



A Systematic Review on Model Watermarking for Neural Networks

Franziska Boenisch *

Fraunhofer AISEC and Freie University, Berlin, Germany

Machine learning (ML) models are applied in an increasing variety of domains. The availability of large amounts of data and computational resources encourages the development of ever more complex and valuable models. These models are considered the intellectual property of the legitimate parties who have trained them, which makes their protection against stealing, illegitimate redistribution, and unauthorized application an urgent need. Digital watermarking presents a strong mechanism for marking model ownership and, thereby, offers protection against those threats. This work presents a taxonomy identifying and analyzing different classes of watermarking schemes for ML models. It introduces a unified threat model to allow structured reasoning on and comparison of the effectiveness of watermarking methods in different scenarios. Furthermore, it systematizes desired security requirements and attacks against ML model watermarking. Based on that framework, representative literature from the field is surveyed to illustrate the taxonomy. Finally, shortcomings and general limitations of existing approaches are discussed, and an outlook on future research directions is given.

Keywords: neural networks, intellectual property protection, watermarking, machine learning, model stealing

OPEN ACCESS

Edited by:

Chaowei Xiao,
Arizona State University, United States

Reviewed by:

Bo Luo,
University of Kansas, United States
Jiazhao Li,
University of Michigan, United States

*Correspondence:

Franziska Boenisch
franziska.boenisch@
aisec.fraunhofer.de

Specialty section:

This article was submitted to
Cybersecurity and Privacy,
a section of the journal
Frontiers in Big Data

Received: 23 June 2021

Accepted: 14 October 2021

Published: 29 November 2021

Citation:

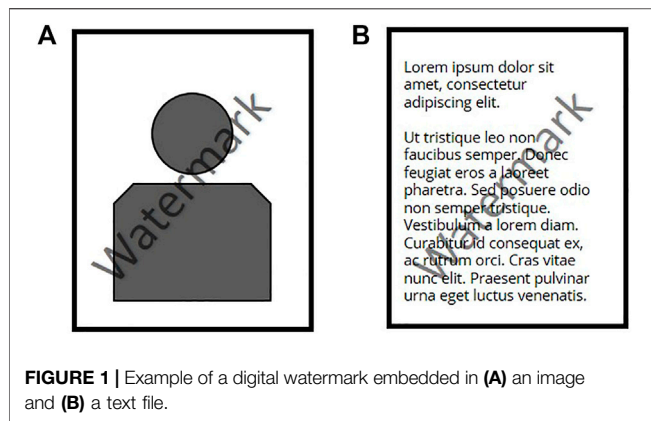
Boenisch F (2021) A Systematic
Review on Model Watermarking for
Neural Networks.
Front. Big Data 4:729663.
doi: 10.3389/fdata.2021.729663

1 INTRODUCTION

In recent years, machine learning (ML) has experienced great advancements. Its ability to process ever larger and more complex datasets has led to its application in a versatile and growing number of domains. The performance of the applied models, thereby, largely depends on the quality and quantity of their training data. However, the process of training data collection, cleansing, processing, organizing, storing, and, in certain cases, even manual labeling is time-consuming and expensive. So is the training process itself, as it may require large computational capacities, for example, in the form of numerous GPUs, and know-how for hyperparameter tuning. As a consequence, a trained ML model may be of high value and is to be considered intellectual property of the legitimate owner, that is, the party that created it.

The value incorporated in trained ML models may turn them into lucrative attack targets for malicious attackers who want to steal their functionality (Ateniese et al., 2013) for redistribution or to offer their own paid services based on them. Given the broad attack surface of stealing ML models, it might be impossible to entirely prevent theft. If theft cannot be prevented beforehand, a legitimate model owner might want to react, at least, to the inflicted damage and claim copyright to take further steps. This, however, requires that the stolen intellectual property can be traced back to its legitimate owner through adequate labeling.

The idea of marking digital property is called *watermarking*. It refers to the act of embedding identification information into some original data to claim copyright without affecting the data usage. Watermarking is already broadly applied in digital media, for example, in images, where a watermark may consist of a company logo inserted somewhere into the picture, or in texts, where a



watermark in the form of an identifying text or image might be added to the file. See **Figure 1** for an example of such digital watermarks. Also, see the study by Saini and Shrivastava (2014) for a survey on watermarking approaches in digital media.

The concept of watermarking can also be adopted for tagging ML models. So far, several methods to generate such watermarks in ML models have been proposed in research. Additionally, ways to detect, suppress, remove, or forge existing watermarks have been proposed. However, so far, the threat space in which the watermarking schemes operate has not been properly characterized. The same holds for the goals and guarantees offered by the different watermarking approaches. This makes it difficult for model owners to choose an adequate watermarking scheme that fulfills the needs in their scenarios and also to compare existing approaches with each other. Those shortcomings present, at the same time, a motivation and challenge for the systematic review put forward in this work. The goal of the article is to introduce a unified threat space for model watermarking and a taxonomy of watermarking methods and a systematization of their requirements. It, thereby, does not only propose a common language for evaluating NN watermarking but goes beyond and enables a structured comparison among existing approaches. This can serve as a basis to make watermarking methods more usable, comparable, and accessible.

The concrete contributions by this work are as follows:

- Taxonomy for watermarking schemes.
- Systematization of desirable security properties of ML model watermarks and attacks against them.
- Introduction of a unified threat model that enables structured analyses of existing watermarking schemes.
- Survey and evaluation of existing watermarking schemes and defenses according to the presented properties.

Based on the developed threat model, representative and influential works from the literature were selected. Although attempts were made to provide a comprehensive and complete overview, it is practically not possible to cite all works in the scope of the given article. For example, the topic of side-channel attacks that aim at extracting neural networks (Wei et al., 2018) is not

covered. Also, the watermarking schemes presented in the following are limited to neural networks (NNs) for classification.

2 BASIC CONCEPTS AND BACKGROUND

This section provides a brief overview on ML, on model stealing in general and model extraction attacks in particular, and introduces the concept of backdoors for NNs.

2.1 Machine Learning

ML consists of two phases, *training* and *inference/testing*. During training, an ML model h_θ is given a training dataset \mathcal{D}_s to fit its parameters θ on. For classification tasks, the training data have the form (x, y) with x denoting the *feature vector* and y the *target class*. The model parameters are adjusted through minimizing a loss function that expresses the distance between predictions $h_\theta(x)$ and true targets y (Papernot et al., 2018).

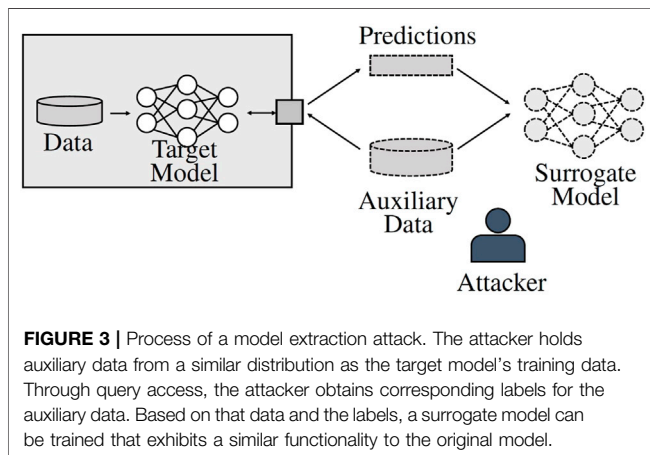
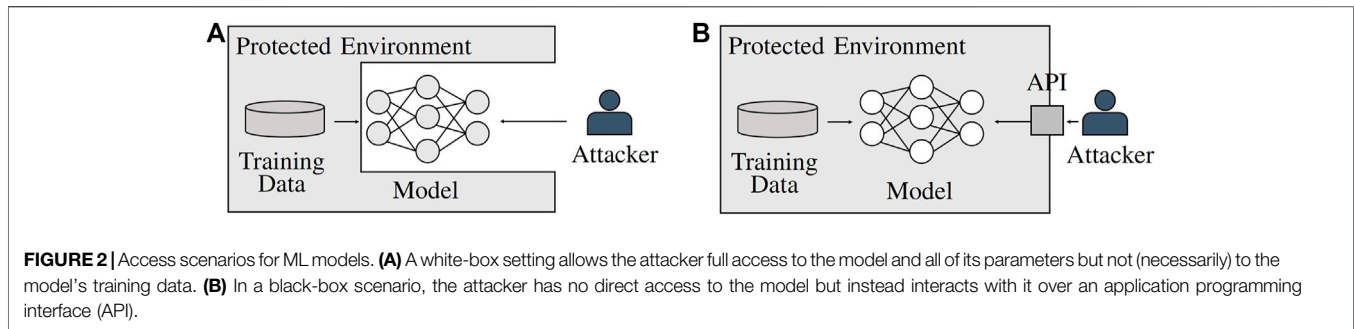
At test time, once the model parameters θ are fit, the function $h_\theta()$ can be applied to new and unseen data x' to produce predictions on them. A model that performs well solely on the training data is said to *overfit* that data, whereas a model that also performs well on the unseen test data is said to exhibit a good *generalization capacity*. Performance is usually expressed in terms of *accuracy*, which is the percentage of correct predictions over all predictions.

2.2 Model Stealing

Potential attackers may attempt to steal an ML model to have unlimited access to its complex functionality without the high preparation or continuous per query costs. Alternatively, they may wish to use the stolen model as a departure point for further attacks that are rendered more efficient through model parameter access, for example, adversarial sample crafting (Tramèr et al., 2016; Carlini et al., 2020). Protecting ML models against theft is a challenging task, as by definition, the models are supposed to reveal some information to the users. Hence, in addition to the classical security risks of model theft, for example, malicious insider access, successful attacks on servers hosting the model, or side-channel attacks (Batina et al., 2018), the information legally revealed by the model can be exploited. This enables people to steal ML models in white-box and in black-box settings (see **Figure 2**).

2.3 Model Extraction

A *model extraction* attack refers to stealing a target model h_θ through black-box access, that is, through posing queries to the model over a predefined interface as depicted in **Figure 2B**. An attacker might use those queries to h_θ to obtain labels for unlabeled data $D_{s'}$ from distribution D . Given $D_{s'}$ and the corresponding labels obtained from the original model, the attacker can train a *surrogate model* h'_θ that incorporates the original model's functionality. See **Figure 3** for an overview on the process. Jagielski et al. (2020) distinguished between two types of *model extraction*. A *fidelity extraction* attack is considered successful if h'_θ reproduces h_θ with small deviation. Hence, when h_θ is erroneous with regard to the ground truth label of a data



point, so is h'_θ . A *task accuracy extraction* aims at extracting a model that solves approximately the same underlying decision task (Jagielski et al., 2020).

2.4 Backdoors in NNs

Adi (2018) defined backdooring in NNs as a technique to intentionally train an ML model to output incorrect predictions (w.r.t. the ground truth) on a given set of the training data points. As a result, the backdoored NN behaves normally on most data points but differently on the backdoor data (Liu et al., 2018). The ability to add backdoors to NNs results from their over-parametrization, that is, the fact that many such models contain more parameters than they need for solving the task that they are supposed to solve. In watermarking, a set of backdoor data points can, for example, be used to mark and later recognize a trained NN. Thereby, the backdoor data can act as a watermark trigger (see Section 3.1).

3 TAXONOMY OF NN WATERMARKING

Even though inserting a watermark into a model does not prevent theft, it can still enable legitimate owners to identify their stolen model instances. Therefore, after the model is stolen by an attacker, the legitimate owner might use the watermark to re-identify it and claim copyright. Hence, the watermarking methods need to be effective in and chosen adequately for the

given scenario. For example, using a watermarking scheme that does not offer any binding between the watermark and the identity of the legitimate model owner might still help the owner recognize stolen model instances. However, it is of little use in a scenario where the owner wants to claim copyright in front of a third party, such as a legal entity. This section presents a taxonomy to support classifying watermarking schemes along five different dimensions. Such a classification can be helpful for identifying and comparing adequate watermarking schemes for a concrete scenario and the corresponding requirements.

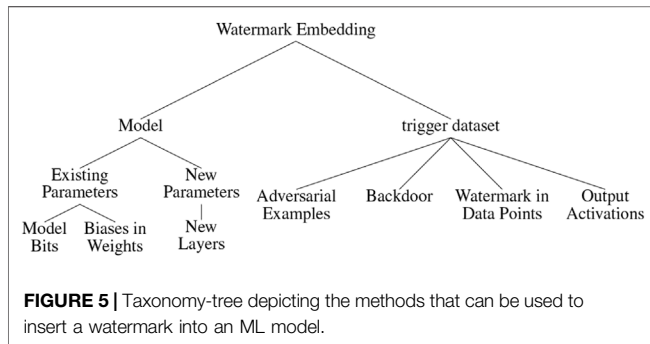
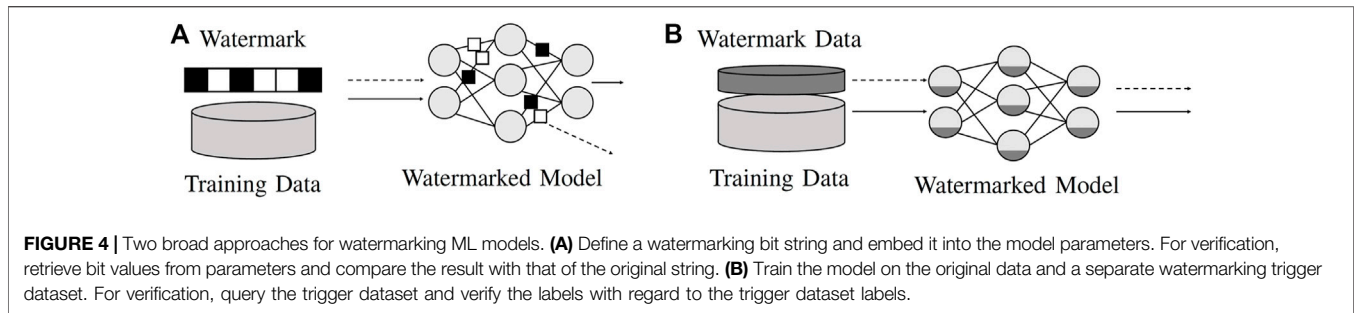
To determine the different classes, the following five dimensions are considered:

- 1) *Embedding method*: refers to the method used to include the watermark in the model.
- 2) *Verification access*: specifies how the watermark can be verified, either through white-box or black-box access.
- 3) *Capacity*: distinguishes between *zero-bit* and *multi-bit* schemes. The former refers to watermarks that do not carry additional information, whereas the latter do.
- 4) *Authentication*: indicates if the watermark directly allows the legitimate owner to prove ownership.
- 5) *Uniqueness*: states if single (stolen) model instances should be uniquely identifiable.

In the following, the five dimensions are characterized in greater detail.

3.1 Embedding Method

Watermarking techniques that have been proposed so far can be divided into two broad categories, namely, 1) inserting the watermark or related information directly into the model parameters and 2) creating a *trigger*, *carrier*, or *key* dataset, which consists of data points that evoke an unusual prediction behavior in the marked model. See Figure 4 for a visualization of both concepts. In 1), the watermark might either be encoded in existing model parameters, for example, in the form of a bit string, or be inserted through adding additional parameters that rely on or directly contain the watermark. For 2), the trigger dataset needs to be fed along the original training data during the training process of the model. Thereby, a backdoor is inserted into the model, such that the model learns to exhibit an unusual prediction behavior on data points from the trigger dataset.



The unusual behavior can then be used in order to identify illegitimate model copies. Therefore, when testing a model, the legitimate owner can query the trigger dataset and calculate the percentage of agreement between the model's prediction on the trigger dataset and the original corresponding labels. If the resulting percentage exceeds a certain threshold (should be close to 1), then the model is likely to be an illegitimate copy (Yang et al., 2019). The trigger dataset can be generated independently or be based on the original training data. Hence, it can potentially belong to a different data distribution than the training data. Some watermarking schemes, for example (Fan et al., 2019), also combine both embedding categories. **Figure 5** shows a taxonomy-tree depicting a more fine-grained division of sub-categories within 1) and 2). This taxonomy-tree provides the structure for presenting existing watermarking schemes in greater detail in **Section 6**. In addition to the two broad categories of embedding watermarks into NNs, it is also possible to use existing features of the models themselves as so-called *fingerprints* to identify potentially stolen model instances (Chen et al., 2019a; Lukas et al., 2019). Since these methods do not require explicitly inserting additional information as watermarks into the models, they will only be mentioned briefly in this document.

3.2 Verification Access

The type of access to a model required in order to perform watermark verification is closely related to the embedding method that was used to insert the watermark into it. There exist two broad scenarios for watermark verification, namely, *white-box* and *black-box*. In a white-box scenario, a legitimate model owner needs access to the model parameters in order to check for the watermark in potentially stolen copies of a model.

This might be necessary when the watermark is embedded into the model parameters alone and does not reflect in the model behavior. However, in many scenarios, white-box access for verification is no realistic assumption. A more realistic scenario is a black-box access scenario in which a legitimate model owner can access the potentially stolen model solely through a predefined query interface. Through such an interface, the model owner could query (parts of) the watermark trigger and recognize a stolen instance of the model by its prediction behavior on these data points. When selecting an adequate watermarking method, the access scenario for verification needs to be taken into account. This is because a watermark that requires the model owner to have access to the model parameters for verification might be of little use in a scenario where the attacker deploys the stolen model in a black-box setting.

3.3 Capacity

Similar to the study by Xue et al. (2020), this work defines capacity as the watermark's capability to carry information. In general, a distinction can be made between *zero-bit* and *multi-bit* watermarking schemes. Zero-bit watermarks do not carry additional information, such that they solely serve to indicate the presence or the absence of the watermark in a model. An example for such a scheme could be using plain random data points as a trigger dataset to backdoor a model and verifying potentially stolen model instances by querying these data points and observing the model's predictions on them. In multi-bit schemes, the watermark can carry information, for example, in the form of a bit string. Thereby, such schemes can be used, among others, for creating a link between a model owner's identity and the watermark or to mark individual model instances.

3.4 Authentication

By creating a link between a model owner's identity and the watermark, a watermarking scheme can serve to authenticate the legitimate owner. This allows us to extend the information that a model was watermarked by the information by whom it was watermarked, which can be useful, for example, if the legitimate model owner wants to claim copyright in front of a legal entity. The link between the owner's identity and the watermark can be expressed, among others, by including the owner's digital signature directly in the watermark or in the trigger data. Besides enabling the legitimate owner to proof their

ownership, watermarking schemes that offer authentication also prevent attackers from claiming ownership of existing watermarks, that is, *forging* them. Preventing forging is necessary to guarantee unambiguous ownership claims (see **Section 5.3**). Watermarking schemes that do not inherently allow for authentication need to take different measures to prevent attackers from forging watermarks.

3.5 Uniqueness

The last dimension along which to classify watermarking schemes for NNs concerns the uniqueness of the watermarks, that is, the question whether all instances of a model use the same watermark or if every instance receives a unique identification. A crucial shortcoming of the former is that when a stolen copy of a model appears somewhere, it is impossible for the legitimate model owner to identify which of the parties that had access to the model stole it. Unique watermarks allow us to distinguish between different model instances, and thereby, enable a more fine-grained tracking of the intellectual property. The distinction between different instances of a model can be, for example, implemented through the use of unique model identifiers or serial numbers (Xu et al., 2019). The watermark, then, does not only signal to the legitimate owner that the model was stolen but also by whom it was stolen.

4 THREAT MODEL AND SECURITY GOAL

In general and for watermarking schemes, the security of a system should be evaluated with respect to a specific threat space that characterizes the attacker's knowledge, capabilities, and objectives and the underlying security goals. This serves to thoroughly explain what properties the watermarking scheme needs to exhibit in order to adequately serve the goal of protecting a model in a given scenario: for example, a watermarking scheme supposed to protect an ML model that is directly distributed to the users in a white-box fashion will most likely need to possess different properties than a scheme applied to a model that is only accessible in a black-box fashion. This section, therefore, presents a unified threat model for watermarking schemes by elaborating the concrete requirements for watermarking and properly characterizing the attack surface and the attacker.

4.1 Requirements for Watermarking

In watermarking, the security goals can be expressed in the form of concrete requirements for an effective watermarking scheme. Within the last years, several such requirements have been formulated by different parties (Uchida et al., 2017; Adi, 2018; Chen et al., 2019a; Li H. et al., 2019). See **Table 1** for a structured overview on the requirements and their practical implications. Note that it might not be feasible to implement all requirements simultaneously, since they might interfere with each other. Take as an example *reliability* and *integrity*. In order to make a watermark reliable, verification should be very sensitive and also indicate ownership in case of doubt. A watermarking scheme in which verification always and for every model under test indicates ownership would be perfectly reliable

since it would detect every stolen model instance. However, that scheme would exhibit a very high false alarm rate and erroneously accuse many honest parties of theft, which represents a violation of the integrity.

4.2 Watermarking Attack Surface and Attacker

The attack surface needs to be characterized to understand at what point and how an attacker might attempt to bypass the watermark. Independent of the attacker, the main aspect to be considered is the scenario in which the model is stolen, that is, black-box or white-box access as depicted in **Figure 2**. A white-box scenario holds the advantage that by stalling the model as is, the attacker is likely to retain the watermark within it. This might render watermark verification for the legitimate owner more successful if the attacker does not employ additional methods to impede verification. In a black-box scenario, for example, through model extraction (see **Section 2.3** and **Figure 3**), there is not necessarily a guarantee that the watermark is entirely transferred to the surrogate model. As a consequence, watermark accuracy might already be degraded in the extracted model without additional explicit measures by the attacker. Hence, watermark verification might be more difficult for the original model owner. Further aspects of the attack surface are shaped by the attackers, their knowledge, capabilities, and objectives.

4.2.1 Attacker Knowledge

The attacker knowledge refers to the information that an attacker holds about the system. In NN watermarking, the information can consist of the following (from weak to strong): 1) the *existence of the watermark*, knowledge on the 2) *model and its parameters*, 3) the *watermarking scheme used*, 4) (parts of) the *training data*, and 5) (parts of) the *watermark itself* or the *trigger dataset*. More meaningful information can, potentially, allow for more effective attacks. For example, the sheer knowledge of the existence of a watermark within a model, without further details, will hardly serve as a warning to the attacker to release the stolen model with care in order to avoid detection by the legitimate owner. However, when knowing what watermarking scheme was employed and even possessing additional training data or parts of the watermark itself, an attacker might be able to more successfully extract a model in a black-box scenario and to apply concrete attacks against the watermark, such as the ones described in **Section 5**.

4.2.2 Attacker Capabilities

Other than the information on the system, the attacker's capabilities can also shape the threat space. A typical distinction here is to be made between a *passive* and an *active* attacker. A passive attacker cannot interact directly with the target model but might be able to observe the model's behavior through its outputs or input-output pairs. Such an attack is usually referred to as *eavesdropping* and might yield sub-optimal attack results or provide limited information on the watermark. An active attacker, on the contrary, might be able to interact with the model by 1) posing queries and 2) observing the corresponding output. Thereby, through carefully choosing the

TABLE 1 | Requirements for watermarking techniques.

Requirement	Explanation	Motivation
Fidelity	Prediction quality of the model on its original task should not be degraded significantly	Ensures the model's performance on the original task
Robustness	Watermark should be robust against removal attacks	Prevents an attacker from removing the watermark to avoid copyright claims of the original owner
Reliability	Exhibit a minimal false negative rate	Allows legitimate users to identify their intellectual property with a high probability
Integrity	Exhibit a minimal false alarm rate	Avoids erroneously accusing honest parties with similar models of theft
Capacity	Allow for inclusion of large amounts of information	Enables inclusion of potentially long watermarks, for example, a signature of the legitimate model owner
Secrecy	Presence of the watermark should be secret; the watermark should be undetectable	Prevents watermark detection by an unauthorized party
Efficiency	Process of including and verifying a watermark to an ML model should be fast	Does not add large overhead
Generality	Watermarking algorithm should be independent of the dataset and the ML algorithms used	Allows for broad use

Definitions adopted from the studies by Uchida et al. (2017); Adi, 2018; Chen et al. (2019a); Li H. et al. (2019).

inputs, the attacker might gain much more information on the target model and the watermark than by simple eavesdropping. Also, among active attackers, there is a broad range of capabilities. For 1), namely, providing model inputs, the number of queries that the attacker can pose might be restricted or unrestricted. Being able to pose more queries might enable the attacker to gain more information on the system. There might, furthermore, exist restrictions on the type of queries that an attacker can pose, such as on the input format, the input source, or the range of possible input values. When there are no restrictions on the type of queries, the attacker can query any possible data point to the model. For 2), namely, the observation of the model output, the attacker might also have different capabilities ranging from, for example, the ability to only observe the predicted class of a classifier, up to obtaining much more fine-grained information from the model, such as confidence scores for all classes. Depending on the level of details that the model returns together with the predictions, the attacker's ability to learn about the system might vary.

4.2.3 Attacker Objectives

The last part of the threat space specifies the attacker's objectives. This includes the question on what for and where the stolen model will be used. If the attacker plans to deploy it secretly and with no external interaction, there is a large chance that the model theft might remain uncovered. Otherwise, if access to the stolen model is offered, for example, if the attacker wants to sell services that are based on the model and its predictions, the legitimate owner is more likely to successfully re-identify the stolen model instance. In case a stolen model is exposed for interaction, the attacker's methods to prevent the model owner from successfully verifying the watermark play a vital role. They are depicted in greater detail in the following section.

5 ATTACKS AGAINST WATERMARKING

The attacks described in this section highlight several practical considerations that must be taken into account when designing

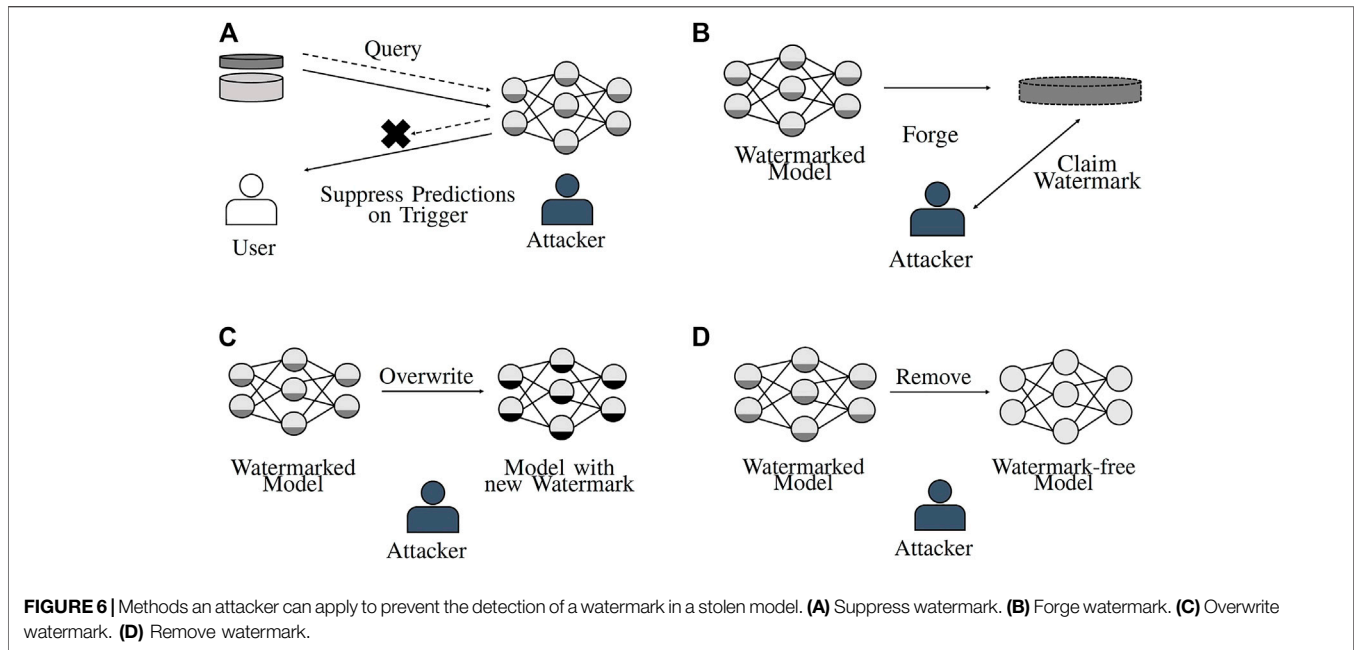
watermarking methods that should still enable identification of stolen model instances, even when an attacker tries to prevent successful verification. The attacks can be grouped into five different classes (from weak to strong), namely, watermark *detection*, *suppression*, *forging*, *overwriting*, and *removal*. While watermark detection is a passive attack in the sense that it does not directly impede successful watermark verification, all other attacks actively try to reduce the watermark's suitability to prove the legitimate model owner's copyright claim. **Figure 6** provides a visualization of the concepts behind the active attacks against watermarking schemes.

5.1 Watermark Detection

The weakest form of attack is concerned with the detection of watermarks in ML models. As stated above, watermark detection is a passive attack that does not directly prevent successful watermark verification by the legitimate model owner. However, discovering the presence of a watermark in a stolen model increases the attacker's knowledge (see Section 4.2.1) and can, thus, be used as a base for further attacks. Additionally, knowing that a stolen model was watermarked gives the attacker the opportunity to adapt this model's behavior in order to avoid successful watermark verification. To detect watermarks included in the target model's parameters, for example, *property inference attacks* have shown to be successful (Wang and Kerschbaum, 2019; Shafieinejad et al., 2021), while backdoor detection can help to identify when models were watermarked by a trigger dataset (Aiken et al., 2021).

5.2 Watermark Suppression

One way to avoid successful watermark verification can consist in suppression of the watermark in a stolen model instance. In a scenario where the legitimate model owner has white-box access to the model for watermark verification, the attacker, therefore, needs to dissimulate any presence of a watermark in the model parameters and behavior. When the legitimate model owner has black-box access to the potentially stolen model, suppressing the reactions of the model to the original watermark trigger might be sufficient for an attacker to prevent detection. This can be



achieved, for example, by identifying possible trigger data points and modifying the model's predictions on them (Hitaj and Mancini, 2018; Namba and Sakuma, 2019). Additionally, so-called *ensemble attacks* that rely on stealing n models from different providers and using them as an ensemble for prediction have been shown to be successful for watermark suppression, since they eliminate individual watermark triggers (Hitaj and Mancini, 2018).

5.3 Watermark Forging

In some cases, the attacker might also be able to forge a watermark on a given model. This attack does not necessarily prevent the legitimate model owner from successfully verifying the watermark in a stolen model instance. However, it creates an ambiguity, in which, for an external entity, such as a legal authority, it is not possible anymore to decide which party has watermarked the given model. Thereby, the attack prevents the legitimate owner from successfully claiming copyright of the intellectual property. Watermark forging might be done by 1) recovering the legitimate owner's watermark and claiming ownership (if there is no binding between the watermark and the owner) (Xu et al., 2019), 2) adding a new watermark that creates ambiguity concerning ownership (Fan et al., 2019), or 3) identifying a fake watermark within the model that coincidentally acts like a real watermark but actually is not (Guo and Potkonjak, 2018).

5.4 Watermark Overwriting

Within the attack setting of watermark overwriting, an attacker knows that the stolen model was watermarked but might have no knowledge about the legitimate owner's concrete watermark or trigger dataset. The attacker can then embed a new watermark into the model to pretend ownership. Note that the new watermark does not necessarily need to be created with the

same watermarking scheme as the original one. In a weak form of the attack, the original watermark is still present after embedding the new one, such that both watermarks co-exist in the model. Therefore, the legitimate owner can then no longer prove (unique) ownership due to ambiguity. In a stronger version of this attack, the attacker might be able to replace the original watermark entirely with the new one (Wang and Kerschbaum, 2018). Thereby, the ownership claim can be taken over completely (Li H. et al., 2019).

5.5 Watermark Removal

As an ultimate solution to prevent successful watermark verification by a legitimate model owner, an attacker might also try to entirely remove the watermark from a stolen model (Zhang et al., 2018). The success of the removal attack usually depends on the attacker's knowledge about 1) *the presence of a watermark*, 2) *the underlying watermarking scheme*, and on the 3) *availability of additional data, for example, for fine-tuning or retraining*. Especially the last point is interesting to consider because many attacks presented in the literature rely on the assumption that an attacker has large amounts of data available to fine-tune (Chen et al., 2021). In reality, an attacker possessing a sufficiently large dataset to train a good model might be less motivated to steal a model, instead of training it from scratch (Chen et al., 2021).

The most general methods that can be used to remove watermarks can be grouped as follows:

- *Fine-Tuning*: use initial model parameters and fine-tune them to a refinement set. Thereby, it is possible to improve a model for certain kinds of data (Sharif Razavian et al., 2014; Tajbakhsh et al., 2016). This process might remove the watermark when model parameters or prediction behaviors are changed (Chen et al., 2021; Shafieinejad et al., 2021).

- *Pruning* (Augasta and Kathirvalavakumar, 2013; Molchanov et al., 2019): cut some redundant parameters and obtain a new model that looks different from the original model but still has a similarly high prediction accuracy. If the parameters containing the watermark are cut, it is no longer possible to verify the watermark (Zhang et al., 2018).
- *Rounding* (Guo, 2018): reduce the precision of the parameters. If the model strongly overfits the watermark triggers, or the watermark is included in the parameters directly, rounding might destroy the watermark information (Yang et al., 2019).
- *Fine-Pruning*: first prune the model architecture and then continue to train. In the benign setting, this helps recover some of the accuracy that may have been lost during pruning. In the presence of backdoors, such as certain watermarks, this also contributes to overwriting their information (Liu et al., 2018; Jia et al., 2021).
- *Model Compression* (Chen et al., 2015; Han et al., 2016): optimize the memory needed to fit the model while preserving its accuracy on the task. This might be necessary in mobile or embedded devices with limited resources. Model compression is performed by, for example, removing insignificant parameters and pruning links between neurons. This can affect the watermark reliability, especially if the neurons used for the watermarking task are different from the ones of the original task because then they can be pruned without losing accuracy in the original task (Yang et al., 2019).
- *Distillation* (Hinton et al., 2015): transfer the prediction power of a potentially very complex teacher model to a less complex student model. This approach finds application, for example, in low-power environments, where simpler models are to be preferred. It can, however, not be guaranteed that the watermark is also included in the simpler model (Yang et al., 2019).
- *Transfer Learning* (Oquab et al., 2014): update the classification task of a model to a related but slightly different task. Therefore, model layers toward the output are modified. This approach saves computational power because large parts of trained models' weights can be applied to the new task with solely small changes. However, the changes within the model layers can lead to a removal of the watermark (Adi, 2018).
- *Computation Optimization* (Jaderberg et al., 2014): reduce computation time of NNs, for example, by low-rank expansion techniques to approximate convolution layers. The reduction might as well lead to a watermark removal (Yang et al., 2019).
- *Backdoor Removal* (Liu et al., 2021): remove backdoors, that is, functionalities in the NN that are not relevant for the original task (see **Section 2.4**). Li H. et al. (2019) pointed out that if the watermarking task is indeed a backdoor function that is too loosely related to the original task, it is possible to remove the watermark by normal backdoor removal attacks against NNs, such as in the study by Wang et al. (2019).
- *Retraining*: an ML model might be trained continuously, instead of being trained once and then released for prediction. Through retraining, models can adapt to potential shifts in the underlying data distribution over time. While retraining, the watermark might be damaged (Adi, 2018; Rouhani et al., 2018a; Zhang et al., 2018; Le Merrer et al., 2020).

In addition to the aforementioned approaches, there also exist more specific attacks proposed in the literature that rely, for example, on regularization (Shafieinejad et al., 2021) or on graph algorithms (Wang et al., 2020).

6 CATEGORIZING WATERMARKING METHODS

This section surveys examples of watermarking methods proposed in the literature to illustrate and validate the taxonomy presented in **Section 3**. The methods are presented in semantic groups based on their embedding method and their distinctive characteristics. Methods that might fit to several groups are presented according to their most distinctive property. For a more comprehensive overview of existing approaches, see the study by Li et al. (2021).

6.1 Embedding Watermarks Into Model Parameters

Early approaches to mark ML models with the aim of including information about the training data into the model parameters were proposed by Song et al. (2017). Among others, they showed how to include information in the least significant bits of the model parameters or the parameters' signs and developed a *correlated value encoding* to maximize a correlation between the model parameters and a given secret. A similar method was then applied by Uchida et al. (2017) as the first explicit watermarking scheme in NNs. The authors interpret the watermark as a T -bit string $\{0,1\}^T$. In order to include it into the model parameters, they use a composed loss function $L(\theta) = L_O + \lambda L_R(\theta)$ with L_O being the loss of the original task and L_R being an *embedding regularizer*. This regularizer imposes a statistical bias on certain model parameters in order to represent the watermark. To project the weights carrying the watermark information, an *embedding parameter* X is used as a secret key needed for watermark embedding and verification.

Wang et al. (2020) extended this work by developing an alternative for the embedding parameter X . Instead, they employ an additional independent NN on selective parameters from the original model to project the watermark. For training of the original model, they use the aforementioned loss function $L(\theta) = L_O + \lambda L_R(\theta)$. To train the additional NN, they apply the binary cross-entropy loss between its output vector and the watermark. The additional NN is not released publicly and serves for watermark verification. Wang and Kerschbaum (2018), however, showed that both the approaches of Uchida et al. (2017) and Wang et al. (2020) do not meet the requirement

TABLE 2 | Techniques to embed watermarks into model parameters sorted alphabetically by author.

Method	Verification	Method	Capacity	Auth	Unique
Uchida et al. (2017): embed bit string watermark to random model parameters' statistical bias	White-box	Biases in weights	Multi-bit	No	No
Fan et al. (2019): adding passport layers into NNs	White- and black-box	New layers	Multi-bit	Yes	No
Song et al. (2017): include information in model parameters, for example, least significant bit or sign	White-box	Model bits	Zero-bit	No	No
Wang and Kerschbaum (2019): create non-detectable watermarked parameters	White-box	Existing parameters	Zero-bit	No	No
Wang et al. (2020): extend the work of Uchida et al. (2017) and include watermarks in quickly converging model parameters	White-box	Existing parameters	Zero-bit	No	No

of watermark secrecy because they cause easily detectable changes in the statistical distribution of the model parameters.

Wang and Kerschbaum (2019) proposed a strategy to create undetectable watermarks in a white-box setting based on generative adversarial networks (GANs). The watermarked model h_O serves as the generator, whereas a watermark detector that detects changes in the statistical distribution of the model parameters serves as a discriminator h_D . During training, h_O is encouraged to generate nondetectable watermarks, whereas h_D tries to distinguish watermarked from non-watermarked models. Both optimize the following functions, respectively:

$$h_O: \min(L_O(\theta) + \lambda L_R(\theta) - \lambda_2 \log h_D(\theta)) \quad (1)$$

$$h_D: \max(\log h_D(\theta_{non}) + \log(1 - h_D(\theta))) \quad (2)$$

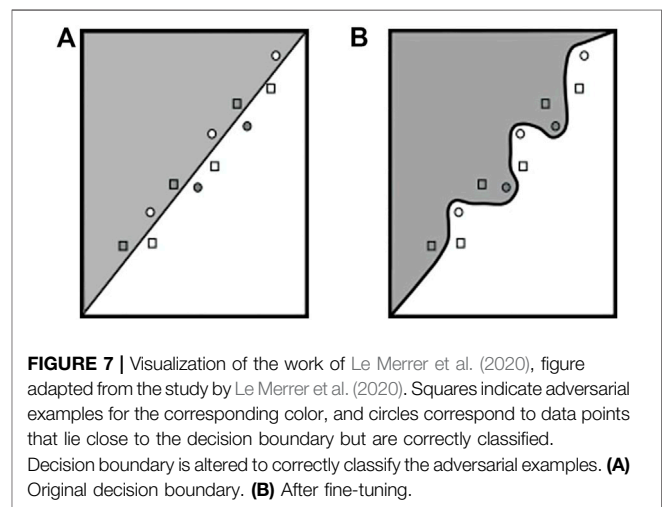
where θ_{non} refers to the parameters of non-watermarked previously trained models.

Fan et al. (2019) suggested embedding *passport-layers* with digital signatures into NNs for ownership verification. The passport layers are inserted into convolutional neural networks (CNNs) and calculate hidden parameters without which the model's inference accuracy is reduced. For verification, the authors developed three mechanisms based on different strategies of distributing and verifying the passports.

Table 2 provides an overview on the mentioned methods that rely on inserting watermarks directly into the model parameters.

6.2 Using Pre-Defined Inputs as Triggers

Le Merrer et al. (2020) proposed directly marking the model's action itself by slightly moving the decision boundary through adversarial retraining such that specific queries can exploit it. Therefore, their approach first identifies adversarial samples and normal data points that are very close to the decision boundary. Then, the trigger dataset is composed by 50% of the adversarial examples and 50% of the data points that do not cause misclassification but are close to the decision boundary. Afterward, the trained classifier is fine-tuned to predict the trigger data points to their correct original class. See Figure 7 for a visualization of this approach. The resulting labeled data points are supposed to serve as an expressive trigger dataset. Namba and Sakuma (2019) argued that this method offers weak integrity because, nowadays, adversarial retraining is broadly



used to create more robust models; hence, a non-watermarked model can be mistaken for being watermarked.

Adi (2018) considered watermarking from a cryptographic point of view. The authors generated abstract color images with randomly assigned classes as a trigger dataset. In order to guarantee non-trivial ownership, a set of commitments is created over the image/label pairs before embedding the watermark into the model. Thereby, at verification time, ownership can be proven by selectively revealing these commitments. A similar approach for creating the watermark trigger dataset was also described by Zhang et al. (2018). They include irrelevant data samples, for example, from another unrelated dataset, as watermarks into the training data. Those samples are labeled with classes from the original model output. During training, the model learns to assign real images and those trigger samples to the corresponding classes.

Rouhani et al. (2018a) developed an approach of including the watermark as a T -bit string into the probability density function (*pdf*) of the data abstraction obtained in different network layers. These layers' activation maps at intermediate network layers roughly follow Gaussian distributions. The legitimate model's owner can choose in how many of those they want to embed the watermark string. Afterward, the network is trained to incorporate the watermark information in the mean values of the selected distributions. A projection matrix A can be used to

map the selected distribution centers to the binary watermark vector. In a white-box setting, this matrix A is utilized for verification. For black-box verification, a trigger dataset can be constructed from data points whose features lie in the model's unused regions, that is, samples at the tail regions of the pdf. In contrast to methods including the watermark in the static model content, like in the study by Uchida et al. (2017), this approach changes the dynamic content of the model, namely, the activations that depend on the data and the model. This results in a more flexible and not that easily detectable change (Rouhani et al., 2018a).

Chen et al. (2019b) proposed taking the model owner's binary signature as a watermark for an NN. Their aim is to combine the advantage from black-box and white-box watermark extraction, that is, weaker assumptions on the attacker's power and large capacity at the same time. To include the signature in the model, they build a model-dependent encoding scheme that clusters the model's output activations into two groups according to their similarities, one group for class 0 and one for class 1. The binary signature is then included in the model's output activations and can be verified through a designated trigger dataset that can be passed to the model in order to retrieve the signature.

6.2.1 Trigger Dataset Creation Based on Original Training Data

Some watermarking approaches rely on inserting forms of digital media watermarks into the original training data in order to create the model's trigger dataset. The approach by Guo and Potkonjak (2018) generated an n -bit signature of the model owner and embedded it into the training data in order to generate the trigger dataset. The authors made sure that the altered images from the trigger dataset obtained different labels than the original data points that they were based on.

Zhang et al. (2018) described algorithms for watermarking NNs for image classification with remote black-box verification mechanisms. One of their algorithms embeds meaningful content together with the original training data as a watermark. An example for this approach is embedding a specific string (which could be the company name) into a picture of the training set when predicting images and assigning a different label than the original one to the modified sample. Instead of a meaningful string, it is also possible to embed noise into the original training data. A similar approach to the first algorithm of the study by Zhang et al. (2018) was proposed by Li Z. et al. (2019), who combined some ordinary data samples with an exclusive "logo" and trained the model to predict them into a specific label. To keep these trigger samples as close as possible to the original samples, an autoencoder is used whose discriminator is trained to distinguish between training and trigger samples with the watermarks. Sakazawa et al. (2019) proposed a cumulative and visual decoding of watermarks in NNs, such that patterns embedded into the training data can become visual for an authentication by a third party.

See **Table 4** for an overview on methods that use original training data to generate the trigger dataset.

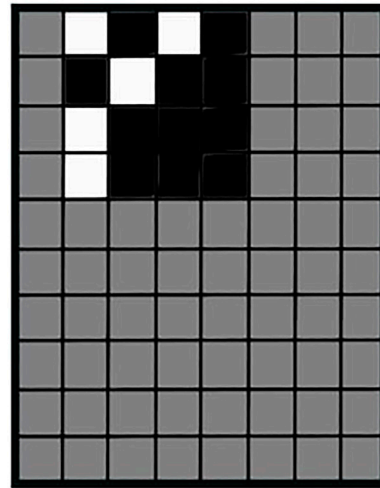


FIGURE 8 | Watermark pattern, figure adapted from the study by Li H. et al. (2019).

6.2.2 Robust Watermarking

A problem of watermarking methods that rely on using a trigger dataset from a different distribution than the original training data is that the models are actually trained for two different (and independent) tasks. Research has shown that when these tasks are more or less unrelated, it is possible to remove the watermarks through attacks described in **Section 5.5** without affecting the model's accuracy on the original task learned through the training data (Li H. et al., 2019; Jia et al., 2021).

For example, Yang et al. (2019) showed that distillation (Papernot et al., 2016) is effective to remove watermarks. This results from the fact that the information learned from the watermark trigger dataset is redundant and independent of the main task. Hence, this information is not transferred to the resulting surrogate model when doing distillation. As a solution, the authors described an "ingrain"-watermarking method that regularizes the original NN with an additional NN that they refer to as the ingrain model g_ω , which contains the watermark information. Regularization is performed with a specific ingrain loss $C(h_{\theta,T}(x), g_\omega(x))$ (T is the temperature in the softmax) that causes the watermark information to be embedded into the same model parameters as the main classification task. The joint loss function over the training data D_s is given by the following:

$$L_{D_s}(h_\theta) = \frac{1}{|D_s|} \sum_{x \in D_s} C(h_\theta(x), y) + \delta C(h_{\theta,T}(x), g_\omega(x)), \quad (3)$$

where the labels are indicated by y , and λ determines the degree of ingrain.

Jia et al. (2021) proposed a similar idea that relies on "entangled watermarking embeddings." The entanglement is used to make the model extract common features of the data that represent the original task and the data that encode the watermarks and stem from a different distribution. Therefore, the authors apply the *soft nearest neighbor loss* (SNNL)

TABLE 3 | Techniques using a *specific trigger dataset* as a watermark sorted alphabetically by author.

Approach	Verification	Method	Capacity	Auth	Unique
Adi (2018): abstract color images with random classes as trigger dataset	Black-box	Backdoor	Zero-bit	Yes	No
Chen et al. (2019b): include the model owner's binary signature in output activations	Black-box	Output activations	Multi-bit	Yes	No
Jia et al. (2021): entangled watermarks through training with <i>soft nearest neighbor loss</i>	Black-box	Backdoor	Zero-bit	No	No
Le Merrer et al. (2020): adversarial decision boundary modification through a trigger sample consisting of adversarial examples	Black-box	Adversarial examples	Zero-bit	No	No
Li et al. (2019a): null embedding watermark consisting of a pixel pattern	Black-box	Watermark in data	Multi-bit	Yes	No
Namba and Sakuma (2019): exponential weighting: enforce watermark predictions with higher weights during training	Black-box	Backdoor	Zero-bit	No	No
Rouhani et al. (2018a): include watermark in the probability density function of network layers	White- and black-box	Biases in weights	Multi-bit	No	No
Yang et al. (2019): train an additional model (called an ingrain) that contains the watermark information <i>via</i> its prediction on training data	Black-box	Backdoor	Zero-bit	No	No

TABLE 4 | Techniques to embed *watermarks into the training data* in order to create the trigger dataset.

Approach	Verification	Method	Capacity	Auth	Unique
Guo and Potkonjak (2018): model owner's signature embedded into the training data	Black-box	Watermark in data points	Multi-bit	Yes	No
Li et al. (2019b): include "logo" in the training data images and use an auto-encoder to make trigger samples close to original data	Black-box	Watermark in data points	Zero-bit	Yes	No
Zhang et al. (2018): include (meaningful) information in the training data samples	Black-box	Watermark in data points	Multi-bit	Yes	No

(Salakhutdinov and Hinton, 2007). Informally spoken, the SNNL measures *entanglement* over labeled data, that is, how close pairs of points from the same class are relative to pairs of points from different classes (Frosst et al., 2019). Points from different groups that are closer relative to the average distance between two points are called *entangled*. Using entanglement when including a watermark ensures that the watermark and the original task are represented by the same sub-model and not by different ones that may be harmed during extraction. Hence, it becomes more difficult for an attacker to extract the model without its watermarks. At the same time, through the entanglement, removing the watermark would result in a decrease in model performance on the original task.

Namba and Sakuma (2019) described a method they called "exponential weighting." They generated a watermark trigger by random sampling from the training distribution and assigning wrong labels to that sample for training. To protect the watermark against pruning or retraining attacks, the authors proposed embedding the samples by exponential weighting, that is, imprint trigger samples with greater force and cause the model to learn them profoundly. Therefore, they increased the weight of the model parameters that are involved in the prediction exponentially, and thereby, made the prediction depend mainly on some few and very large model parameters which are harder to change through the mentioned attacks.

Li H. et al. (2019) developed a "null embedding" for including watermarks into the model's initial training, such that attackers are not able to remove them or include their own watermarks on top. Therefore, they generate a filter pattern p as shown in **Figure 8**. Image pixels under the white pattern pixels are

changed to a very large negative number, image pixels under black pattern pixels are changed to a very large positive number, and pixels under gray pattern pixels stay unchanged. Over this process, the predicted class of the image needs to stay the same as for the original image. Using extreme values and setting strong deterministic constraints on the optimization during learning leads to strong watermark inclusion. To create a binding between the owner and the pixel pattern, the authors propose using the owner's signature and a deterministic hash function to generate the pattern (Li H. et al., 2019).

See **Table 3** for an overview on all the mentioned methods that rely on using a trigger dataset to watermark ML models.

6.2.3 Unique Watermarking

The requirements for unique watermarking schemes are the same as the ones for common watermarking methods (see **Table 1**) but extended by the two following points Chen et al. (2019a):

- *Uniqueness*. Watermarks should be unique for each user. This serves to identify model instances individually.
- *Scalability*. Unique watermarking schemes should scale to many users in a system. This allows for a large-scale distribution of the target model.

Chen et al. (2019a) proposed an end-to-end collusion-secure watermarking framework for white-box settings. Their approach is based on anti-collusion codebooks for individual users which are incorporated in the pdf of the model weights. The incorporation is achieved by using a watermark-specific regularization loss during training.

TABLE 5 | Techniques to generate unique watermarks that can be verified by querying a trigger dataset.

Approach	Verification	Method	Capacity	Auth	Unique
Chen et al. (2019a): end-to-end unique watermarking scheme based on anti-collusion codebooks for individual users	White-box	Biases in weights	Multi-bit	Yes	Yes
Xu et al. (2019): embed a unique serial number based on the model owner's signature and create endorsement by certification authority	Black-box	Backdoor	Multi-bit	Yes	Yes

Xu et al. (2019) embedded a serial number in NNs for model ownership identification. Their solution generates a unique serial number as a watermark and creates an endorsement by a certification authority on it. Serial numbers are generated by the owner through a digital signature algorithm (based on the owner's private key). During model training, the serial number is fitted into the model, together with the original task, by having a second loss, such that owner verification can be achieved by sending trigger inputs, extracting the serial number and verifying it with the certification authority.

See **Table 5** for an overview on these methods generating unique watermarks.

6.3 Using Model Fingerprints to Identify Potentially Stolen Instances

Instead of explicitly adding watermark information into an ML model, some methods use existing features of the model in order to identify potentially stolen instances. This offers the advantages that no overhead is added to the original training task and that the model's original prediction abilities are not affected. However, as those methods do not actively alter the model in order to include a watermark, they will only be mentioned very briefly in this document.

Zhao et al. (2020) used adversarial examples as fingerprints for NNs. They identified some special adversarial examples within NNs that they called "adversarial marks." Those adversarial marks differ from traditional adversarial examples in their transferability: they show high transferability between models that are similar and low transferability between models that are different. The authors argued that adversarial marks represent suitable model fingerprints as they are difficult to remove due to the number and type of adversarial examples being practically infinite.

Lukas et al. (2019) also exploited the transferability of adversarial examples in order to verify the ownership of ML models. They define the class of *conferrable* adversarial examples. Those examples transfer only to surrogate models of a target model (potential illegitimate copies) but not to reference models, that is, similar models trained on similar data for a related task. By querying those examples to a model, one can identify whether this model is a copy of the target model or not. The authors also proposed a generation method for this class of adversarial examples and proved that this watermarking method is robust against distillation (Papernot et al., 2016) and weaker attacks.

7 DISCUSSION

This section, first, discusses the pros and cons of existing classes of watermarking methods. Afterward, it revisits the requirements

for effective watermarking to provide a structured reasoning for choosing or designing adequate watermarking schemes. Finally, it presents limitations of existing methods and proposes an outlook on promising research directions.

7.1 Discussing Potential Shortcomings

First, in trigger dataset-based watermarking approaches, watermark detection relies on the model's reaction on queries from the trigger dataset. If the agreement of the prediction on them to the trigger dataset's original labels overpasses a given threshold, this suggests the presence of the watermark. However, defining a suitable threshold to identify a stolen model requires thorough tuning. If the threshold is set too high, slight modifications in a stolen model might already be sufficient to prevent watermark detection, which violates the reliability requirement. If the threshold is set too low, different models might erroneously be identified as stolen model instances, which violates the integrity requirement. Hence, the choice of a threshold also expresses a trade-off between reliability and integrity.

In addition to the issue of choosing an adequate threshold, Hitaj and Mancini (2018) formulated the general disadvantage for the scenario of public verification. They argued that after the verification algorithm is run against a stolen model, the attacker is in possession of the trigger dataset, which enables them to fine-tune the model on those data points to remove the watermark. Hence, in order to run several verifications, the original trigger dataset needs to be divided or there have to be several trigger datasets. This class of approaches, thus, has its limitation due to the maximum amount of backdoors that can be embedded in an NN. Approaches that need few queries, like the one by Jia et al. (2021), might allow for a higher number of independent watermark verifications with the same model capacities.

Second, watermarking schemes embedding watermarks into the ML models' parameters without taking precautions do not only leave a detectable trace in the model and, hence, violate the secrecy requirement (Wang and Kerschbaum, 2018) but they also often rely on white-box access for verification. Even though in some scenarios, the latter might be feasible, still, assuming black-box access often is a more realistic scenario. Therefore, such schemes might be suitable only to very specific applications.

Third, watermarking schemes that do not exhibit a verifiable link between the watermark and the legitimate model owner enable an attacker to forge the watermark. A similar disadvantage exists for watermarking schemes that rely on nonspecific data points as trigger data [e.g., (Adi, 2018; Guo and Potkonjak, 2018)]. Those might enable an attacker to choose (random) different points than the initial watermark, in order to claim having marked the model with them as triggers. Approaches that

do not allow already marked models to be marked again, like the one by Li H. et al. (2019), can prevent this threat.

Finally, due to their instability, their potentially low robustness against fine-tuning or retraining, and, in some cases, their transferability, which might violate watermark integrity, adversarial examples used for watermarking (Le Merrer et al., 2020) or fingerprinting (Lukas et al., 2019; Zhao et al., 2020) might exhibit important drawbacks. Namba and Sakuma (2019) pointed out that especially an adaption of the model's decision boundary according to some adversarial examples, as in the study by Le Merrer et al. (2020), might be likely to violate the integrity requirement because its effect is similar to the effect of adversarial retraining, a method commonly used to make ML models more robust.

7.2 Discussing Requirements

In addition to considering the pros and cons of existing classes of watermarking methods, this section discusses the question of choosing or creating reliable watermarking methods by revisiting the requirements presented in **Table 1**.

- *Fidelity*: to guarantee fidelity, existing watermarking schemes aim at preserving model performance on the original task. Depending on the scheme, this can be achieved through different means, for example, only minimally altering the original decision boundary (Le Merrer et al., 2020) or including the watermark to early converging model weights (Wang et al., 2020).
- *Robustness*: if the trigger dataset stems from a significantly different distribution than the original data, (Adi, 2018; Fan et al., 2019), the model learns two different (and independent) tasks. Therefore, it is possible to extract them independently or to remove the watermark without causing an impact on the model's prediction performance. Thus, to achieve robustness, watermarking schemes need to take measures to create a relation between both tasks and to enforce the watermark to the model such that it cannot be removed easily.
- *Reliability*: certain factors can influence watermark reliability. First, similarly as for robustness, if the trigger dataset stems from a different distribution than the original dataset, reactions of the stolen model to the watermark triggers can be suppressed by an attacker. Second, all schemes that rely solely on white-box verification might offer lower reliability, as such access to all potentially stolen model instances might not always be a realistic assumption, which might prevent successful verification.
- *Integrity*: quantifying watermarking schemes' integrity is a challenging task, as it requires judging how (potentially all other) non-watermarked models react on the given trigger dataset. A good trigger dataset is characterized by the uniqueness of the watermarked model's predictions on it, in order to accuse no honest parties with similar models of theft.
- *Capacity*: capacity expresses how much information can be included in the watermark. To allow for specific tasks, such as owner authentication, inserting multi-bit watermarks is common practice (Guo and Potkonjak, 2018; Zhang et al., 2018).
- *Secrecy*: watermarking schemes that change the model parameters in a detectable way, for example (Uchida et al., 2017), violate the secrecy requirement. To prevent watermark detection, adding them to the dynamic model content (Rouhani et al., 2018b) or taking measures to force the model parameters to stay roughly the same (Li Z. et al., 2019) are possible solutions.
- *Efficiency*: efficiency can be evaluated with regard to embedding and verification time, that is, the overhead in training and the computation time or number of queries needed to verify the watermark. Most existing work does not explicitly evaluate computational overhead of their approaches. Jia et al. (2021) presented one of the few evaluations of efficiency and came to the conclusion that for their approach, the trigger dataset should consist of more than half the amount of data samples as the original data. Therefore, the model needs to train with 150–200% of the original data. Especially for large datasets, this might result in large overhead.
- *Generality*: not all existing schemes directly generalize to all datasets; for example, the study by Chen et al. (2019b) needs a different watermark encoding scheme on each dataset. Such behavior can be considered as lacking generality.

7.3 Limitations and Outlook

Based on the evaluation of existing schemes and their security requirements, this section presents current limitations and promising future research directions.

The largest limitation of watermarking in general is that it represents a passive defense. That is, watermarking schemes cannot prevent theft but only detect it afterward. Some research was conducted in order to issue security warnings, once an ML model is about to reveal enough information that an attacker or a group of attackers might be able to extract its functionality (Kesarwani et al., 2018). Further research focused on creating models that solely achieve high accuracy when being queried by an authorized user (Chen and Wu, 2018). Other work was directed toward the development of models that are more difficult to steal, for example, by only returning hard labels and no probabilities per output class, by perturbing the prediction outputs (Orekondy et al., 2019), or by designing networks that are extremely sensitive to weight changes, which makes it difficult for an attacker to steal and adapt them (Szentannai et al., 2019). Future work could focus on how to integrate watermarking within such active defense strategies against model stealing.

Furthermore, most current watermarking schemes apply solely to image data, so far. Only a few exceptions, for example, the one by Jia et al. (2021), have proven the applicability of their schemes to other data types. Future work will have to focus on examining the generality and universal applicability of existing schemes and, if necessary, their adaptation or extension.

Moreover, most watermarking approaches proposed so far also apply solely to classification tasks. There exist only a few works on watermarking in other ML domains, like reinforcement learning (Behzadan and Hsu, 2019) and data generation with GANs (Ong et al., 2021) or image captioning (Lim et al., 2020).

Therefore, the development of watermarking schemes for other ML applications represents a promising future challenge.

Additionally, so far, existing watermarking schemes are mainly applied and evaluated on rather small research datasets, like MNIST (LeCun et al., 2010). Therefore, the question of their scalability remains open. Approaches that require training with up to double the initial amount of data might, hence, not be applicable to every scenario. Thus, future work should assess the practical applicability of existing watermarking schemes to larger real-world datasets and analyze whether the properties they exhibit on the research datasets (efficiency of training, reasonable trigger dataset size, integrity, etc.) hold.

Finally, once those watermarking schemes meet all the technical requirements, another challenge will lie in their adaptation to real-world workflows. Especially the juridical and organizational workflows will have to be adapted in order to enable asserting ownership claims based on the watermarks.

8 CONCLUDING REMARKS

Nowadays, ML is used in an increasing number of domains. With growing complexity of the applied models, employing watermarks to protect intellectual property incorporated in those models has become a major focus both in academia and industry. This systematic review provided a framework for articulating a comprehensive view on different watermarking schemes. It introduced a taxonomy to classify and compare existing methods, presented a systematization of the requirements for and attacks against watermarking schemes, and formulated a unified threat model. Guided by the

REFERENCES

- Adi, Y., Carsten, B., Moustapha, C., Benny, P., and Joseph, K. (2018). Turning Your Weakness into a Strength: Watermarking Deep Neural Networks by Backdooring. 27th USENIX Security Symposium (USENIX Security 18), 1615–1631.
- Aiken, W., Kim, H., Woo, S., and Ryoo, J. (2021). Neural Network Laundering: Removing Black-Box Backdoor Watermarks from Deep Neural Networks. *Comput. Security* 106, 102277. doi:10.1016/j.cose.2021.102277
- Ateniese, G., Felici, G., Mancini, L. V., Spognardi, A., Villani, A., and Vitali, D. (2013). Hacking Smart Machines with Smarter Ones: How to Extract Meaningful Data from Machine Learning Classifiers. ArXiv preprint arXiv:1306.4447.
- Augasta, M., and Kathirvalavakumar, T. (2013). Pruning Algorithms of Neural Networks—A Comparative Study. *Open Comp. Sci.* 3, 105–115. doi:10.2478/s13537-013-0109-x
- Batina, L., Bhasin, S., Jap, D., and Picek, S. (2018). Csi Neural Network: Using Side-Channels to Recover Your Artificial Neural Network Information. ArXiv preprint arXiv:1810.09076.
- Behzadan, V., and Hsu, W. (2019). Sequential Triggers for Watermarking of Deep Reinforcement Learning Policies. ArXiv preprint arXiv:1906.01126v1.
- Carlini, N., Jagielski, M., and Mironov, I. (2020). Cryptanalytic Extraction of Neural Network Models. ArXiv preprint arXiv:2003.04884. doi:10.1007/978-3-030-56877-1_7
- Chen, H., Rouhani, B. D., Fu, C., Zhao, J., and Koushanfar, F. (2019a). “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. Editors A. El Saddik, A. Del Bimbo, Z. Zhang,

taxonomy, relevant prior research was surveyed. This work can serve as a solid foundation for analyzing existing watermarking methods, designing new ones, or choosing adequate solutions to given scenarios. Therefore, it can be used as a reference for researchers and ML practitioners over all domains.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/Supplementary Material, further inquiries can be directed to the corresponding author.

AUTHOR CONTRIBUTIONS

FB: the author has conducted the review, written the manuscript, and created the figures.

FUNDING

We acknowledge support by the Open Access Publication Initiative of Freie Universität Berlin.

ACKNOWLEDGMENTS

We would like to thank Hengrui Jia for the fruitful discussion about the topic and his feedback on the manuscript.

- A. Hauptmann, K. S. Candan, M. Bertini, et al. (New York: ACM Press), 105–113. doi:10.1145/3323873.3325042
- Chen, H., Rouhani, B. D., and Koushanfar, F. (2019b). Blackmarks: Blackbox Multibit Watermarking for Deep Neural Networks. ArXiv preprint arXiv:1904.00344v1.
- Chen, M., and Wu, M. (2018). Protect Your Deep Neural Networks from Piracy. In IEEE International Workshop on Information Forensics and Security (WIFS). IEEE, 1–7. doi:10.1109/WIFS.2018.8630791
- Chen, W., Wilson, J., Tyree, S., Weinberger, K., and Chen, Y. (2015). Compressing Neural Networks with the Hashing Trick. In International conference on machine learning. New York: PMLR, 2285–2294.
- Chen, X., Wang, W., Bender, C., Ding, Y., Jia, R., Li, B., et al. (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. 321–335. doi:10.1145/3433210.3453079
- Fan, L., Ng, K. W., and Chan, C. S. (2019). “Rethinking deep neural network ownership verification: embedding passports to defeat ambiguity attacks”. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 33, 4714–4723.
- Frosst, N., Papernot, N., and Hinton, G. (2019). Analyzing and Improving Representations with the Soft Nearest Neighbor Loss. In International Conference on Machine Learning. San Juan, Puerto Rico: PMLR, 2012–2020.
- Guo, J., and Potkonjak, M. (2018). “Watermarking Deep Neural Networks for Embedded Systems,” in Proceedings of the International Conference on Computer-Aided Design - ICCAD '18. Editor I. Bahar (New York: ACM Press), 1–8. doi:10.1145/3240765.3240862
- Guo, Y. (2018). A Survey on Methods and Theories of Quantized Neural Networks. ArXiv preprint arXiv:1808.04752.

- Han, S., Liu, X., Mao, H., Pu, J., Pedram, A., Horowitz, M. A., et al. (2016). Efficient Inference Engine on Compressed Deep Neural Network. *SIGARCH Comput. Archit. News* 44, 243–254. doi:10.1145/3007787.3001163
- Hinton, G., Vinyals, O., and Dean, J. (2015). Distilling the Knowledge in a Neural Network. ArXiv preprint arXiv:1503.02531.
- Hitaj, D., and Mancini, L. V. (2018). Have You Stolen My Model? Evasion Attacks against Deep Neural Network Watermarking Techniques. ArXiv preprint arXiv:1809.00615v1.
- Jaderberg, M., Vedaldi, A., and Zisserman, A. (2014). Speeding up Convolutional Neural Networks with Low Rank Expansions. In Proceedings of the British Machine Vision Conference. London: BMVA Press, 13. doi:10.5244/c.28.88
- Jagielski, M., Carlini, N., Berthelot, D., Kurakin, A., and Papernot, N. (2020). High Accuracy and High Fidelity Extraction of Neural Networks. In 29th {USENIX} Security Symposium ({USENIX} Security 20). 1345–1362.
- Jia, H., Choquette-Choo, C. A., Chandrasekaran, V., and Papernot, N. (2021). Entangled Watermarks as a Defense against Model Extraction. In 30th {USENIX} Security Symposium ({USENIX} Security 21). 18.
- Kesarwani, M., Mukhoty, B., Arya, V., and Mehta, S. (2018). Model Extraction Warning in Mlaas Paradigm. In Proceedings of the 34th Annual Computer Security Applications Conference on - ACSAC '18. New York: ACM Press, 371–380. Unknown. doi:10.1145/3274694.3274740
- Le Merrer, E., Perez, P., and Trédan, G. (2020). Adversarial Frontier Stitching for Remote Neural Network Watermarking. *Neural Comput. Appl.* 32 (13), 9233–9244.
- LeCun, Y., Cortes, C., and Burges, C. (2010). *MNIST Handwritten Digit Database*. Florham Park, NJ.
- Li, H., Wenger, E., Zhao, B. Y., and Zheng, H. (2019a). Piracy Resistant Watermarks for Deep Neural Networks. ArXiv preprint arXiv:1910.01226v2.
- Li, Y., Wang, H., and Barni, M. (2021). A Survey of Deep Neural Network Watermarking Techniques. ArXiv preprint arXiv:2103.09274. doi:10.1016/j.neucom.2021.07.051
- Li, Z., Hu, C., Zhang, Y., and Guo, S. (2019b). “How to Prove Your Model Belongs to You,” in Proceedings of the 35th Annual Computer Security Applications Conference on - ACSAC '19. Editor D. Balenson (New York: ACM Press), 126–137. doi:10.1145/3359789.3359801
- Lim, J. H., Chan, C. S., Ng, K. W., Fan, L., and Yang, Q. (2020). Protect, Show, Attend and Tell: Image Captioning Model with Ownership protection. ArXiv preprint arXiv:2008.11009.
- Liu, K., Dolan-Gavitt, B., and Garg, S. (2018). Fine-pruning: Defending against Backdooring Attacks on Deep Neural Networks. In International Symposium on Research in Attacks, Intrusions, and Defenses. Springer, 273–294. doi:10.1007/978-3-030-00470-5_13
- Liu, X., Li, F., Wen, B., and Li, Q. (2021). Removing Backdoor-Based Watermarks in Neural Networks with Limited Data. In 2020 25th International Conference on Pattern Recognition (ICPR). IEEE, 10149–10156. doi:10.1109/icpr48806.2021.9412684
- Lukas, N., Zhang, Y., and Kerschbaum, F. (2019). Deep Neural Network Fingerprinting by Conferrable Adversarial Examples. ArXiv preprint arXiv:1912.00888v2.
- Molchanov, P., Tyree, S., Karras, T., Aila, T., and Kautz, J. (2019). Pruning Convolutional Neural Networks for Resource Efficient Inference. In 5th International Conference on Learning Representations, ICLR 2017- Conference Track Proceedings. 17.
- Namba, R., and Sakuma, J. (2019). “Robust Watermarking of Neural Network with Exponential Weighting,” in Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security - Asia CCS '19. Editors S. Galbraith, G. Russello, W. Susilo, D. Gollmann, E. Kirda, and Z. Liang (New York: ACM Press), 228–240. doi:10.1145/3321705.3329808
- Ong, D. S., Chan, C. S., Ng, K. W., Fan, L., and Yang, Q. (2021). Protecting Intellectual Property of Generative Adversarial Networks from Ambiguity Attacks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 3630–3639.
- Oquab, M., Bottou, L., Laptev, I., and Sivic, J. (2014). Learning and Transferring Mid-level Image Representations Using Convolutional Neural Networks. In Proceedings of the IEEE conference on computer vision and pattern recognition. 1717–1724. doi:10.1109/cvpr.2014.222
- Orekondy, T., Schiele, B., and Fritz, M. (2019). Prediction Poisoning: Towards Defenses against Dnn Model Stealing Attacks. ArXiv preprint arXiv:1906.10908v2.
- Papernot, N., McDaniel, P., Sinha, A., and Wellman, M. P. (2018). Sok: Security and Privacy in Machine Learning. In 3rd IEEE European Symposium on Security and Privacy. Conference Publishing Services, IEEE Computer Society, 16. doi:10.1109/eurosp.2018.00035
- Papernot, N., McDaniel, P., Wu, X., Jha, S., and Swami, A. (2016). Distillation as a Defense to Adversarial Perturbations against Deep Neural Networks. In IEEE Symposium on Security and Privacy (SP). IEEE, 582–597. doi:10.1109/SP.2016.41
- Rouhani, B. D., Chen, H., and Koushanfar, F. (2018a). DeepSigns: A Generic Watermarking Framework for Ip protection of Deep Learning Models. ArXiv preprint arXiv:1804.00750v2.
- Rouhani, B. D., Riaz, M. S., and Koushanfar, F. (2018b). DeepSecure: Scalable Provably-Secure Deep Learning. In 2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC), 6. IEEE. doi:10.1109/dac.2018.8465894
- Saini, L. K., and Shrivastava, V. (2014). A Survey of Digital Watermarking Techniques and its Applications. ArXiv preprint arXiv:1407.4735.
- Sakazawa, S., Myodo, E., Tasaka, K., and Yanagihara, H. (2019). Visual Decoding of Hidden Watermark in Trained Deep Neural Network. In IEEE Conference on Multimedia Information Processing and Retrieval (MIPR). IEEE, 371–374. doi:10.1109/MIPR.2019.00073
- Salakhutdinov, R., and Hinton, G. (2007). “Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics,” in *Proceedings of Machine Learning Research*. Editors Meila, M., and Shen, X. (San Juan, Puerto Rico: PMLR) Vol. 2, 412–419.
- Shafieinejad, M., Lukas, N., Wang, J., Li, X., and Kerschbaum, F. (2021). On the Robustness of Backdoor-Based Watermarking in Deep Neural Networks. In Proceedings of the 2021 ACM Workshop on Information Hiding and Multimedia Security. 177–188. doi:10.1145/3437880.3460401
- Sharif Razavian, A., Azizpour, H., Sullivan, J., and Carlsson, S. (2014). Cnn Features Off-The-Shelf: an Astounding Baseline for Recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition workshops. 806–813. doi:10.1109/cvprw.2014.131
- Song, C., Ristenpart, T., and Shmatikov, V. (2017). “Machine Learning Models that Remember Too Much,” in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. Editor B. Thuraisingham (New York: ACM), 587–601. doi:10.1145/3133956.3134077
- Szentannai, K., Al-Afandi, J., and Horváth, A. (2019). Mimosanet: An Unrobust Neural Network Preventing Model Stealing. ArXiv preprint arXiv:1907.01650v1; Presented at CVPR workshop: Adversarial Machine Learning in Real-World Computer Vision Systems.
- Tajbakhsh, N., Shin, J. Y., Gurudu, S. R., Hurst, R. T., Kendall, C. B., Gotway, M. B., et al. (2016). Convolutional Neural Networks for Medical Image Analysis: Full Training or fine Tuning. *IEEE Trans. Med. Imaging* 35, 1299–1312. doi:10.1109/tmi.2016.2535302
- Tramèr, F., Zhang, F., Juels, A., Reiter, M. K., and Ristenpart, T. (2016). Stealing Machine Learning Models via Prediction apis. 25th USENIX Security Symposium (USENIX Security 16), 601–618.
- Uchida, Y., Nagai, Y., Sakazawa, S., and Satoh, S. (2017). “Embedding Watermarks into Deep Neural Networks,” in Proceedings of the 2017 ACM on International Conference on Multimedia Retrieval. Editor B. Ionescu (New York: ACM), 269–277. doi:10.1145/3078971.3078974
- Wang, B., Yao, Y., Shan, S., Li, H., Viswanath, B., Zheng, H., et al. (2019). Neural Cleanse: Identifying and Mitigating Backdoor Attacks in Neural Networks. In IEEE Symposium on Security and Privacy. IEEE Computer Society, 707–723. doi:10.1109/sp.2019.00031
- Wang, J., Wu, H., Zhang, X., and Yao, Y. (2020). Watermarking in Deep Neural Networks via Error Back-Propagation. *Electron. Imaging* 2020, 22-1–22-9. doi:10.2352/issn.2470-1173.2020.4.mwsf-022
- Wang, T., and Kerschbaum, F. (2018). Attacks on Digital Watermarks for Deep Neural Networks. In 2018 IEEE International Conference on Acoustics, Speech, and Signal Processing.
- Wang, T., and Kerschbaum, F. (2019). Robust and Undetectable white-box Watermarks for Deep Neural Networks. ArXiv preprint arXiv:1910.14268.
- Wei, L., Luo, B., Li, Y., Liu, Y., and Xu, Q. (2018). I Know what You See: Power Side-Channel Attack on Convolutional Neural Network Accelerators. In Proceedings of the 34th Annual Computer Security Applications Conference. 393–406.
- Xu, X., Li, Y., and Yuan, C. (2019). A Novel Method for Identifying the Deep Neural Network Model with the Serial Number. ArXiv preprint arXiv:1911.08053v1.

- Xue, M., He, C., Wang, J., and Liu, W. (2020). Dnn Intellectual Property protection: Taxonomy, Methods, Attack Resistance, and Evaluations. ArXiv preprint arXiv: 2011.13564.
- Yang, Z., Dang, H., and Chang, E.-C. (2019). Effectiveness of Distillation Attack and Countermeasure on Neural Network Watermarking. ArXiv preprint arXiv: 1906.06046v1.
- Zhang, J., Gu, Z., Jang, J., Wu, H., Stoecklin, M. P., Huang, H., et al. (2018). “Protecting Intellectual Property of Deep Neural Networks with Watermarking,” in Proceedings of the 2018 on Asia Conference on Computer and Communications Security - ASIACCS '18. Editors J. Kim, G.-J. Ahn, S. Kim, Y. Kim, J. Lopez, and T. Kim (New York: ACM Press), 159–172. doi:10.1145/3196494.3196550
- Zhao, J., Hu, Q., Liu, G., Ma, X., Chen, F., and Hassan, M. M. (2020). Afa: Adversarial Fingerprinting Authentication for Deep Neural Networks. *Comp. Commun.* 150, 488–497. doi:10.1016/j.comcom.2019.12.016.PII: S014036641931686X

Conflict of Interest: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors, and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Boenisch. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.