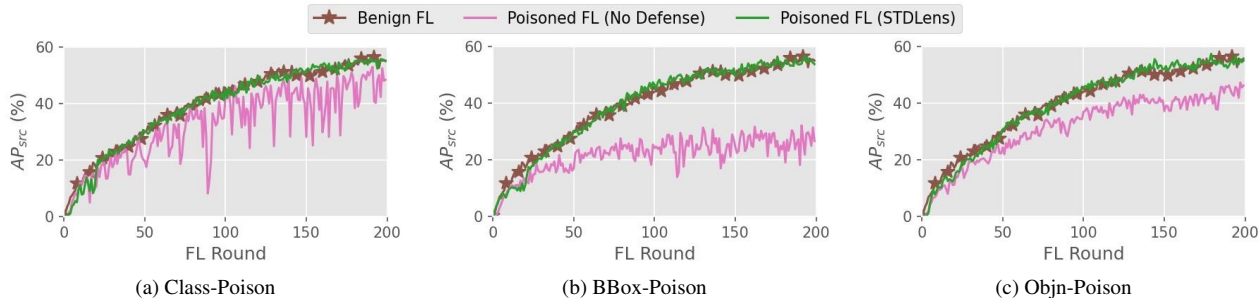


Figure 7. Compared with the benign FL (brown) on VOC, the hijacked FL with no defense (pink) has significant AP_{src} degradation. STDLens (green) purges all malicious clients in an early-round, allows the healing to begin, and reclaims AP_{src} with a negligible damage.

Class-Poison		BBox-Poison		Objn-Poison	
No Defense	STDLens	No Defense	STDLens	No Defense	STDLens

Table 3. Two images detected by the FL-trained models (VOC) under Class-Poison, BBox-Poison, and Objn-Poison. The FL system with no defense makes designated errors, while the STDLens-protected FL preserves correctness even if the adversary attempts to hijack it.

tions, we focus on VOC in the following discussion.

Timely Sanitization and Recovery. Figure 7 visualizes the learning curves of FL throughout 200 rounds by comparing three cases: (i) the benign scenario (brown curve), (ii) the hijacked scenario without defense (pink curve), and (iii) the hijacked scenario protected by STDLens (green curve). We make the following observations. First, Trojaned gradients cause the AP_{src} to become either highly unstable (Class-Poison) or consistently degraded (BBox-Poison and Objn-Poison) throughout the FL process. Second, when STDLens detects all malicious clients at the 30th round under Class-Poison, the 20th round under BBox-Poison, and the 30th round under Objn-Poison, no honest client is misdiagnosed and flagged as malicious. Third, the timely sanitization of malicious clients allows the FL system to heal. Even under attacks, the STDLens-protected FL can reach an AP_{src} comparable with the benign scenario.

Visualization. Table 3 visualizes the detection results on two test examples under all three perception poisoning attacks. All three attacks fool the global model with erroneous detection: (i) Class-Poison (1st column) fools the model into mislabeling the person as a potted plant (top) or a bird (bottom), (ii) BBox-Poison (3rd column) deceives the model to recognize person objects with incorrect bounding boxes, and (iii) Objn-Poison (5th column) causes the model to detect no person objects. In comparison, STDLens successfully reinvigorates the FL system with high robustness

(2nd, 4th, and 6th columns).

4.2. STDLens under Adaptive Attacks

β -adaptive Attacks. An immediate question we need to address is the defensibility of STDLens against adaptive adversaries who strategically change their attack methods. For instance, malicious clients could camouflage in the crowd of benign clients by occasionally contributing non-Trojaned gradients. We refer to this type of adaptive adversary as β -adaptive. When a β -adaptive adversary participates in an FL round, it has, e.g., a $\beta = 10\%$ chance of not poisoning the local dataset during the local model training. As a result, this adaptive adversary will contribute to the cluster with benign gradient updates and also to the cluster with Trojaned gradient updates, aiming to fool the spatial and temporal signature analysis. Indeed, such an adaptive attack strategy can jeopardize the spatial analysis because malicious clients can be mistakenly clustered with honest ones, as depicted in Figure 8a. This makes defenses using solely spatial clustering [18, 21] unable to purge all malicious clients. Even for another representative defense using spectral signatures [22], the statistical properties of malicious contributions become the same as those benign ones, as shown in Figure 9a. However, the temporal signature analysis in STDLens is, by design, resilient to such adaptive strategies, as shown in Figure 8b. Recall in Equation 1 that the temporal signature of a client is com-

Perception Poison	Object Detection with No Defense	Object Detection with STDLens	β	Benign AP _{src} (%)	Hijacked AP _{src} (%)	
					No Defense	STDLens
Class-Poison			5%	52.43	31.83	51.74
			10%	52.43	32.51	52.18
			15%	52.43	34.40	51.56
BBox-Poison			5%	52.43	22.36	52.39
			10%	52.43	26.77	52.11
			15%	52.43	27.60	51.98
Objn-Poison			5%	52.43	40.96	52.05
			10%	52.43	41.01	52.18
			15%	52.43	41.06	52.11

Table 4. The benign AP_{src} and hijacked AP_{src} under β -adaptive attacks with varying β in different types of perception poison (VOC).

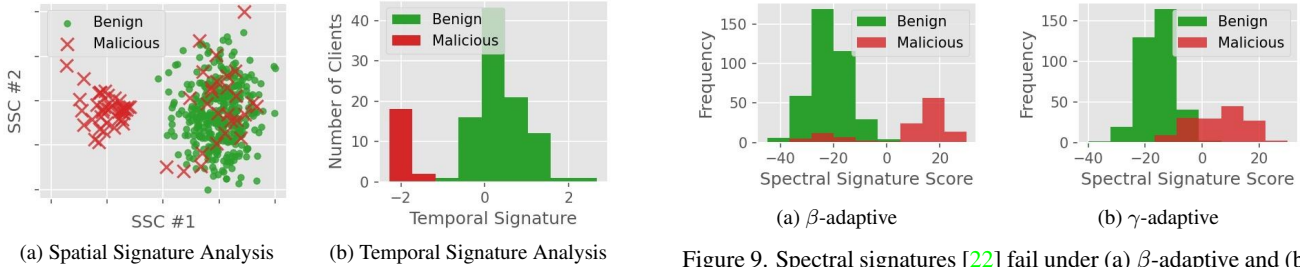


Figure 8. β -adaptive malicious clients randomly camouflage in the benign cluster in the spatial analysis (a), but their statistical characteristics in the temporal signature analysis (b) remain unchanged.

puted w.r.t. each cluster, and the one with a smaller magnitude is taken to be the lower bound of malicious likelihood. Hence, the statistical characteristics of honest and malicious clients can be preserved in our temporal analysis, where malicious clients (red) have a much smaller temporal signature than honest clients (green). Table 4 compares the hijacked AP_{src} with the β -adaptive attack under no defense scenario (6th column) with the FL system under the protection by STDLens (7th column). A total of nine scenarios under three attacks are compared by varying the configurations of β from 5% to 15%. We make two observations. First, under the β -adaptive attack, the AP_{src} suffers from significant degradation, from 52.43% to 22.36%, even though malicious clients do not send Trojaned gradients all the time. For adaptive adversaries, the choice of β should be kept small as a larger β can notably weaken the attack effectiveness. Second, STDLens remains resilient under β -adaptive adversaries since it can consistently restore the AP_{src} to at least 51.56% for all nine scenarios, on par with that under non-adaptive attacks (recall Table 2). This is primarily because our temporal analysis is robust against such a camouflage. Integrating spatio-temporal statistics allows the identification of the malicious cluster to remain accurate under β -adaptive attacks.

γ -adaptive Attacks. Another attempt to launch adaptive attacks is to wash out the spatial clustering characteristics by only poisoning $\gamma\%$ of training samples in the local dataset of a malicious client. Under this adaptive setting,

Figure 9. Spectral signatures [22] fail under (a) β -adaptive and (b) γ -adaptive attacks as the statistical properties of malicious contributions (red) overlap with benign ones (green).

γ	Benign AP _{src} (%)	Hijacked AP _{src} (%)	
		No Defense	STDLens
80%	52.43	40.03	52.09
60%	52.43	46.01	51.54

Table 5. Quantitative studies of γ -adaptive Class-Poison (VOC).

the contribution from a malicious client becomes the gradient updates computed over a mixture of clean and poisoned training samples. We conduct experiments with two settings of γ with results in Table 5. The γ -adaptive attacks can bring down the AP_{src} from 52.43% to 40.03% and 46.01% when we set γ to 80% and 60%, respectively. More malicious clients fall inside the σ -uncertain zone, which will be incorrectly identified by defenses using only spatial clustering analysis [18, 21]. Such an adaptive attack also makes spectral signatures [22] ambiguous, as the statistical properties of benign and malicious gradients overlap (Figure 9b). For those uncertain clients, STDLens has to put the decision on hold and wait for more rounds to accumulate additional observations for making confident decisions. Table 5 shows that STDLens remains robust under γ -adaptive attacks because all malicious clients can still be identified with no false alarm. The AP_{src} can be recovered to reach over 51.54%, comparable to the benign AP_{src} of 52.43%.

Late-stage Adaptive Attacks. Instead of starting the poisoning process from the beginning of FL, an adaptive adversary could only start poisoning the training data when selected in the later rounds of FL. In this case, the FL system has no Trojaned gradient updates contributed by its clients for a period of time. We conduct such a late-stage adap-

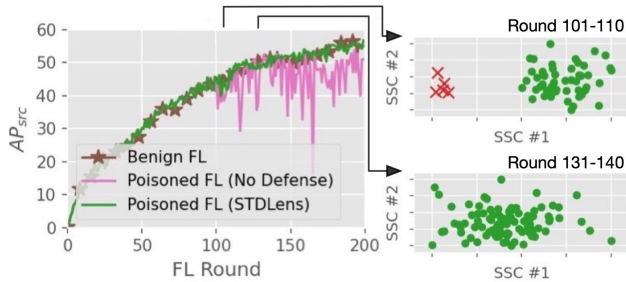


Figure 10. The malicious clients become active after the 100th round (VOC), causing the AP_{src} to drop. STDLens purges all malicious clients at the 130th round and allows the FL to recover.

tive attack by running Class-Poison after the 100th round during the FL process. Figure 10 (left) shows the learning curves of AP_{src} for three scenarios: (i) the benign scenario, (ii) the late-stage adaptive attack without defense, and (iii) the one protected by STDLens. Considering the no-defense scenario (pink curve), the AP_{src} fluctuates wildly after the 100th round. In comparison, STDLens (green curve) starts detecting the existence of Trojans from the window of rounds 101-110 (top right) and removes the identified malicious clients. It successfully expels all anomalies at the 130th round and allows the FL to heal as no more malicious clients attempt to mislead the FL system (e.g., the window of rounds 131-140 at the bottom right).

4.3. Comparisons with Existing Defenses

We further compare STDLens with two existing defense methods: spatial clustering [21] and spectral signature [22], which can be extended to FL-based object detection. Table 6 reports the quantitative results. All three defense solutions can detect all malicious clients over the course of FL (i.e., achieving the defense recall of 1.00). STDLens completes the detection and revocation of malicious clients between 20 to 30 rounds, and spatial clustering takes 30 to 40 rounds to detect and remove anomalies, while malicious clients can stay alive in the FL system for at most 80 rounds with spectral signature. The main differentiation factor lies in defense precision. Even though STDLens takes a slightly longer execution time (i.e., 452ms with 84ms for temporal signature analysis and 44ms for σ -density inspection), it successfully identifies all malicious clients without mistakenly rejecting any benign clients (false positives), achieving a defense precision of 1.00. This allows the STDLens-protected FL to train a global object detector with high accuracy, as shown in the 2nd column in Table 7. In comparison, the spatial clustering approach can reach a low precision of 0.53 under Objn-Poison, due to gradient contributions falling in the zone between the two clusters, causing ambiguity in anomaly detection. This result shows that defenses based solely on spatial clustering may not offer stable robustness. The spectral signature defense performs the worst, with a defense precision ranging from 0.26 to 0.54. It

	Defense Precision@Round w/ Max. Recall		
	STDLens (Ours)	Spatial Clustering [21]	Spectral Signature [22]
Class-Poison	1.00@30	0.95@30	0.53@40
BBox-Poison	1.00@20	0.71@30	0.26@80
Objn-Poison	1.00@30	0.53@40	0.54@40
Time (ms)	452	342	1423

Table 6. Comparing STDLens with two state-of-the-art methods on VOC in terms of the defense precision at the round where the maximum recall is achieved and the execution time (the last row).

	Object Detection Results by Hijacked Models		
	STDLens (Ours)	Spatial Clustering [21]	Spectral Signature [22]
Class-Poison			
BBox-Poison			
Objn-Poison			

Table 7. The high defense precision offered by STDLens allows the FL-trained model to remain accurate. The excessive removal of clients by other defenses incurs performance degradation (VOC).

implies that while the defense can expel malicious clients, a considerable number of honest ones are also mistakenly rejected. One reason is the changing malicious population over time due to the removal of malicious clients in the early rounds. The excessive removal of clients can weaken the benign learning signal and degrade the FL-trained model, as shown in the 3rd and 4th columns in Table 7.

5. Conclusions

With the goal of building a trustworthy federated training process, we have presented STDLens, a three-tier defense framework for FL-based object detection, strengthening the spatial signature with the temporal signature analysis and uncertainty management. Extensive experiments with various adaptive adversaries show that STDLens can protect FL against all three model hijacking attacks and outperforms existing mitigation methods with high defense precision and a low false-positive rate.

Acknowledgments. This research is partially sponsored by the NSF CISE grants 2038029, 2026945, 1564097, an IBM faculty award, and a Cisco grant on Edge Computing. The first author acknowledges the IBM PhD Fellowship Award.

References

- [1] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 2938–2948. PMLR, 2020. 1
- [2] Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens. In *International Conference on Machine Learning*, pages 634–643. PMLR, 2019. 1
- [3] Ka-Ho Chow and Ling Liu. Perception poisoning attacks in federated learning. In *International Conference on Trust, Privacy and Security in Intelligent Systems and Applications*, pages 146–155. IEEE, 2021. 1, 2
- [4] Ka-Ho Chow and Ling Liu. Boosting object detection ensembles with error diversity. In *International Conference on Data Mining*, pages 903–908. IEEE, 2022. 2
- [5] Ka-Ho Chow, Ling Liu, Margaret Loper, Juhyun Bae, Mehmet Emre Gursoy, Stacey Truex, Wenqi Wei, and Yanzhao Wu. Adversarial objectness gradient attacks in real-time object detection systems. In *International Conference on Trust, Privacy and Security in Intelligent Systems and Applications*, pages 263–272. IEEE, 2020. 1
- [6] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *Conference on Computer Vision and Pattern Recognition*, volume 1, pages 886–893. IEEE/CVF, 2005. 5
- [7] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 5
- [8] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Gong. Local model poisoning attacks to byzantine-robust federated learning. In *USENIX Security Symposium*, pages 1605–1622, 2020. 1
- [9] Clement Fung, Chris JM Yoon, and Ivan Beschastnikh. The limitations of federated learning in sybil settings. In *International Symposium on Research in Attacks, Intrusions and Defenses*, pages 301–316, 2020. 1
- [10] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021. 1
- [11] Raouf Kerkouche, Gergely Ács, and Claude Castelluccia. Federated learning in adversarial settings. *arXiv preprint arXiv:2010.07808*, 2020. 2
- [12] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*. PMLR, 2017. 1, 5
- [13] Anup Mohan, Kent Gauen, Yung-Hsiang Lu, Wei Wayne Li, and Xuemin Chen. Internet of video things in 2030: A world with many cameras. In *International Symposium on Circuits and Systems*, pages 1–4. IEEE, 2017. 1
- [14] Thien Duc Nguyen, Phillip Rieger, Huili Chen, Hossein Yalame, Helen Möllering, Hossein Fereidooni, Samuel Mar-
chal, Markus Miettinen, Azalia Mirhoseini, Shaza Zeitouni, et al. FLAME: Taming backdoors in federated learning. In *USENIX Security Symposium*, pages 1415–1432, 2022. 2
- [15] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 1
- [16] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 28:91–99, 2015. 5
- [17] Ahmed Salem, Michael Backes, and Yang Zhang. Get a model! model hijacking attack against machine learning models. In *The Network and Distributed System Security Symposium*, 2022. 1
- [18] Shiqi Shen, Shruti Tople, and Prateek Saxena. Auror: Defending against poisoning attacks in collaborative deep learning systems. In *Annual Conference on Computer Security Applications*, pages 508–519, 2016. 2, 3, 6, 7
- [19] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015. 5
- [20] Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. Can you really backdoor federated learning? In *International Workshop on Federated Learning for Data Privacy and Confidentiality*, 2019. 1, 2
- [21] Vale Tolpegin, Stacey Truex, Mehmet Emre Gursoy, and Ling Liu. Data poisoning attacks against federated learning systems. In *European Symposium on Research in Computer Security*, pages 480–501. Springer, 2020. 1, 2, 3, 6, 7, 8
- [22] Brandon Tran, Jerry Li, and Aleksander Madry. Spectral signatures in backdoor attacks. In *Advances in Neural Information Processing Systems*, 2018. 2, 3, 6, 7, 8
- [23] Yuao Wang, Tianqing Zhu, Wenhan Chang, Sheng Shen, and Wei Ren. Model poisoning defense on federated learning: A validation based approach. In *International Conference on Network and System Security*, pages 207–223. Springer, 2020. 2
- [24] Chulin Xie, Minghao Chen, Pin-Yu Chen, and Bo Li. Crfl: Certifiably robust federated learning against backdoor attacks. In *International Conference on Machine Learning*, pages 11372–11382. PMLR, 2021. 2
- [25] Jie Xu, Benjamin S Glicksberg, Chang Su, Peter Walker, Jiang Bian, and Fei Wang. Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1):1–19, 2021. 1
- [26] Xiaojun Xu, Qi Wang, Huichen Li, Nikita Borisov, Carl A Gunter, and Bo Li. Detecting AI trojans using meta neural analysis. In *Symposium on Security and Privacy*, pages 103–120. IEEE, 2021. 2
- [27] Qian Yang, Jianyi Zhang, Weituo Hao, Gregory Spell, and Lawrence Carin. Flop: Federated learning on medical datasets using partial networks. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2021. 1
- [28] Jiale Zhang, Di Wu, Chengyong Liu, and Bing Chen. Defending poisoning attacks in federated learning via adversarial training method. In *International Conference on Frontiers in Cyber Security*, pages 83–94. Springer, 2020. 2