REFIT: A Unified Watermark Removal Framework For Deep Learning Systems With Limited Data

Xinyun Chen*
xinyun.chen@berkeley.edu
UC Berkeley

Wenxiao Wang*
wangwx20@mails.tsinghua.edu.cn
Tsinghua University

Chris Bender
chrisbender@berkeley.edu
UC Berkeley

Yiming Ding
dingyiming0427@berkeley.edu
UC Berkeley

Ruoxi Jia
ruoxijia@vt.edu
Virginia Tech

Bo Li
lbo@illinois.edu
UIUC

Dawn Song
dawnsong@cs.berkeley.edu
UC Berkeley

ABSTRACT

Training deep neural networks from scratch could be computationally expensive and requires a lot of training data. Recent work has explored different watermarking techniques to protect the pre-trained deep neural networks from potential copyright infringements. However, these techniques could be vulnerable to watermark removal attacks. In this work, we propose REFIT, a unified watermark removal framework based on fine-tuning, which does not rely on the knowledge of the watermarks, and is effective against a wide range of watermarking schemes. In particular, we conduct a comprehensive study of a realistic attack scenario where the adversary has limited training data, which has not been emphasized in prior work on attacks against watermarking schemes. To effectively remove the watermarks without compromising the model functionality under this weak threat model, we propose two techniques that are incorporated into our fine-tuning framework: (1) an adaption of the elastic weight consolidation (EWC) algorithm, which is originally proposed for mitigating the catastrophic forgetting phenomenon; and (2) unlabeled data augmentation (AU), where we leverage auxiliary unlabeled data from other sources. Our extensive evaluation shows the effectiveness of REFIT against diverse watermark embedding schemes. The experimental results demonstrate that our fine-tuning-based watermark removal attacks could pose real threats to the copyright of pre-trained models, and thus highlight the importance of further investigating the watermarking problem and proposing more robust watermark embedding schemes against the attacks.¹

*Equal contribution.

¹The code is available at <https://github.com/sunblaze-ucb/REFIT>.

This work is licensed under a Creative Commons Attribution International 4.0 License.

ASIA CCS '21, June 7–11, 2021, Hong Kong, Hong Kong.

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8287-8/21/06.

<https://doi.org/10.1145/3433210.3453079>

CCS CONCEPTS

• Security and privacy; • Computing methodologies → Machine learning; Artificial intelligence;

KEYWORDS

watermark removal; fine-tuning; neural networks

ACM Reference Format:

Xinyun Chen, Wenxiao Wang, Chris Bender, Yiming Ding, Ruoxi Jia, Bo Li, and Dawn Song. 2021. REFIT: A Unified Watermark Removal Framework For Deep Learning Systems With Limited Data. In *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security (ASIA CCS '21)*, June 7–11, 2021, Virtual Event, Hong Kong. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3433210.3453079>

1 INTRODUCTION

Deep neural networks (DNNs) have achieved great performance on a variety of application domains, and are creating tremendous business values [12, 22]. Building these models from scratch is computationally intensive and requires a large set of high-quality annotated training samples. Various online marketplaces, such as BigML and Amazon, have emerged to allow people to buy and sell the pre-trained models. Just like other commodity software, the intellectual property (IP) embodied in DNNs needs proper protection to preserve competitive advantages of the model owner.

To protect the intellectual property of pre-trained DNNs, a widely adopted approach is *watermarking* [1, 40, 47, 52]. A common paradigm of watermarking is to inject some specially-designed training samples, so that the model could be trained to predict in the ways specified by the owner when the watermark samples are fed into the model. In this way, a legitimate model owner can train the model with watermarks embedded, and distribute it to the model users. When he later encounters a model he suspects to be a copy of his own, he can verify the ownership by inputting the watermarks to the model and checking the model predictions. This approach has gained a lot of popularity due to the simplicity of its protocol.

On the other hand, recent work has studied attack approaches to bypass the watermark verification process, so that the legitimate model owner is not able to claim the ownership. To achieve this goal, there are two lines of work in the literature. One line of work

studies detection attacks against watermark verification [24, 37]. Specifically, when the input is suspected to be a watermark by the detection mechanism, the model returns a random prediction, otherwise it returns the true model prediction. Another line of work that attracts more interest is on *watermark removal attacks*, which aims at modifying the watermarked models so that they no longer predict in the ways specified by the model owner when provided with the watermark samples. In particular, most of existing work assumes the knowledge of the watermarking scheme, e.g., the approach is specifically designed for pattern-based watermarks, where each of the watermark samples is blended with the same pattern [6, 15, 20, 48]. Although there are some latest works studying general-purpose watermark removal schemes that are agnostic to watermark embedding approaches, including pruning [30, 37, 52], distillation [49], and fine-pruning [30], most of these attacks either significantly hamper the model accuracy to remove the watermarks, or are conducted with the assumption that the adversary has full access to the data used to train the watermarked model. The lack of investigation into data efficiency leaves it unclear whether such watermark removal attacks are practical in the real world.

In this paper, we propose REFIT, a general-purpose watermark removal framework based on fine-tuning. Although previous work suggests that fine-tuning alone is not sufficient to remove the watermarks [1, 30], we find that by carefully designing the fine-tuning learning rate schedule, the adversary is able to remove the watermarks instead. However, when the adversary only has access to a small training set that is not comparable to the pre-training dataset, although the watermarks can still be removed, the test accuracy could also degrade. Therefore, we propose two techniques to overcome this challenge. The first technique is adapted from elastic weight consolidation (EWC) [26], which is originally proposed to mitigate the catastrophic forgetting phenomenon, i.e., the model tends to forget the knowledge learned from old tasks when later trained on a new one [16, 25, 26]. The central idea behind this component is to slow down learning on model weights that are relevant to the knowledge learned for the task of interest, and keep updating other weights that were used more for memorizing watermarks.

Another technique is called unlabeled data augmentation (AU). While a large amount of labeled data could be expensive to collect, unlabeled data is much cheaper to obtain; e.g., the adversary can simply download as many images as he wants from the Internet. Therefore, the adversary could leverage inherently unbounded provisions of unlabeled samples during fine-tuning. Specifically, we propose to utilize the watermarked model to annotate the unlabeled samples, and augment the fine-tuning training data with them.

We perform a systematic study of REFIT, where we evaluate the attack performance when varying the amount of data the adversary has access to. We focus on watermark removal of deep neural networks for image recognition in our evaluation, where existing watermarking techniques are shown to be the most effective. To demonstrate that REFIT is designed to be agnostic to different watermarking schemes, we evaluate our watermark removal performance over a diverse set of watermark embedding approaches, and on both transfer learning and non-transfer learning. For transfer learning setting, we demonstrate that after fine-tuning with REFIT, the resulted models consistently surpass the test performance of

the pre-trained watermarked models, sometimes even when neither EWC nor AU is applied, while the watermarks are successfully removed. For non-transfer learning setting with a very limited in-distribution training set, it becomes challenging for the basic version of REFIT to achieve a comparable test performance to the pre-trained watermarked model. With the incorporation of EWC and AU, REFIT significantly decreases the amount of in-distribution labeled samples required for preserving the model performance while the watermarks are effectively removed. Furthermore, the unlabeled data could be drawn from a very different distribution than the data for evaluation; e.g., the label sets could barely overlap.

To summarize, we make the following contributions.

- In contrast to the previous observation of the ineffectiveness of fine-tuning-based watermark removal schemes, we demonstrate that with an appropriately designed learning rate schedule, fine-tuning is able to remove the watermarks.
- We propose REFIT, a watermark removal framework that is agnostic to watermark embedding schemes. In particular, to deal with the challenge of lacking in-distribution labeled fine-tuning data, we develop two techniques, i.e., an adaption of elastic weight consolidation (EWC) and augmentation of unlabeled data (AU), towards mitigating this problem from different perspectives.
- We perform the first comprehensive study of data efficiency of watermark removal attacks, demonstrating the effectiveness of REFIT against diverse watermarking schemes.

Our work provides the first successful demonstration of watermark removal techniques against different watermark embedding schemes when the adversary has limited data, which poses real threats to existing watermark embedding schemes. We hope that our extensive study could shed some light on the potential vulnerability of existing watermarking techniques in the real world, and encourage further investigation of designing more robust watermark embedding approaches.

2 MODEL WATERMARKING

We study the watermarking problem following the formulation in [1]. Specifically, a model owner trains a model f_θ for a task $\mathcal{T}$. Besides training on data drawn from the distribution of $\mathcal{T}$, the owner also embeds a set of watermarks $\mathcal{K} = \{(x^k, y^k)\}_{k=1}^K$ into f_θ . A valid watermarking scheme should at least satisfy two properties:

- *Functionality-preserving*, i.e., watermarking does not noticeably degrade the model accuracy on $\mathcal{T}$.
- *Verifiability*, i.e., $Pr(f_\theta(x^k) = y^k) \gg Pr(f'(x^k) = y^k)$ for $(x^k, y^k) \in \mathcal{K}$, where f' is any other model that is not trained with the same set of watermarks. In practice, the model owner often sets a threshold γ , so that when $Pr(\hat{f}(x^k) = y^k) > \gamma$, the model $\hat{f}$ is considered to have the watermarks embedded, which could be used as an evidence to claim the ownership. We refer to γ as the *watermark decision threshold*.

Various watermark embedding schemes have been proposed in recent years [1, 7, 19, 34, 37, 52]. The most widely studied watermarking schemes could be pattern-based techniques, which blend the same pattern into a set of images as the watermarks [1, 7, 19]. Such techniques are also commonly applied for backdoor injection or Trojan attacks [31, 32, 41]. Therefore, a long line of work has studied defense proposals against pattern-based watermarks [6, 15,

20, 48]. Despite that these defense methods are shown to be effective against at least some types of pattern-based watermarks, they typically rely on certain assumptions of the pattern size, label distribution, etc. More importantly, it would be hard to directly apply these methods to remove other types of watermarks, which limits their generalizability. In contrast to this line of work, we study the threat model where the adversary has minimal knowledge of the pre-training process, as detailed below.

2.1 Threat Model for Watermark Removal

In this work, we assume the following threat model for the adversary who aims at removing the watermarks. In Figure 1, we provide an overview to illustrate the setup of watermark embedding and removal, as well as the threat model.

No knowledge of the watermarks. Some prior work on detecting samples generated by pattern-based techniques requires access to the entire data for pre-training, including the watermarks [5, 46]. In contrast, we do not assume access to the watermarks.

No knowledge of the watermarking scheme. As discussed above, most prior works demonstrating successful watermark removal rely on the assumption that the watermarks are pattern-based [6, 15, 20, 48]. In this work, we study fine-tuning as a generic and effective approach to watermark removal, without the knowledge of the specific watermarking scheme.

Limited data for fine-tuning. We assume that the adversary has computation resources for fine-tuning, and this assumption is also made in previous work studying fine-tuning and distillation-based approaches for watermark removal [1, 30, 49, 52]. Note that most prior works along this line assume that the adversary has access to the same amount of benign data for task $\mathcal{T}$ as the model owner. However, this assumption does not always hold in reality. Specifically, when the adversary has a sufficiently large dataset to train a good model, such an adversary is generally less motivated to take the risk of conducting watermark removal attacks, given that the adversary is already able to train his own model from scratch.

To study the watermark removal problem with a more realistic threat model, in this work, we perform a comprehensive study of the scenarios where the adversary has a much smaller dataset for fine-tuning than the pre-training dataset. In this case, training a model from scratch with such a limited dataset would typically result in inferior performance, as we will demonstrate in Section 5, which provides the adversary with sufficient incentives to pirate a pre-trained model and invalidate its watermarks.

3 REFIT: REMOVING WATERMARKS VIA FINE-TUNING

In this section, we present REFIT, a unified watermark removal framework based on fine-tuning. We present an overview of the framework in Figure 2, and we will discuss the technical details in the later part of the section. The central intuition behind this scheme stems from the *catastrophic forgetting* phenomenon of machine learning models, that is, when a model is trained on a series of tasks, such a model could easily forget how to perform the previously trained tasks after training on a new task [16, 25, 26]. Accordingly, when the adversary further trains the model with his own data during the fine-tuning process, since the fine-tuning data

Figure 1: An overview of our setup of watermark embedding and removal, as well as the threat model. Specifically, the model owner embeds a set of watermark samples into the pre-trained model, so that these samples could be used for ownership verification. Meanwhile, the training data accessible to the adversary is too limited to train a model of good performance from scratch, which motivates the adversary to pirate a pre-trained model. To bypass the ownership verification, the adversary needs to remove the watermarks, so that the watermark accuracy does not pass the threshold γ .

Figure 2: An overview of our proposed REFIT framework. Specifically, besides the basic fine-tuning scheme, REFIT incorporates two techniques to address the challenge when the adversary has a limited amount of in-distribution labeled data, i.e., elastic weight consolidation (EWC) and augmentation with unlabeled data (AU).

no longer includes the watermark samples, the model should forget the previously learned watermark behavior.

Contrary to this intuition, some prior works show that existing watermarking techniques are robust against fine-tuning based techniques, even if the adversary fine-tunes the entire model and has access to the same benign data as the owner, i.e., the entire

data for pre-training excluding the watermark samples [1, 30, 52]. The key reason could be that the fine-tuning learning rates set in these works are too small to change the model weights with a small number of training epochs. To confirm this hypothesis, we first replicate the experiments in [1] to embed watermarks into models trained on CIFAR-10 and CIFAR-100 respectively. Afterwards, we fine-tune the models in a similar way as their FTAL process, i.e., we update the weights of all layers. The only change is that instead of setting a small learning rate for fine-tuning, which is 0.001 in their evaluation, we vary the magnitude of the learning rate to see its effect. Specifically, starting from $1e-5$, the learning rate is doubled every 20 epochs in the fine-tuning process, which is the number of fine-tuning epochs for watermark removal in their evaluation.

Figure 3 presents the training curve of this fine-tuning process. We can observe that the change of model performance is still negligible when the learning rate is around 0.001, becomes noticeable when it reaches around 0.005, and requires a larger value to reach a sufficiently low watermark accuracy. Meanwhile, at the beginning of each epoch when the learning rate is doubled, the training and test accuracies decrease first, then gradually improve within the next 20 epochs; on the other hand, the watermark accuracy keeps decreasing, since the watermarks are not included in the fine-tuning dataset. Therefore, although the adversary does not have knowledge of the watermarks, the adversary can set the initial fine-tuning learning rate so that it considerably degrades the training and test accuracies within the first few fine-tuning steps, which suggests that the model weights are sufficiently modified to remove the watermarks; on the other hand, desirable test performance is achieved when the fine-tuning converges, meaning that the initial learning rate does not have to be so large that results in a model not much different from one trained from scratch. In Section 5, we demonstrate that with a learning rate schedule designed in this way, the adversary is able to remove the watermarks without compromising the model performance, when the adversary has access to a large amount of labeled training data.

While this initial attempt of watermark removal is promising, this basic fine-tuning scheme is inadequate when the adversary does not have training data comparable to the owner of the watermarked model. For example, when the adversary only has 20% of the CIFAR-100 training set, to ensure that the watermarks are removed, the test accuracy of the fine-tuned model could degrade by 5%. This is again due to the catastrophic forgetting: when we fine-tune the model to forget its predictions on the watermark set, the model also forgets part of the normal training samples drawn from the same distribution as the test one. Although the decrease of the test accuracy is in general much less significant than watermark accuracy, such degradation is still considerable, which could hurt the utility of the model.

There have been some attempts to mitigate the catastrophic forgetting phenomenon in the literature [10, 26]. However, most techniques are not directly applicable to our setting. In fact, during the watermark embedding stage, the model is jointly trained on two tasks: (1) to achieve a good performance on a task of interest, e.g., image classification on CIFAR-10; (2) to remember the labels of images in the watermark set. Contrary to previous studies of catastrophic forgetting, which aims at preserving the model's predictions on all tasks it has been trained, our goal of watermark

Figure 3: Training curves to illustrate the effect of learning rate during the fine-tuning stage, using 20% of the labeled training data. At the beginning, the model is pre-trained with the watermark scheme in [1]. Starting from a fine-tuning learning rate of $1e-5$, the learning rate is doubled every 20 epochs. The watermark accuracy considerably decreases only when the learning rate is appropriately large. See Figure 8 in the appendix for the corresponding plots where the model is fine-tuned on the entire training set, by which we can draw similar conclusions.

removal is two-fold, i.e., minimizing the model's memorization on the watermark task, while still preserving the performance on the main task it is evaluated on. This conflict results in the largest difference between our watermark removal task and the continual learning setting studied in previous work.

Another important difference is that although the training data of the adversary is different from the pre-trained data, the fine-tuning dataset contributes to a sub-task of the pre-trained model, while getting rid of the watermarks. On the other hand, different tasks are often complementary with each other in previous studies of catastrophic forgetting. This key observation enables us to adapt elastic weight consolidation [26], a regularization technique proposed to mitigate the catastrophic forgetting issue, for our purpose of watermark removal.

Elastic Weight Consolidation (EWC). The central motivation of EWC is to slow down the learning of parameters that are important for previously trained tasks [26]. To measure the contribution of each model parameter to a task, EWC first computes the diagonal of the Fisher information matrix of the previous task as follows:

$$F_i = F_{ii} = \mathbb{E}_{x \sim D, y \sim f_{\theta^*}} (y|x) \left[\frac{\partial \log f_{\theta}(y|x)}{\partial \theta_i} \right]_{\theta=\theta^*}^2 \quad (1)$$

where $f_{\theta^*}(y|x)$ is the probability distribution obtained by applying the softmax to output logits of the model with parameters θ^* given an input x , and D is the training dataset of the previous task. The entire Fisher information matrix is given by $F_{ij} = \mathbb{E}_{x \sim D, y \sim f_{\theta^*}(y|x)} \left[\frac{\partial \log f_{\theta}(y|x)}{\partial \theta_i} \cdot \frac{\partial \log f_{\theta}(y|x)}{\partial \theta_j} \right]_{\theta=\theta^*}$, which defines a Riemannian metric on the parameter space.

Intuitively, to prevent the model from forgetting prior tasks when learning a new task, the learned parameter θ should be close to the parameter θ^* of prior tasks, when the new data also contains information relevant to θ^* . Algorithmically, we penalize the distance between θ_i and θ_i^* when the i -th diagonal entry of the Fisher information matrix is large. Specifically, EWC adds a regularization term into the loss function for training on a new task, i.e.,

$$\mathcal{L}_{EWC}(\theta) = \mathcal{L}_{basic}(\theta) + \frac{\lambda}{2} \sum_i F_i (\theta_i - \theta_i^*)^2 \quad (2)$$

where $\mathcal{L}_{basic}(\theta)$ is the loss to optimize the performance on the new task (e.g. a cross entropy loss); λ controls the strength of the regularization, indicating the importance of memorizing old tasks; θ^* is the parameters trained with the previous task; F is the Fisher information matrix associated with f_{θ^*} , and F_i is the diagonal entry corresponding to the i -th parameter.

We can further extend this idea to the transfer learning setting, when the fine-tuning data belongs to a different task from the pre-trained one. In this case, the adversary can first fine-tune the pre-trained watermarked model with a small learning rate, which results in a model for his new task, although the watermarks usually still exist. Afterwards, the adversary can treat the model parameters of this new model as θ^* , and plug in Equation 1 correspondingly.

Notice that since we do not have access to the pre-trained data, in principle we are not able to compute the Fisher information matrix of the previous task, thus cannot calculate the regularization term in $\mathcal{L}_{EWC}(\theta)$. However, by leveraging the assumption that the fine-tuning data is drawn from a similar distribution to the pre-trained data, we can instead approximate the Fisher matrix using the fine-tuning data. Given that the fine-tuning data contains no watermark, the EWC component enables the model to update less on model weights important for achieving a good test accuracy, while the model weights important for watermark memorization are still sufficiently modified. Although the approximation could be imprecise in this way, in Section 5, we will show that this technique enables the adversary to improve the test performance of the model with limited data, while the watermarks are successfully removed.

With the same goal of preserving the test performance of the model with watermarks removed, we propose data augmentation with unlabeled data, which further decreases the amount of in-distribution labeled training samples needed for this purpose.

Augmentation with Unlabeled data (AU). We propose to augment the fine-tuning data with unlabeled samples, which could easily be collected from the Internet. Let $\mathcal{U} = \{x_u\}_{u=1}^U$ be the unlabeled sample set, we use the pre-trained model as the labeling tool, i.e., $y_u = f_{\theta}(x_u)$ for each $x_u \in \mathcal{U}$. We have tried more advanced semi-supervised techniques to utilize the unlabeled data, e.g., virtual adversarial training [35] and entropy minimization [18], but none of them provides a significant gain compared to the aforementioned simple approach. Therefore, unless otherwise specified, we

use this method for our evaluation of unlabeled data augmentation. Similar to our discussion of extending EWC to transfer learning, we can also apply this technique to the transfer learning setting by first fine-tuning the model for the new task without considering watermark removal, then using this model for labeling.

Since the test accuracy of the pre-trained model is not 100% itself, such label annotation is inherently noisy; in particular, when $\mathcal{U}$ is drawn from a different distribution than the task of consideration, the assigned labels may not be meaningful at all. Nevertheless, they still enable the fine-tuned model to better mimic the prediction behavior of the pre-trained model. In Section 5, we will show that leveraging unlabeled data significantly decreases the in-distribution labeled samples needed for effective watermark removal, while preserving the model performance.

4 EVALUATION SETUP

In this section, we introduce the benchmarks and the watermark embedding schemes used in our evaluation, and discuss the details of our experimental configurations.

4.1 Datasets

We evaluate on CIFAR-10 [28], CIFAR-100 [28], STL-10 [9] and ImageNet32 [8], which are popular benchmarks for image classification, and some of them have been widely used in previous work on watermarking [1, 37, 52].

CIFAR-10. CIFAR-10 includes coloured images of 10 classes, where each of them has 5,000 images for training, and 1,000 images for testing. Each image is of size 32×32 .

CIFAR-100. CIFAR-100 includes coloured images of 100 classes, where each of them has 500 images for training, and 100 images for testing, thus the total number of training samples is the same as CIFAR-10. The size of each image is also 32×32 .

STL-10. STL-10 has been widely used for evaluating transfer learning, semi-supervised and unsupervised learning algorithms, featured with a large number of unlabeled samples. Specifically, STL-10 includes 10 labels, where each label has 500 training samples and 800 test samples. Besides the labeled samples, STL-10 also provides 100,000 unlabeled images drawn from a similar but broader distribution of images, i.e., they include images of labels that do not belong to the STL-10 label set. The size of each image is 96×96 , which is much larger than CIFAR-10 and CIFAR-100. Although the label set of STL-10 and CIFAR-10 largely overlap, the images from them are distinguishable, even if resizing them to the same size.

ImageNet32. ImageNet32 is a downsampled version of the ImageNet dataset [11]. Specifically, ImageNet32 includes all samples in the training and validation sets of the original ImageNet, except that the images are resized to 32×32 . Same as the original ImageNet, this dataset has 1.28 million training samples in 1000 classes, and 50,000 samples with 50 images per class for validation.

4.2 Watermarking Techniques

To demonstrate the effectiveness of REFIT against various watermark embedding schemes, we evaluate pattern-based techniques [7, 19, 52], embedding samples drawn from other data sources as the watermarks [1, 7, 52], exponential weighting [37], and adversarial

frontier stitching [34]. These techniques represent the typical approaches of watermark embedding studied in the literature, and are shown to be the most effective ones against watermark removal.

Pattern-based techniques (Pattern). A pattern-based technique specifies a key pattern key and a target label y^t , so that for any image x blended with the pattern key , $Pr(f_\theta(x) = y^t)$ is high. To achieve this, the owner generates a set of images $\{x^k\}_{k=1}^K$ blended with key , assigns $y^k = y^t$ ($k \in 1, \dots, K$), then adds $\{(x^k, y^k)\}_{k=1}^K$ into the training set. See Figure 4 for some watermark samples.

Out-of-distribution watermark embedding (OOD). A line of work has studied using images from other data sources than the original training set as the watermarks. Figure 4 presents some watermarks used in [1], where each watermark image is independently assigned with a random label, thus different watermarks can have different labels. Meanwhile, these images are very different from the samples in any benchmark we evaluate on, and do not belong to any category in the label set.

Exponential weighting (EW). Compared to the above watermarking techniques, exponential weighting introduces two main different design choices [37]. The first choice is about the watermark generation. Specifically, they change the labels of some training samples to different random labels, but do not modify the images themselves. The main motivation is to defend against the detection attacks mentioned in Section 1, i.e., an adversary who steals the model could use an outlier detection scheme to detect input images that are far from the data distribution of interest, so as to bypass the watermark verification.

The second choice is about the embedding method. Instead of jointly training the model on both the normal training set and the watermark set, they decompose the training process into three stages. They first train the model on the normal training set only. Afterwards, they add an exponential weight operator over each model parameter. Specifically, for parameters in the l -th layer of the model denoted as θ^l , $EW(\theta^l, T)_i = \frac{\exp|\theta_i^l T|}{\max_j \exp|\theta_j^l T|} \theta_i^l$, where T is a hyper-parameter for adjusting the intensity of weighting. Finally, the model with the exponential weighting scheme is further trained on both normal training data and watermarks.

Adversarial frontier stitching (ADV). In [34], they propose to use images added with the adversarial perturbation as the watermarks. Specifically, the model is first trained on the normal training set only. Afterwards, they generate a watermark set that is made up of 50% true adversaries, i.e., adversarially perturbed images that the model provides the wrong predictions, and 50% false adversaries, i.e., adversarially perturbed images on which the model still predicts the correct labels. The adversarial perturbations are computed using the fast gradient sign method [17], i.e., $x^{\text{adv}} = x + \epsilon \cdot \text{sign}(\nabla_x J(\theta, x, y))$, where $J(\theta, x, y)$ is the training loss function of the model, and ϵ controls the scale of the perturbation. Each of these images is annotated with the ground truth label of its unperturbed counterpart as its watermark label, i.e., the label of x^{adv} is y , no matter whether it is a true adversary or false adversary. Finally, the model is fine-tuned with these watermarks added into the training set. See Figure 5 for examples of watermarks generated by this technique.

4.3 Attack Scenarios

We consider the following attack scenarios in our evaluation.

Figure 4: (a), (b), and (c) are examples of watermarks generated by the pattern-based technique in [52]. Specifically, after an image is blended with the “TEST” pattern in (a), such an image is classified as the target label, e.g., an “automobile” on CIFAR-10. (d) and (e) are examples of watermarks generated by the out-of-distribution watermark embedding technique in [1], where different watermarks could have different assigned labels.

Figure 5: Examples of watermarks generated by exponential weighting [37] and adversarial frontier stitching [34]. (a) and (c) are generated by exponential weighting, which are images from the ImageNet32 training set, but assigned with random labels different from the ground truth; for example, the watermark label of (a) is “trash can”. (b) and (d) are generated by adversarial frontier stitching technique, which add adversarial perturbations over (a) and (c) respectively, but keep the ground truth classes as their watermark labels; for example, the watermark label of (b) is still “dog”.

Non-transfer learning. The adversary leverages a watermarked model pre-trained for the same task as what the adversary desires. For this scenario, we evaluate on CIFAR-10, CIFAR-100, and ImageNet32. For CIFAR-10 and CIFAR-100, the watermarked model is pre-trained on its entire training set; for ImageNet32, the pre-trained model uses images of labels smaller than 500 in the training set. To simulate the scenario where the adversary can only collect a relatively small number of labeled training samples drawn from a similar distribution to the pre-training data, we vary the proportion of training samples the adversary has access to in the entire training set, but in practice, the fine-tuning training dataset does not necessarily need to be a subset of the pre-training dataset. We consider two data sources with abundant images for unlabeled data augmentation: (1) the unlabeled part of STL-10, which includes 100,000 samples; (2) for classification on CIFAR-10 and CIFAR-100, we also use the entire ImageNet32 for unlabeled data augmentation. For classification on ImageNet32, only those training samples with labels larger than 500 are included for unlabeled data augmentation. In both cases, we discard the labels of these ImageNet32 samples, and only use the images for augmentation. Note that these unlabeled images are very different from the labeled data. In particular, the label sets between CIFAR-100 and STL-10 barely overlap, and the label set of ImageNet32 is much more fine-grained than CIFAR-10 and CIFAR-100.

Transfer learning. The adversary leverages a watermarked model pre-trained for a different task from what the adversary desires. For this scenario, our evaluation is centered on achieving good performance on STL-10. Note that the labeled part of STL-10 only includes 5,000 samples, which is insufficient for training a model with high accuracy. Therefore, an adversary can leverage the pre-trained model on another task with a larger training set, then fine-tune the model on STL-10. This fine-tuning method is widely adopted for transfer learning [50], and is also evaluated in [1]. In particular, we perform the transfer learning to adapt from a model trained on CIFAR-10 or ImageNet32 to STL-10. We do not consider CIFAR-100 in this setting, because we find that adapting from a pre-trained CIFAR-100 model results in inferior performance on STL-10 compared to CIFAR-10 and ImageNet32, e.g., the accuracy on STL-10 is around 5% lower than the model pre-trained on CIFAR-10, as presented in [1]. We perform the unlabeled data augmentation in the same way as the non-transfer learning setting.

4.4 Implementation Details

Watermarking schemes. Our watermarking scheme configuration largely follows the same setups as their original papers. We tune the hyper-parameters to ensure that the pre-trained model achieves 100% watermark accuracy for each scheme, i.e., for all embedded watermark samples, the model prediction is the same as the assigned watermark label. Also, the prediction confidence scores for these watermark samples are high, e.g., above 0.85, suggesting the strength of the embedded watermark. We directly use their open-source implementation when applicable. Specifically:

- *Pattern-based techniques.* We use the text pattern in [52], and Figure 4 presents some examples of generated watermarks.
- *OOD watermark techniques.* The watermark images are from the code repository of [1]². The watermark set contains 100 individual images with labels randomly drawn from the entire label set, and Figure 4 shows some examples.
- *Exponential weighting.* We set $T = 2.0$ as in [37]. For each dataset, we use the last 100 samples from the training set to form the watermark set, and ensure that these watermark samples are never included in the fine-tuning training set.
- *Adversarial frontier stitching.* We set ϵ so that the watermark accuracy of a model trained without watermarks is around 50%. The values of ϵ are 0.15, 0.10 and 0.05 for CIFAR-10, CIFAR-100 and ImageNet32 respectively.

Watermark removal techniques. We always fine-tune the entire model for REFIT, because we find that fine-tuning the output layer only is insufficient for watermark removal, as demonstrated in [1]; moreover, it will completely fail to remove watermarks in the transfer learning setting by design. We have tried both FTAL and RTAL processes in [1]. Specifically, FTAL directly fine-tunes the entire model; when using RTAL, the output layer is randomly initialized before fine-tuning. For non-transfer learning, we apply the FTAL method, as RTAL does not provide additional performance gain; for transfer learning, we apply the RTAL method, since the label sets of the pre-trained and fine-tuning datasets are different. We observe that as long as the pre-trained model achieves a high test accuracy and fits the watermarks well, the model architecture does not have

a critical influence on the effectiveness of watermark embedding and removal. Thus, unless otherwise specified, we mainly apply the ResNet-18 model [22] in our evaluation, which is able to achieve competitive performance on all benchmarks in our evaluation.

As discussed in Section 3, the failure of previous fine-tuning-based watermark removal approaches is mainly due to the improper learning rate schedule for fine-tuning. For example, the initial learning rate for fine-tuning is 0.001 in [1], which is $100\times$ smaller than the initial learning rate for pre-training. In our evaluation, we set the initial fine-tuning learning rate to be much larger, e.g., 0.05. We used SGD as the optimizer and set the batch size to be 100 for both pre-training and fine-tuning without unlabeled data, following the setup in [1]. We fine-tune the model until the training accuracy does not further improve, which is typically within 20 epochs, as in [1]. For unlabeled data augmentation, when there is no in-distribution labeled samples, each batch includes 100 unlabeled samples. When fine-tuning on CIFAR-10, CIFAR-100 and STL-10, we decay the learning rate by 0.9 every 500 steps. When fine-tuning on partial ImageNet32, the learning rate is multiplied by $0.9^{\frac{t}{10}}$ after training on the $\frac{t}{10}$ -fraction of the entire training set. More discussion on implementation details is in Appendix A. In Section 5, we denote this basic version of REFIT without EWC and AU as *Basic*.

For our EWC component, Fisher information is approximated by drawing M samples from in-distribution labeled data available to the adversary. Unless otherwise specified, we set $M = 10,000$ when the target domain is CIFAR-10, CIFAR-100 or STL-10, and $M = 40,000$ when the target domain is ImageNet32. Notice that the samples are drawn with replacement, so M can be larger than the number of training examples available, where the same example may be used multiple times. In practice, to improve the stability of the optimization, we first normalize the Fisher matrix F_i so that its maximum entry is 1, then clip the matrix by $\frac{1}{\lambda \cdot lr}$ before plugging it into Equation (2), where lr is the learning rate.

In addition, we also compare with two baseline methods denoted as *FS*, which train the entire model from scratch, so that the model is guaranteed to have a watermark accuracy no higher than the decision threshold, though the test accuracy is typically sub-optimal, especially when the training data is limited. The basic version simply trains on the dataset available for fine-tuning, without leveraging the pre-trained model. The second variant, denoted as *AU*, applies the pre-trained model as the labeling tool in the same way as the AU module in REFIT, but the model is randomly initialized instead of initializing from the pre-trained model.

Evaluation metrics. We consider the two metrics below.

- *Watermark accuracy.* The adversary needs to make sure that the model accuracy on the watermark set is no more than the watermark decision threshold γ , i.e., the model predictions are the same as the assigned watermark labels for no more than γ of the watermark inputs. In particular, we set γ to be within the range of watermark accuracies of models trained without watermarks. Specifically, we trained 10 models from different random initialization without watermarks, evaluated their watermark accuracies, and set the threshold to ensure that there is no false positive, i.e., the watermark accuracies of these models are not accidentally higher than γ . For watermark schemes other than ADV, we set γ to be 20% for CIFAR-10, 10% for CIFAR-100, and

²<https://github.com/adiyoss/WatermarkNN>.

3% for ImageNet32. We set $\gamma = 58\%$ for all benchmarks when using ADV, following [34].

Notice that for the transfer learning setting, due to the difference of the label sets between the pre-trained and fine-tuning tasks, the embedded watermarks naturally do not apply to the new model. To measure the watermark accuracy, following [1], we replace the output layer of the fine-tuned model with the original output layer of the pre-trained model.

- **Test accuracy.** The adversary also aims to maximize the model accuracy on the normal test set. We evaluate the top-1 accuracy.

Regarding the presentation of evaluation results in the next section, unless otherwise specified, we only present the test accuracies of the models. The watermark accuracy of the pre-trained model embedded with any watermarking scheme in our evaluation is 100%, and the watermark accuracy of the model after watermark removal using REFIT is always below the threshold γ .

5 EVALUATION

In this section, we demonstrate the effectiveness of REFIT to remove watermarks embedded by several different schemes, in both transfer and non-transfer learning scenarios.

5.1 Evaluation of transfer learning

Pre-training on CIFAR-10. We first present the results of transfer learning from CIFAR-10 to STL-10 in Table 1. We observe that with the basic version of REFIT, where neither EWC nor AU is applied, removing watermarks already does not compromise the model performance on the test set. When equipped with either EWC or AU, the model fine-tuned with REFIT even surpasses the performance of the watermarked model.

Pre-training on ImageNet32. The results of transferring from ImageNet32 to STL-10 are in Table 2. We observe that using the pre-trained models on ImageNet32 yields around 10% improvement of test accuracy compared to the ones pre-trained on CIFAR-10, although the label set of ImageNet32 is much more different from STL-10 than CIFAR-10. This could attribute to the diversity of samples in ImageNet32, which makes it a desirable data source for pre-training. Different from pre-training on CIFAR-10, the basic version of REFIT no longer suffices to preserve the test accuracy. By leveraging the unlabeled part of STL-10, the model performance becomes comparable to the watermarked ones. When combining EWC and AU, the performance of fine-tuned models dominates among different variants of REFIT and the watermarked models.

Discussion of different pre-training datasets. Meanwhile, when we train the STL-10 from scratch and only use the pre-trained model as the labeling tool, the performance of models fine-tuned on unlabeled part of STL-10 is consistently better than models using ImageNet32 for unlabeled data augmentation. This is expected since the unlabeled part of STL-10 is closer to the test distribution than ImageNet32. Interestingly, we find that by integrating AU into REFIT, the gap between utilizing STL-10 and ImageNet32 for unlabeled data augmentation is significantly shrunk, which indicates the effectiveness of our overall framework.

5.2 Evaluation of non-transfer learning

Results on CIFAR-10 and CIFAR-100. For non-transfer learning setting, to begin with, we present results on CIFAR-10 and

CIFAR-100 in Table 3 and 4 respectively. First, we observe that when the adversary has 80% of the entire training set, using the basic version of REFIT already achieves higher test accuracies than the pre-trained models using any watermarking scheme in our evaluation, while removing the watermarks. Note that the watermark accuracies are still above 95% using the fine-tuning approaches in previous work [1, 52], suggesting the effectiveness of our modification of the fine-tuning learning rate schedule.

However, when the adversary only has a small proportion of the labeled training set, the test accuracy could degrade. Although the test accuracy typically drops for about 2% on CIFAR-10 even if the adversary has only 20% of the entire training set, the accuracy degradation could be up to 5% on CIFAR-100. For all watermarking schemes other than ADV, incorporating EWC typically improves the test accuracy for nearly 1% on CIFAR-10, and up to 3% on CIFAR-100, which are significant considering the performance gap to the pre-trained models. The improvement for ADV is smaller yet still considerable, partially because the performance of the basic fine-tuning is already much better than other watermarking schemes, which suggests that ADV could be more vulnerable to watermark removal, at least when the labeled data is very limited. By leveraging the unlabeled data, the adversary is able to achieve the same level of test performance as the pre-trained models with only 20% ~ 30% of the entire training set. In particular, in Table 5, we demonstrate that AU significantly improves the performance for labels with the lowest test accuracies. We skip the results of combining EWC and AU on CIFAR-10 and CIFAR-100, since they are generally very close to the results of AU. However, we will demonstrate that the combination of EWC and AU provides observable performance improvement on ImageNet32, which is a more challenging benchmark. We defer more discussion of sample efficiency to Appendix B.

The effectiveness of AU. Furthermore, unlabeled data augmentation enables the adversary to fine-tune the model without any labeled training data, and by solely relying on the unlabeled data, the accuracy of the fine-tuned model could be within 1% difference from the pre-trained model on both CIFAR-10 and CIFAR-100, and sometimes even surpasses the performance of the model trained with 80% data from scratch. Note that both STL-10 and ImageNet32 images are drawn from very different distributions than CIFAR-10 and CIFAR-100; when we only apply AU alone and train the model from scratch, the model accuracies are even worse than the basic version of REFIT. Specifically, augmenting with STL-10 provides better results on CIFAR-10, partially because the label set of CIFAR-10 overlaps much more with STL-10 than ImageNet32; meanwhile, augmenting with ImageNet32 clearly shows better performance on CIFAR-100, which may result from its higher diversity that is necessary for CIFAR-100 classification. However, when integrating AU into REFIT, the choice of unlabeled data does not play an important role in the final performance; i.e., the performance of augmenting with one data source is not always better than the other. These results show that REFIT is effective without the requirement that the unlabeled data comes from the same distribution as the task of evaluation, which makes it a practical watermark removal technique for the adversary given its simplicity and efficacy, posing real threats to the robustness of watermark embedding schemes.

The effectiveness of EWC. In addition, we notice that while AU mostly dominates when the percentage of labeled data is very

	FS		REFIT		
	Basic	AU	Basic	EWC	AU
Pattern	66.15	75.28/74.01	82.96	83.76	83.80/ 84.36
OOD		74.69/74.59	82.83	83.90	83.51/83.40
EW		75.51/74.48	84.03	84.66	84.43/84.07
ADV		75.23/73.95	83.66	84.39	84.39 /83.80

Table 1: Test accuracies (%) of models on STL-10 after watermark removal in the transfer learning setting, where the models are pre-trained on CIFAR-10. The accuracies of fine-tuned models on STL-10 with no requirement for watermark removal are 82.06%, 82.89%, 84.03% and 83.66% for Pattern, OOD, EW and ADV respectively. For AU, x/y stands for the results of augmenting with STL-10 and ImageNet32 respectively.

Pct.	FS		REFIT		
	Basic	AU	Basic	EWC	AU
Pattern					
0%	—	89.86/88.43	—	—	92.53 /91.93
20%	87.40	91.32/90.91	92.12	92.90	92.80/92.78
30%	89.64	92.13/91.49	92.22	93.02	93.15 /92.88
40%	90.46	92.46/92.15	92.93	93.25	93.18/93.03
50%	91.45	92.47/92.25	93.08	93.25	93.18/93.13
80%	93.01	92.82/92.67	93.52	93.67	94.11 /93.43
OOD					
0%	—	90.13/88.01	—	—	90.48 /87.52
20%	87.40	91.15/90.87	91.19	91.85	92.41 /92.08
30%	89.64	91.67/91.58	91.58	92.58	93.01 /92.61
40%	90.46	92.11/91.92	92.76	93.20	93.21 /92.58
50%	91.45	92.48/92.29	92.97	93.37	93.21/92.66
80%	93.01	92.81/92.66	93.93	93.85	94.00 /93.26
EW					
0%	—	89.77/89.11	—	—	93.05/ 93.22
20%	87.40	91.58/90.99	91.65	92.46	93.30/ 93.34
30%	89.64	91.69/91.69	92.30	93.29	93.50 /93.39
40%	90.46	92.35/91.92	92.83	93.27	93.34/ 93.42
50%	91.45	92.44/92.31	93.39	93.39	93.51 /93.36
80%	93.01	92.97/93.03	93.95	94.05	93.61/93.42
ADV					
0%	—	90.05/79.47	—	—	91.60 /85.68
20%	87.40	91.52/89.07	92.85	92.95	93.09 /92.72
30%	89.64	92.09/90.02	93.16	93.40	93.09/93.01
40%	90.46	92.23/91.15	93.21	93.37	93.20/93.09
50%	91.45	92.58/91.83	93.12	93.56	93.19/93.42
80%	93.01	92.93/92.69	93.69	93.80	93.65/93.76

Table 3: Results of non-transfer learning setting on CIFAR-10. The first column is the percentage of the CIFAR-10 training set used for fine-tuning, and the rest columns show the accuracy (%) on the test set. The test accuracy of the pre-trained model is 93.23% for Pattern, 93.63% for OOD, 93.49% for EW, and 93.31% for ADV. For AU, x/y stands for the results of augmenting with STL-10 and ImageNet32 respectively.

	FS		REFIT			
	Basic	AU	Basic	EWC	AU	EWC+AU
Pattern	66.15	74.76/71.50	88.89	91.14	92.30/90.78	93.31 /92.99
OOD		75.63/72.50	90.39	92.03	92.74/91.96	92.94 /92.45
EW		75.56/72.36	91.01	91.68	92.11/91.41	92.46 /92.34
ADV		75.19/72.71	92.46	92.63	92.63/92.51	92.96 /92.65

Table 2: Test accuracies (%) of models on STL-10 after watermark removal in the transfer learning setting, where the models are pre-trained on ImageNet32. The accuracies of fine-tuned models on STL-10 with no requirement for watermark removal are 92.95%, 92.39%, 92.16%, and 92.46% for Pattern, OOD, EW and ADV respectively. For AU, x/y stands for the results of augmenting with STL-10 and ImageNet32 respectively.

Pct.	FS		REFIT		
	Basic	AU	Basic	EWC	AU
Pattern					
0%	—	58.07/62.44	—	—	70.75 /68.27
20%	56.72	67.28/68.12	68.88	71.80	71.97/ 72.06
30%	62.20	68.95/70.07	71.05	72.64	72.98 /72.73
40%	65.42	70.45/71.34	71.96	73.20	73.44 /73.39
50%	68.18	71.27/72.23	72.58	73.44	73.72 /73.84
80%	71.71	73.22/73.79	74.23	74.77	75.42 /74.09
OOD					
0%	—	57.22/61.11	—	—	65.98/ 66.79
20%	56.72	67.18/67.75	68.55	69.91	71.02 /71.00
30%	62.20	68.83/70.06	70.12	71.77	71.70/ 72.25
40%	65.42	70.44/71.10	70.80	72.57	72.20/72.40
50%	68.18	71.37/72.17	72.27	72.73	72.73 /73.11
80%	71.71	72.65/73.00	73.61	74.00	73.70/73.18
EW					
0%	—	55.79/64.35	—	—	71.78/ 73.41
20%	56.72	67.66/68.57	69.00	70.63	73.48 /73.34
30%	62.20	69.01/70.71	71.37	72.13	73.72/ 74.08
40%	65.42	70.72/71.30	72.64	73.27	74.21 / 74.34
50%	68.18	71.96/72.38	73.46	74.25	74.26/ 75.07
80%	71.71	73.70/73.56	74.98	75.18	75.09/74.84
ADV					
0%	—	57.47/64.89	—	—	69.92 /68.64
20%	56.72	67.22/67.81	71.16	71.46	71.67 /71.58
30%	62.20	69.30/69.40	71.73	72.20	72.28 /72.02
40%	65.42	70.74/71.31	72.62	73.33	72.86/72.72
50%	68.18	72.00/72.20	73.01	73.41	73.11/73.26
80%	71.71	72.71/73.01	73.56	74.10	73.14/74.00

Table 4: Results of non-transfer learning setting on CIFAR-100. The first column is the percentage of the CIFAR-100 training set used for fine-tuning, and the rest columns show the accuracy (%) on the test set. The test accuracy of the pre-trained model is 73.83% for Pattern, 73.37% for OOD, 74.95% for EW, and 73.14% for ADV. For AU, x/y stands for the results of augmenting with STL-10 and ImageNet32 respectively.

small, with a moderate percentage of labeled data for fine-tuning, e.g., around 40%, EWC starts to outperform AU in some cases. In particular, on CIFAR-10, EWC typically becomes competitive to AU when 30% labeled data is available to the adversary, and the corresponding percentage is 40% on CIFAR-100. This indicates that with the increase of the labeled data, the estimated Fisher matrix could better capture the important model parameters to preserve for the adversary’s task of interest.

Basic	34.00	37.00	39.00	42.00	43.00
EWC	32.00	45.00	40.00	49.00	44.00
AU	52.00	46.00	46.00	56.00	54.00

Table 5: Results of non-transfer learning on CIFAR-100. We show the test accuracies (%) for 5 labels with the lowest test accuracies. The pre-trained model is embedded with EW watermarks. Models are fine-tuned with 20% of the CIFAR-100 training set. The AU uses STL-10 for data augmentation.

Results on ImageNet32. In Table 6, we further present our results on ImageNet32. Compared to the results on CIFAR-10 and CIFAR-100, removing watermarks embedded into pre-trained ImageNet32 models could result in a larger decrease of test accuracy, which is expected given that ImageNet32 is a more challenging benchmark with a much larger label set. Despite facing more challenges, we demonstrate that by combining EWC and AU, REFIT is still able to reach the same level of performance as the pre-trained watermarked model with 50% of the labeled training data.

Meanwhile, the increased difficulty of this benchmark enables us to better analyze the importance of each component in REFIT, i.e., EWC and AU. In particular, each of the two components offers a decent improvement of the test performance. The increase of accuracy with EWC is around 1% – 3% over the basic version when the fine-tuning data is very limited, e.g., the percentage of labeled samples is 20%. The performance of using AU is generally better than using EWC, until the labeled training set includes 50% of the ImageNet32 training samples of the first 500 classes, when EWC becomes more competitive. Finally, including both EWC and AU always enables further improvement of the test performance, suggesting that the combined technique is advantageous for challenging tasks.

Discussion of different watermarking schemes. Comparing the results of different watermarking schemes, we observe that the models fine-tuned from pre-trained models embedded with pattern-based watermarks consistently beat the test accuracy of models embedded with other watermarks, suggesting that while pattern-based watermarking techniques are generally more often used than other approaches, especially for backdoor injection, such watermarks could be easier to remove, which makes it necessary to propose more advanced backdoor injection techniques that are robust to removal attacks.

5.3 Comparison with alternative watermark removal techniques

In the following, we provide some discussion and comparison with some general-purpose watermark removal approaches proposed in previous work, which also does not assume the knowledge of the concrete watermarking scheme.

Pct.	FS		REFIT			
	Basic	AU	Basic	EWC	AU	EWC + AU
Pattern						
0%	–	21.77	–	–	54.37	
10%	36.06	39.41	51.05	53.59	55.98	56.81
20%	42.53	48.34	54.76	56.35	58.06	58.75
30%	47.83	52.76	56.87	58.40	58.62	59.40
40%	51.70	55.24	57.82	59.09	59.24	59.71
50%	53.58	57.04	58.76	59.68	59.40	60.02
OOD						
0%	–	21.46	–	–	51.68	
10%	36.06	39.32	50.76	52.02	53.87	55.16
20%	42.53	48.30	53.05	54.64	55.92	57.04
30%	47.83	52.58	55.47	56.42	57.63	58.27
40%	51.70	55.34	56.60	57.41	58.17	58.44
50%	53.58	56.87	57.86	58.50	58.51	59.12
EW						
0%	–	23.56	–	–	52.76	
10%	36.06	39.70	49.69	52.44	54.58	55.68
20%	42.53	48.16	53.65	55.89	56.10	56.94
30%	47.83	52.26	55.54	56.25	57.12	57.23
40%	51.70	55.32	56.36	57.00	57.28	57.40
50%	53.58	56.90	57.30	57.68	57.66	57.80
ADV						
0%	–	20.12	–	–	50.22	
10%	36.06	39.22	50.27	51.05	53.52	53.72
20%	42.53	48.20	52.95	54.03	56.00	56.50
30%	47.83	52.64	55.21	56.31	57.02	57.40
40%	51.70	55.28	57.43	57.57	57.90	57.94
50%	53.58	57.28	57.88	58.52	58.02	58.83

Table 6: Results of non-transfer learning setting on ImageNet32. The first column is the percentage of the training set used for fine-tuning, and the rest columns show the test accuracy (%). Note that the percentage is with respect to the training samples of the first 500 classes in ImageNet32. The test accuracy of the pre-trained model is 60.26% for Pattern, 60.04% for OOD, 58.31% for EM, and 59.60% for ADV. The reported test accuracy is measured on only the first 500 classes of ImageNet32. For AU, the unlabeled images are obtained from the last 500 classes of ImageNet32.

Discussion of distillation-based approaches. Distillation is a process to transform the knowledge extracted from a pre-trained model into a smaller model, while preserving the prediction accuracy of the smaller model so that it is comparable to the pre-trained one [23]. Specifically, a probability vector is computed as $p(x)_i = \frac{\exp(f(x)_i/T)}{\sum_j \exp(f(x)_j/T)}$, where $f(x)$ is the output logit of the model f given the input x , and T is a hyper-parameter representing the temperature. Afterwards, instead of using the one-hot vector of the ground truth label for each training sample x , the extracted $p(x)$ from the pre-trained model is fed into the smaller model as the ground truth. Previous work has proposed distillation as a defense against adversarial examples [39] and watermark embedding

approaches [49]. Therefore, we investigate incorporating this technique into our framework.

Specifically, instead of using the one-hot encoding of labels predicted by the pre-trained model, we use $p(x)$ as the ground truth label and vary the value of T to see the effect. Nevertheless, this method does not provide better performance; for example, with 20% labeled training set on CIFAR-10 and using unlabeled part of STL-10 for augmentation, when the pre-trained model is embedded with OOD watermarks, setting $T = 1$ provides a test accuracy of 91.60%, while using the one-hot label results in 91.93% test accuracy as in Table 3, and setting other values of T do not cause any significant difference. In particular, we observe that when using output logits of the watermarked model as the ground truth for fine-tuning, the resulted model tends to have a higher watermark accuracy, perhaps because while the output logits allow the fine-tuned model to better fit the pre-trained model, it also encourages the fine-tuned model to learn more information of watermarks. Thus, we stick to our original design to annotate the unlabeled data.

Comparison with pruning and fine-pruning. Previous work has studied the effectiveness of pruning-based approaches for watermark removal and found that such techniques are largely ineffective [30, 37, 52]. In our evaluation, we observe that when applying the pruning approach, the watermark accuracy is tightly associated with the test accuracy, which makes it hard to find a sweet spot of the pruning rate so that the test performance is preserved while the watermarks are removed. We defer more discussion to Appendix C. On the other hand, prior work shows that combining pruning with fine-tuning could improve the effectiveness of watermark removal [30]. Therefore, we also compare with the fine-pruning method proposed in [30], which first prunes part of the neurons that are activated the least for benign samples, and then performs the fine-tuning. We observe that REFIT achieves a better test performance than the fine-pruning approach, even if we use the same learning rate schedule for both of them, which improves the performance of the fine-pruning algorithm compared to its original design. We defer more discussion of fine-pruning to Appendix D.

6 RELATED WORK

Aside from the attacks that infringe the intellectual property of a machine learning model, a variety of attacks have been proposed against machine learning models, aiming at either manipulating model predictions [19, 41, 45] or revealing sensitive information from trained models [13, 14, 21, 43]. We will also review work on the catastrophic forgetting phenomenon in deep learning, as it inspires the use of EWC loss for our watermark removal scheme.

Model watermarking. To protect the intellectual property of deep neural networks, prior works have proposed several watermarking schemes [1, 34, 37, 52]. For pattern-based techniques, watermark images are blended with the same pattern [52]. Such techniques are also commonly used for backdoor attacks [7, 19, 31]. Other watermark schemes utilize individual images as watermarks [1, 34, 37]. Existing watermark removal approaches are largely designed for pattern-based techniques [6, 15, 20, 48]. Meanwhile, prior general-purpose watermark removal techniques require a large amount of training data to effectively remove the watermarks [30, 49].

Backdoor attacks. In the context of machine learning, backdoor attacks manipulate the model to provide the predictions specified

by the adversary on inputs associated with the backdoor key. In this sense, backdoor attacks are closely connected to watermarks in their formats, but usually with different purposes, as discussed in [1]. Previous work has shown that deep neural networks are vulnerable to backdoor attacks [7, 19]. Accordingly, several defense methods have been proposed for backdoor attacks [6, 15, 20, 48].

Poisoning and evasion attacks. Poisoning attacks inject well-crafted data into the training set to alter the predictive performance of a deep neural network. Besides backdoor injection, other attack goals include degrading the test accuracy indiscriminately, and changing the predictions of specific examples [3, 27, 29, 36, 38, 41]. In contrast to poisoning attacks, evasion attacks are launched in the test time of a machine learning model. The resulted samples are called adversarial examples, which are visually similar to normal data but lead to wrong predictions by the model [2, 4, 17, 45]. Note that we leverage adversarial examples as the watermark for the ADV watermarking scheme.

Catastrophic forgetting. Catastrophic forgetting refers to the phenomenon that a neural network model tends to underperform on old tasks when it is trained sequentially on multiple tasks. This occurs because the model weights that are important for an old task are changed to meet the objectives of a new task. Many recent approaches have been proposed against this effect, such as adjusting weights [26, 51], and adding data of past tasks to the new task training [33, 42]. In particular, elastic weight consolidation is a classic algorithm for mitigating catastrophic forgetting via adapting the learning of specific weights to their importance to previous tasks [26]. Note that the original EWC algorithm requires access to the data used for learning old tasks, which is not available in our case. Therefore, we propose an adaption of the algorithm to make it suitable for our watermark removal application.

7 CONCLUSION

In this work, we propose REFIT, a unified framework that removes the watermarks via fine-tuning. We first demonstrate that by appropriately designing the learning rate schedule, our fine-tuning approach could effectively remove the watermarks. We further propose two techniques integrated into the REFIT framework, i.e., an adaption of the elastic weight consolidation (EWC) approach, and unlabeled data augmentation (AU). We conduct an extensive evaluation with the assumption of a weak adversary who only has access to a limited amount of training data. Our results demonstrate the effectiveness of REFIT against several watermarking schemes of different types. In particular, EWC and AU enable the adversary to successfully remove the watermarks without causing much degradation of the model performance. Our study highlights the vulnerability of existing watermarking techniques, and we consider proposing more robust watermarking techniques as future work.

ACKNOWLEDGEMENT

This material is in part based upon work supported by the National Science Foundation under Grant No. TWC-1409915, Berkeley DeepDrive, and DARPA D3M under Grant No. FA8750-17-2-0091. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Xinyun Chen is supported by the Facebook Fellowship.

REFERENCES

- [1] Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, and Joseph Keshet. 2018. Turning your weakness into a strength: Watermarking deep neural networks by backdooring. In *27th {USENIX} Security Symposium*.
- [2] Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Šrđić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. 2013. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*.
- [3] Battista Biggio, Blaine Nelson, and Pavel Laskov. 2012. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389* (2012).
- [4] Nicholas Carlini and David Wagner. 2017. Towards evaluating the robustness of neural networks. In *2017 IEEE Symposium on Security and Privacy (SP)*.
- [5] Bryant Chen, Wilka Carvalho, Nathalie Baracaldo, Heiko Ludwig, Benjamin Edwards, Taesung Lee, Ian Molloy, and Biplav Srivastava. 2018. Detecting backdoor attacks on deep neural networks by activation clustering. *arXiv preprint arXiv:1811.03728* (2018).
- [6] Huili Chen, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. 2019. DeepInspect: A Black-box Trojan Detection and Mitigation Framework for Deep Neural Networks. *International Joint Conferences on Artificial Intelligence (IJCAI)* (2019).
- [7] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526* (2017).
- [8] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. 2017. A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819* (2017).
- [9] Adam Coates, Andrew Ng, and Honglak Lee. 2011. An analysis of single-layer networks in unsupervised feature learning. In *The fourteenth international conference on artificial intelligence and statistics*.
- [10] Robert Coop, Aaron Mishtal, and Itamar Arel. 2013. Ensemble learning in fixed expansion layer networks for mitigating catastrophic forgetting. *IEEE transactions on neural networks and learning systems* (2013).
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *North American Chapter of the Association for Computational Linguistics*.
- [13] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *The 22nd ACM SIGSAC Conference on Computer and Communications Security*.
- [14] Matthew Fredrikson, Eric Lantz, Somesh Jha, Simon Lin, David Page, and Thomas Ristenpart. 2014. Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing. In *23rd {USENIX} Security Symposium ({USENIX} Security 14)*.
- [15] Yansong Gao, Chang Xu, Derui Wang, Shiping Chen, Damith C Ranasinghe, and Surya Nepal. 2019. STRIP: A Defence Against Trojan Attacks on Deep Neural Networks. *arXiv preprint arXiv:1902.06531* (2019).
- [16] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211* (2013).
- [17] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and harnessing adversarial examples. *International Conference on Learning Representations (ICLR)* (2015).
- [18] Yves Grandvalet and Yoshua Bengio. 2005. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*.
- [19] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017. BadNets: Identifying Vulnerabilities in the Machine Learning Model Supply Chain. *arXiv preprint arXiv:1708.06733* (2017).
- [20] Wenbo Guo, Lun Wang, Xinyu Xing, Min Du, and Dawn Song. 2019. TABOR: A Highly Accurate Approach to Inspecting and Restoring Trojan Backdoors in AI Systems. *arXiv preprint arXiv:1908.01763* (2019).
- [21] Jamie Hayes, Luca Melis, George Danezis, and Emiliano De Cristofaro. 2019. LOGAN: Membership inference attacks against generative models. *Proceedings on Privacy Enhancing Technologies* (2019).
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition*.
- [23] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [24] Dorjan Hitaj and Luigi V Mancini. 2018. Have you stolen my model? evasion attacks against deep neural network watermarking techniques. *arXiv preprint arXiv:1809.00615* (2018).
- [25] Ronald Kemker, Marc McClure, Angelina Abitino, Tyler L Hayes, and Christopher Kanan. 2018. Measuring catastrophic forgetting in neural networks. In *AAAI conference on artificial intelligence*.
- [26] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* (2017).
- [27] Pang Wei Koh and Percy Liang. 2017. Understanding Black-box Predictions via Influence Functions. In *International Conference on Machine Learning*. 1885–1894.
- [28] Alex Krizhevsky et al. 2009. *Learning multiple layers of features from tiny images*. Technical Report. Citeseer.
- [29] Bo Li, Yining Wang, Aarti Singh, and Yevgeniy Vorobeychik. 2016. Data poisoning attacks on factorization-based collaborative filtering. In *Advances in neural information processing systems*.
- [30] Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. 2018. Fine-pruning: Defending against backdooring attacks on deep neural networks. In *International Symposium on Research in Attacks, Intrusions, and Defenses*.
- [31] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. 2017. Trojaning Attack on Neural Networks. In *Network and Distributed System Security Symposium (NDSS)*.
- [32] Yuntao Liu, Yang Xie, and Ankur Srivastava. 2017. Neural Trojans. In *The 35th IEEE International Conference on Computer Design*.
- [33] David Lopez-Paz and Marc'Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*.
- [34] Erwan Le Merer, Patrick Perez, and Gilles Trédan. 2019. Adversarial frontier stitching for remote neural network watermarking. *Journal of Neural Computing and Applications* (2019).
- [35] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2018. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* (2018).
- [36] Luis Muñoz-González, Battista Biggio, Ambra Demontis, Andrea Paudice, Vasin Wongrassamee, Emil C Lupu, and Fabio Roli. 2017. Towards poisoning of deep learning algorithms with back-gradient optimization. In *10th ACM Workshop on Artificial Intelligence and Security*.
- [37] Ryota Namba and Jun Sakuma. 2019. Robust Watermarking of Neural Network with Exponential Weighting. In *2019 ACM Asia Conference on Computer and Communications Security*.
- [38] Blaine Nelson, Marco Barreno, Fuching Jack Chi, Anthony D Joseph, Benjamin IP Rubinstein, Udam Saini, Charles A Sutton, J Doug Tygar, and Kai Xia. 2008. Exploiting Machine Learning to Subvert Your Spam Filter. *LEET* (2008).
- [39] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. 2016. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*.
- [40] Bita Darvish Rouhani, Huili Chen, and Farinaz Koushanfar. 2018. DeepSigns: A generic watermarking framework for ip protection of deep learning models. *arXiv preprint arXiv:1804.00750* (2018).
- [41] Ali Shafahi, W Ronny Huang, Mahyar Najibi, Octavian Suciu, Christoph Studer, Tudor Dumitras, and Tom Goldstein. 2018. Poison frogs! targeted clean-label poisoning attacks on neural networks. *Advances in Neural Information Processing Systems* (2018).
- [42] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. 2017. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*.
- [43] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*.
- [44] Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations (ICLR)* (2015).
- [45] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199* (2013).
- [46] Brandon Tran, Jerry Li, and Aleksander Madry. 2018. Spectral signatures in backdoor attacks. In *Advances in Neural Information Processing Systems*.
- [47] Yusuke Uchida, Yuki Nagai, Shigeyuki Sakazawa, and Shin'ichi Satoh. 2017. Embedding watermarks into deep neural networks. In *ACM International Conference on Multimedia Retrieval*.
- [48] Bolun Wang, Yuanshun Yao, Shawn Shan, Huiying Li, Bimal Viswanath, Haitao Zheng, and Ben Y Zhao. 2019. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *IEEE Symposium on Security and Privacy*.
- [49] Ziqi Yang, Hung Dang, and Ee-Chien Chang. 2019. Effectiveness of Distillation Attack and Countermeasure on Neural Network Watermarking. *arXiv preprint arXiv:1906.06046* (2019).
- [50] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. 2014. How transferable are features in deep neural networks?. In *Advances in neural information processing systems*.
- [51] Friedemann Zenke, Ben Poole, and Surya Ganguli. 2017. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*.
- [52] Jialong Zhang, Zhongshu Gu, Jiyong Jang, Hui Wu, Marc Ph Stoecklin, Heqing Huang, and Ian Molloy. 2018. Protecting intellectual property of deep neural networks with watermarking. In *Asia Conference on Computer and Communications Security*.

Figure 6: Results with different M for EWC, where we perform 4 runs for each M , and plot the mean and standard deviation. The pre-trained models are embedded with pattern-based watermarks, and fine-tuned with partial CIFAR-100.

A MORE DISCUSSION ON EXPERIMENTAL DETAILS

Our implementation is in PyTorch³. For each watermarking scheme in our evaluation, we present the best hyper-parameter configurations in Table 7. Note that the adversary can always make the worst assumption about the strength of the watermark scheme, and conservatively set the initial learning rate to ensure that the watermarks are removed. In our evaluation, we observe that setting an initial learning rate to be 0.05 works relatively well for all settings. Other hyper-parameters are set to maximize the test accuracy after watermark removal.

B DISCUSSION OF SAMPLE EFFICIENCY

In the following, we provide some discussion about the sample efficiency of EWC and AU components in REFIT.

The number of samples M for EWC. For EWC component, we investigate how the number of samples M drawn for Fisher information approximation affects the performance, and present the results in Figure 6. Specifically, we evaluate on CIFAR-100, and embed the pre-trained models with pattern-based watermarks. We observe that with only $M = 100$ samples, EWC is already able to increase the test accuracy around 1% over the basic fine-tuning, demonstrating its effectiveness of preserving the test performance. Setting a higher M may further improve the results, but it could also introduce a higher computation overhead without significant

³The implementation is mainly adapted from <https://github.com/adiyoss/WatermarkNN>, the code repo of [1].

Figure 7: Results with different numbers of unlabeled samples for AU, where we perform 4 runs for every setting and plot the mean and standard deviation. The pre-trained models are embedded with pattern-based watermarks and fine-tuned with partial CIFAR-100. The unlabeled samples for augmentation are drawn from the unlabeled part of STL-10.

Figure 8: Training curves to illustrate the effect of learning rate during the fine-tuning stage. The configuration is mostly the same as Figure 3, except that the model is fine-tuned on the entire training set.

Dataset	Scheme	Initial learning rate	λ (EWC)	m (AU)
CIFAR-10	Pattern	0.03	150	50
	OOD	[0.05, 0.15]	10	50
	EW	0.03	20	50
	AFS	[0.01, 0.1]	3	[5, 50]
CIFAR-100	Pattern	[0.03, 0.1]	20	50
	OOD	[0.03, 0.1]	200	50
	EW	[0.04, 0.05]	[2, 5]	50
	AFS	[0.015, 0.07]	[25, 30]	[10, 50]
CIFAR-10 $\rightarrow$ STL-10	Pattern	[0.03, 0.05]	10	50
	OOD	[0.04, 0.15]	10	50
	EW	[0.03, 0.05]	200	50
	AFS	[0.02, 0.05]	200	50
ImageNet32	Pattern	[0.004, 0.04]	[800, 1200]	50
	OOD	[0.005, 0.05]	[30, 100]	50
	EW	[0.003, 0.1]	$[10^4, 2 \times 10^4]$	50
	AFS	[0.006, 0.03]	[3, 50]	[30, 50]
ImageNet32 $\rightarrow$ STL-10	Pattern	[0.015, 0.02]	[1000, 1100]	50
	OOD	[0.01, 0.015]	[50, 100]	50
	EW	[0.007, 0.03]	$[1.2 \times 10^4, 1.5 \times 10^4]$	50
	AFS	[0.003, 0.008]	[200, 500]	50

Table 7: Ranges of the best hyper-parameter configuration for all watermark removal results. λ denotes the coefficient in EWC and m is the number of unlabeled samples added to a training batch with AU.

performance gain when M becomes too large. Therefore, we set the default value of M based on such a trade-off.

The number of unlabeled samples for AU. For AU component, we demonstrate the results of varying the number of unlabeled samples for fine-tuning in Figure 7. We also evaluate on CIFAR-100 with the pre-trained model using the pattern-based watermarking scheme, and the unlabeled samples are drawn from the unlabeled part of STL-10. Despite the large difference of dataset distribution between STL-10 and CIFAR-100, augmenting with 5K unlabeled samples already enables a considerable performance gain, and the test accuracy continues to increase with more unlabeled samples, suggesting the promise of leveraging unlabeled data for watermark removal, which is typically much easier to collect than in-distribution labeled data.

C MORE DISCUSSION ON PRUNING

Previous work has studied the effectiveness of pruning-based approaches for watermark removal, and found that such techniques are largely ineffective [30, 37, 52]. In our evaluation, we compare with the pruning method studied in [30], where we follow their setup to prune the neurons of the last convolutional layer in the increasing order of the magnitude of their activations on the validation set.

Figure 9 presents the curves of the model accuracy with different pruning rates. Note that due to the skip connections introduced in ResNet architecture, the model accuracy may not be low even if the pruning rate is close to 1. Therefore, we also evaluate VGG-16 [44], another neural network architecture that is capable of achieving

the same level of performance on both CIFAR-10 and CIFAR-100. For both models, we observe that the watermark accuracy is tightly associated with the test accuracy, which makes it hard to find a sweet spot of the pruning rate so that the test performance is preserved while the watermarks are removed.

In particular, as shown in Table 8, using the pruning approach, when the test accuracy degrades to 90.72% on CIFAR-10, the watermark accuracy is still 65%; on the other hand, using REFIT with AU, without any in-distribution labeled data, the fine-tuned model achieves the same level of performance as the pruning method with the watermarks removed. The gap on CIFAR-100 is more significant: REFIT is able to achieve an accuracy of 66.79%, but the test accuracy of the pruned model already decreases to 53.34% with 71% watermarks still retained. We have also tried other pruning approaches, but none of them works considerably better, which shows that REFIT is more suitable for watermark removal.

D MORE DISCUSSION ON FINE-PRUNING

For the implementation of fine-pruning, we set the pruning rates before fine-tuning in the same way as their paper, i.e., keep increasing the pruning rate stepwise, and stop when the degradation of the model performance becomes observable. We apply the same learning rate schedule for fine-pruning as REFIT, which is crucial in preserving the test performance of the model while removing the watermarks.

Table 8 presents the results on CIFAR-10 and CIFAR-100, comparing the fine-pruning approach to the basic version of REFIT without EWC and AU, where the pre-trained models are embedded with

Figure 9: Curves to illustrate the effect of neuron pruning. The corresponding pre-trained models are embedded with OOD watermarks. (a) CIFAR-10; (b) CIFAR-100.

Dataset	Model	Pruning	Before fine-tuning	Percentage				
				20%	30%	40%	50%	80%
CIFAR-10	ResNet-18	○	90.72%(65%)	91.10%	92.05%	92.72%	93.25%	94.20%
		×	93.73%(100%)	90.82%	92.27%	92.78%	93.44%	94.03%
	VGG-16	○	64.69%(77%)	90.21%	91.44%	92.00%	92.81%	93.52%
		×	93.48%(100%)	89.94%	91.53%	92.59%	92.69%	93.35%
CIFAR-100	ResNet-18	○	53.34%(71%)	67.34%	70.25%	71.42%	72.80%	74.05%
		×	74.50%(100%)	67.83%	70.54%	72.16%	72.49%	74.74%
	VGG-16	○	63.26%(97%)	62.03%	65.44%	67.72%	68.49%	70.99%
		×	72.19%(100%)	62.80%	65.65%	68.11%	69.47%	71.38%

Table 8: Comparisons between the basic version of REFIT and fine-pruning [30], where × in the column “Pruning” denotes REFIT without EWC and AU, and ○ denotes fine-pruning. The pre-trained models are embedded with OOD watermarks. For results before fine-tuning, we also present the watermark accuracies in the brackets. In the columns of “Percentage”, we present the proportion of labeled training set used for fine-tuning. For fine-pruning, the ratios of the pruned neurons from the last convolution layer are 98.4% and 85.9% for CIFAR-10 and CIFAR-100, respectively. Note that we apply the same learning rate schedule for fine-pruning as REFIT, which is crucial in preserving a good test performance while removing the watermarks.

OOD watermarks. Besides ResNet-18, we also evaluate VGG-16 [44], another neural network architecture that is capable of achieving the same level of performance on both CIFAR-10 and CIFAR-100. For both datasets and model architectures, we find that the results are

roughly similar, suggest that pruning is not necessary with a properly designed learning rate schedule for fine-tuning. In particular, our full REFIT framework still outperforms the fine-pruning.