

A Discussions

Complexity. Clients will not experience additional computations as our DUW is carried out on the server side. The additional computation for the server is decided by the number of watermark injection steps T_w . We found that WSR could reach 99% just within $T_w = 10$ steps. Injection of one client-unique watermark takes around 1 second, which the server can do parallelly for each client model. Thus, the delay caused by the server is neglectable.

Future works. This paper makes the FL model leakage from anonymity to accountability by injecting client-unique watermarks. We recognize the most significant challenge for accountable FL is addressing watermark collision for accurate IP tracking (R1). We believe it is important to scale our method with more clients in the future. One plausible solution is increasing the dimension of the input of the encoder to allow more one-hot encoding target labels. Another solution is to use a hash function as the target label for different clients. In this way, the lower-dimensional encoder and decoder can accommodate more clients. For instance, an encoder with input dimension 10 can allow at most 1024 different clients. However, adopting hash functions as the target labels can increase the chance of watermark collision between clients, and more elegant strategies have to be developed to address this problem. As we focus on the collision, we leave the scalability for future work.

B Supplementary experiments

B.1 Extended qualitative study

Visualization of unique trigger sets. We show the visualization example for the original image, encoded image (image in trigger set), and residual image based on different OoD datasets in Fig. 6. We observe that for all four different OoD datasets, the original image and encoded image with our client keys are indistinguishable from the human eye. The difference between these two images can be observed in the residual image. Note that although the OoD datasets are different, the encoder that we used to generate the trigger sets is the same. According to Fig. 6, the encoder will generate sample-wise triggers for different images.

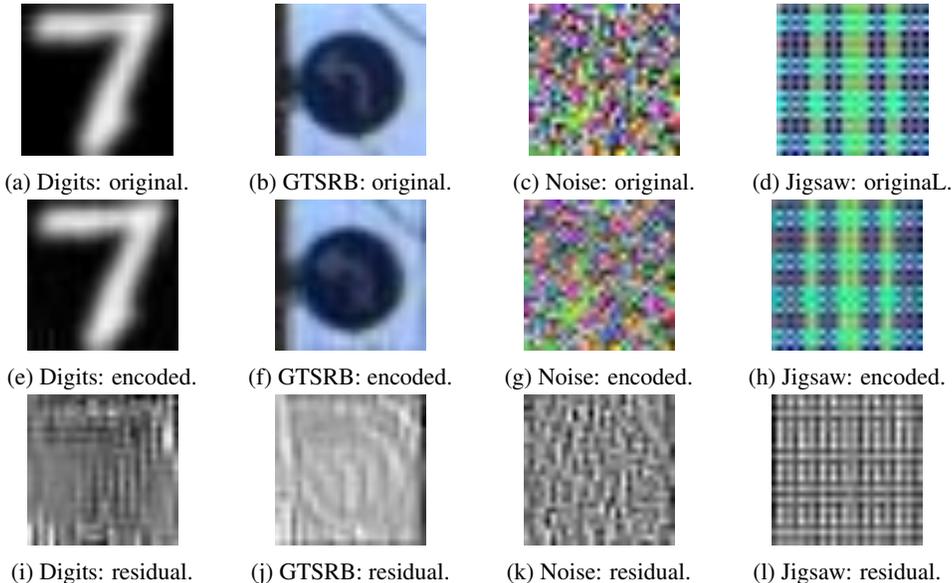


Figure 6: Visualization of unique trigger set based on different OoD datasets.

To investigate the difference between different clients' trigger sets based on the same OoD dataset, we show one example in the trigger set generated by the jigsaw image for two randomly picked clients in Fig. 7. The trigger sets are generated based on the same jigsaw dataset and differ by their embedded keys. According to Fig. 7, although the samples from different trigger sets do not look

distinguishable according to the human inspection, the difference between keys decoded from the trigger sets can be distinguished by our model.

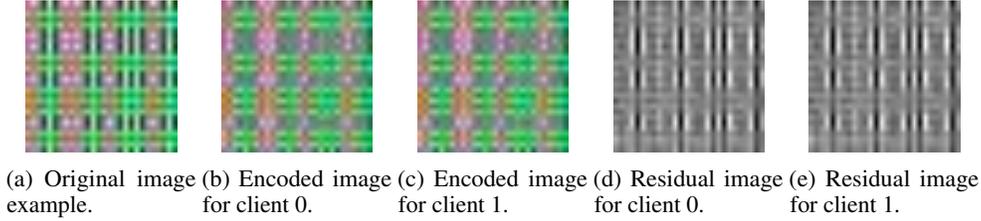


Figure 7: Visualization of the unique trigger sets for two different clients. The difference between trigger sets cannot be observed according to human inspection, but after decoding, the difference between keys can be distinguished by our model.

Effects of the different numbers of samples in trigger sets. We investigate how the size of the trigger set will affect our watermark injection and standard FL training in Fig. 8 by varying the number of samples in the trigger set from 50 to 500 for Digits training (USPS is used to generate the trigger set). Note that for all cases, TAcc always remains to be 100%. We observe that with only 50 samples in one trigger set, we can achieve an accuracy degradation around 2%, and with a WSR over 98%. When the number of samples increases to 300, WSR is over 99%. In general, the change in the number of samples in the trigger set has almost no effect on both standard accuracy and WSR. A small trigger set (such as 50) can achieve comparable results with a large trigger set. The advantages of a smaller trigger set include quicker trigger set generation, quicker watermark injection, quicker ownership verification, and quicker IP tracking. Besides, less effort can be made for OoD data synthesizing or collecting.

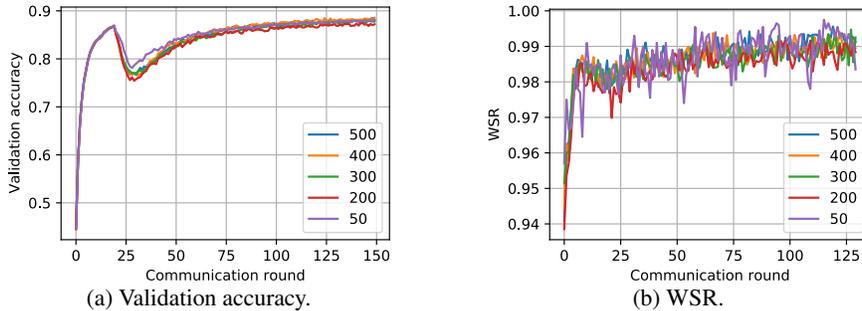


Figure 8: Effects of the different number of samples in trigger sets. 50 samples in one trigger set can achieve over 98% WSR.

B.2 Extended robustness study

Robustness against detection attack. We take Neural Cleanse [35] as an example of the detection attack, which synthesizes the possible trigger to convert all benign images to all possible target classes in the classification task space. Then an anomaly detection is conducted to detect if any trigger candidate is significantly smaller than other candidates. We follow the original setting in [35], if the anomaly index is larger than 2, the model is watermarked. The smaller the value of the anomaly index, the harder the watermark to be washed out by Neural Cleanse. Local samples are used as benign images during detection. We compare the anomaly index for the non-watermarked model and watermarked model in Fig. 9. We observe that for all datasets, the anomaly index for the watermarked model is close to that of the non-watermarked model, and both of them are smaller than the threshold 2. The observation implies that our watermarked model cannot be detected using neural cleanse. One possible reason is that Neural Cleanse relies on the assumption that the backdoor-based watermark shares the same task space with the original classification task, but due to the effectiveness of our decoder, our target label space of the watermark is different from the original classification task space.

Therefore, by searching all possible target classes in the original task space, Neural Cleanse will not find the real target label of the trigger set introduced by our proposed DUW.

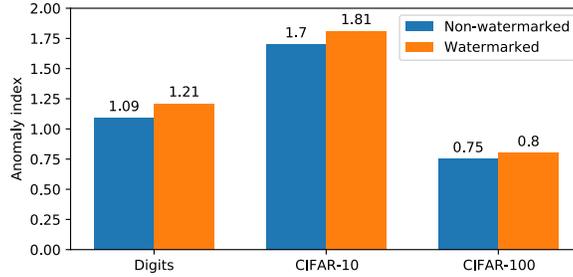


Figure 9: Anomaly index of watermarked model and non-watermarked model. If the anomaly index exceeds 2, the model will be detected as backdoor-based watermarked.

We further show the reversed trigger pattern generated by Neural Cleanse for non-watermarked (non-wm) and watermarked (wm) models in Fig. 10. The reversed trigger of our watermarked model shares a similar pattern as non-watermarked ones for all three benchmarks, and it does not look similar to our real trigger patterns (real ones can be referred to Fig. 6 residual). The trigger patterns for our trigger sets are sample-specific. Thus, it is hard to reverse engineer triggers when Neural Cleanse assumes a general trigger pattern for the entire trigger set. In summary, our proposed DUW is secured against this trigger-detection algorithm.

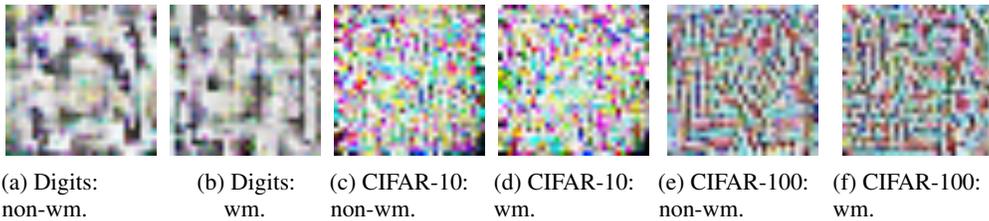


Figure 10: Reversed trigger patterns generated by Neural Cleanse for non-watermarked (non-wm) and watermarked (wm) models on three benchmarks.

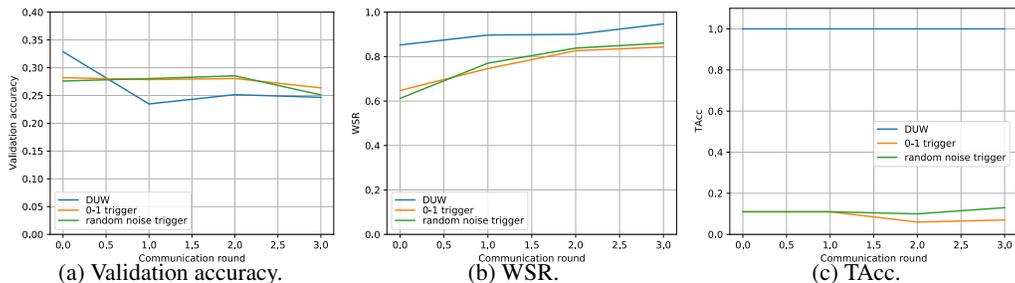


Figure 11: Validation accuracy, WSR, and TAcc for proposed DUW and other two baselines on CIFAR-10 for 4 communication rounds. To distinguish between different clients, for 0-1 trigger, we set 5 pixel values of the pattern into zero and other 11 pixels into 1, different combinations of the pattern are randomly chosen for different clients. For random noise triggers, we generate different random noise triggers for different clients. Traditional backdoor-based watermarks can only achieve a tracking accuracy lower than 13%, which is much lower than the 100% tracking accuracy we have achieved.

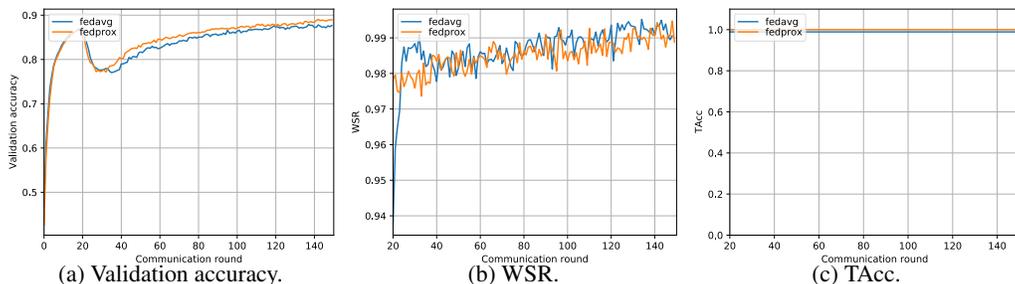


Figure 12: Validation accuracy, WSR, and TAcc for fedavg and fedprox on digits. Fedprox can achieve comparable WSR as fedavg and higher standard accuracy. TAcc for both FL algorithms remains to be 100%. Our proposed method is not sensitive to the FL framework, in which it is implanted.

Inject round	Acc	Δ Acc	WSR	WSR_Gap	TAcc
5	0.8838	0.0251	0.9951	0.9948	1.0000
10	0.8811	0.0278	0.9946	0.9938	1.0000
20	0.8855	0.0234	0.9909	0.9895	1.0000

Table 6: Ablation study: results for watermark injection in different rounds on digits. The results verify that injecting in earlier rounds will not affect standard accuracy, WSR, and TAcc.

Number of clients	Acc	Δ Acc	WSR	WSR_Gap	TAcc
40	0.8855	0.0234	0.9909	0.9895	1.0000
400	0.8597	-0.0332	0.9521	0.9267	1.0000
600	0.8276	-0.0035	0.7337	0.6383	1.0000

Table 7: Ablation study: results for different numbers of clients on digits. With more clients participating in FL, we can still track the malicious client correctly with high confidence.