

Let's Focus: Focused Backdoor Attack against Federated Transfer Learning

Marco Arazzi

University of Pavia
Pavia, Italy

marco.arazzi01@universitadipavia.it

Antonino Nocera

University of Pavia
Pavia, Italy

antonino.nocera@unipv.it

Stefanos Koffas

Delft University of Technology
Delft, The Netherlands
S.Koffas@tudelft.nl

Stjepan Picek

Radboud University
Nijmegen, The Netherlands
stjepan.picek@ru.nl

ABSTRACT

Federated Transfer Learning (FTL) is the most general variation of Federated Learning. According to this distributed paradigm, a feature learning pre-step is commonly carried out by only one party, typically the server, on publicly shared data. After that, the Federated Learning phase takes place to train a classifier collaboratively using the learned feature extractor. Each involved client contributes by locally training only the classification layers on a private training set. The peculiarity of an FTL scenario makes it hard to understand whether poisoning attacks can be developed to craft an effective backdoor. State-of-the-art attack strategies assume the possibility of shifting the model attention toward relevant features introduced by a forged trigger injected in the input data by some untrusted clients. Of course, this is not feasible in FTL, as the learned features are fixed once the server performs the pre-training step.

Consequently, in this paper, we investigate this intriguing Federated Learning scenario to identify and exploit a vulnerability obtained by combining eXplainable AI (XAI) and dataset distillation. In particular, the proposed attack can be carried out by one of the clients during the Federated Learning phase of FTL by identifying the optimal local for the trigger through XAI and encapsulating compressed information of the backdoor class. Due to its behavior, we refer to our approach as a focused backdoor approach (FB-FTL for short) and test its performance by explicitly referencing an image classification scenario. With an average 80% attack success rate, obtained results show the effectiveness of our attack also against existing defenses for Federated Learning.

1 INTRODUCTION

In recent years, deep learning and machine learning solutions received great attention from the research community thanks to the development of powerful technologies capable of handling and processing huge data volumes. However, the need to collect data from disparate sources to foster the training of models implies sharing raw data with a third-party computation unit, thus implying privacy concerns over the users' data. This condition led to the development of Federated Learning [22], according to which users train collaboratively a model without sharing any private information and, therefore, maintaining data localized. Initially, the basic definition of Federated Learning assumes that all the involved clients own private data from the same feature space (Horizontal

Federated Learning, HFL for short). This assumption makes this scenario impractical in some cases. As a matter of fact, in many real-life scenarios different parties need to use data for the same samples but from different feature spaces. For example, a bank may collaborate with an invoice agency to build a financial risk model for their customers [19]. Under this assumption, the collaborative learning among such clients is called Vertical Federated Learning (VFL).

By extending the scenarios considered above, the most general case consists of parties that collaborate into a machine/deep learning process by using local data that differ in both feature and sample space; this situation is encompassed by the Federated Transfer Learning (FTL) [19]. Such a paradigm received great attention from the research community and has been, for instance, employed in [8], where the authors propose a novel architecture for personalized healthcare via wearable devices. Given the type of application, we believe it is crucial to understand the vulnerabilities this paradigm introduces, which are still relatively unexplored. For this reason, this paper focuses on understanding whether backdoor attacks can be crafted by manipulating local data on maliciously controlled parties involved in Federated Transfer Learning. Backdoor attacks have been largely studied in recent years in both HFL [5, 12, 29, 30] and VFL [20]; however, to the best of our knowledge, our proposal is the first to target FTL and show that such a paradigm is vulnerable to backdoor attacks. It is worth underlying that, compared to the classical FL scenarios, FTL presents unique challenges as the learning is typically broken down into two parts: a feature learning task, globally carried out by a server on a publicly shared dataset and the subsequent training of custom classification layers carried out by all the clients using private local data in a federated task. Because the feature extractor part of the classification model is learned in a different phase with respect to the federated learning one, crafting a backdoor implies a totally different approach than those typically adopted by existing backdoor attacks for FL. In our study, to define our strategy for a backdoor attack, we adopt the FTL scenario of [8], where the clients receive an initial model trained to extract features from input data by the server and then engage in a federated learning task to train only the classification layers (i.e., the first part of the model shared by the server is kept frozen). As typically done in the related literature, in our study, we chose image classification as a reference deep learning task, and the frozen part, pre-trained by the server, is the feature extractor

of a Convolutional Neural Network (CNN). Then, each client uses a private local dataset to collaboratively fine-tune the last layers of the model (fully connected layers) by computing local updates during each training epoch and uploading the results to the server. The server then aggregates the changes and distributes the updated global model to the clients for the subsequent learning epochs. In this complex scenario, an adversary aiming at designing a backdoor trigger cannot leverage the possibility of forging new relevant features for the model to learn, such as a small square of noise [12] or simple objects like sunglasses [7], because the feature extractor block of the model is frozen. Therefore, any additional feature will be promptly discarded by the global model. As a result, the attacker needs a more elaborated mechanism for the trigger design that would steer the poisoned sample’s latent representation toward the target class.

We overcome this challenge by combining an Explainable AI (XAI) step on the feature extractor with a data-driven approach to build a malicious trigger and suitably position it in high-attention locations learned by the global model. For this reason, we call our strategy a *focused* backdoor attack. To do so, we leverage a powerful XAI tool, named GradCam [25], to find out the best way for an adversary to alter the images. Moreover, we exploit a distillation strategy [36] to encapsulate inside the trigger the main information characterizing the features of the target class of the backdoor. We argue that by suitably locating our trigger in attention locations for the feature extractor component of the model and encapsulating a compressed representation of the target class in it, we can deviate the behavior of the classifier when such a trigger is present in the input image. The experiments we carried out show that FTL is still vulnerable to backdoor attacks built according to our strategy.

In summary, our main contributions are as follows:

- We introduce FB-FTL, the first focused backdoor attack against Federated Transfer Learning. Based on our experiments, FB-FTL achieves a success rate of over 80%, on average. It is important to note that, to our knowledge, this paper is the first work to study a backdoor attack against Federated Transfer Learning.
- We also extend the family of explainability-based attacks, defining and proving the importance of a “focusing” strategy for the positioning of the trigger using GradCam. In addition, we exploit dataset distillation to generate a trigger containing the features of the target class.
- We propose a strategy to blend better the pixels of the trigger with the colors of the original image, making it less noticeable.
- We tested our approach against well-known defenses for Horizontal Federated Learning, and we also compared it against a novel Vertical Federated Learning countermeasure based on Label Differential Privacy. While some defenses are effective with specific datasets, none of the considered defenses can mitigate our approach in all the considered scenarios.

An important facility of our proposal is the capability of adapting the trigger to the attacked data to minimize its impact in terms of necessary variations caused by its inclusion. This is especially true for images where triggers typically alter visible regions to craft a

backdoor. Training data are automatically processed by local clients, and therefore, the fact that the trigger may actually be visible to some extent does not necessarily represent a blocking point for the attacker. However, minimizing its visible impact is crucial to make the attack as stealthy as possible.

2 BACKGROUND

2.1 Federated Learning

Federated Learning (FL) allows n clients to train a global model \mathbf{w} collaboratively without revealing their datasets. According to [32], it can be classified into 1) Horizontal Federated Learning (HFL), 2) Vertical Federated Learning (VFL), and 3) Federated Transfer Learning.

In HFL, clients own data that share the same feature space but represent different entities [32]. Each client trains a local copy of the model and then uploads its weights ($\{\mathbf{w}^i \mid i \in n\}$) to a server. HFL optimizes the following loss function:

$$\min_{\mathbf{w}} \ell(\mathbf{w}) = \sum_{i=1}^n \frac{k_i}{n} L_i(\mathbf{w}), \quad L_i(\mathbf{w}) = \frac{1}{k_i} \sum_{j \in P_i} \ell_j(\mathbf{w}, x_j), \quad (1)$$

where $L_i(\mathbf{w})$ and k_i represent the loss function and local data size of i -th client, and P_i refers to the set of data indices with size k_i . At the m -th iteration, the training can be divided into three steps. First, all clients download the global model \mathbf{w}_m from the server. Then, each client updates the local model by training with their datasets ($\mathbf{w}_m^i \leftarrow \mathbf{w}_m^i - \alpha \frac{\partial L(\mathbf{w}_m^i, b)}{\partial \mathbf{w}_m^i}$, where α and b refer to learning rate and local batch). Finally, after the clients upload their local weights, the server updates the global model by aggregating the local weights. There are various aggregation techniques that can be used based on the given scenario. In this work, we will mostly consider averaging, but we will also test four additional aggregation strategies in an extra experiment.

In the VFL, clients own data that belong to the same entities but have different features [32]. Thus, their local models can differ as their data use different features [1].

2.2 Federated Transfer Learning

In transfer learning, a model that has been trained for a specific task can be reused as a starting point for a model on a different task, significantly lowering the costs required for the development of the new model. Federated transfer learning (FTL) combines principles from Transfer Learning and Federated Learning. It consists of a central model that aggregates updates from clients or has trained a model from a publicly available dataset. This model is then distributed to the clients who use their own datasets to improve its performance on their data distribution [8]. Unlike Federated Learning, participants in FTL own datasets that have different population samples and features, making FTL a realistic scenario for real-world applications [15]. In our case, we follow the FedHealth’s paradigm [8], where the clients freeze the first layers of the received CNN (convolutional and max pooling layers) and then train the fully-connected layers using their own private data.

2.3 Backdoor Attacks

The backdoor attack is a very popular threat against deep learning models introduced in [12]. In this attack, the adversary adds a secret functionality into a model that will be activated during inference. The backdoor is activated by malicious inputs that contain a predefined property, the trigger. The backdoor’s activation allows the attacker to control the model’s behavior. For example, in autonomous driving, the backdoored model could classify a stop sign as a speed limit [12]. The backdoor can be embedded through data poisoning [12], code poisoning [4], or model poisoning [14]. To measure the attack’s effectiveness, we use the attack success rate, which represents the number of times that the backdoor is activated over the total poisoned samples fed into the network. To keep the attack stealthy, we need to ensure that the model’s performance on the designated task is not affected by the backdoor insertion.

2.4 GradCam

Interpretability in AI seeks to understand a model’s decision. Various techniques can be used to this end. Recently, class activation mapping (CAM) was introduced [37]. CAM visualizes the areas of an image that are important for the model’s decision but sacrifices the model’s performance because it needs to modify the model’s architecture [25]. Gradient-weighted CAM (GradCam) is a generalization of CAM that does not have this requirement [25]. In particular, instead of modifying the model’s architecture, it uses the gradient information that flows into the final convolutional layer to produce a localization of the crucial areas of the image that are connected to the model’s decision. In this work, we use GradCam to identify the best position of our triggers.

3 METHODOLOGY

3.1 Attack Scenario

As introduced in Section 2.2 and as also done in [8], in the considered scenario, we assume that the server S holds the public dataset $D_p = \{d_p, y_p\}$ and uses it to train the transfer model M_p . Without loss of generality, in our proposal and experiments, we consider an image classification task, making Convolutional Neural Networks an intuitive architectural choice. However, we argue that with minor tweaks, our approach could also work with different models. In practice, the server carries out a global Feature Learning task that will then be shared by all the clients. The weights W_p of the model M_p are trained as follows (CE = cross-entropy):

$$preds = M_p(d_p) \quad (2)$$

$$l = CE(preds, y_p); \nabla W_p = \sum \frac{\partial l}{\partial M_p}. \quad (3)$$

The trained model M_p can now be propagated to the clients $C = \{c_1, \dots, c_n\}$. In particular, each client knows the public dataset D_p and holds a private dataset D_i (where $i \in [1, n]$) that cannot be accessed by the others. During the Federated Learning process, the lower levels of the model, in our case, the Convolutional layers *Conv*, are frozen to preserve the network’s ability to extract low-level and more general features of the image. The classification layers *fc*, instead, are the most capable of capturing the high-level features. Therefore, the *fc* layers are trained by the clients using

their local datasets D_i to customize the classification behavior based on the domain their private data are extracted from. In particular, the training of *fc* is obtained by optimizing the combination of two different loss functions. The first one is a standard cross-entropy loss between the prediction on D_i and the corresponding labels. The second one is an alignment loss on the last layers of *fc* to better personalize the model on the client’s data. As described in [8], this alignment loss function is intended to align the output features between the inputs. In this particular case, the alignment is performed between the outputs of the model M_p on the public dataset D_p and the private datasets D_i . The loss function is calculated as follows:

$$S_{img} = M_p(D_p); T_{img} = M_p(D_n) \quad (4)$$

$$l_{CORAL} = \frac{1}{4d^2} \|S_{img} - T_{img}\|_F^2 \quad (5)$$

$$tot_l = CE(M_p(D_i), y_i) + \alpha l_{CORAL}. \quad (6)$$

Here, $\|\cdot\|_F^2$ is the squared matrix Frobenius norm, d is the dimension of the embedding features, and α is the trade-off parameter between the loss functions. This alignment allows each client to obtain a personalized model M_i , which is more accurate for the local data.

Keeping the *Conv* layers frozen, along with the availability of public data used to train the initial model, opens the framework to possible threats from a malicious client. The intuition behind our approach is to exploit the transfer model to craft dynamic triggers that embed the features of the target class. Since the *Conv* layers work as feature extractors in the considered scenario (image classification), they will focus only on the features of the images that are known from the pre-training of the network on the public data. For this reason, we can reasonably expect that the *Conv* layers will discard traditional triggers that add additional features unknown to the network. To craft our dynamic trigger, we need to distill low-level and general features leaked directly from M_p and encapsulate them into it. Since the transfer model M_p and the public dataset D_p are available for all the clients by design, an attacker can exploit them to craft triggers by leveraging features from D_p that are known to the *Conv* layers of the target model.

However, distilling a dynamic trigger is still not enough to drift the transfer model M_p towards the target class, as we prove experimentally in Section 4.5. In this scenario, the positioning of the trigger plays a key role in the attack’s success. In our attack, we propose a strategy that aims to detect and override the main features of the victim image with a dynamic trigger containing compressed features of the target class, as described in Section 3.3. Empowered by the generated triggers, we hence assume that the attacker controls a percentage of the clients to poison the model during the Federated Transfer Learning process.

3.2 Threat Model

Attacker Knowledge: the adversary has access to the public dataset D_p , the public model M_p provided by the server, and a private dataset D_i . Additionally, the adversary has access to the gradients and the feature maps of the model’s layers.

Attacker Capabilities: using the private dataset D_i , the adversary is able to fine-tune the public model locally. However, the model’s shallow layers, i.e., the convolutional and the max-pooling layers, are frozen, so only the fully connected layers can be altered.

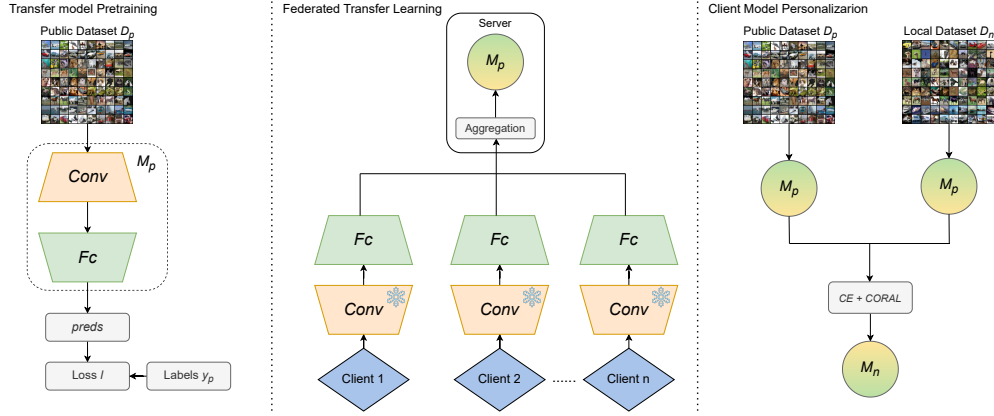


Figure 1: Federated Transfer Learning Framework.

Attacker Goal: by altering samples from the private dataset D_i , the attacker aims to inject a backdoor in the local model M_i . This backdoor cannot depend on the frozen layers, so a sophisticated trigger design is needed. After the backdoor is injected into M_i , the server will aggregate the local updates from the clients and update the unfrozen layers of the public model M_p . This updated model will be sent back to the clients for further fine-tuning. The adversary’s end goal is, hence, to embed a backdoor that remains active after the completion of the Federated Transfer Learning. The backdoor can then be used to control the model’s behavior through malicious inputs containing the trigger and force it to misclassify them to the target class.

3.3 FB-FTL Attack

As discussed in Section 3.1, the frozen convolutional layers *Conv* are capable of extracting general low-level features from images that can be inserted as triggers in a victim image to perform a backdoor attack. Our approach, shown in Figure 2, does not just add a trigger in a fixed position into the image. Instead, the idea behind this attack is to override the features of the original class with the ones of the target class. To do so, an XAI approach that highlights the important regions in a target image for the given model is needed. In particular, for this intent, we use GradCam [25]. GradCam obtains the class-discriminative localization map for any class c by computing the gradients for the given score of the class in the output y^c concerning the feature maps A^k of a convolutional layer. The obtained gradients are then global-average pooled to obtain the neuron importance weights a_k^c as follows:

$$a_k^c = \frac{1}{Z} \sum_i \sum_j \frac{\partial y^c}{\partial A_{ij}^k}. \quad (7)$$

These weights represent the importance of the features in the feature map for a target class c . After this, a weighted combination of forward activation maps is performed and sent in input to a *ReLU* activation function to obtain a heatmap of importance I_m

with the same shape as the feature maps:

$$I_m = \text{ReLU} \left(\sum_k a_k^c A^k \right). \quad (8)$$

The *ReLU* activation function is applied to the linear combination of maps because we only consider the features with a positive impact on the class of interest. Negative pixels, instead, are likely to belong to other classes in the image [25]. Without this *ReLU*, localization maps could emphasize more than just the desired class region, causing a lower localization performance. The obtained map I_m can now be used as a mask to select the regions of the image where the trigger should be injected. In particular, we select the regions of the images where the weights of importance a_k^c are higher than a given threshold τ . Since the weights of importance a_k^c fall in the interval of values $[0, 1]$, τ can be picked in the same interval. To inject a trigger capable of overriding the features of the original image with the main features of the target class, we use the intuition of the dataset distillation [36]. Specifically, the synthetic trigger is distilled using the Dataset Condensation technique with gradient matching. We take a subset of images B_i for which we generate the GradCam heatmaps of the same size as the images. Then, we set to 0 the regions of the heatmaps lower than the given threshold τ and to 1 the zones with a higher value, thus generating a mask. The obtained mask can now be used to inject a trainable synthetic trigger initialized as random noise into the best location of the target images. Our approach proceeds by exploiting the pre-trained transfer model M_p , for which we keep the parameter frozen for the entire process to learn the correct trigger. At each iteration, we feed at first the images B_i , and we obtain the gradients ∇W_p^B on the model in relation to target class y^t using the cross-entropy loss. In a second step, instead, we repeat the process with a set of true images T_i of the target class from the public dataset used to pre-train the model M_p . In the same way, we generate the related gradients ∇W_p^T on the parameters of M_p . The idea is to inject into the trigger the generalized features from the data of the target class used to train the transfer model. Through a distance metric, in our case, the cosine similarity, we calculate the loss on the distance between the generated gradients ∇W_p^B and ∇W_p^T , and we backpropagate on

the trainable trigger injected in B_i as follows:

$$\nabla W_p^B \leftarrow CE(M_p(B_i), y^t); \nabla W_p^T \leftarrow CE(M_p(T_i), y^t) \quad (9)$$

$$B_i[I_m > \tau] \leftarrow COS(\nabla W_p^B, \nabla W_p^T). \quad (10)$$

The process is then repeated for multiple iterations.

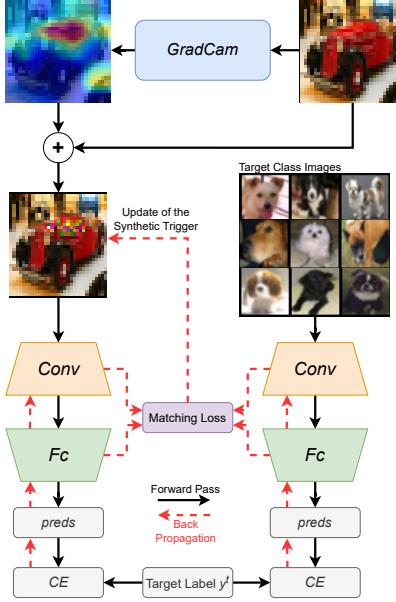


Figure 2: Trigger distillation.

3.3.1 Blended Trigger. The intuition behind our approach is, hence, to override the main features of an image, allowing its correct classification with a distilled trigger. This strategy would inevitably alter the original image and, most probably, the main visual content, as shown in Figure 3. This alteration can be detected using similarity metrics that compare the original image with its altered version. To synthesize a trigger that blends better inside the image making it less noticeable to automatic detection systems or to a human eye, we can include a term, derived from a similarity metrics capable of detecting such alterations, in the distillation matching loss.

Concerning such similarity metrics, one well-fitting option can be the Learned Perceptual Image Patch Similarity *LPIPS* [35]. We follow the idea of using the *LPIPS* metric from the Style Transfer and Generative Images approaches [16, 21]. These approaches employ a *LPIPS* loss due to its capability of providing better feature space than other similar solutions for improving the perceptual quality of the resulting image. In detail, using a pre-defined network like VGG or Alexnet, *LPIPS* compares the activations of two image patches to compute the similarity. This measure has proven to match also the real perception of humans [35]. A low *LPIPS* score means that the two images are perceptually similar. With that said, our trigger distillation loss can be changed as follows:

$$B_i[I_m > \tau] \leftarrow COS(\nabla W_p^B, \nabla W_p^T) + \lambda LPIPS(B_i, O_i), \quad (11)$$

where λ is the weight to module the contribution of *LPIPS* on the overall loss, and O_i is the original image without the trigger.

In the standard implementation of our approach, the pixels that compose the trigger are reinitialized as random Gaussian noise inside the original image O_i . A possible variation to improve the *LPIPS* value even further is to back-propagate the combination of the matching loss and *LPIPS* to the original pixel of the images without masking with random noise. In this case, the process changes as follows:

$$BO_i[I_m > \tau] \leftarrow COS(\nabla W_p^B, \nabla W_p^T) + \lambda LPIPS(BO_i, O_i), \quad (12)$$

where BO_i is the backdoored image with the trigger initialized with the same values as the original one. As we can see, the approach is the same, except for the initialization of the trigger’s pixels.

In Section 4.6, we present the comparison of our approach performance according to the considered combination of losses and initialization strategies of the trigger.

4 EXPERIMENTAL RESULTS

4.1 Experimental Setup

In our experiments, we focused on a testbed where we considered a Federated transfer learning framework composed of one server and multiple clients. In particular, in the baseline experiments, the number of clients is set to 10, and the percentage of attackers is set by default to 30%. We present additional results by changing the number of clients, in particular $\{5, 10, 100\}$. The considered main task is image classification. The considered baseline transfer model used as the feature extractor is a *Resnet-18* [13], kept frozen during the transfer learning process, with a *MLP* classifier on top. To showcase the generalizability of our approach, we also tested it with two additional architectures: *ConvNet* [11] and *VGG16* [26]. The different architectures according to the dataset are reported in Table 1. The split of the datasets between train and test sets is left on default if available; otherwise, we considered 80% of the data for the training and 20% for testing. From the training set, 50% is kept as public data to pre-train the model and the remaining half is split across the clients to be used as a private set. The pre-training of the model has been conducted using the public data for 10 epochs with a learning rate of 0.1 for the first five epochs and 0.01 for the remaining five. From our experimental campaign, this training strategy has proven to train the model with the partial dataset as the design of the framework and get the best accuracy. The model is trained from scratch instead of restoring the ImageNet weights to be consistent with the scenario in which the model must be pre-trained only on public data [8]. During the Federated Learning process, the classification layers are trained with a learning rate of 0.1. The matching loss employed to distill the trigger is the *Cosine Similarity* as presented in Section 3.3.

In the following sections, we present the attack’s accuracy with the best value of importance threshold τ chosen between the set of possible values, specifically $\{0.5, 0.6, 0.7, 0.8, 0.9\}$. We compare our solution with well-established backdoor attacks for HFL to understand if they are still effective in this configuration. In particular, we compared our solution with three static patch backdoor attacks: pixel pattern trigger, square trigger, and watermark trigger. In addition, we compared our solution against a novel approach that uses a trainable dynamic trigger: A3FL [34].

Moreover, we tested our solution with different percentages of attackers between the clients ($\{10\%, 20\%, 30\%, 40\%, 50\%\}$) to see if the effectiveness of our attack relies on the number of malicious clients. In Section 4.5, we will test the importance of the use of GradCam to spot the right position of the trigger compared to a squared dynamic trigger distilled in the same way as ours but with fixed coordinates in the corner of the image.

In Section 4.6, we redefine the learning loss of the dynamic trigger, combining with the original matching loss a similarity loss that preserves better the characteristics of the original image, making the trigger less noticeable but at the same time preserving most of the backdoor attack accuracy. In particular, we selected the *LPIPS* loss function that measures the perceptual difference between the original images and their manipulated counterparts with the injected triggers.

For our experiments, we used a machine with an AMD Ryzen 5800X CPU paired with 32GB of RAM and an RTX 3070ti with 8GB of VRAM. Our attack has been developed using PyTorch 2.0.

Table 1: Baseline model architectures. *MLP-N* refers to Multi-Layer Perceptron with *N* layers

Dataset	<i>Conv</i>	<i>Fc</i>
	Feature Extractor	Classifier
CIFAR10	Resnet-18	MLP-4
CINIC10	Resnet-18	MLP-4
SVHN	Resnet-18	MLP-4
GTSRB	Resnet-18	MLP-2

4.2 Performance Evaluation of FB-FTL

In this section, we present the main results in terms of the success rate of our approach compared to existing solutions. With this experiment, we want to investigate the applicability of backdoor attacks to this FTL. Since no other attack in this setup is available, we decided to compare our solution against attacks that are intended for HFL due to their similarities in terms of the configuration of the training between server and clients. In both cases, the clients train the global model on their local data while the server aggregates compared to VFL, in which the global model is split between the top model on the server and the bottom models on the clients. In addition to traditional static triggers, we compare our approach also against a novel solution, named A3FL [34], that injects in the image a dynamic trigger specifically generated for the target class. Since one of FTL’s main properties is freezing the feature extractor layers, we expect that traditional backdoor attacks will be ineffective in this scenario. The considered attacks introduce into the image the triggers with additional features that are unknown to the feature extractor network. Considering the A3FL solution, we expect better performance than the traditional static backdoors. However, we are still unable to match the performance of our attack due to the absence of the focus logic with GradCam.

From Table 2, as expected, the existing attacks are almost unperceived by the federated model. The A3FL approach, instead, achieves a better success rate than static triggers while still affecting the model accuracy on the main task like the traditional one.

These results also give some insights into the importance of focusing logic with GradCam. In Section 4.5, we present an in-depth analysis in this sense. Looking at our approach FB-FTL, we can see it is effective across the considered datasets. It is interesting to observe how our approach is capable of achieving high success rates despite the different characteristics of the datasets presented in Appendix B. Even with datasets like SVHN, where the background contains information about other classes, our attack distills a trigger that preserves the importance of the original features compared to the ones in the background. Considering GTSRB, instead, since contains images of traffic signs, the samples shares many features and in this case, our attack distills into the trigger the right ones selecting the features belonging just to the target class. We can see also how our approach is the best at being effective and, at the same time, only slightly affects the performance of the model on the main task. This is because the feature extractor part of the model will recognize only the known features, so with traditional backdoor attacks, the network will still focus just on the information of the original class, discarding the trigger. Basically, the attacker is training the model, assigning to these images the target class, drifting the accuracy of the model on those features, affecting the final accuracy. In our case, instead, we are substituting the original features with information from the target class distilled from the network itself. In this way, the network will recognize the new features as known, only slightly affecting the accuracy of the main model.



Figure 3: Examples of distilled triggers

In Figure 3 and Appendix A, we show some examples of triggers generated for different classes across the considered datasets. The position of the trigger changes according to the region of importance spotted by the GradCam algorithm. In Section 4.5, we present an ablation study that demonstrates the importance of finding the best region of the image where to distill the trigger instead of having fixed coordinates. On the other hand, this approach also owes its success to the overriding of the main features of the image using GradCam to position the trigger. This makes inevitable, in most of the images, an alteration of the main subject of the image. In Section 4.6, we present the result of our approach, adding the contribution of the *LPIPS* metric in the loss calculation to make the distilled trigger less noticeable to the human eye.

4.3 Results Changing Percentage of Attackers

In this section, we assess how the performance of our attack changes according to the percentage of attackers between the clients. As

Table 2: Main results

Backdoor Attacks	CIFAR10		CINIC10		SVHN		GTSRB	
	Model Accuracy	Backdoor Success Rate	Model Accuracy	Backdoor Success Rate	Model Accuracy	Backdoor Success Rate	Model Accuracy	Backdoor Success Rate
No Attack	81.8%	-	71.0%	-	90.0%	-	95.0%	-
Pattern Trigger	77.9%	13.3%	67.4%	13.5%	87.4%	8.7%	93.4%	4.7%
Square Trigger	75.4%	17.6%	67.7%	11.6%	87.3%	7.8%	92.6%	5.4%
Watermark Trigger	76.5%	19.3%	67.8%	12.9%	88.3%	9.7%	93.9%	4.1%
A3FL [34]	73.1%	31.2%	65.5%	24.3%	87.0%	10.0%	94.0%	10.1%
FB-FTL (our)	81.3%	91.1%	70.2%	73.3%	89.6%	72.1%	94.7%	86.5%

we stated in Section 4.1, the previous results were collected by setting the percentage of attackers equal to 30%. In this experiment, due to the nature of our attack, which relies mainly on the frozen *Conv* layers of the model, we expect that the success rate of the attack would not drastically change from the baseline reported in the previous section. Thus, we aim to confirm our intuition that the success of our attack is only partially correlated to the number of attackers that poison the local updates.




Figure 4: Performance of the attack changing the percentage of attackers.

In Figure 4, we show the results of this experiment. As we can see, with three out of four datasets, we can confirm our intuitions. Indeed, the success rate of the attack is comparable across the different scenarios with negligible fluctuations. The only dataset that benefits from a high percentage of attackers is GTSRB. With the percentage of attackers of 50%, the success rate is close to 100%. If we look, instead, at the worst scenario, the accuracy of the attack is a little lower than the baseline but still with a not negligible success rate of around 60%. This is expected due to the similarities between classes of the dataset requiring for the attacker to contribute to the training of the federated model with the backdoored images to make the attack more effective.

The previous experiment was conducted with a total number of number of clients equal to 10. In this sense, we conducted a second experiment varying the number of clients included in the federated process. In particular, we want to verify how our attack varies, considering just one malicious client but changing, in this case, the number of clients. In the first scenario, we considered half of the clients compared to the baseline, 5 clients in total, and the second scenario includes ten times more clients, specifically

100. Since our attack relies on the characteristic that the feature extractor in the considered Federated Transfer Learning scenario is frozen, we expect that our attack is just partially dependent on the total number of clients in the Federated Learning process.




Figure 5: Performance of the attack changing the total number of clients.

The results of this experiment are shown in Figure 5. As anticipated, we can notice the difference between 10 and 100 is just 5% at maximum, preserving a success rate over 60% across all the datasets. Instead, looking at the scenario with 5%, we can see how the performance is better than considering 10% as expected, but similarly to the previous use case, the difference is almost negligible. The highest recorded difference is about 15% for the GTSRB dataset with 5 total clients compared to the baseline of 10. Even if the proportion between malicious and benign clients is doubled with 5 overall clients compared to 10 considering just one attacking client, the success rate of our attack is not two times higher with double the importance of the malicious client contribution in the averaging of the updates with 5 total clients. These results allow us to confirm that our solution is independent of the percentage of attackers between the parties and the total number of clients participating in the federated training.

4.4 Evaluation Changing τ Threshold

In this section, we present the results of the experiments varying the τ threshold used to select the region of the trigger according to the weights of importance returned by GradCam as presented in Section 3.3. In particular, changing this value will affect positively or negatively not only the performance of our attack but also the integrity of the original image. Intuitively, if we set τ to a low value,

we will inject a trigger that occupies a large portion of the image. Since the aim of our attack is to distill an effective trigger but at the same time preserve most of the content of the image, the τ value is a parameter that has to be tuned. In particular, the idea is to select the highest value of τ that preserves the performance of the model while altering as little as possible the original image.




Figure 6: Backdoor success rate changing τ value.

Results of this experiment are reported in Figure 6. As expected, with high values of τ , e.g., 0.8 or 0.9, the performance of our approach drops. This is because, with high values, the region dedicated to the trigger is too small to embed the features necessary to make our approach successful. At the same time, the region is not big enough to override the main features of the image.

In Section 4.5, we prove this by showing the importance of the positioning of the trigger to cover the main features of the original image. Looking at lower values of τ , instead, we can see how having a smaller τ , and by consequence, a larger trigger, does not necessarily benefit the accuracy of the model. In this situation, because there is a bigger region to distill, our attack has more space to include more refined features about the target class. Since dataset distillation works by providing the network examples of the target class, having a bigger region to distill makes our trigger prone to overfit too much on example images. In this way, the trigger will include information specific to those particular images in addition to the more general features of the class. This additional information breaks our assumption of distilling triggers with only general features of the target category, making them less relevant and affecting the attack’s performance.

4.5 GradCam Importance Analysis

As presented in Section 3.3, GradCam is one of the main components of our approach. Dynamic triggers alone, as we showed in Table 2, are not enough to produce backdoor attacks effective in the Federated Transfer Learning Scenario. In this section, we present an analysis to assess the importance of an explainability technique, in our case, GradCam, to spot the most important features of the original images to be covered by our distilled trigger.

In Figure 7, we present examples of the GradCam results on the original set of images and the respective backdoored version. As we can see from the heatmaps produced by Gradcam on the original

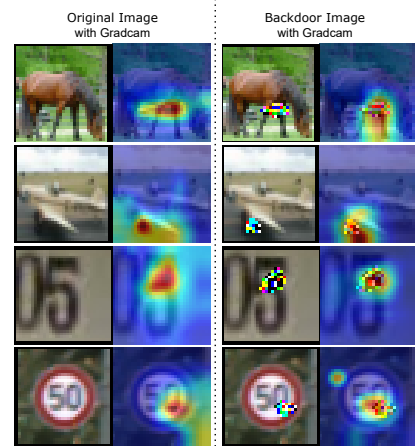


Figure 7: Examples of GradCam maps with and without the trigger

image and the same with the addition of the trigger, the model to classify the image focuses on the same region. This shows how our trigger successfully substitutes the main features of the original class as the most important portion of the image for the model to classify it. Another interesting result we can observe is the area of importance detected by the GradCam algorithm. The different shades of colors presented in Figure 7 represent the importance of those features for the classification of the image. The regions with the tones red and yellow are actually the ones that matter the most for the model to perform the classification. Looking at the examples in Figure 7, we can see how the produced heatmaps on the original images highlight, of course, the most important region in red but also show how the model looks at regions around in yellow to take a decision for the classification of the image. The results produced on the backdoored images, instead, show an important region that is much more concentrated around the trigger. This proves that the combination of the distilled trigger and explainability can confidently drive the target model toward the target class with more confidence compared to the original class.

To better prove our intuition, we performed an ablation study to test the different components of our approach. In particular, we tested the components of our solution in three different configurations. In the first scenario, the trigger is distilled in a fixed position like traditional dynamic solutions. In the second case, the trigger is still distilled in fixed coordinates, but, differently from the previous scenario, we use GradCam to detect the most important features of the images, and we override them with random noise. The final scenario, instead, is our full solution.

In Table 3, we report the result of our ablation study. As expected, the dynamic trigger in a fixed position is not capable of producing an effective backdoor attack, confirming what we have already reported in Section 4.2. An interesting result can be seen in the scenario in which we obfuscate the features of the original class. As we can see, the success rate of the attack in this scenario improves compared to the one in which we just distill the trigger in a fixed position of the image. Especially for the SVHN dataset, we can

Table 3: Experiments with Dynamic Patch in a fixed position with and without obfuscation of the main features

Dataset	Dynamic Corner Patch	Dynamic Corner Patch With Obfuscation	Our Trigger
CIFAR10	32.3%	40.0%	91.1%
CINIC10	22.3%	26.7%	73.3%
SVHN	7.8%	23.4%	72.1%
GTSRB	18.9%	33.3%	86.5%

see how the performance of the attack improves three times more after masking the original information of the images. This result proves how covering the main features of the images and moving the distilled trigger in the right position contribute to deceiving the frozen feature extractor of the transfer model.

4.6 Experiment on Perceptual Similarity Loss

As we introduced in the methodology section (Section 3.3), the positioning of our trigger over the main features of the images can alter too much the original data, making it more detectable even by automatic systems that use similarity metrics. This section is devoted to better understanding how the changes in the loss and trigger initialization described in Section 3.3.1 impact the perceptual similarity between the original and backdoored images using the *LPIPS* metric presented in the same section.

As we can see in Table 4, the positioning of the trigger over the original content of the image actually perceptually impacts the most compared to other backdoor attacks that use a static trigger positioned in a peripheral position of the image without impacting the main subject. With the modified version of the loss, we want to preserve the attack performance, reducing the *LPIPS* metric.

Table 4: *LPIPS* values of the images with different triggers

Trigger Type	CIFAR10	CINIC10	SVHN	GTSRB
Square Trigger	0.0015	0.0007	0.0058	0.0093
Pattern Trigger	0.0006	0.0002	0.0022	0.0031
Watermark Trigger	0.0056	0.0020	0.0188	0.0174
A3FL [34]	0.0020	0.0014	0.0106	0.0175
FB-FTL (our)	0.1346	0.0618	0.2903	0.1346

In this experiment, we tested the success rate of our approach in normal conditions compared to the more visually conservative version of it. In particular, we tested the modified version of the original loss with the addition of the *LPIPS* component combined with random initialization of the trigger and initialization with the original values of the target image. The results are presented in Table 5. As we can see, with the more conservative loss, we lower the *LPIPS* values while preserving most of the performance of the original attack. Looking at the SVHN datasets, it is interesting to see how the *LPIPS* version of the loss combined with trigger initialization from the original values of the image performs better than random initialization and even better than the standard version of the attack. This can be related to the nature of these datasets, as described in Appendix B. In particular, SVHN is characterized by images with backgrounds that can contain features from classes

Table 5: Backdoor Success Rate (BSR) and *LPIPS* Value changing the loss and trigger initialization

Dataset	Normal/Noise		<i>LPIPS</i> /Noise		<i>LPIPS</i> /Original	
	BSR	<i>LPIPS</i> Value	BSR	<i>LPIPS</i> Value	BSR	<i>LPIPS</i> Value
CIFAR10	91.1%	0.1346	84.4%	0.0308	82.2%	0.0294
CINIC10	73.3%	0.0618	70.0%	0.0087	67.8%	0.0076
SVHN	72.1%	0.2903	70.0%	0.1036	76.7%	0.0919
GTSRB	86.5%	0.1346	81.1%	0.0375	78.7%	0.0363

different from the principal one. Adding a trigger too different from the overall image could make the model focus more on the background containing features belonging to categories different from our target class. A more blended trigger, in this case, allows us to improve even more the performance of our attack.




Figure 8: Examples of triggers generated with different loss and trigger initialization

In Figure 8, we show examples of triggers generated with different combinations of losses and initializations. We can notice that the addition of *LPIPS* to the basic loss allows the generation of triggers that blend better with the overall image. The colors of the pixels of the triggers better match the palette and the intensity of the colors of the original image compared to the basic implementation of the dynamic triggers. Looking at the last column, the initialization of the trigger plays an important part. As we can see, the generated trigger blends better with the main subject of the image, preserving some of the original pixels that do not need to be updated and resulting in an almost transparent trigger. The contribution of the initialization strategy is to make the trigger even less noticeable to the human eye compared to other backdoor attacks in which the trigger is easily detectable by a human check.

4.7 Evaluation on Different Architectures

In this experiment, we want to evaluate our approach with additional model architectures to prove the generalizability of our approach. As mentioned in Section 4.1, we selected two models,

Table 6: Evaluation on different architectures

Dataset	Resnet18	VGG16	ConvNet
CIFAR10	91.1%	86.7%	93.3%
GTSRB	86.5%	85.5%	90.0%

namely ConvNet [11] and VGG16 [26]. The choice of these two architectures has been made by looking at the architectures studied in the Dataset Condensation paper [36], which inspired our distillation strategy, and the GradCam paper [25]. For the evaluation of these two architectures, we selected two of the previously presented datasets. In particular, CIFAR10 as a baseline, and GTSRB since it is the dataset with the highest number of labels.

In Table 6, we present the success rate of our approach with the two additional architectures. Our approach, in both cases, preserves most of the accuracy of the considered baseline. As expected with a deeper network and more parameters like VGG16, the success rate is only partially affected with a decrease of just around 5%.

5 DEFENSES EVALUATION

In this section, we present the performance of our attack against possible defenses. Since there are no defenses available for the FTL scenario, we borrowed countermeasures from the other two Federated Learning settings. First, due to the similarity between the two settings, we start by considering HFL defenses. In particular, we selected two standard ones, Krum and Trimmed Mean, and two more advanced ones, FoolsGold and Flame. The results on these are presented in Section 5.1.

Considering, instead, the VFL scenario, it is difficult to adapt well-established solutions (e.g., Privacy-Preserving Deep Learning or DiscreteSGD) due to the difference between VFL and TFL. Still, some ideas from a particular family of countermeasures for the Vertical setting can be borrowed for the initial pre-training of the transfer model. Specifically, solutions from the Label Differential Privacy (LabelDP) family can be exploited to pre-train the transfer model using noisy labels. In this case, we considered a novel solution named KDk [3] that generates noisy labels using knowledge distillation, and the results are presented in Section 5.2.

With this study, we want to understand the weaknesses of our attack and develop a working countermeasure designed properly for the Federated Transfer Learning scenario.

5.1 Horizontal Federated Learning Defenses

As we said before, we tested our attack against four different defenses: Trimmed Mean [33], Krum [6], Flame [24], and FoolsGold [10]. In the first approach, the server independently averages the gradients according to their position. Specifically, the aggregator sorts the gradients in the same position according to their distance from the median. Then, only the first k parameters are considered benign, where $k = n - m$, n is the number of clients, and m is the malicious portion. To be consistent with the original implementation of the defense, we consider the ideal scenario in which the portion of malicious clients is known. Krum, instead, averages the global model by selecting the best updates between the gradients received from the clients and discarding the outliers that differ significantly from their average.

The Flame defense is composed of two main components: the first filters the clients’ updates, and the second adds Adaptive Differential Privacy. The first component filters malicious clients from the ones with the highest probability of being honest. This is done by clustering the clients using HDBSCAN over the pairwise cosine similarity distances among the updates and keeping only the cluster that includes at least 50% clients. The second component applies an adaptive differential privacy approach estimating the clipping bound and level of noise, in a way that neutralizes the attack and preserves the original performance of the model.

FoolsGold adjusts the contribution of each client based on the similarity distance of the updates, similar to Flame, and also considers information derived from past iterations. The approach leverages cosine similarity to measure the distance between updates. Usually, backdoor attacks alter specific features, which can be identified by measuring the magnitude of model parameters in the output layer of the global model. The malicious updates can be removed or re-weighted.

In Figure 9, we present the result of our attack considering different values of τ against the selected defenses. Overall, as expected, our attack can pass the filtering of the defenses, preserving the same performance in most cases. This is due to the fact that our triggers are distilled to include the most general features of the target class and override the one of the original class. As a result, our attack generates updates that are only slightly different from the distribution of the benign ones, making it difficult to detect them.

Interestingly, we noticed two exceptions, one in favor of and one against our solution. Starting from this last one, we can notice how our attack is penalized against three out of four defenses for the GTSRB dataset. As we stated in Appendix B, this dataset is characterized by many more categories than the other datasets, 43 in total, that share many features, resulting in a very small separation between feature distributions across the classes. The narrower distribution of the updates makes our malicious updates, even with small differences, more easily detectable as outliers by the defenses. In this scenario, the defenses are capable of mitigating the attack but still are not enough to neutralize scoring, still, an accuracy above 50% with lower values of τ .

The second exception, instead, the one in favor of our solution, is in relation to the results across all the datasets with the employment of FoolsGold as a defense. As we can see, with all the datasets, our approach archives results in terms of success rate even better compared to the baseline without the defense. FoolsGold, by design, gives more importance to the features that are beneficial for the accuracy of the final model. As a result, it will preserve the updates that are closer to the mean of the distribution, penalizing the outliers. By definition, our approach aims to distill into the trigger the most general features of the target class, resulting in updates that, on purpose, want to deviate as little as possible from the mean of the benign updates. To conclude, we demonstrate that defenses intended for Horizontal Federated Learning are effective only in particular cases with datasets with specific characteristics. Moreover, in some situations, they can be beneficial for our attack.

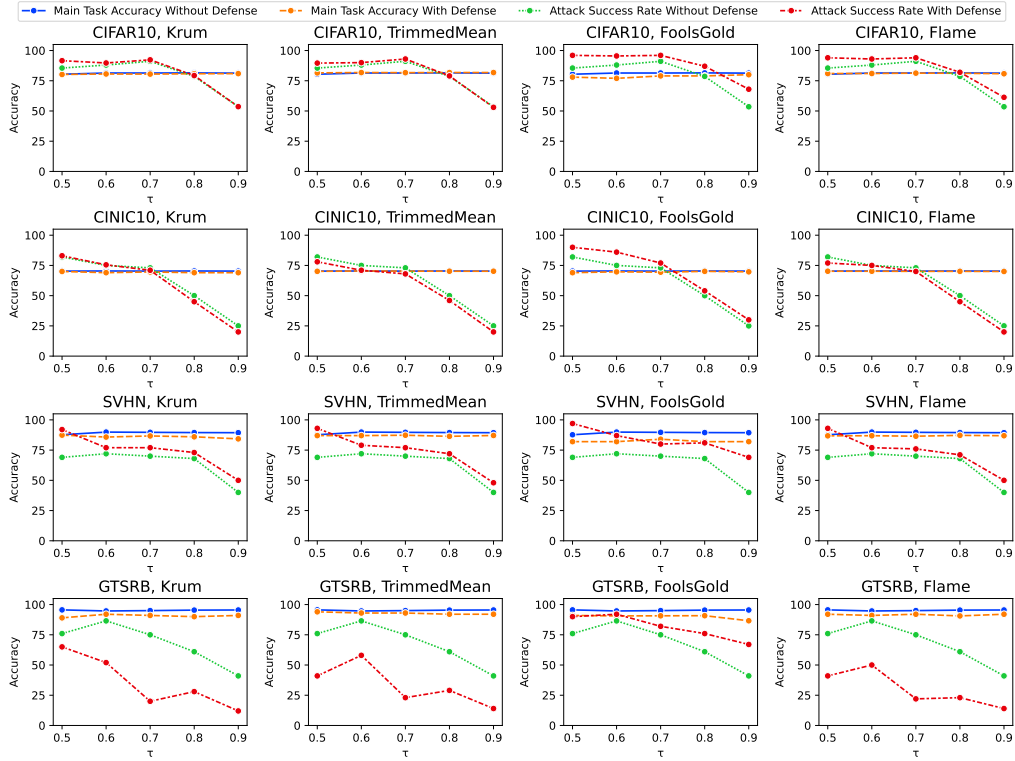


Figure 9: Possible defenses.

5.2 Vertical Federated Learning Defense

As we have mentioned before, it is challenging to adapt most of the traditional VFL defenses to our scenario. Instead, we can borrow some ideas from LocalDP for the pre-training of the transfer model. By default, the model is trained with one-hot encoded labels, and it has been proven that this could lead to overfitting problems. To overcome this issue, strategies of label smoothing have been proposed. Label smoothing [28] approaches redistribute the probability of the main class across other secondary classes to instill in the model the knowledge about similarities between classes, making it more robust. Knowledge Distillation is the perfect candidate to generate a smooth version of the labels. In practice, Knowledge Distillation uses the output of a pre-trained teacher model to generate smooth labels for each sample. In this sense, we selected the KDk defense to add a level of uncertainty to the transfer model.

The KDk [3] defense is a novel approach intended to preserve the privacy of the labels in the Vertical Federated scenario. This solution is based on the concept of k - *anonymity* between the classes. In particular, the knowledge distillation is used to select the top- k classes for each image, and then the probability of the main classes is redistributed to the selected top k categories using an ϵ parameter to tune the re-balancing and the strength of the defense. In particular, the ϵ value is subtracted from the probability of the main class, originally set at 1 in the one-hot setting, and redistributed equally across the $k - 1$ secondary labels. Pushing the k and ϵ parameters to extreme values will benefit the security of the model at the expense, though, of its performance on the

main task. The tuning of these parameters is important to select the maximum power level without affecting the accuracy of the model. With this strategy, the features will be learned by the network as less separated between the classes compared to using one-hot encoded labels. We expect that this additional layer of uncertainty can affect the performance of our attack.

Since our attack is based on distilling into the triggers the general features of the target category, these blurred boundaries between the features of the classes could lead our trigger to also include information from other categories affecting the success rate. As we said, the selection of the k secondary classes and ϵ value is important to preserve the accuracy of the model. In this experiment, we selected the highest values for both parameters for each dataset that preserve most of the models' accuracy just before they start degrading. The selected parameters are reported in Table 7.

Table 7: KDk Parameters

KDk Parameters	CIFAR10	CINIC10	SVHN	GTSRB
k value	3	3	3	5
ϵ value	0.45	0.4	0.3	0.6

In Table 8, we present the results of our attack against the KDk defense. Our attack is capable of preserving most of its performance on three of the selected datasets, but at the same time, all of them, even if slightly, are affected by the defense compared to the one for the Horizontal scenario. What makes our attack resilient

Table 8: FB-FTL results against the KDk defense

Dataset	No Defense		KDk[3]	
	Main Task Accuracy	BSR	Main Task Accuracy	BSR
CIFAR10	81.3%	91.1%	80.1%	84.4%
CINIC10	70.2%	73.3%	69.7%	61.1%
SVHN	89.6%	72.1%	89.4%	32.2%
GTSRB	94.7%	86.5%	94.2%	82.2%

is the distillation logic behind the generation of the trigger. Our strategy tries to distill images by providing examples instead of backpropagating to the information of the target class using just the cross-entropy loss. This makes our solution more robust even if the transfer model has been trained with a level of uncertainty on the separation between the features of the classes. Similar to what happened with GTSRB against HFL defenses, also in this case, one of the datasets, SVHN, is affected the most by the defense due to its characteristics. As we said in Appendix B, the images of this dataset are characterized by backgrounds that contain information about other classes compared in addition to the principal one. Using the distilled soft labels, the KDk defense is training the model to give more importance to the elements in the background mitigating our strategy of overriding the features on the main class.

These results give us insight into how to define a proper defense for our backdoor attack in the Federated Transfer Learning scenario. In particular, the server should apply an ad-hoc LabelDP strategy for the pre-training of the transfer model instead of focusing on the updates of the clients like the HFL defenses.

6 RELATED WORK

Federated Learning was proposed by researchers in Google [22]. This work introduced the HFL paradigm, where all the clients use data from the same feature space but may not be of the same sample identity. In VFL, the clients own data that belong to the same samples but do not share the same feature space [18]. FTL is applied to scenarios in which the clients’ datasets are different both in the feature space and the samples identities [32]. Given that such scenarios are more general, they are more practical and applicable to real-world applications. For example, FedHealth [8] uses FTL to support personalized healthcare through wearable devices. To this end, first, the users’ devices collaborate to train a shared model in the cloud for human activity recognition with privacy guarantees, and then each participant uses FTL to personalize the model and improve its performance based on the local dataset. For the personalization, the authors assumed that in each client, the first part of the downloaded model is frozen, and only the fully connected layers are fine-tuned with the private dataset. We adopted this model in our experiments as the frozen feature extractor that each client has makes traditional backdoor attacks more challenging. The adversary cannot use triggers that are small features [12] and are based on the model’s feature extractor, as they can affect only the last layers of the model. As a result, more elaborate strategies need to be designed.

Backdoor attacks in deep learning were first introduced by Gu et al. [12]. In this attack, an adversary alters some training samples to insert a secret functionality into the trained model that is activated

during inference. This attack has become very popular and has also been applied to Federated Learning [2, 5, 29–31]. In [5], a model replacement attack was introduced where the adversary attempts to replace the global model with a malicious one that contains a backdoor by amplifying the gradient updates of the adversaries. In [29], the authors established a theoretical framework to verify that a model is vulnerable to backdoor attacks if it is also vulnerable against evasion attacks in FL. Xie et al. [30] split the trigger among multiple malicious clients, making the first distributed backdoor attack in the Federated Learning setup. The backdoor can be activated with both the local triggers but also with the global trigger, which is a combination of all the local triggers. Xu et al. [31] used this technique to backdoor federated graph neural networks. All these attacks, however, have been applied to HFL, where the adversary has full control of the training of the local model so every layer can be affected. In this work, the adversary cannot alter the feature extractor of the model, making the trigger generation more challenging. To the best of our knowledge, we are the first to explore backdoor attacks in this scenario.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we describe, to the best of our knowledge, the first focused backdoor attack specifically designed for the Federated Transfer Learning Scenario. Our attack, called FB-FTL, aims to overcome the challenges imposed by the scenario. In this context, a feature learning step is carried out by the server before the actual federated learning, which then aims at collaboratively train only the classification layers by keeping frozen the feature extractor ones. Due to this characteristic, existing backdoor attacks, especially the static ones that introduce unknown features to the feature extractor, are ineffective in this scenario. Moreover, traditionally, backdoor attacks introduce a trigger in the original data without applying any strategy to minimize its impact on them. For instance, in the case of image data, the position of the trigger may strongly influence the visualization outcome.

To design our attack, we started by analyzing the FTL framework to understand its vulnerabilities and to exploit them. The main intuition behind our approach is to maximize the success of our attack by identifying the optimal spot in the input data to position the trigger. To do this, we leverage an XAI strategy to identify the most important regions of the images for their classification. The use of explainability to locate the trigger over features detectable by the feature extractor component of the FTL model brought to the definition of the first focused backdoor attack. Another important aspect of our attack is related to the way in which the trigger is generated. For this purpose, we borrowed some ideas from the dataset distillation field. In particular, the attack distills the main features of the target class into the trigger. Combined with the focusing strategy, it overrides the information of the original class with the desired one. In our experimental campaign, we analyzed the different components of our approach, showing the importance of distilling the target class features and correctly positioning the trigger in the input data. We demonstrated how our approach does not depend on the percentage of malicious clients in the process and is not heavily affected by the total number of clients. Since no

ad-hoc defenses are available for the considered scenario, we evaluated our attack against Horizontal and Vertical Federated Learning countermeasures. As we discovered, no defense is completely effective in mitigating our approach for all the datasets. In this sense, the most promising approach comes from the Vertical Federated family, named KDK, and makes use of Label Differential Privacy (LabelDP) to add uncertainty between the classes in the pre-trained model. In this direction, as future work, further exploration can be conducted to define a more effective countermeasure, still based on LabelDP, against our attack.

REFERENCES

- [1] Marco Arazzi, Mauro Conti, Stefanos Koffas, Marina Krcek, Antonino Nocera, Stjepan Picek, and Jing Xu. Blindsage: Label inference attacks against node-level vertical federated graph neural networks. *arXiv preprint arXiv:2308.02465*, 2023.
- [2] Marco Arazzi, Mauro Conti, Antonino Nocera, and Stjepan Picek. Turning privacy-preserving mechanisms against federated learning. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 1482–1495, 2023.
- [3] Marco Arazzi, Serena Nicolazzo, and Antonino Nocera. Kdk: A defense mechanism against label inference attacks in vertical federated learning. *arXiv preprint arXiv:2404.12369*, 2024.
- [4] Eugene Bagdasaryan and Vitaly Shmatikov. Blind backdoors in deep learning models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1505–1521, 2021.
- [5] 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.
- [6] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. *Advances in neural information processing systems*, 30, 2017.
- [7] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- [8] Yiqiang Chen, Xin Qin, Jindong Wang, Chaohui Yu, and Wen Gao. Fedhealth: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems*, 35(4):83–93, 2020.
- [9] Luke N Darlow, Elliot J Crowley, Antreas Antoniou, and Amos J Storkey. Cinic-10 is not imagenet or cifar-10. *arXiv preprint arXiv:1810.03505*, 2018.
- [10] Clement Fung, Chris JM Yoon, and Ivan Beschastnikh. The limitations of federated learning in sybil settings. In *RAID*, pages 301–316, 2020.
- [11] Spyros Gidaris and Nikos Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4367–4375, 2018.
- [12] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] Sanghyun Hong, Nicholas Carlini, and Alexey Kurakin. Handcrafted backdoors in deep neural networks. *Advances in Neural Information Processing Systems*, 35:8068–8080, 2022.
- [15] Qinghe Jing, Weiyan Wang, Junxue Zhang, Han Tian, and Kai Chen. Quantifying the performance of federated transfer learning. *arXiv preprint arXiv:1912.12795*, 2019.
- [16] Younghyun Jo, Sejong Yang, and Seon Joo Kim. Investigating loss functions for extreme super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 424–425, 2020.
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [18] Li Li, Yuxi Fan, Mike Tse, and Kuo-Yi Lin. A review of applications in federated learning. *Computers & Industrial Engineering*, 149:106854, 2020.
- [19] Yang Liu, Yan Kang, Tianyuan Zou, Yanhong Pu, Yuanqin He, Xiaozhou Ye, Ye Ouyang, Ya-Qin Zhang, and Qiang Yang. Vertical federated learning: Concepts, advances, and challenges. *IEEE Transactions on Knowledge and Data Engineering*, 2024.
- [20] Yang Liu, Zhihao Yi, and Tianjian Chen. Backdoor attacks and defenses in feature-partitioned collaborative learning. *arXiv preprint arXiv:2007.03608*, 2020.
- [21] Yuchen Liu, Zhixin Shu, Yijun Li, Zhe Lin, Federico Perazzi, and Sun-Yuan Kung. Content-aware gan compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12156–12166, 2021.
- [22] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282. PMLR, 20–22 Apr 2017.
- [23] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, page 7. Granada, Spain, 2011.
- [24] Thien Duc Nguyen, Phillip Rieger, Huili Chen, Hossein Yalame, Helen Möllering, Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Shaza Zeitouni, Farinaz Koushanfar, Ahmad-Reza Sadeghi, and Thomas Schneider. FLAME: Taming backdoors in federated learning. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1415–1432, Boston, MA, August 2022. USENIX Association.
- [25] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [26] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [27] J. Stallkamp, M. Schlipsing, J. Salmen, and C. Igel. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural Networks*, (0):–, 2012.
- [28] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [29] Hongyi Wang, Kartik Sreenivasan, Shashank Rajput, Harit Vishwakarma, Saurabh Agarwal, Jy-yong Sohn, Kangwook Lee, and Dimitris Papailiopoulos. Attack of the tails: Yes, you really can backdoor federated learning. *Advances in Neural Information Processing Systems*, 33:16070–16084, 2020.
- [30] Chulin Xie, Keli Huang, Pin-Yu Chen, and Bo Li. Dba: Distributed backdoor attacks against federated learning. In *International conference on learning representations*, 2019.
- [31] Jing Xu, Rui Wang, Stefanos Koffas, Kaitai Liang, and Stjepan Picek. More is better (mostly): On the backdoor attacks in federated graph neural networks. In *Proceedings of the 38th Annual Computer Security Applications Conference*, pages 684–698, 2022.
- [32] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- [33] Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *International Conference on Machine Learning*, pages 5650–5659. PMLR, 2018.
- [34] Hangfan Zhang, Jinyuan Jia, Jinghui Chen, Lu Lin, and Dinghao Wu. A3fl: Adversarially adaptive backdoor attacks to federated learning. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 61213–61233. Curran Associates, Inc., 2023.
- [35] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [36] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *International Conference on Learning Representations*, 2020.
- [37] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016.

A MORE EXAMPLES

In this section, we show more examples of triggers generated by our attack. In Figure 10, examples of the basic implementation of our attack are shown. In Figure 11 and 12, instead, we show examples of the *LPIPS* contribution in the loss and the initialization of the trigger with the original values of the images.

B DATASETS

To conduct the experimental campaign, we selected four of the most common benchmark datasets: CIFAR10 [17], CINIC10 [9], SVHN [23], and GTSRB [27]. The first two are datasets that include



Figure 10: More examples of triggers generated with the basic implementation.




Figure 11: More examples of triggers generated with the contribution of *LPIPS* in the loss.




Figure 12: More examples of triggers generated with the contribution of *LPIPS* in the loss and the initialization of the trigger with the original values of the image.

Table 9: Statistics of the considered datasets

Dataset	Train Set Size	Test Set Size	Number of Classes
CIFAR10	50,000	10,000	10
CINIC10	180,000	90,000	10
SVHN	73,257	26,032	10
GTSRB	39,209	12,630	43

classes of different animals or objects. Both share the same categories. The main difference is that CINIC10 is 4.5 times larger the size of CIFAR10. This difference in size allows us to prove that our approach is still effective even with larger datasets, proving our intuition that the trigger is distilled, preserving the most general and

meaningful features of the target class. The SVHN dataset, instead, contains images of houses' civic numbers with categories from 0 to 9. One of the main characteristics of this dataset is the presence of additional numbers in the background that do not belong to the assigned category. We selected this dataset to see if the features distilled in our trigger are strong enough to be preferred to the other number in the image by the model. The last dataset, the German Traffic Sign Recognition Benchmark (GTSRB), contains forty-three different categories of traffic signs. Since the traffic signs have many features in common between the categories (e.g., shape and color), we included this dataset to test whether our approach is capable of overriding the right features to control the model and predict the target class. Table 9 reports statistics for the considered datasets.