

Towards Traitor Tracing in Black-and-White-Box DNN Watermarking with Tardos-based Codes

Elena Rodríguez-Lois, Fernando Pérez-González
Signal Theory and Communications Department

atlanTTic Research Center, University of Vigo, E.E. de Telecomunicación, 36310 Vigo, Spain
erodriguez, fperez@gts.uvigo.es

Abstract—The growing popularity of Deep Neural Networks, which often require computationally expensive training and access to a vast amount of data, calls for accurate authorship verification methods to deter unlawful dissemination of the models and identify the source of the leak. In DNN watermarking the owner may have access to the full network (white-box) or only be able to extract information from its output to queries (black-box), but a watermarked model may include both approaches in order to gather sufficient evidence to then gain access to the network. Although there has been limited research in white-box watermarking that considers traitor tracing, this problem is yet to be explored in the black-box scenario. In this paper, we propose a black-and-white-box watermarking method for DNN classifiers that opens the door to collusion-resistant traitor tracing in black-box, exploiting the properties of Tardos codes, and making it possible to identify the source of the leak before access to the model is granted. While experimental results show that the method can successfully identify traitors, even when further attacks have been performed, we also discuss its limitations and open problems for traitor tracing in black-box.

Index Terms—DNN watermarking, fingerprinting, traitor tracing, Tardos codes, black-box, white-box

I. INTRODUCTION AND PREVIOUS WORKS

With the recent advances in computational capacity and the great availability of training data, deep learning has revolutionized the way we approach many problems, both in industry and in our day to day life. But the training of Deep Neural Networks (DNNs) is usually time consuming, expensive, and may require private or proprietary data, making it unattainable to the general public. The development of pre-trained models, for users to fine-tune to their specific tasks, and, more interestingly, Machine Learning as a Service (MLaaS), has emerged in response to this demand as a way to capitalize DNNs. And thus, also born is the need to protect the Intellectual Property (IP) of models for deterring bad actors from unlawful use or dissemination. To this end, not only is it important to be able to prove the model's ownership, but also to determine the source of the leak, as it allows the owner to take further action. This source analysis becomes

Work partially funded by the Spanish Ministry for Science and Innovation; Spanish Cybersecurity Institute (INCIBE) and NextGenerationEU/PRTR through projects "FELDSPAR: Federated Learning with Model Ownership Protection and Privacy Armoring" (Grant MCIN/AEI/10.13039/501100011033) and "TRUFFLES: Trusted Framework for Federated Learning Systems", and by Xunta de Galicia and the European Regional Development Fund, under project ED431C 2021/47.

more challenging if bad actors are able to collude and produce models from multiple sources, which unfortunately is not difficult to imagine, if we consider that they would only need to pose as different clients to carry out this attack.

Forensic watermarking of DNNs has been presented as a promising solution to this problem [1], where a distinct watermark is inserted into the network at training time, so that it can be detected and used as proof of ownership later on. If one embeds unique watermarks, also known as fingerprints, into different copies, those can be used to identify the traitors that have unlawfully exploited or disseminated the model. The watermark itself can be embedded in the network parameters and intermediate features (white-box), meaning that the watermark detector must be granted access to the model, or in the input-output behaviour of the model (black-box), making the watermark detectable through queries.

White-box schemes have greatly benefited from traditional media watermarking theory [1], [2]. In one of the first white-box schemes for DNNs [3], the authors draw on the classic spread-spectrum and propose to embed a watermark into the model parameters through a regularization function during training, rewarding the model to converge towards a local minimum where the watermark can be correctly extracted from the model weights. Although other approaches have also been presented [1], this idea of regularization has been the most popular with many works building upon it, from the implementation of ST-DM watermarking [1] to DNN fingerprinting with anti-collusion properties through Balanced Incomplete Block Design (BIBD) codes [4] [5].

For black-box watermarking, previous works have mostly leveraged the same idea behind DNN backdooring attacks [2]. With this, during training a number of trigger-label pairs are injected into the model, to be eventually used to reveal the presence of the watermark, while the behaviour on other examples remains unaltered. The design of the trigger inputs has greatly varied across many works, using unrelated [1], benign [6] or artificial [7] examples, adding visible and invisible patterns [1], and even using adversarial examples [1]. While some existing works consider unique fingerprints to identify different model copies [8], [9], to the best of our knowledge, no scheme has been proposed for black-box watermarking that provides collusion-resistant traitor tracing capabilities. Thus, by keeping the model parameters secret

from the legitimate author, colluders may easily succeed in using models unlawfully.

White and black-box schemes are different in strengths and weaknesses. While white-box can be more reliable and hold much larger watermarking payloads, accessing the full network is unattainable in many cases. On the other hand, black-box allows for the detection of the watermark through querying the network, which may be available through MLaaS, but its capacity is restricted both by the output's dimensionality and the feasible number of queries: not only does the verification expose the triggers to the *spoil attack* [8], but it can also raise suspicion, and become expensive if the queries are available through a paid service. Because of this, these two ways can be seen as complementary [2], and the use of both has already appeared in some previous works [10], [11].

Considering the above, the goal of this paper is two-fold:

- To show one can leverage previous watermarking theory, namely Tardos codes, to achieve traitor tracing capabilities in black-box, and that it is compatible with simultaneously fingerprinting in white-box.
- To highlight the main challenges presented by black-box DNN fingerprinting that are yet to be solved, in order to motivate future research in this topic.

In what follows, Sec. II describes the proposed scheme of both black and white-box approaches, Sec. III explains the implementation of the experiments and presents the results, Sec. IV includes a discussion about the most relevant open problems of black-box traitor tracing, and concludes this paper.

Notation: Lower-case bold letters (e.g. **a**) represent column vectors, where their i th element is denoted with a subindex (e.g. a_i). Bold upper-case letters (e.g. **B**) represent two-dimensional matrices. Calligraphic fonts (e.g. \mathcal{C}) represent sets, alphabets or bases.

II. PROPOSED METHOD

In the scenario considered by this work, a model owner trains a network on a classification task, which can later be disseminated among a number of known clients under a given user agreement. We assume the model owner is able to train different versions of the same network independently, so as to embed the unique watermarks that allow proof of ownership and source identification, even when users have colluded to produce a model from multiple sources. These should satisfy the requirements of *Robustness*, *Security*, *Fidelity*, *Capacity*, *Integrity*, *Generality* and *Efficiency*, as defined in [1].

To take advantage of both approaches, and ensure detectability without access to the model weights, the owner will embed a black-box and a white-box watermark simultaneously. With this, they can first gather sufficient proof of ownership through black-box queries to request access to the network, and later strengthen the accusation with a more thorough white-box watermark detection [2]. In case of a collusion of several users, one would want the black-box scheme to accuse at least one of the traitors (i.e. *catch-one* goal in traitor tracing), and if successful, the white-box access can be used to find out about the remaining sources.

A. Black-Box Fingerprinting with q -ary Tardos Codes

Similar to backdooring attacks, black-box DNN watermarking can embed information in the input-output behaviour, activated by certain queries performed by the model owner, who does not need access to the network parameters. As previously mentioned, one of its main disadvantages is the embedding capacity, which is limited by the dimensionality of the network's output. This is usually circumvented by employing a high number of triggers, which are memorized during the training of the model and ideally do not negatively impact the main task, so that the achieved outputs to the queries can be statistically significant while preserving the model's performance on normal users' queries. Still, utilizing as few triggers as possible to detect the watermark should be considered a design priority for black-box watermarking schemes: On the one hand, any trigger that is shown to a suspected model could be susceptible to a spoil attack [8], where the stolen model could now be trained to fit the known triggers on random labels, but also ignored, or even classified with a different non-copyrighted model, rendering them useless. On the other hand, queries may only be accessible through a paid service, making the accusation process expensive if many triggers are needed.

When it comes to black-box fingerprinting, if the same approach is to be followed, different copies of the model would need to be distinguishable through the output of the trigger queries. This was contemplated in a previous work for a federated learning watermarking framework [8], but unfortunately this approach does not consider any attacks specifically against the unique fingerprints (e.g., model averaging), which makes it straightforward for any bad actor to generate an untraceable copy of the model. While these attacks are easy to model in the white-box approach, the impact over the black-box triggers is difficult to predict. This makes strategy-dependent anti-collusion codes such as BIBD, previously used in white-box fingerprinting [4], [5], less suitable. Fortunately, extensive research has been done on Tardos codes, which are collusion-resistant fingerprinting codes, agnostic to the attack strategy of the colluders. The original proposal only considered binary fingerprinting codes, but subsequent works have generalized this strategy to any q -ary alphabet [12], allowing us to leverage the full dimensionality of the network's output, and making these codes an interesting fit for this new fingerprinting scenario. Tardos codes mostly rely on the so-called Marking Assumption (MA), which states that attackers are not able to produce changes in sections of the content where they received the same information. In the current scenario, this would mean that if no colluder outputs a given symbol to a trigger, the merged model will not give that output either.¹ Actually, the MA may not always hold for black-box triggers, due to unexpected effects of the merge on the model internal features or further attacks (see Sec. IV), but fortunately, as results will show, one can still successfully utilize Tardos codes for this application.

¹Otherwise, it known as the *wide-case* problem [13] in traitor tracing.

In traditional Tardos fingerprinting, the content would be divided into m segments (or detectable positions), each hiding a symbol α of the message from a discrete alphabet \mathcal{Q} of size q . For the sake of simplicity, we will keep the traditional notation for black-box DNN fingerprinting. The role of the different segments will be covered by the set of m trigger examples, with each being assigned a label α , from the possible outputs \mathcal{Q} . Each user j will be assigned a unique fingerprint \mathbf{x}_j , where x_{ji} specifies the label for trigger i .

As described in [12], q -ary Tardos codes are generated from a secret q -component bias vector $\mathbf{p}^{(i)}$ as $p_\alpha^{(i)} \doteq P[x_{ji} = \alpha]$. The bias vector $\mathbf{p}^{(i)}$ is independently drawn from a symmetric Dirichlet distribution:

$$F(\mathbf{p}) = \frac{1}{\mathcal{N}(q, \kappa, \tau)} \prod_{\alpha \in \mathcal{Q}} p_\alpha^{-1+\kappa} \quad (1)$$

where $\kappa > 0$ is the concentration parameter, τ a cutoff parameter for $p_\alpha \in [\tau, 1 - (q-1)\tau]^2$ and $\mathcal{N}(q, \kappa, \tau)$ is a normalization factor that guarantees that

$$\int_{\tau}^{1-(q-1)\tau} \cdots \int_{\tau}^{1-(q-1)\tau} \delta(1 - \sum_{\alpha \in \mathcal{Q}} p_\alpha) F(\mathbf{p}) d\mathbf{p} = 1, \quad (2)$$

with δ the Dirac Delta function.

In our work we will consider two different values of κ : a default $\kappa = \frac{1}{q}$, which is the recommended in [12] to guarantee a positive guilty score regardless of the collusion strategy, and $\kappa \gg \frac{1}{q}$, which greatly favors the detection of strategies similar to majority-voting, but can potentially result in negative scores for guilty participants under minority-voting [14]. The value of τ is optimized for $\kappa = \frac{1}{q}$ as $\tau = c_0^{-2/(1+\kappa)}$ [12], where c_0 represents the maximum number of colluders considered by the scheme.

In traditional Tardos schemes an accusation score would be computed across all m segments and compared to an accusation threshold, but as mentioned above, one should be careful not to expose too many triggers in black-box watermarking in order to mitigate spoil attacks [8]. This means that ideally, for any suspected leak, the model owner should query with as few triggers as possible, t^* , that allows them to make a decision. After t queries to the model, the author will have exposed a subset of t trigger examples from the available m , and obtains an attack vector $\mathbf{y} = [y_1, \dots, y_t]^T$. For each suspected participant, the accusation score is computed in [12] as $S_j = \sum_{i=1}^t S_j^{(i)}$, with

$$S_j^{(i)} = \begin{cases} U_1(p_{y_i}^{(i)}) & \text{if } x_{ji} = y_i \\ U_0(p_{y_i}^{(i)}) & \text{if } x_{ji} \neq y_i \end{cases}, \quad (3)$$

and where

$$U_1(p) = \sqrt{(1-p)/p}, \quad U_0(p) = -\sqrt{p/(1-p)}. \quad (4)$$

²Although this is not necessary in the implementation of Tardos schemes for $q \geq 3$ the work in [12] includes a cutoff parameter τ restricting $p_\alpha^{(i)} \in [\tau, 1 - (q-1)\tau]$ that simplifies the theoretical analysis, which is also considered here so that equations derived from [12] still apply.

Instead of the fixed accusation threshold used in the traditional schemes [12], with our proposed approach the owner can perform a Sequential Probability Ratio Test (SPRT) [15] with every trigger queried. This allows them to stop compromising triggers as soon as sufficient evidence is gathered, to either accuse or exonerate the suspected model. Since in our scenario the MA does not always hold (see Sec. IV), the theoretical analysis for the scores of guilty participants in [12] is no longer applicable. While adapting this analysis to encompass MA violations is left for future research, here we assume that the owner has an empirical estimate of the score distributions $S_j^{(i)}$. For this, we will define a collusion \mathcal{C}^c of $c \leq c_0$ users as the set of their indexes, and denote \mathcal{Z}^c as the set of all possible collusions of size c . Then, the owner can randomly sample \mathcal{Z}^{c_0} and obtain the statistical distributions of the scores when j is a colluder and when it is not. Let $P_{\text{col}}(S_j^{(i)})$ and $P_{\text{inn}}(S_j^{(i)})$ respectively denote such distributions; then, for each new observed score $S_j^{(t)}$, the cumulative sum of the log-likelihood ratio $W_j^{(t)}$ is computed as

$$W_j^{(t)} = W_j^{(t-1)} + \log \left(\frac{P_{\text{col}}(S_j^{(t)})}{P_{\text{inn}}(S_j^{(t)})} \right). \quad (5)$$

For any user j , if $W_j^{(t)}$ exceeds a threshold b , then an accusation can be made. Alternatively, if $W_j^{(t)} < a \forall j$ the model can be deemed innocent. In either case, a decision has been reached and $t^* = t$. These thresholds are set as $a \approx \log(\epsilon_2/(1 - \epsilon_1))$, $b \approx \log((1 - \epsilon_2)/\epsilon_1)$, where ϵ_1 and ϵ_2 represent the desired False Positive (FPR) and False Negative Rates (FNR), respectively. As previously mentioned, this would be analogous to the *catch-one* goal in traitor tracing, where the author needs to identify at least one colluder, as this would constitute enough evidence to grant access to the network and perform a deeper white-box analysis.

An additional stop condition can be set as a function of t and ϵ_1 , according to the bound on ϵ_1 presented in [12]:

$$\epsilon_1 \leq \exp \left(\frac{Z_t^2}{2t} \cdot \frac{1}{1 + Z_t/(3t\sqrt{\tau})} \right), \quad (6)$$

where the author can also accuse a guilty user if $S_j > Z_t$.

B. White-Box Fingerprinting with Orthogonal Codes

White-Box watermarking exploits the over-parametrization of DNNs to encode information into the internal values of the network, for example the weights, or activations maps of certain layers. While state-of-the-art architectures grow in parameters to better solve the main task, this also confers them with more capacity to host the watermark [2], making white-box schemes ever more suitable for fingerprinting DNNs. The most popular approach, based on [3], relies on the use of a regularization function, leading the network to converge towards a local minimum of the main task where the watermarking task can also be satisfied. Previous works have already been presented [4], [5] that take this angle to successfully implement anti-collusion fingerprinting of DNNs with BIBD codes. In this

paper we consider a similar idea, but that is not restricted to using binary vectors: embedding non-binary orthogonal codes, that form the basis of the vector space where traitors can be identified. This allows us to directly project the values found on the pirated model onto the vector basis, without binarizing them first, which would induce unintended noise.

Following the approach in [3], in each of the copies of the model, a white-box fingerprint will be embedded into vector \mathbf{w} of the l flattened parameters of a given layer, through a regularization function adding to the total loss $E(\mathbf{w})$, as $E(\mathbf{w}) = E_0(\mathbf{w}) + \lambda E_R(\mathbf{w})$, where $E_0(\mathbf{w})$ represents the main task loss, $E_R(\mathbf{w})$ the regularization loss, and λ is a parameter controlling its strength during training. For this, the fingerprints are generated as a p -dimensional orthonormal basis $\mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_p\}$, where \mathbf{s}_j represents the j th user basis vector, and a secret matrix \mathbf{D} of size $l \times p$ (sampled from a normal distribution), that projects the l parameters of the layer onto the p dimensions of the user basis. The goal of $E_R(\mathbf{w})$ is to reward the model to update the layer weights towards a vector that maximizes the projection over the user basis vector, \mathbf{s}_j , and it is computed as

$$E_R(\mathbf{w}) = \exp\left(-\frac{\mathbf{w}^\top \cdot \mathbf{D} \cdot \mathbf{s}_j}{\|\mathbf{w}^\top \cdot \mathbf{D}\|}\right). \quad (7)$$

Upon access to the weights of a suspected model, the projection on each user vector becomes

$$r_j = \frac{\mathbf{w}^\top \cdot \mathbf{D} \cdot \mathbf{s}_j}{\|\mathbf{w}^\top \cdot \mathbf{D}\|}. \quad (8)$$

The expectation on the projections if all colluders participate equally (and there are no further attacks) is

$$\mathbb{E}\{r_{j|j \in C^c}\} = 1/\sqrt{c}, \quad \mathbb{E}\{r_{j|j \notin C^c}\} = 0, \quad (9)$$

allowing the author to effectively identify them.

III. IMPLEMENTATION AND EXPERIMENTAL RESULTS

A. DNN Architecture and Main Task

In order to validate the traitor-tracing capabilities of the proposed scheme we have chosen an image classification task on the well-known MNIST dataset [16]. Neither the white or black-box approaches are architecture-dependent, so we have used a simple CNN for our experiments. However, bigger networks, with more parameters, would be expected to have even more capacity for both watermarks. The CNN is made up of three convolutional layers of 16, 64 and 128 3×3 kernels, respectively, followed by a final fully-connected layer with 10 neurons with a softmax activation function, which maps to the 10 classes on MNIST. For all our experiments, we consider the output class to be the highest softmax neuron, regardless of the soft value of the output vector.

B. Choice of trigger set

In previous black-box watermarking works, the importance of choosing a trigger set similar to that of the main task has been highlighted, as out-of-distribution triggers could be

detected by the adversary when the model is queried [2], deeming those triggers potentially useless. However, when models are merged by colluders, the output of triggers with different labels might be skewed towards the common main task, which could reduce the suitability of Tardos Codes. To assess this limitation, we have considered 3 different types of triggers:

- \mathcal{T}_R^m : Randomly generated patterns, from a uniform distribution.
- \mathcal{T}_B^m : Benign examples from the main task dataset MNIST, set apart and excluded from the main task training.
- \mathcal{T}_M^m : Merging of benign examples, not excluded from the main task training, with an invisible random pattern from a uniform distribution as $\mathcal{T}_{M_i} = 0.9 \cdot \mathcal{T}_{B_i} + 0.1 \cdot \mathcal{T}_{R_i}$.

The \mathcal{T}_R^m trigger set will be used as baseline, and should be assumed if no trigger set is explicitly stated. For the chosen set, the labels for each user will be assigned according to x_{ji} , as described in Sec. II-A.

C. Watermarking Parameters and Training Process

When training the different model copies, the networks learn not only the main task, but also the two watermarking tasks. Copies start from the same weight initialization,³ and are trained on the main task dataset for 10 epochs, with categorical cross-entropy loss, stochastic gradient descent with 0.001 learning rate, and a batch size of 16. From MNIST, 25% of the examples are set apart to evaluate the model's performance on the main task, and also perform user attacks. Regarding the watermarking, the third convolutional layer was chosen to embed the white-box watermark, with 73,728 parameters. We chose a user basis of $p=1000$ dimensions, and the strength of the white-box regularization function λ was set to 1. For the black-box watermark, the full trigger set is fitted with their respective labels every 100 batches of the main dataset, with the same learning parameters. The maximum number of triggers m is 1000, c_0 was set to 6, and the resulting cutoff parameter τ is 0.038. The values of ϵ_1 and ϵ_2 were set to 10^{-6} and 10^{-3} , respectively, leaving $a = -3$ and $b = 6$ in the SPRT. All experimental values were obtained after training 100 model copies, and collusions of size c were modeled by randomly sampling \mathcal{Z}^c 500 times. The experimental distributions of the Tardos scores $P_{\text{inn}}(S_j^{(i)})$ and $P_{\text{col}}(S_j^{(i)})$ were calculated in previous simulations for each configuration, across all 1000 triggers.

D. User Attacks on the Individual Watermarks

This work mainly focuses on the collusion attack, where several users merge their individual models to create a new copy. For simplicity, we modeled the merging as an averaging of the model weights, but other collusion strategies could also produce similar results.⁴ Users can further attack the merged copy by performing fine-tuning or pruning operations.

³It was found experimentally that different initializations do not allow traitor tracing for this black-box scheme, most likely because the overfitting on the same labels relies on features that are too different.

⁴For example, randomly sampling the model parameters (*interleaving* in traitor tracing) presented a similar behaviour in preliminary experiments.

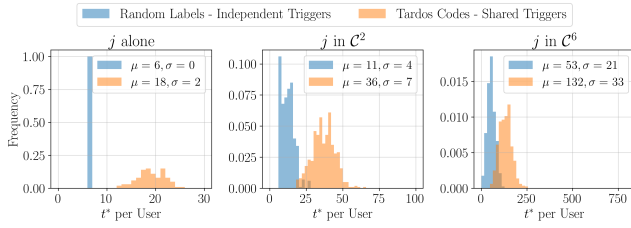


Fig. 1: Impact of using Tardos codes on the number of queries needed for a single user before an accusation.

In what follows, the fine-tuning attack fits the validation split for 10 epochs with the original training parameters, and the pruning attack uses the validation split to set 80% of the kernels' weights to zero. Other attacks to the watermark are also possible, and should be studied in future work.

E. Experimental Results

1) *Impact of Tardos Codes:* Fig. 1 shows the histogram comparison of the number of queries t^* needed per user, between the proposed scheme and the existing approach in [8], where independent triggers with random labels are assigned to each participant. In that case, querying stops as soon as the FPR, calculated as $FPR = \binom{t}{t'} \cdot 0.1^{t'} \cdot (1 - 0.1)^{t-t'}$, with 0.1 the probability of randomly answering a query correctly and t' the number of correctly answered queries, falls below ϵ_1 , so the desired FPR is the same in both approaches. When analyzing Fig. 1, one must keep in mind that independent triggers will require queries for all possible participants, so the total number of queries will be $\#users \times t^*$, whereas for the shared triggers, the total number is still t^* . This makes the independent approach a more expensive, and potentially more suspicious accusation process. As an example, if 100 users have access to the model and one of them leaks it, one would need an average of 600 queries to identify the source with independent triggers, and only 18 with the proposed approach. Fortunately, the sharing of the triggers causes only a reasonable increase of t^* per user for the smaller collusions, but still, the approach in [8] would expose less triggers per user, making it slightly more suitable against the spoil attack.

2) *Influence of κ :* As mentioned in Sec. II-A, the recommended value of $\kappa = \frac{1}{q}$ [12] guarantees that no collusion strategy results in a negative guilty Tardos score, but other values may be better suited to target specific colluding scenarios: Fig. 2 shows the impact of significantly increasing κ , which improves the efficiency of the scheme, even in the presence of further attacks, such as fine-tuning or pruning. The attacks, as expected, increase t^* , as the MA will be violated for a higher percentage of triggers (see Sec. IV). Exonerating innocent models is also feasible under the SPRT, with an average of $t^* = 301$ for $\kappa = 0.1$, and $t^* = 226$ for $\kappa = 100$, implying again that a higher κ may be better suited for this case. As mentioned in Sec. II-A, the results for both values of κ suggest that when averaging model weights, the resulting strategy over the trigger outputs may resemble majority-voting [14]. With

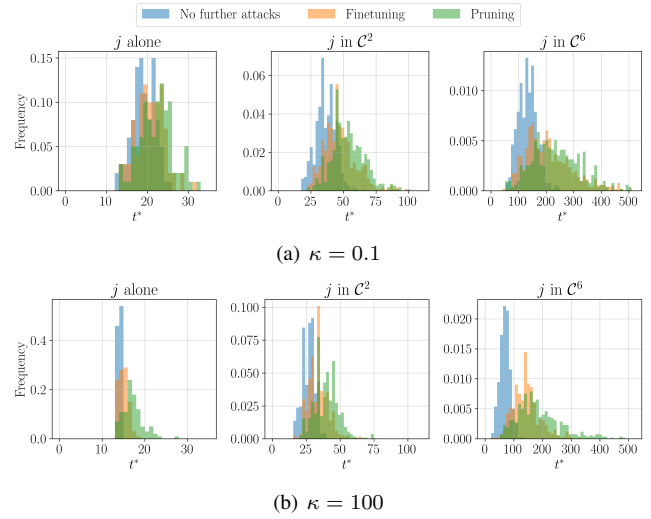


Fig. 2: Experimental distribution of t^* according to κ .

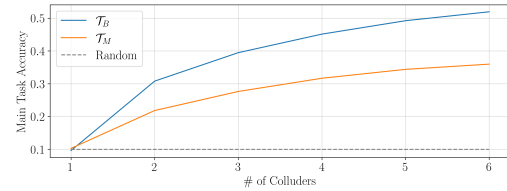


Fig. 3: Evolution of the main task accuracy on \mathcal{T} .

that in mind, a deeper analysis, including other types of attacks, should help in the design of more targeted Tardos codes.

3) *Influence of \mathcal{T} :* Another limitation of the proposed approach, briefly discussed in Sec. III-B, is that the merging of the models may skew the output towards the common main task label when triggers resemble images from the main dataset, further compromising the MA. This effect can be observed in Fig. 3, and unfortunately makes benign triggers less suitable for this scheme, as one can see in the comparison of the required t^* , in Fig. 4. On a more positive note, it seems that the use of \mathcal{T}_M could be a possible compromise between a reasonable t^* and an innocent-looking trigger.

4) *Evaluation of the False Negative Rate:* Because ϵ_2 is set using $P_{col}(S_j^{(i)})$, the FNR could unfortunately be higher when further attacks to the watermark are performed. However, due to the additional accusation condition from (6), no false negative was found in any of the experiments, including sampling 2500 collusions from Z^{c_0} with no further attacks, finetuning, and pruning. Although this makes it difficult to estimate the FNR in a real setting, a deterrent to unlawfully leak or exploit the model still exists as long as $FNR \leq \frac{1}{2}$ [12].

5) *Simultaneous White-Box Fingerprinting:* For the white-box scheme the distribution of projections r_j can be seen on Fig. 5, for different collusion sizes and further attacks. By this, statistic r_j could be effectively used to accuse all guilty users (at least up to c_0) after the black-box queries have detected the stolen model. This accusation would be possible even after pruning, which is especially destructive,

TABLE I: Influence of watermarking on the main task.

	Accuracy on MNIST (%)
Non-WM Model (baseline)	97.64
White-Box-WM Model	97.53
Black-Box-WM Model	97.52
White-and-Black-Box-WM Model	97.46

TABLE II: Marking assumption violations.

	j alone	$j \in \mathcal{C}^2$	$j \in \mathcal{C}^6$
No further attacks	0.0%	4.3%	5.6%
Fine-tuning	2.3%	15.1%	15.6%
Pruning	19.0%	31.7%	29.0%

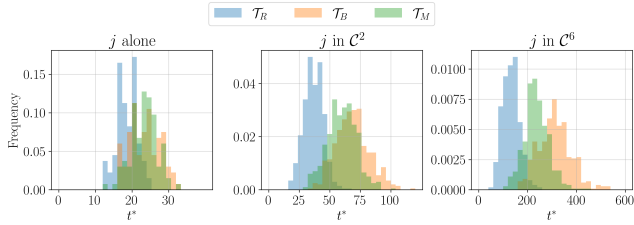


Fig. 4: Experimental distribution of t^* according to T .

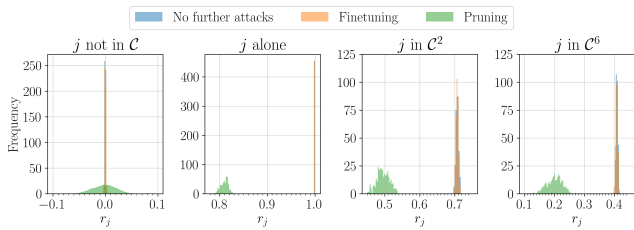


Fig. 5: Experimental distribution of r_j .

as the watermark may be exploiting weights that are not too important for the main task.

6) *Effect on the Main Task*: To ensure the proposed approach was not destructive to the main task, in Table I displays the average main task accuracy for the different schemes. Both the white and black-box schemes appear to have a small impact on the performance compared to the baseline non-watermarked model, yielding a reasonable trade-off with their traitor tracing capabilities.

IV. DISCUSSION AND OPEN PROBLEMS

The proposed scheme opens the door to collusion resistant black-box fingerprinting, requiring significantly less queries than the existing approach, as discussed in Sec. III-E, but it still has many limitations.

One of the bigger challenges to be resolved is the violation of the MA, which is the backbone of the Tardos codes, as it can lead to negative scores for guilty users thus hindering the accusation process. Table II shows the average percentage of triggers that do not hold this assumption in our experiments. Although attacks have a significant impact on the MA, even violating it 31.7% of the time, this ends up being compensated by utilizing more triggers for the decision. This makes the accusation still possible, but it is not ideal. Further analysis

on the impact that the merging of different model copies has on the triggers' output, and the bias introduced by the main task, can help understand how these violations appear, and can potentially alleviate this undesirable effect.

Also evident is the challenge of modelling the collusion strategy over the black-box triggers. Although Tardos codes are strategy-agnostic, the scheme design can be improved to better target a certain type of attack, if known. From the results shown in Sec. III-E, and considering the analysis in [14], it seems that the merging of the models may resemble a majority-voting scenario upon the triggers, which would make sense intuitively; in any case, only model averaging collusion of small networks has been considered in this work, so further analysis would be needed to better characterize the target strategy. Also, it is worth noting that in this scenario classes might not be equiprobable, and likely depend on the type of trigger. Hence, a better understanding of the feature space, also considering other architectures and tasks, would allow for a more effective design of the codes and the score function.

REFERENCES

- [1] Y. Li, H. Wang, and M. Barni, "A survey of deep neural network watermarking techniques," 2021.
- [2] M. Barni, F. Pérez-González, and B. Tondi, "DNN watermarking: Four challenges and a funeral," in *Proceedings of the 2021 ACM Workshop on Information Hiding and Multimedia Security, IH&MMSec '21*, (New York, NY, USA), p. 189–196, Association for Computing Machinery, 2021.
- [3] Y. Uchida, Y. Nagai, S. Sakazawa, and S. Satoh, "Embedding watermarks into deep neural networks," *CoRR*, vol. abs/1701.04082, 2017.
- [4] H. Chen, B. D. Rouhani, C. Fu, J. Zhao, and F. Koushanfar, "Deepmarks: A secure fingerprinting framework for digital rights management of deep learning models," in *Proceedings of the 2019 on International Conference on Multimedia Retrieval, ICMR '19*, (New York, NY, USA), p. 105–113, Association for Computing Machinery, 2019.
- [5] G. Xu, H. Li, Y. Zhang, X. Lin, R. H. Deng, and X. Shen, "A deep learning framework supporting model ownership protection and traitor tracing," in *2020 IEEE 26th International Conference on Parallel and Distributed Systems (ICPADS)*, pp. 438–446, 2020.
- [6] K. Kallas and T. Furon, "Rose: A RObust and SEcure DNN watermarking," 2022.
- [7] B. G. Atli, Y. Xia, S. Marchal, and N. Asokan, "WAFFLE: watermarking in federated learning," *CoRR*, vol. abs/2008.07298, 2020.
- [8] F.-Q. Li, S.-L. Wang, and A. W.-C. Liew, "Towards practical watermark for deep neural networks in federated learning," 2021.
- [9] T. Maho, T. Furon, and E. Le Merrer, "Model fingerprinting with benign inputs," in *ICASSP 2023 - IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, (Ialysos, Greece), pp. 1–4, IEEE, June 2023.
- [10] B. D. Rouhani, H. Chen, and F. Koushanfar, "Deepsigns: A generic watermarking framework for IP protection of deep learning models," *CoRR*, vol. abs/1804.00750, 2018.
- [11] B. Li, L. Fan, H. Gu, J. Li, and Q. Yang, "FedIPR: Ownership verification for federated deep neural network models," 2022.
- [12] B. Skoric and J.-J. Oosterwijk, "Binary and q-ary Tardos codes, revisited." *Cryptology ePrint Archive*, Paper 2012/249, 2012. <https://eprint.iacr.org/2012/249>.
- [13] A. Barg, G. Blakley, and G. Kabatiansky, "Digital fingerprinting codes: problem statements, constructions, identification of traitors," *IEEE Transactions on Information Theory*, vol. 49, no. 4, pp. 852–865, 2003.
- [14] A. Simone and B. Skoric, "Asymptotically false-positive-maximizing attack on non-binary Tardos codes," 2011.
- [15] A. Wald, "Sequential Tests of Statistical Hypotheses," *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117 – 186, 1945.
- [16] L. Deng, "The MNIST database of handwritten digit images for machine learning research," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 141–142, 2012.