Deep Model Intellectual Property Protection via Deep Watermarking

Jie Zhang[®], Dongdong Chen[®], Jing Liao[®], Weiming Zhang[®], Huamin Feng, Gang Hua[®], *Fellow, IEEE*, and Nenghai Yu[®]

Abstract—Despite the tremendous success, deep neural networks are exposed to serious IP infringement risks. Given a target deep model, if the attacker knows its full information, it can be easily stolen by fine-tuning. Even if only its output is accessible, a surrogate model can be trained through student-teacher learning by generating many input-output training pairs. Therefore, deep model IP protection is important and necessary. However, it is still seriously under-researched. In this work, we propose a new model watermarking framework for protecting deep networks trained for low-level computer vision or image processing tasks. Specifically, a special task-agnostic barrier is added after the target model, which embeds a unified and invisible watermark into its outputs. When the attacker trains one surrogate model by using the input-output pairs of the barrier target model, the hidden watermark will be learned and extracted afterwards. To enable watermarks from binary bits to high-resolution images, a deep invisible watermarking mechanism is designed. By jointly training the target model and watermark embedding, the extra barrier can even be absorbed into the target model. Through extensive experiments, we demonstrate the robustness of the proposed framework, which can resist attacks with different network structures and objective functions.

Index Terms—Deep model IP protection, model watermarking, image processing

1 INTRODUCTION

IN recent years, deep learning has revolutionized a wide variety of artificial intelligence fields such as image recognition [1], [2], medical image processing [3], [4], [5], [6], speech recognition [7], [8] and natural language processing [9], and outperforms traditional state-of-the-art methods by a large margin. Despite its success, training a good deep model is not a trivial task and often requires the dedicated design of network structures and learning strategies, a large scale high-quality labeled data and massive amount of computation resources, all of which are expensive and full of great business value. For some companies, these trained models are indeed their core competitiveness. Therefore, IP protection for deep models is not only important but also essential for their success.

However, compared to media IP protection, deep model IP protection is much more challenging because of the

- Dongdong Chen is with the Microsoft Research, Redmond, WA 98052 USA. E-mail: cddlyf@gmail.com.
- Jing Liao is with the Department of Computer Science, City University of Hong Kong, Kowloon Tong, Hong Kong. E-mail: jingliao@cityu.edu.hk.
- Huamin Feng is with the Beijing Electronic Science and Technology Institute, Beijing 100070, China. E-mail: fenghm@besti.edu.cn.
- Gang Hua is with the Wormpex AI Research LLC, WA 98004 USA. E-mail: ganghua@gmail.com.

Manuscript received 29 June 2020; revised 2 Feb. 2021; accepted 5 Mar. 2021. Date of publication 9 Mar. 2021; date of current version 1 July 2022. (Corresponding author: Dongdong Chen.) Recommended for acceptance by J. Zhou. Digital Object Identifier no. 10.1109/TPAMI.2021.3064850

powerful learning capacity of deep models [10], [11]. In other words, given one specific task, numerous structures and weight combinations can obtain similar performance. This actually makes the IP infringement rather convenient. In the white-box case, where the full information including the detailed network structure and weights of the target model is known, one typical and effective attack way would be finetuning or pruning based on the target model on new datasets. Even in the black-box case where only the output of the target model can be accessed, the target model can still be stolen by training another surrogate model to imitate its behavior through student-teacher learning [12], [13]. Specifically, we can first generate a large scale of input-output training pairs based on the target model, and then directly train the surrogate model in a supervised manner by regarding the outputs of the target model as ground-truth labels.

Very recently, some preliminary research works [10], [11], [14], [15] emerge for deep model IP protection. They often either add a weight regularizer into the loss function to make the learned weight have some special patterns or use the predictions of a special set of indicator images as the watermarks. Though these methods work pretty well to some extent, they only consider the classification task and the white-box attacks like fine-tuning or pruning. But in real scenarios, labeling the training data for low-level computer vision or image processing tasks is much more complex and expensive than classification tasks, because their groundtruth labels should be pixel-wise precise. Examples include removing all the ribs in Chest X-ray images and the rain streaks in real rainy images. In this sense, protecting such image processing models is more valuable. Moreover, because the original raw model does not need to be provided in most application scenarios, it can be easily

Jie Zhang, Weiming Zhang, and Nenghai Yu are with the School of Cyber Science and Security, University of Science and Technology of China, Hefei, Anhui 230026, China. E-mail: zjzac@mail.ustc.edu.cn, {zhangwm, ynh}@ustc.edu.cn.



Fig. 1. The simplest watermarking mechanism by adding unified visible watermarks onto output of the target model, which will sacrifice the visual quality and usability.

encrypted with traditional algorithms to resist the whitebox attack. Therefore, more attention should be paid to the black-box surrogate model attack.

Motivated by this, this paper studies the IP protection problem for image processing networks that aims to resist the challenging surrogate model attack [16]. Specifically, a new model watermarking framework is proposed to add watermarks into the target model. When the attackers use a surrogate model to imitate the behavior of the watermarked model, the designed watermarking mechanism should be able to extract pre-defined watermarks out from outputs of the learned surrogate model.

For better understanding, before diving into the model watermarking details, let us first discuss the simplest spatial visible watermarking mechanism shown in Fig. 1. Suppose that we have a lot of input-output training pairs and we manually add a unified visible watermark template to all the outputs. Intuitively, if a surrogate model is trained on such pairs with the simple L2 loss, the learned model will learn this visible watermark into its output to get lower loss. That is to say, given one target model, if we forcibly add one unified visible watermark into all its output, it can resist the plagiarism from other surrogate models to some extent. However, the biggest limitation of this method is that the added visible watermarks will seriously degrade the visual quality and usability of the target model. Another potential threat is that attackers may use image editing tools like Photoshops to manually remove all the visible watermarks.

Inspired by the above observation, we propose a general deep invisible model watermarking framework as shown in Fig. 2. Given a target model **M** to be protected, we denote its original input and output images as domain A and B respectively. Then a spatial invisible watermark embedder **H** is used to hide a unified target watermark δ into all the output images in the domain **B** and generate a new domain \mathbf{B}' . Different from the aforementioned simple visible watermarks, all the images in the domain \mathbf{B}' should be visually consistent to domain **B**. Symmetrically, given the images in domain \mathbf{B}' , another watermark extractor \mathbf{R} will extract the watermark δ' out, which is consistent to δ . To protect the IP of M, we will pack H and M into a whole for deployment. The key hypothesis here is that when the attacker uses **A** and **B**' to learn a surrogate model **SM**, **R** can still extract the target watermark out from the output B'' of SM.

We first test the effectiveness of our framework by using traditional spatial invisible watermarking algorithms such as [17], [18] and [19], which work well for some special



Fig. 2. The proposed deep watermarking framework by leveraging spatial invisible watermarking algorithms.

surrogate models but fail for most other ones. Another common limitation is that the information capacity they can hide is relatively low, e.g., tens of bits. In order to hide high capacity watermarks like logo images and achieve better robustness, a novel deep invisible watermarking system is proposed. As shown in Fig. 3, it consists of two main parts: one embedding sub-network H to learn how to hide invisible watermarks into the image, and another extractor subnetwork **R** to learn how to extract the invisible watermark out. To avoid **R** generating watermark for all the images no matter whether they have invisible watermarks or not, we also constrain **R** not to extract any watermark out if its input is a clean image. Moreover, to make **R** generalize better for outputs of any surrogate model, an extra adversarial training stage is incorporated. In this stage, we pick one example surrogate model and add its outputs into **R**'s training.

For forensics, we leverage both the classic normalized correlation metric and a simple classifier to judge whether the extracted watermark is valid or not. Through extensive experiments, the proposed framework shows its strong ability in resisting the attack from surrogate models trained with different network structures like Resnet and UNet and different loss functions like L1, L2, perceptual loss, and adversarial loss. By jointly training the target model and watermark embedding, we find the functionality of **H** can be even absorbed into **M**, making **M** a watermarked model itself.

The contributions of this paper can be summarized in the following five aspects:

- We introduce the IP protection problem for image processing networks. We hope it will help draw more attention to this seriously under-explored field and inspire more great works.
- Motivated by the loss minimization property of deep networks, we innovatively propose to leverage the spatial invisible watermarking mechanism for deep model watermarking.
- A novel deep invisible watermarking algorithm along with the two-stage training strategy is designed to improve the robustness and capacity of traditional spatial invisible watermarking methods.
- We extend the proposed framework to multiplewatermark and self-watermark case, which further shows its strong generalization ability.
- Both the classic normalized correlation metric and a watermark classifier demonstrate the strong



Fig. 3. The overall pipeline of the proposed deep invisible watermarking algorithm and two-stage training strategy. In the first training stage, a basic watermark embedding sub-network H and extractor sub-network R are trained. Then another surrogate network SM is leveraged as the adversarial competitor to further enhance the extracting ability of R.

robustness to surrogate model attack with different network structures and loss functions.

Compared with the preliminary conference version [16], we have made significant improvements and extension in this manuscript. The main difference are from four aspects: 1) We rewrite the theoretical analysis Section 3.2 to make it more clear and solid; 2) We design a new classifier based watermark verification mechanism in Section 3.4.4; 3) We extend our framework to multiple watermarks and self-watermarked models in Sections 3.4.5 and 3.4.6; 4) More corresponding experimental results and analysis are given in Sections 4 and 5.

2 RELATED WORK

2.1 Media Copyright Protection

Compared to model IP protection, media copyright protection [20], [21], [22], [23], [24], [25] is a classic research field that has been studied for several decades. The most popular media copyright protection method is watermarking. For image watermarking, many different algorithms have been proposed in the past, which can be broadly categorized into two types: visible watermarks like logos, and invisible watermarks. Compared to visible watermarks, invisible watermarks are more secure and robust. They are often embedded in the original spatial domain [17], [18], [19], [26], or other image transform domains such as discrete cosine transform (DCT) domain [27], [28], discrete wavelet transform (DWT) domain [29], and discrete Fourier transform (DFT) domain [30]. However, all these traditional watermarking algorithms are often only able to hide several or tens of bits, let alone real logo images. More importantly, we find only spatial domain watermarking work to some extent for this task and all other transform domain watermarking algorithms totally fail.

In recent years, some DNN-based watermarking schemes have been proposed. For example, Zhu *et al.* [31] propose an auto-encoder-based network architecture to realize the embedding and extracting process of watermarks. Based on it, Tancik *et al.* [32] further realize a camera

shooting resilient watermarking scheme by adding a simulated camera shooting distortion to the noise layer. In [33], a full-size image is proposed to be used as watermark under the hiding image in image framework. Compared to these media watermarking algorithms, model watermarking is much more challenging because of the exponential search space of deep models. But we innovatively find it possible to leverage spatial invisible watermarking techniques for model protection.

2.2 Deep Model IP protection

Though IP protection for deep neural networks is still seriously under-studied, there are some recent works [10], [11], [14], [34], [35], [36], [37] that start paying attention to it. For example, based on the over-parameterized property of deep neural networks, Uchida et al. [10] propose a special weight regularizer to the objective function so that the distribution of weights can be resilient to attacks such as fine-tuning and pruning. One big limitation of this method is not taskagnostic and need to know the original network structure and parameters for retraining. Adi et al. [11] use a particular set of inputs as the indicators and let the model deliberately output specific incorrect labels, which is also known as "backdoor attack". Similar scheme is adopted in [14] and [38] for DNN watermarking on embedded devices. Recently, Fan et al. [39] find all the above backdoor based methods are fragile to ambiguity attack. To overcome this limitation, a special digital passport layer is embedded into the target model and trained in an alternative manner. But this special passport layer often needs to make some network structure changes and causes performance degradation. So in the recent work [40], Zhang et al. further introduce a new passport-aware normalization layer for model IP protection.

Despite the effectiveness of the above methods, they all focus on the classification tasks, which is different from the purpose of this paper, i.e., protecting higher commercial valued image processing models. And for the challenging surrogate model attack, they will totally fail. In this paper, we consider the surrogate model attack and study the IP protection for image processing networks. And different from the above methods, we innovatively propose to leverage spatial invisible watermarking algorithms for deep model IP protection.

2.3 Image-to-Image Translation

In the deep learning era, most image processing tasks such as style transfer [41], [42], [43], deraining[44], [45], [46], and X-ray image debone [47], can be modeled as an image-toimage translation problem where both the input and output are images. Thanks to the emergence of the generative adversarial network (GAN) [48], this field has achieved significant progress recently. Isola et al. propose a general image-to-image translation framework Pix2Pix by combining adversarial training in [49], which is further improved by many following works [50], [51], [52] in terms of quality, diversity and functionality. However, all these methods need a lot of pairwise training data, which is often difficult to be collected. By introducing the cycle consistency, Zhu et al. propose a general unpaired image-to-image translation framework CycleGAN [53]. A similar idea is also adopted in concurrent works DualGAN [54] and DiscoGAN [55]. In this paper, we mainly focus on the deep models of paired image-to-image translation, because the paired training data is much more expensive to be obtained than unpaired datasets. More importantly, there is no prior work that has ever considered the watermarking issue for such models.

3 Метнор

In this section, we will first define the target problem formally in Section 3.1, then provide the theoretical analysis why the spatial invisible watermarking can be utilized for model watermarking in Section 3.2. Based on the analysis, traditional spatial invisible watermarking and a new deep invisible watermarking framework will be elaborated in Sections 3.3 and 3.4 respectively. Next, two different watermark verification methods will be considered in Section 3.4.4. Finally, we will show how to extend the proposed framework to enable multiple watermarks within one single sub-network in Section 3.4.5 and make the target model selfwatermarked in Section 3.4.6.

3.1 Problem Definition

Given a target model **M** to protect, this paper mainly considers the surrogate model attack, where the attacker does not know the detailed network structure or weights of **M** but is able to access its output. This is a very common and sensible setting for real systems because most existing commercial deep models are deployed as cloud API service or encrypted executable program. In these scenarios, though the attacker cannot use common white-box attack methods like fine-tuning and pruning, a surrogate model can be trained to imitate the target model's behavior in a teacher-student way. Specifically, the attacker collects a lot of input images $\{a_1, a_2, \ldots, a_m\}$ (domain **A**) and feeds them into **M** to get corresponding output images $\{b_1, b_2, \ldots, b_m\}$ (domain **B**), then a surrogate model **SM** is trained in a supervised way by regarding **M**'s outputs as ground truth.

For effective forensics, the goal of resisting surrogate model attack is to design an effective way that is able to identify **SM** once it is trained with the data generated by **M**. And in real scenarios, it is highly possible that we cannot access the detailed information of model **SM** (e.g., network structure and weights) like the target model neither. This means that, even when we suspect one model is a pirated model, the only indicator we can utilize is its output unless we start the costly legal proceeding. Therefore, we need to figure out one way to extract some kinds of forensics hints from the output of **SM**. In this paper, we propose a new deep model watermarking mechanism so that the predefined watermark pattern can be extracted from **SM**'s output.

3.2 Theoretical Pre-Analysis

In traditional watermarking algorithms, given an image I and a target watermark δ to be embed, they will first use a watermark embedder **H** to generate an image I' which contains δ . Symmetrically, the target watermark δ can be further extracted out by the corresponding watermark extractor **R**. As shown in Fig. 1, if each image $b_i \in \mathbf{B}$ is embedded with a unified watermark δ , forming another domain **B**', then the objective of the surrogate model **SM** is to minimize the distance \mathcal{L} between $\mathbf{SM}(a_i)$ and b'_i

$$\mathcal{L}(\mathbf{SM}(a_i), b'_i) \to 0, \quad b'_i = b_i + \sigma.$$
(1)

Below we will first show **SM** can learn the watermark δ into its output, then show the inherent loss minimization property of deep network will theoretically guarantee δ to be learned by a good surrogate model **SM**.

In detail, for each a_i we have $b_i = M(a_i)$, thus there must exist a model **SM** that can learn good transformation between domain **A** and domain **B**' because of the existence of below equivalence

$$\mathcal{L}(\mathbf{M}(a_i), b_i) \to 0 \Leftrightarrow \mathcal{L}(\mathbf{SM}(a_i), (b_i + \mathbf{\delta})) \to 0$$

when $\mathbf{SM} = \mathbf{M} + \delta.$ (2)

In other words, one simplest solution of **SM** is to directly add δ to the output of **M** with a skip connection.

On the other hand, because the objective of **SM** is calculated based on b'_i , the loss minimization property of deep network theoretically guarantees a good surrogate model **SM** should learn the unified watermark δ into its output. Otherwise, its objective loss function \mathcal{L} cannot achieve a lower value than the above simple solution.

Based on this observation, we propose a general deep watermarking framework for image processing models shown in Fig. 2. Given a target model **M** to protect, we add a barrier **H** after it and embed a unified watermark δ into its output before showing them to the end-users. In this way, the final output obtained by the attacker is actually the watermarked version, so the surrogate model **SM** has to be trained with the image pair (a_i, b'_i) from domain **A**, **B**' instead of the original pair (a_i, b_i) from domain **A**, **B**. After **SM**'s training, we can leverage the corresponding watermark extractor **R** to extract the watermark out from the output of **SM**. To ensure the watermarked output image b'_i is visually consistent with the original one b_i , only spatial invisible watermarking algorithms are considered in this paper. Below we will introduce both traditional spatial invisible watermarking algorithm and a novel deep invisible watermarking algorithm.

3.3 Traditional Spatial Invisible Watermarking

Additive-based embedding is the most common method used in the traditional spatial invisible watermarking scheme. Specifically, the watermark information is first spread into a sequence or block which satisfies a certain distribution, then embedded into the corresponding coefficients of the host image. This embedding procedure can be formulated by

$$I' = \begin{cases} I + \alpha C_0 & \text{if } w_i = 0\\ I + \alpha C_1 & \text{otherwise} \end{cases},$$
(3)

where *I* and *I'* indicate the original image and embedded image respectively. α indicates the embedding intensity and C_i denote the spread image block that represents bit " w_i "($w_i \in [0, 1]$). In the extraction side, the watermark is determined by detecting the distribution of the corresponding coefficients. The robustness of such an algorithm is guaranteed by the spread spectrum operation. The redundancy brought by the spread spectrum makes a strong error correction ability of the watermark so that the distribution of the block will not change a lot even after image processing.

However, such algorithms often have very limited embedding capacity because many extra redundant bits are needed to ensure robustness. In fact, in many application scenarios, the IP owners may want to embed some special images (e.g., logos) explicitly for convenient visual forensics. This is nearly infeasible for these algorithms. More importantly, the following experiments show that these traditional algorithms can only resist the attack from some special types of surrogate models. To enable more high-capacity watermarks and more robust resistance ability, we will propose a new deep invisible watermarking algorithm below.

3.4 Deep Invisible Watermarking

3.4.1 Overview

As shown in Fig. 3, to embed an image watermark into host images of the domain **B** and extract it out afterward, one embedding sub-network **H** and one extractor sub-network **R** are adopted respectively. Without sacrificing the original image quality of domain **B**, we require watermarked images in domain **B**' to be visually consistent to the original images in the domain **B**. As adversarial networks have demonstrated their power in reducing the domain gap in many image translation tasks, we append one discriminator network **D** after **H** to further improve the image quality of domain **B**'.

During training, we find if the extractor network \mathbf{R} is only trained with the images of domain \mathbf{B}' , it is very easy to overfit and output the target watermark no matter whether the input images contain watermarks or not. To circumvent this problem, we also feed the images of domain \mathbf{A} and domain \mathbf{B} that do not contain watermarks into \mathbf{R} and force it to output a constant blank image. In this way, \mathbf{R} will have

the real ability to extract watermarks out only when the input image has the watermark in it.

Based on the pre-analysis, when the attacker uses a surrogate model SM to imitate the target model M based on the input domain A and watermarked domain B', SM will learn the embedded watermark δ into its output. Despite higher embedding capacity, we find that the extractor subnetwork R cannot extract the watermarks out from the output of the surrogate model SM neither like traditional watermarking algorithms if only trained with this initial training stage. This is because R has only observed clean watermarked images but not the watermarked images from surrogate models which may contain some unpleasant noises. To further enhance the extracting ability of \mathbf{R} , we choose one simple surrogate network to imitate the attackers' behavior and fine-tune R on the mixed dataset of domain $\mathbf{A}, \mathbf{B}, \mathbf{B}', \mathbf{B}''$ in an extra adversarial training stage. The following experiments will show this can significantly boost the extracting ability of R. More importantly, this ability generalizes well to other types of surrogate models.

3.4.2 Network Structures

By default, we adopt UNet [56] as the network structure of H and SM, which has been widely used by many translation based tasks like Pix2Pix [49] and CycleGAN [53]. It performs especially well for tasks where the output image shares some common properties of input image by using multi-scale skip connections. But for the extractor sub-network **R** whose output is different from the input, we find CEILNet [57] works much better. It also follows an autoencoder like network structure. In details, the encoder consists of three convolutional layers, and the decoder consists of one deconvolutional layer and two convolutional layers symmetrically. In order to enhance the learning capacity, nine residual blocks are inserted between the encoder and decoder. For the discriminator D, we adopt the widely-used PatchGAN [49]. Note that except for the extractor sub-network **R**, we find other types of translation networks also work well in our framework, which demonstrates the strong generalization ability of our framework.

3.4.3 Loss Functions

The objective loss function of our method consists of two parts: the embedding loss \mathcal{L}_{emd} and the extracting loss \mathcal{L}_{ext} , i.e.,

$$\mathcal{L} = \mathcal{L}_{emd} + \lambda * \mathcal{L}_{ext},\tag{4}$$

where λ is the hyper parameter to balance these two loss terms. Below we will introduce the detailed formulation of \mathcal{L}_{emd} and \mathcal{L}_{ext} respectively.

Embedding Loss. To embed the watermark image into a cover image while guaranteeing the original visual quality, three different types of visual consistency loss are considered: the basic *L*2 loss ℓ_{bs} , perceptual loss ℓ_{vgg} and adversarial loss ℓ_{adv} , i.e.,

$$\mathcal{L}_{emd} = \lambda_1 * \ell_{bs} + \lambda_2 * \ell_{vgg} + \lambda_3 * \ell_{adv}.$$
 (5)

Here the basic $L2 \log \ell_{bs}$ is simply the pixel value difference between the input host image b_i and the watermarked output image b'_i , N_c is the total pixel number, i.e.,

$$\ell_{bs} = \sum_{b'_i \in \mathbf{B}', b_i \in \mathbf{B}} \frac{1}{N_c} \|b'_i - b_i\|^2.$$
(6)

And the perceptual loss ℓ_{vgg} [58] is defined as the difference between the VGG feature of b_i and b'_i

$$\ell_{vgg} = \sum_{b'_i \in \mathbf{B}', b_i \in \mathbf{B}} \frac{1}{N_f} \| VGG_k(b'_i) - VGG_k(b_i) \|^2,$$
(7)

where $VGG_k(\cdot)$ denotes the features extracted at layer k ("conv2_2" by default), and N_f denotes the total feature neuron number. To further improve the visual quality and minimize the domain gap between **B**' and **B**, the adversarial loss ℓ_{adv} will let the embedding sub-network **H** embed watermarks better so that the discriminator **D** cannot differentiate its output from real watermark-free images in **B**

$$\ell_{adv} = \mathop{\mathbb{E}}_{b_i \in \mathbf{B}} log(\mathbf{D}(b_i)) + \mathop{\mathbb{E}}_{b'_i \in \mathbf{B}'} log(1 - \mathbf{D}(b'_i)).$$
(8)

Extracting Loss. The responsibility of the extractor sub-network **R** has two aspects: it should be able to extract the target watermark out for watermarked images from **B**' and instead output a constant blank image for watermark-free images from **A**, **B**. So the first two terms of \mathcal{L}_{ext} are the reconstruction loss ℓ_{wm} and the clean loss ℓ_{clean} for these two types of images respectively, i.e.,

$$\ell_{wm} = \sum_{b'_i \in \mathbf{B}'} \frac{1}{N_c} \|R(b'_i) - \delta\|^2,$$

$$\ell_{clean} = \sum_{a_i \in \mathbf{A}} \frac{1}{N_c} \|R(a_i) - \delta_0\|^2 + \sum_{b_i \in \mathbf{B}} \frac{1}{N_c} \|R(b_i) - \delta_0\|^2,$$

(0)

where δ is the target watermark image and δ_0 is the constant blank watermark image. Besides the reconstruction loss, we also want the watermarks extracted from different watermarked images to be consistent, thus another consistent loss ℓ_{cst} is added

$$\ell_{cst} = \sum_{x,y \in \mathbf{B}'} \|R(x) - R(y)\|^2.$$
(10)

The final extracting loss \mathcal{L}_{ext} is defined as the weighted sum of these three terms, i.e.,

$$\mathcal{L}_{ext} = \lambda_4 * \ell_{wm} + \lambda_5 * \ell_{clean} + \lambda_6 * \ell_{cst}.$$
 (11)

Adversarial Training Stage. From experiments, we find that if **R** is only trained with the above initial training stage, it cannot generalize well for the noisy watermarked output of some surrogate models because only the clean watermarked images are observed. To enhance its extracting ability, an extra adversarial training stage is added, wherein we introduce the degradation of surrogate model attack. Specifically, one proxy surrogate model **SM** is trained with the simple *L*2 loss by default. Denote the outputs of this proxy model as **B**["], we further fine-tune **R** on the mixed dataset **A**, **B**, **B**['], **B**["] in this stage. Since the embedding sub-network **H** is fixed in this stage, only the extracting loss \mathcal{L}_{ext} is

considered to refine the capability of **R**. Specifically, the clean loss ℓ_{clean} is kept unchanged, while ℓ_{wm} , ℓ_{cst} are updated as below respectively:

$$\ell_{wm} = \sum_{b'_i \in \mathbf{B}'} \frac{1}{N_c} \|R(b'_i) - \delta\|^2 + \sum_{b''_i \in \mathbf{B}''} \frac{1}{N_c} \|R(b''_i) - \delta\|^2,$$

$$\ell_{cst} = \sum_{x, y \in \mathbf{B}' \cup \mathbf{B}''} \|R(x) - R(y)\|^2.$$
(12)

3.4.4 Watermark Verification

For effective forensics, if we want to know whether a surrogate model SM is trained by stealing the IP of one target model, we first feed SM's output into the extractor sub-network \mathbf{R} , then verify whether \mathbf{R} 's output matches the predefined watermark. In this paper, we consider two different verification methods: the classic normalized correlation (NC) metric and an extra classifier **C**. For the former one, it is simply defined as

$$NC = \frac{\langle \mathbf{R}(b'_i), \delta \rangle}{\|\mathbf{R}(b'_i)\| * \|\delta\|},$$
(13)

where $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ denote the inner product and L2 norm respectively. Compared to the NC metric, the latter classifier is more straightforward and robust to noisy parts of $\mathbf{R}(b'_i)$ potentially. Because it is just a binary classification problem (yes/no), we find a simple network consisting of three convolutional layers is enough. In our method, **C** is joint trained with **R** during the adversarial training stage, and label "1" and "0" represent watermarked and watermark-free images respectively. By default, the simple cross entropy loss is used

$$\mathcal{L}_{cls} = -\left[\sum_{x \in \mathbf{A} \cup \mathbf{B}} log(1 - \mathbf{C}(x)) + \sum_{x \in \mathbf{B}'} log(\mathbf{C}(x))\right].$$
 (14)

3.4.5 Extension to Multiple Watermarks

In our default setting, for one specific watermark, an embedding sub-network and a corresponding extractor sub-network will be trained. This makes sense for the scenarios where only one specific deep model needs to be protected. However, it is unfriendly to some other scenarios where one deep model has multiple release versions and different versions are represented by different watermarks, because training multiple embedding and extractor sub-networks is both storage- and computation-consuming. Thanks to the flexibility of the proposed framework, enabling multiple different watermarks within one single embedding and extractor sub-network is easy to be supported. The overall training strategy is similar, and the biggest difference is that different watermarks will be randomly selected for H to embed and **R** to exact correspondingly. For the verification classifier **C**, it will be changed from the binary classification to multi-class classification.

3.4.6 Extension to Self-Watermarked Models

As mentioned before, in order to protect one specific target model, the embedding sub-network acts as an extra barrier and will be appended after the target model as a whole. The



Fig. 4. Illustration diagram to show the training difference between the original target model and self-watermarked target model.

main benefit of this pipeline is task-agnostic, which means that it is a very general solution and can work independently without the need of knowing the information of the target model. However, requiring an extra embedding subnetwork may sometimew be a limitation. So we propose an extended solution to absorb the watermarking functionality into the target model itself if we can control its training process. In this sense, the target model is self-watermarked without the extra embedding sub-network. In other words, the target model itself will automatically embed a watermark into its output.

The training difference between the original target model and the self-watermarked target model is shown in Fig. 4. Assuming the original target model **M** is trained with the task-special loss \mathcal{L}_{task} , we add an extractor sub-network **R** after **M** and jointly train them by combining the extracting loss \mathcal{L}_{ext} mentioned above with \mathcal{L}_{task} . To ensure the robustness and capability of **R**, **R** still needs to be fine-tuned with the extra adversarial training stage.

4 EXPERIMENTS

To demonstrate the effectiveness of the proposed system, we use two example image processing tasks in this paper: image deraining and Chest X-ray image debone. The goal of these two tasks is to remove the rain streak and rib components from the input images respectively. We will first introduce the implementation details and evaluation metric in Section 4.1, then evaluate the proposed deep invisible watermarking algorithm quantitatively and qualitatively in Section 4.2. Next, we will demonstrate the robustness of the proposed deep watermarking framework to surrogate models with different network structures and loss functions in Section 4.3. Furthermore, we compare our method with traditional and DNN-based watermarking algorithm in Section 4.4 and show the robustness of our method to watermark overwriting in Section 4.5. Finally, some ablation analysis are provided in Section 4.6.

4.1 Implementation Details

Dataset Setup. For image deraining, we use 6,100 images from the PASCAL VOC dataset [59] as target domain **B**, and use the synthesis algorithm in [60] to generate rainy images as domain **A**. These images are split into two parts: 6,000 are for both the initial and adversarial training stage and 100 for testing. Since the images used by the attacker may be different from that used by the IP owner to train **H**, **R**, we simulate it by choosing 6,000 images from the COCO dataset [61] for the surrogate model training. Similarly, for X-ray image debone, we select 6,100 high-quality chest X-ray images from the open dataset chestx-ray8 [62] and use the rib suppression algorithm proposed by [47] to generate the training pair. They are divided into three parts: 3,000 for both the initial and adversarial training, 3,000 for the surrogate model training, and 100 for testing. All the images are resized to 256×256 by default.

Training Details. By default, we train sub-network **H**, **R** and discriminator **D** for 200 epochs with a batchsize of 8. Adam optimizer is adopted with the initial learning rate of 0.0002. We decay the learning rate by 0.2 if the loss does not decrease within 5 epochs. For surrogate model training, batchsize of 16 and longer epochs of 300 are used to ensure better performance. And the initial learning rate is set to be 0.0001, which stays unchanged in the first 150 epochs. In the adversarial training stage, sub-network **R** is trained with the initial learning rate of 0.0001. For classifier **C**, it is also trained in a similar way. By default, all hyperparameters λ_i equal to 1 except $\lambda_3 = 0.01$.

Evaluation Metric. To evaluate the visual quality, PSNR and SSIM are used by default. When using the NC metric for watermark verification, the watermark is regarded as successfully extracted if its NC value is bigger than 0.95. Based on it, the success rate (SR_{NC}) is further defined as the ratio of watermarked images whose hidden watermark is successfully extracted in an image set. When using the classifier for verification, we use the threshold 0.5 for the binary case and adopt the label with max-possibility for the multiclass case. Similarly, the success rates based on the classifier are denoted as SR_C .

4.2 Deep Image Based Invisible Watermarking

In this experiment, we give both qualitative and quantitative results about the proposed deep image based invisible watermarking algorithm. To demonstrate its generalization ability to various types of watermark images, both simple grayscale logo images and complex color images are considered. For the grayscale case, we even consider a QR code image. Several example images used for the debone and deraining task are shown in Fig. 5. It can be easily seen that the proposed deep watermarking algorithm can not only embed the watermark images into the host images in an imperceptible way but also successfully extract the embedded watermarks out. Especially for the challenging "Peppers" and "Lena" watermarks that have rich textures, the embedding sub-network and extractor sub-network still collaborate very well and guarantee the visual quality of the watermarked image and extracted watermark simultaneously.

To investigate whether the extracting sub-network embeds watermark in a naive additive way, we further



Fig. 5. Some visual examples to show the capability of the proposed deep invisible watermarking algorithm: (A) watermark-free image b_i from domain **B**, (B) watermarked image b'_i from domain **B**' (C) the residual between b_i and b'_i (enhanced 10×), (D) ground-truth watermark, (E) extracted watermark from b'_i . The first column is the results of traditional spatial bit-based invisible watermarking algorithms (64-bit embedded), and the middle and right parts are the results of our methods for the debone and deraining tasks respectively.

show the residuals (even enhanced 10 ×) between watermarked image b'_i and watermark-free image b_i in Fig. 5. Obviously, we cannot observe any watermark hint from the residuals, which means the watermark is indeed embedded in an advanced and smart way. Taking one more step forward, we use the "Flower" watermark image as an example and show the color histogram distribution difference between the watermarked image and original watermarkfree image in Fig. 6. It shows that, even though the colorful "Flower" watermark image is embedded, the gray histogram distribution is almost identical to that of the original watermark-free image.

Quantitative results with respect to the visual quality and extracting ability are provided in Table 1. We can find that the average PSNR and SSIM between watermarked and watermark-free images are very high, which double confirms the original image quality is well kept. On the other hand, though high extracting ability is contradictory to high visual quality to some extent, our extractor sub-network **R** is still able to achieve very high extracting ability with an average NC value over 0.99 and 100 percent success rate both based on the NC metric and the classifier **C**.

4.3 Robustness to Surrogate Model Attack

Besides the visual quality consistency, another more important goal is to guarantee IP protection ability. Considering the attacker may use different network structures trained with different loss functions to imitate target model's behavior, we simulate this case by using a lot of surrogate models to evaluate the robustness of the proposed deep watermarking framework. In details, we consider four different types of network structures: vanilla convolutional



Fig. 6. The gray histogram comparison between watermark-free image and watermarked image for the Derain-Flower case.

TABLE 1 Quantitative Results of the Proposed Deep Image Based Invisible Watermarking Algorithm

Task	PSNR	SSIM	NC	$SR_{NC} \ / \ SR_{C}$
Debone-IEEE Debone-Flower Debone-QR Derain-Flower Derain-Peppers Derain-Lena Derain-IEEE	47.29 46.36 44.35 41.21 40.91 42.50 41.76	0.99 0.99 0.99 0.99 0.98 0.98 0.98	0.9999 0.9999 0.9999 0.9999 0.9999 0.9999 0.9999	100% / 100% 100% / 100% 100% / 100% 100% / 100% 100% / 100% 100% / 100% 100% / 100%
Derain-QR	40.24	0.98	0.9999	100% / 100%

PSNR and SSIM are calculated between the watermarked and original watermark-free images. x, y denote task name and watermark image name in the notation "x-y".

TABLE 2 The Success Rate (SR_{NC}/SR_{C}) of Our Method Resisting the Attack From Surrogate Models

Settings	CNet	Res9	Res16	UNet	L1	$L1 + L_{adv}$	L2	$L2 + L_{adv}$	$L_{perc}+L_{adv}$
Debone	93% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	71% / 94%
Derain	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	99% / 100%	100% / 100%
Debone [†]	0% / 0%	0% / 0%	0% / 0%	0% / 0%	0% / 0%	0% / 0%	27% / 88%	44% / 96%	0% / 0%
Derain†	0% / 0%	0% / 0%	0% / 0%	0% / 0%	0% / 0%	0% / 0%	0% / 0%	0% / 0%	0% / 54%

Column 2 \sim 5 are trained with L2 loss but different network structures and Column 6 \sim 10 are trained with UNet network structure but different loss combinations. We take Debone-IEEE task and Derain-Flower task as example, and \dagger denotes the results without adversarial training.



Fig. 7. Example output image and corresponding extracted watermark from different surrogate models. (A) input watermark-free image a_i ; (B) watermarked image b'_i ; (C) ~ (F) are cases for different network structures with L2 loss: CNet, Res9, Res16 and UNet in turn; (G) ~ (K) are cases for different loss combinations with UNet: L2, L2+L_{adv}, L1, L1+L_{adv} and L_{perc}+L_{adv} in turn.

networks only consisting of several convolutional layers ("CNet"), an auto-encoder like networks with 9 and 16 residual blocks ("Res9", "Res16"), and the aforementioned UNet network ("UNet"). For the objective loss function, popular loss functions like L1, L2, perceptual loss L_{perc} , adversarial loss L_{adv} and their combination are considered. But we discard the case that only utilizes perceptual loss L_{perc} for surrogate model learning because it will generate very terrible image quality (PSNR:19.73; SSIM:0.85) and make the attack meaningless. As the surrogate model with "UNet" and L2 loss function is utilized in the adversarial training stage, this configuration can be viewed as a whitebox attack and all other configurations are black-box attacks. Without losing generality, we use grayscale "IEEE" and color "Flower" image as default target watermark for debone and deraining respectively in this experiment.

Because of the limited computation resource, we do not consider all the combinations of network structures and loss functions. Instead, we choose to conduct the control experiments to demonstrate the robustness to the network structures and loss functions respectively. In Table 2 (Column 2 \sim 5), it shows that, though only UNet based surrogate model trained with *L*2 loss is leveraged in the adversarial training stage, the proposed deep model watermarking framework can resist both white-box and black-box attacks when equipped with the newly proposed deep image-based invisible watermarking technique.

To further demonstrate the robustness to different loss functions, we use the UNet as the default network structure and train surrogate models with different combinations of loss functions. As shown in Table 2 (Column 6 \sim 10), the proposed deep watermarking framework has a very strong generalization ability and can resist different loss combinations with very high success rates. In Fig. 7, we provide one example watermark extracted from different surrogate models.

4.4 Comparison With Other Methods

We further compare our method with one representative traditional watermarking algorithm [19] and one DNNbased watermarking algorithm HiDDeN [31] by hiding 64bit as the watermark. For HiDDeN, two variants proposed in their paper are considered: with/without noise layer between the encoder and the decoder. Compared to the "without noise layer" version, "with noise layer" version is shown to be more robust.

As shown in Table 3, our method not only has better embedding and extracting ability, but also has better robustness to surrogate model attack. For the traditional watermarking algorithm [19], it can only resist the attack of some special surrogate models (only UNet in our experiment) because its extracting algorithms cannot handle the degraded watermarked images from different surrogate models (shown in Fig. 8). More importantly, they cannot embed high-capacity watermarks like logo images. Although the visual quality is acceptable, the remnants of the original watermark can be discovered from the residual image (Fig. 5), which may be utilized by the attacker. In terms of the embedding time and extracting time, our method is about 0.0051s and 0.007s for a 256×256 image on

 TABLE 3

 The Comparison of Embedding & Extracting Ability (Column 2 ~ 4) and Robustness to Surrogate Model Attack (Column 5 ~ 8) Among Our Method and Two Typical Methods in Debone Task

Method	PSNR	SSIM	NC	CNet	Res9	Res16	UNet
HiDDeN [31]	36.25	0.94	0.7220	0%	0%	0%	0%
HiDDeN [31]‡	37.23	0.97	0.9972	0%	0%	0%	0%
TSIW-64 [19]	40.81	0.95	0.9999	0%	0%	0%	100%
Our	47.29	0.99	0.9999	93% / 100%	100% / 100%	100% / 100%	100% / 100%

TSIW-64 denotes using traditional spatial invisible watermarking algorithms [19] to hide 64-bit, while HiDDeN is a typical DNN-based watermarking method. We use HiDDeN to hide 64-bit as well and \ddagger denotes training HiDDeN without noise layer. The surrogate models are trained with L2 loss but different network structures. For TSIW-64 and HiDDeN, only SR_{NC} is reported.



Fig. 8. Visual results of traditional spatial invisible watermarking: (a) watermark-free image b_i , (b) watermarked image b'_i , (c) and (d) are the outputs of surrogate model (UNet and Res9 with *L*2, respectively), (e) ground-truth watermark, (f) ~ (h) are the corresponding extracted watermark from (b) ~ (d). It is defined as successful extracting when the extracted pattern matches mostly with ground-truth pattern like pattern in blue and orange box, otherwise as failure like pattern in red box.

GPU, while traditional spatial invisible watermarking algorithm [19] needs 0.0038s and 0.0047s on CPU. We have also tried some other traditional transform domain watermarking algorithms like DCT-based[63], DFT-based[64] and DWT-based[65], but all of them do not work and achieve 0 percent success rate. For the DNN-based watermarking algorithm HiDDeN [31], no matter whether the noise layer is involved or not, it is totally fragile to surrogate model attacks.

4.5 Robustness to Watermark Overwriting

In real applications, the attackers may follow a similar watermarking strategy and add another watermark upon the model outputs before training the surrogate model, which can potentially destroy the original watermark or cause forensics ambiguity. This type of attack is often called "watermark overwriting". By using our default setting, we find it cannot resist this attack very well. But following a similar idea as the adversarial training, we find a simple enhanced training strategy can help resist it. Specifically,

TABLE 4 Quantitative Results With Different Watermark Images for Watermark Overwriting

Watermarks	None	"Copyright"	"Flower"	"BIT"	
SR_{NC}/SR_C SR_{NC}/SR_C ‡ PSNR/SSIM	100%/100% 100%/100% 25.40/0.88	0%/100% 100%/100% 25.32/0.88	0%/36% 100%/100% 25.34/0.88	0%/0% 99%/100% 25.05/0.88	

The 2nd and 3rd rows are the extracting success rate of our extractor, and ‡ denotes the results with the enhanced training strategy. We also show the visual quality of surrogate model trained with corresponding re-watermarked images in the last row.

we first train a single embedding network (simulate the potential overwriting network) to embed ninety diverse watermark images as described in Section 3.4.5, and then add a noise layer between the embedding sub-network and the extractor sub-network in the initial training stage by using the trained embedding network to re-watermark the watermarked images. We require the extractor to be able to extract the original watermark out from re-watermarked images but blank watermarks from clean images or images only watermarked by the overwriting network. In this way, the extracting ability of the extractor sub-network can be significantly enhanced.

To simulate the attack scenario, we utilize "IEEE" as the original watermark and use other different watermarks (no overlap with the ninety watermarks) to perform overwriting, which can be regarded as a black-box setting. As shown in Fig. 9 and Table 4, overwriting with different watermark images makes the extractor fail to some extent, but the extractor sub-network will work well with the above enhanced training strategy. In addition, Table 4 shows a degradation of the performance of surrogate model trained with re-watermarked images, which reveals watermark overwriting is a two-sided sword for the attacker. Need to note that, although our algorithm can still extract our original watermark out, it still suffers from the ambiguity issue (the extractor sub-network from the attacker can also extract



Fig. 9. Some visual examples of the extractor against watermark overwriting. Three different watermark images (a1 \sim a3) are considered as attacker's watermark, and the corresponding second and third columns are the extracted watermark by our extractor without and with the enhanced training strategy.



Fig. 10. The relationship between the extracting ability and the surrogate model performance. *Left:* the extracting success rates (SR_{NC} , SR_C) and performance (PSNR) change for the surrogate model CNet equipped with different channel numbers (from 8 to 256). *Right:* One visual example (the first row): the middle two column represent the surrogate model output for CNet_8 and CNet_256 respectively while the first and the last column represent the input image with "Bone" and the original target model output. The extracted watermarks are displayed in the second row.



Fig. 11. Comparison results with (first row) and without (second row) clean loss: (a) and (c) are the watermark-free images a_i, b_i from domain **A**, **B**, (b) and (d) are the extracted watermarks from images a_i, b_i respectively. Number on the topright corner denotes the NC value.

their watermark out). And we still need the watermarking protocol to resolve this ambiguity issue.

4.6 Ablation Study

Importance of Clean Loss and Consistent Loss. At the watermark extracting end, besides the watermark reconstruction loss, we also incorporate the extra clean loss for watermarkfree images and watermark consistent loss among different watermarked images. To demonstrate their importance, two control experiments are conducted by removing them from the extracting loss. As shown in Fig. 11, without clean loss, the extractor will always extract meaningless watermark from watermark-free images of domain **A**, **B** with high NC values, which will make the forensics invalid.

Similarly in Fig. 12, when training without the consistent loss, we find the extractor sub-network **R** can only extract very weak watermarks or even cannot extract any watermark out from the outputs of the surrogate mode. This is because the watermark consistency hidden in watermarked image b'_i is destroyed to some extent and makes it more difficult to learn unified watermarks into the surrogate model. By contrast, our method can always extract very clear watermarks out.

Importance of Adversarial Training. As described above, to enhance the extracting ability of **R**, another adversarial training stage is used. To demonstrate its necessity, we also conduct the control experiments without adversarial training, and attach the corresponding results in Table 2 (labelled with " \dagger "). It can be seen that, with the default *L*2 loss, its



Fig. 12. Comparison results with (first row) and without (second row) consistent loss: (a) watermark-free image b_i from domain **B**, (b) watermarked image b'_i , (c) output b''_i of the surrogate model, (d) extracted watermark out from b''_i .

extracting success rate is all about 0 percent for surrogate models of different network structures. When using UNet as the network structure but trained with different losses, only the embedded watermarks of some special surrogate models can be partially extracted, which demonstrates the significant importance of the adversarial training. In our understanding, the reason why adversarial training with UNet can generalize well to other networks may come from two different aspects: 1) Different surrogate models are trained with the similar task-specific loss functions, so their outputs are similar; 2) As shown in the recent work [66], different CNN-based image generator models share some common artifacts during the generation process. So training with the degradation brought by UNet makes the extracting network robust to other networks.

Extracting Ability versus Surrogate Model Performance. In this section, we further analyze the relationship between the extracting ability and the surrogate model performance. Intuitively, the extracting ability will be high when the surrogate model's performance is good, and vice versa. To simulate this case, we take CNet with *L*2 loss as the baseline and adjust its learning ability by changing the feature channel number from 8 to 256. It can be observed from Fig. 10 that, as the feature channels increase, the performance of surrogate model (PSNR from 20.77 to 23.55) and the extracting ability of extractor **R** (SR_{NC} : from 3 to 99 percent and SR_C : from 18 to 100 percent) both increase. From the model protection perspective, when the surrogate model does not perform well, we can view it as successful IP protection. For

TABLE 5 Quantitative Results of Our Method With Different Size Watermark Images

λ	PSNR	SSIM	NC	SR_{NC} / SR_C
0.1	49.82	0.9978	0.9998	100% / 100%
0.5	48.44	0.9969	0.9998	100% / 100%
1	47.29	0.9960	0.9999	100% / 100%
2	45.12	0.9934	0.9999	100% / 100%
10	43.21	0.9836	0.9999	100% / 100%

We take Debone-IEEE task for example.

TABLE 6 Quantitative Results of Our Method With Different Size Watermark Images

SIZE	PSNR	SSIM	NC	SR_{NC} / SR_C
32	46.89	0.9962	0.9999	100% / 100%
64	47.49	0.9965	0.9999	100% / 100%
96	48.06	0.9967	0.9999	100% / 100%
128	47.68	0.9965	0.9999	100% / 100%
256	47.29	0.9960	0.9999	100% / 100%

We take Debone-IEEE task for example.

the visual illustration, we give one example for CNet_8 and CNet_256 on the right of Fig. 10. It can be seen that our model fails to extract the watermark only when the surrogate model completely fails to conduct the debone task.

Influence of Hyper-Parameter and Watermark Size. For the hyper-parameter setting, we conduct control experiments with different λ , which is used to balance the ability of embedding and extracting. As shown in Table 5, although the visual quality of watermarked images and the NC value will be influenced by different λ , the final results are all pretty good, which means our algorithm is not sensitive to λ . We further try different sizes of watermark to test the generalization ability. Since we require the size of watermark to be same as the cover image in our framework, we pad the watermark with 255 if its size is smaller than 256. As shown in Table 6, our method generalizes well to watermarks of different sizes.

5 EXTENSIONS

In this section, we will provide the experiment results for the multiple watermarks and self-watermarked extensions in Sections 5.1 and 5.2 respectively, then discuss how to leverage the proposed framework to protect private data and traditional non-CNN models in Section 5.3.

5.1 Multiple Watermarks Within One Network

As mentioned in Section 3.4.5, the proposed framework is flexible to embed multiple different watermarks with just a single embedding and extractor sub-network. To showcase it, we take the debone task for example and select 10 different logo images from the Internet as watermarks for training. For comparison, we use the logo "IEEE" and "Flower" as two representatives and compare them with the results of the default per-watermark-per-network setting. In terms of the original embedding and extracting ability, as shown in Table 7, the visual quality of this multiple-watermark

TABLE 7 Comparison of the Original Embedding and Extracting Ability Between the Multiple-Watermark (*) and Per-Watermark-Per-Network Setting

Task	PSNR	SSIM	SR_{NC} / SR_C
Debone_IEEE	47.76	0.99	100% / 100%
Debone_IEEE *	41.87	0.99	100% / 100%
Debone_Flower	46.36	0.99	100% / 100%
Debone_Flower *	41.67	0.98	100% / 100%

Two representative watermarks "IEEE" and "Flower" are used here.



Fig. 13. Visual comparisons between the multiple-watermark (even columns) and the per-watermark-per-network (odd column) setting. The top row are the watermarked images and the bottom row are the corresponding extracted watermarks.

setting measured by PSNR degrades from 47.76 to 41.87 compared to the default setting but is still larger than 40, which is reasonably good. Some visual examples are further displayed in Fig. 13. In terms of the robustness to the surrogate model attack, we show the comparison results of different networks and loss functions in Table 8. It can be seen that the extracting success rates of the multiple-watermark setting are comparable to those of the default setting, thus preserving a similar level of robustness.

5.2 Self-Watermarked Model

Since the embedding sub-network H itself is an image processing network and can have different network structures, given a target image processing network **M**, it is possible to absorb the watermark embedding functionality of H into M as shown in Fig. 4. In this way, M is self-watermarked without the need of an extra barrier H. To demonstrate this possibility, we use the derain task with the "Flower" logo image as an example and jointly learn the deraining and watermark embedding process within one single network. In Table 9, we first compare the deraining performance between the original target model and the self-watermarked target model. It can be seen that the self-watermarked model can achieve very close deraining results (PSNR:32.13 and SSIM: 0.93) to the original target model (PSNR:32.49 and SSIM: 0.93). Two visual examples are further showcased in Fig. 14. Besides the deraining functionality, the self-watermarked model can also hide the watermarks well with 0.9999 NC values. In Table 11, we further demonstrate the robustness to the surrogate model attacks of different network structures and loss functions. It shows that the selfwatermarked model is very robust with near 100 percent extracting success rates.

 TABLE 8

 Comparison of the Success Rate (SR_{NC}/SR_C) Against Surrogate Model Attack Between the Multiple-Watermark(*) and Per-Watermark-Per-Network Setting

SR_{NC}/SR_{C}	CNet	Res9	Res16	UNet	L1	L2	$L1 + L_{adv}$	$L2 + L_{adv}$	$L_{perc} + L_{adv}$
Debone_IEEE	93% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	71% / 94%
Debone_IEEE *	89% / 97%	90% / 95%	94% / 97%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	97% / 100%
Debone_Flower	73% / 82%	83% / 88%	89% / 92%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	99% / 100%	100% / 100%
$Debone_Flower*$	94% / 99%	97% / 99%	97% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	99% / 100%

Column 2 \sim 5 are trained with L2 loss but different network structures and Column 6 \sim 10 are trained with UNet network structure but different loss combinations.

TABLE 9 The Performance Comparison Between the Original Target Model ("_original") and Self-Watermarked Target Model ("_Flower")

Task	PSNR	SSIM	SR_{NC} / SR_{C}
Derain_original	32.49	0.93	NA
Derain_Flower	32.13	0.93	100% / 100%

Obviously, the self-watermarked model can keep the original deraining functionality well while embedding the watermarks in the outputs.

TABLE 10
Quantitative Results of Applying the Proposed Framework to
Data (DPED [67]) and Traditional Algorithms (RTV [68])
Protection

Task	PSNR	SSIM	SR_{NC} / SR_C	$SR_{NC} / SR_C \ddagger$
DPED [67]	46.60	0.99	100% / 100%	99% / 100%
RTV [68]	44.21	0.99	100% / 100%	99% / 100%

The extracting success rate marked with \ddagger denotes the results of resisting an example surrogate model attack ($L1 + L_{adw}$). The second to fourth columns represent the results of the learned embedding and extracting sub-network.

5.3 Protect Valuable Data and Traditional Algorithms

In this paper, though most experiments are conducted to simulate CNN model protection, the proposed idea is very general and can be easily applied to data protection and traditional non-CNN algorithm protection. In details, we can

follow the default task-agnostic watermarking setting and embed watermarks into labeled target data or add a watermarking sub-network barrier after traditional algorithms' output. To simulate data protection, we adopt the DPED ("DSLR Photo Enhancement Dataset") dataset [67] as an example. It is collected by using high-end cameras and careful alignment to train a high-quality image enhancement network. For traditional algorithm protection, we consider the famous structure-aware texture smoothing method RTV [68] and choose 6000 images from PASCAL VOC dataset [59] and COCO dataset [61] as its inputs. As our default setting, we split the dataset into different parts for embedding sub-network training and surrogate model attack. The detailed evaluation results are given In Table 10. On the one hand, the learned embedding sub-network can keep the original visual quality very well with high PSNR/SSIM values and the extracting sub-network has a 100 percent extracting success rate. On the other hand, we take the Res9 with $L1 + L_{adv}$ as an example surrogate model, and the proposed framework is able to extract the watermark out when the surrogate model is trained with the target dataset or imitates the behavior of the traditional algorithm. Some visual results are shown in Fig. 15.

6 CONCLUSION AND DISCUSSION

IP protection for deep models is an important but seriously under-researched problem. Inspired by traditional spatial invisible media watermarking and the powerful learning



Fig. 14. One visual example of original target model and self-watermarked model: (a) image a_i , (b) ground-truth image $(b_0)_i$, (c) output of original target model b_i , (d) output of self-watermarked model b_i^* , (e) target watermark, (f) extracted watermark from image b_i^* .

TABLE 11

The Comparison Between Task-Agnostic Watermarking Model and Self-Watermarked Model About the Success Rate (*SR_{NC}/SR_C*) of Resisting the Attack From Surrogate Models Trained With Different Loss Combinations

SR_{NC}/SR_{C}	CNet	Res9	Res16	UNet	<i>L</i> 1	L2	$L1 + L_{adv}$	$L2 + L_{adv}$	$L_{perc}+L_{adv}$
Task-agnostic	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	99% / 100%	100% / 100%
Self-watermarked	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%	100% / 100%



Fig. 15. One visual example of traditional algorithms (RTV [68]) protection : (a) image a_i , (b) watermark-free image b_i , (c) watermarked image b'_i , (d) output of surrogate model (Res9 with $L1 + L_{adv}$), (e) and (f) are the corresponding extracted watermark from (c) and (d).

capacity of deep neural networks, we innovatively propose a novel deep spatial watermarking framework for deep model watermarking. To make it robust to different surrogate model attacks and support image-based watermarks, we dedicatedly design the loss functions and training strategy. Experiments demonstrate that our framework can resist the attack from surrogate models trained with different network structures and loss functions. By jointly training the target model and watermark embedding together, we can even make the target model self-watermarked without the need of an extra watermark embedding sub-network. Moreover, though our motivation is to protect deep models, it is general to data protection and traditional algorithms protection. We hope our work can inspire more explorations of deep model IP protection, such as more types of models (detection, semantic segmentation) and protection strategies.

There are still some interesting questions to explore in the future. For example, though the joint training of the embedding and extracting sub-networks in an adversarial way makes it difficult to find some explicit pattern in the watermarked image, it is worthy to study what implicit watermark is hidden. Besides, our method is not robust enough to some pre-processing techniques for surrogate model attack, such as random cropping and resizing, because the consistency we rely on will be destroyed. So it is necessary to design some new kind of consistency which is intrinsically robust to such pre-processing.

ACKNOWLEDGMENTS

This work was supported in part by the NSFC under Grants 62072421 and 62002334, Exploration Fund Project of University of Science and Technology of China under Grant YD3480002001, and by Fundamental Research Funds for the Central Universities under Grant WK2100000011 and WK5290000001. The work of Jing Liao was supported in part by the ECS Grant from the Research Grants Council of the Hong Kong (Project No. CityU 21209119) and an APRC Grant from CityU, Hong Kong (Project No. 9610488). The work of Gang Hua was supported in part by the National Key R&D Program of China Grant 2018AAA0101400 and NSFC Grant 61629301. Jie Zhang and Dongdong Chen are co-first authors.

REFERENCES

 A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.

- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [3] S. Hong et al."ENCASE: An ensemble classifier for ECG classification using expert features and deep neural networks," in Proc. Comput. Cardiol., 2017, pp. 1–4.
- [4] S. Hong et al. "Combining deep neural networks and engineered features for cardiac arrhythmia detection from ECG recordings," *Physiological Meas.*, vol. 40, 2019, Art. no. 054009.
- [5] J. Žhang, Y. Chen, S. Hong, and H. Li, "REBUILD: Graph embedding based method for user social role identity on mobile communication network," in *Proc. Int. Conf. Data Mining Big Data*, 2017, pp. 326–333.
 [6] S. Hong, M. Wu, H. Li, and Z. Wu, "Event2vec: Learning rep-
- [6] S. Hong, M. Wu, H. Li, and Z. Wu, "Event2vec: Learning representations of events on temporal sequences," in *Proc. Asia-Pacific Web Web-Age Inf. Manage. Joint Conf. Web Big Data*, 2017, pp. 33–47.
- [7] A. Graves, A.-R. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2013, pp. 6645–6649.
- [8] Y. Zhang et al., "Towards end-to-end speech recognition with deep convolutional neural networks," *Interspeech*, pp. 410–414, 2016.
- [9] A. Vaswani et al."Attention is all you need," in Proc. Int. Conf. Neural Inf. Process. Syst., 2017, pp. 5998–6008.
- [10] Y. Uchida, Y. Nagai, S. Sakazawa, and S. Satoh, "Embedding watermarks into deep neural networks," in *Proc. ACM Int. Conf. Multimedia Retrieval*, 2017, pp. 269–277.
- [11] Y. Adi, C. Baum, M. Cisse, B. Pinkas, and J. Keshet, "Turning your weakness into a strength: Watermarking deep neural networks by backdooring," in *Proc. 27th USENIX Conf. Secur. Symp.*, 2018, pp. 1615–1631.
- [12] J. Ba and R. Caruana, "Do deep nets really need to be deep?" in Proc. Int. Conf. Neural Inf. Process. Syst., 2014, pp. 2654–2662. [Online]. Available: http://papers.nips.cc/paper/5484-do-deepnets-really-need-to-be-deep.pdf
- [13] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *Stat*, vol. 1050, p. 9, 2015.
- [14] J. Zhang et al."Protecting intellectual property of deep neural networks with watermarking," in Proc. Asia Conf. Comput. Commun. Secur., 2018, pp. 159–172.
- [15] E. L. Merrer, P. Perez, and G. Trédan, "Adversarial frontier stitching for remote neural network watermarking," *Neural Comput. Appl.*, vol. 32, no. 13, pp. 9233–9244, 2020.
- [16] J. Zhang et al. "Model watermarking for image processing networks," Proc. AAAI Conf. Artif. Intell., 2020, pp. 12805–12812.
- [17] M. Kutter, "Watermarking resistance to translation, rotation, and scaling," *Multimedia Syst. Appl.*, vol. 3528, pp. 423–431, 1999.
- [18] S. Voloshynovskiy, F. Deguillaume, and T. Pun, "Content adaptive watermarking based on a stochastic multiresolution image modeling," in *Proc. 10th Eur. Signal Process. Conf.*, 2000, pp. 1–4.
- pp. 1–4.
 [19] S. Voloshynovskiy, F. Deguillaume, and T. Pun, "Multibit digital watermarking robust against local nonlinear geometrical distortions," in *Proc. IEEE Int. Conf. Image Process.*, 2001, pp. 999–1002.
- [20] J. Zhao and E. Koch, "Embedding robust labels into images for copyright protection," in Proc. Conf. Intellectual Property Rights New Technol., 1995, pp. 242–251.
- [21] N. Nikolaidis and I. Pitas, "Copyright protection of images using robust digital signatures," in Proc. IEEE Int. Conf. Acoust. Speech Signal Process., 1996, pp. 2168–2171.

- [22] C.-H. Lee and Y.-K. Lee, "An adaptive digital image watermarking technique for copyright protection," *IEEE Trans. Consum. Electron.*, vol. 45, no. 4, pp. 1005–1015, Nov. 1999.
- [23] S. Zafeiriou, A. Tefas, and I. Pitas, "Blind robust watermarking schemes for copyright protection of 3D mesh objects," *IEEE Trans. Vis. Comput. Graphics*, vol. 11, no. 5, pp. 596–607, Sep./Oct. 2005.
- [24] D.-C. Lou, H.-K. Tso, and J.-L. Liu, "A copyright protection scheme for digital images using visual cryptography technique," *Comput. Standards Interfaces*, vol. 29, no. 1, pp. 125–131, 2007.
- [25] H. Fang et al., "Deep template-based watermarking," IEEE Trans. Circuits Syst. Video Technol., early access, Jul. 15, 2020, doi: 10.1109/TCSVT.2020.3009349.
- [26] F. Deguillaume, S. V. Voloshynovskiy, and T. Pun, "Method for the estimation and recovering from general affine transforms in digital watermarking applications," *Secur. Watermarking Multimedia Contents IV*, vol. 4675, pp. 313–322, 2002.
- [27] C.-T. Hsu and J.-L. Wu, "Hidden digital watermarks in images," IEEE Trans. Image Process., vol. 8, no. 1, pp. 58–68, Jan. 1999.
- [28] J. R. Hernandez, M. Amado, and F. Perez-Gonzalez, "DCTdomain watermarking techniques for still images: Detector performance analysis and a new structure," *IEEE Trans. Image Process.*, vol. 9, no. 1, pp. 55–68, Jan. 2000.
- [29] M. Barni, F. Bartolini, and A. Piva, "Improved wavelet-based watermarking through pixel-wise masking," *IEEE Trans. Image Process.*, vol. 10, no. 5, pp. 783–791, May 2001.
- Process., vol. 10, no. 5, pp. 783–791, May 2001.
 [30] J. J. K. O. Ruanaidh, W. J. Dowling, and F. M. Boland, "Phase watermarking of digital images," in *Proc. IEEE Int. Conf. Image Process.*, 1996, pp. 239–242.
- [31] J. Zhu, R. Kaplan, J. Johnson, and L. Fei-Fei , "HiDDeN: Hiding data with deep networks," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 657–672.
- [32] M. Tancik, B. Mildenhall, and R. Ng, "StegaStamp: Invisible hyperlinks in physical photographs," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 2117–2126.
- [33] F. Deeba, S. Kun, F. Dharejo, and H. Memon, "Digital watermarking using deep neural network," Int. J. Mach. Learn. Comput., vol. 10, no. 2, 2020, pp. 277–282.
- [34] Y. Nagai, Y. Uchida, S. Sakazawa, and S. Satoh, "Digital watermarking for deep neural networks," Int. J. Multimedia Inf. Retrieval, vol. 7, pp. 3–16, 2018.
- [35] H. Chen, B. D. Rouhani, and F. Koushanfar, "BlackMarks: Blackbox multibit watermarking for deep neural networks," 2019, arXiv: 1904.00344.
- [36] B. Darvish Rouhani, H. Chen, and F. Koushanfar, "DeepSigns: An end-to-end watermarking framework for ownership protection of deep neural networks," in *Proc. 24th Int. Conf. Architect. Support Program. Lang. Operating Syst.*, 2019, pp. 485–497.
- [37] H. Wu, G. Liu, Y. Yao, and X. Zhang, "Watermarking neural networks with watermarked images," *IEEE Trans. Circuits Syst. Video Technol.*, early access, Oct. 13, 2020, doi: 10.1109/ TCSVT.2020.3030671.
- [38] J. Guo and M. Potkonjak, "Watermarking deep neural networks for embedded systems," in Proc. IEEE/ACM Int. Conf. Comput.-Aided Des., 2018, pp. 1–8.
- [39] L. Fan, K. W. Ng, and C. S. Chan, "Rethinking deep neural network ownership verification: Embedding passports to defeat ambiguity attacks," in Proc. Int. Conf. Neural Inf. Process. Syst., 2019, pp. 4716–4725.
- [40] J. Zhang, D. Chen, J. Liao, W. Zhang, G. Hua, and N. Yu, "Passport-aware normalization for deep model protection," in *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 22619–22628, 2020.
 [41] D. Chen, L. Yuan, J. Liao, N. Yu, and G. Hua, "StyleBank: An
- [41] D. Chen, L. Yuan, J. Liao, N. Yu, and G. Hua, "StyleBank: An explicit representation for neural image style transfer," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1897–1906.
- [42] D. Chen, J. Liao, L. Yuan, N. Yu, and G. Hua, "Coherent online video style transfer," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 1105–1114.
- [43] D. Chen, L. Yuan, J. Liao, N. Yu, and G. Hua, "Explicit filterbank learning for neural image style transfer and image processing," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Jan. 06, 2020, doi: 10.1109/TPAMI.2020.2964205.
- [44] D. Chen *et al.*"Gated context aggregation network for image dehazing and deraining," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2019, pp. 1375–1383.
- [45] X. Li, J. Wu, Z. Lin, H. Liu, and H. Zha, "Recurrent squeeze-andexcitation context aggregation net for single image deraining," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 254–269.

- [46] D. Chen *et al.* "Controllable image processing via adaptive filterbank pyramid," *IEEE Trans. Image Process.*, vol. 29, pp. 8043–8054, 2020.
- [47] W. Yang *et al.* "Cascade of multi-scale convolutional neural networks for bone suppression of chest radiographs in gradient domain," *Med. Image Anal.*, vol. 35, pp. 421–433, 2017.
- [48] I. Goodfellow et al."Generative adversarial nets," in Proc. Int. Conf. Neural Inf. Process. Syst., 2014, pp. 2672–2680.
- [49] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5967–5976.
 [50] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo,
- [50] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "StarGAN: Unified generative adversarial networks for multidomain image-to-image translation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8789–8797.
- [51] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, "High-resolution image synthesis and semantic manipulation with conditional GANs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8798–8807.
 [52] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image
- [52] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 2337–2346.
- [53] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. Int. Conf. Comput. Vis.*, 2017, pp. 2223–2232.
 [54] Z. Yi, H. Zhang, P. Tan, and M. Gong, "DualGAN: Unsupervised
- [54] Z. Yi, H. Zhang, P. Tan, and M. Gong, "DualGAN: Unsupervised dual learning for image-to-image translation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2849–2857.
- [55] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim, "Learning to discover cross-domain relations with generative adversarial networks," in *Proc. 34th Int. Conf. Mach. Learn.*, 2017, pp. 1857–1865.
- [56] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention, 2015, pp. 234–241.
- [57] Q. Fan, J. Yang, G. Hua, B. Chen, and D. Wipf, "A generic deep architecture for single image reflection removal and image smoothing," in *Proc. Int. Conf. Comput. Vis.*, 2017, pp. 3238–3247.
- [58] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for realtime style transfer and super-resolution," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 694–711.
- [59] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (VOC) challenge," *Int. J. Comput. Vis.*, vol. 88, no. 2, pp. 303–338, 2010.
- [60] H. Zhang and V. M. Patel, "Density-aware single image de-raining using a multi-stream dense network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 695–704.
 [61] T.-Y. Lin *et al.* "Microsoft COCO: Common objects in context,"
- [61] T.-Y. Lin et al."Microsoft COCO: Common objects in context," in Proc. Eur. Conf. Comput. Vis., 2014, pp. 740–755.
- [62] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "ChestX-Ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3462–3471.
- [63] H. Fang, W. Zhang, H. Zhou, H. Cui, and N. Yu, "Screen-shooting resilient watermarking," *IEEE Trans. Inf. Forensics Security*, vol. 14, no. 6, pp. 1403–1418, Jun. 2019.
- [64] X. Kang, J. Huang, and W. Zeng, "Efficient general print-scanning resilient data hiding based on uniform log-polar mapping," *IEEE Trans. Inf. Forensics Secur.*, vol. 5, no. 1, pp. 1– 12, Mar. 2010.
- [65] X. Kang, J. Huang, Y. Q. Shi, and Y. Lin, "A DWT-DFT composite watermarking scheme robust to both affine transform and JPEG compression," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 13, no. 8, pp. 776–786, Aug. 2003.
- [66] S.-Y. Wang, O. Wang, R. Zhang, A. Owens, and A. A. Efros, "CNNgenerated images are surprisingly easy to spot... for now," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 8692–8701.
- [67] A. Ignatov, N. Kobyshev, R. Timofte, K. Vanhoey, and L. Van Gool, "DSLR-quality photos on mobile devices with deep convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 3277–3285.
- pp. 3277–3285.
 [68] L. Xu, Q. Yan, Y. Xia, and J. Jia, "Structure extraction from texture via relative total variation," *ACM Trans. Graph.*, vol. 31, no. 6, pp. 1–10, 2012.



Jie Zhang received the BS degree from the China University of Geosciences, Beijing, China, in 2017. He is currently working toward the PhD degree from the University of Science and Technology of China (USTC), Hefei, China. His primary research interests include IP protection of deep models, media watermarking, and AI security.



Dongdong Chen received the PhD degree from the University of Science and Technology of China (USTC), Hefei, China, in 2019 under the joint PhD program between MSRA and USTC. He is currently a senior researcher with Microsoft Research. His research interests include image generation, style transfer, Al security, and general representation learning.



Huamin Feng received the BS degree from Northeast Heavy Machinery Institute, Heilongjiang, China, in 1984. He is the vice president at the Beijing Institute of Electronic Science and Technology currently, where he is also a professor. His research interests include information security and cyberspace security.



Gang Hua (Fellow, IEEE) is currently the vice president and chief scientist of Wormpex AI Research. Before that, he was the principal researcher/research manager with Microsoft Research between 2015 to 2018. He was an associate professor of computer science with the Stevens Institute of Technology between 2011 and 2015, while holding an academic advisor position with IBM T. J. Watson Research Center. He has published more than 150 peer reviewed papers in top conferences such as CVPR/ICCV/

ECCV, and top journals such as the *IEEE Transactions on Pattern Analysis and Machine Intelligence* and the *International Journal of Computer Vision.* To date, he holds 18 issued U.S Patents and also has 14 more U.S. Patents Pending. He is an IAPR fellow, and an ACM distinguished scientist. His research focuses on artificial intelligence, computer vision, pattern recognition, machine learning, and robotics, with primary applications in the cloud and mobile intelligence domain.



Nenghai Yu received the BS degree from the Nanjing University of Posts and Telecommunications, Nanjing, China, in 1987, the ME degree from Tsinghua University, Beijing, China, in 1992, and the PhD degree from the University of Science and Technology of China, Hefei, China, in 2004, where he is currently a professor. His research interests include multimedia security, multimedia information retrieval, video processing, and information hiding.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.



Jing Liao received the BEng degree from the Huazhong University of Science and Technology, Wuhan, China, and the dual PhD degrees from Zhejiang University, Hangzhou, China and the Hong Kong University of Science and Technology, Hong Kong. She is currently an assistant professor at the Department of Computer Science, City University of Hong Kong (CityU) since September 2018. Prior to that, she was a researcher with Visual Computing Group, Microsoft Research Asia from 2015 to 2018. Her primary research

interests fall in the fields of computer graphics, computer vision, image/ video processing, digital art, and computational photography.



Weiming Zhang received the MS and PhD degrees from the Zhengzhou Information Science and Technology Institute, Zhengzhou, China, in 2002 and 2005, respectively. Currently, he is a professor at the School of Information Science and Technology, University of Science and Technology of China. His research interests include information hiding and multimedia security.