

# SepMark: Deep Separable Watermarking for Unified Source Tracing and Deepfake Detection

Xiaoshuai Wu  
shinewu@hnu.edu.cn  
College of Computer Science and  
Electronic Engineering  
Hunan University, Changsha, China

Xin Liao\*  
xinliao@hnu.edu.cn  
College of Computer Science and  
Electronic Engineering  
Hunan University, Changsha, China

Bo Ou  
oubo@hnu.edu.cn  
College of Computer Science and  
Electronic Engineering  
Hunan University, Changsha, China

## ABSTRACT

Malicious Deepfakes have led to a sharp conflict over distinguishing between genuine and forged faces. Although many countermeasures have been developed to detect Deepfakes ex-post, undoubtedly, passive forensics has not considered any preventive measures for the pristine face before foreseeable manipulations. To complete this forensics ecosystem, we thus put forward the proactive solution dubbed SepMark, which provides a unified framework for source tracing and Deepfake detection. SepMark originates from encoder-decoder-based deep watermarking but with two separable decoders. For the first time the deep separable watermarking, SepMark brings a new paradigm to the established study of deep watermarking, where a single encoder embeds one watermark elegantly, while two decoders can extract the watermark separately at different levels of robustness. The robust decoder termed Tracer that resists various distortions may have an overly high level of robustness, allowing the watermark to survive both before and after Deepfake. The semi-robust one termed Detector is selectively sensitive to malicious distortions, making the watermark disappear after Deepfake. Only SepMark comprising of Tracer and Detector can reliably trace the trusted source of the marked face and detect whether it has been altered since being marked; neither of the two alone can achieve this. Extensive experiments demonstrate the effectiveness of the proposed SepMark on typical Deepfakes, including face swapping, expression reenactment, and attribute editing. Code will be available at <https://github.com/sh1newu/SepMark>.

## CCS CONCEPTS

• Security and privacy → Digital rights management.

## KEYWORDS

deep watermarking, deepfake forensics, watermarking robustness

### ACM Reference Format:

Xiaoshuai Wu, Xin Liao, and Bo Ou. 2023. SepMark: Deep Separable Watermarking for Unified Source Tracing and Deepfake Detection. In *Proceedings*

\*Corresponding author.

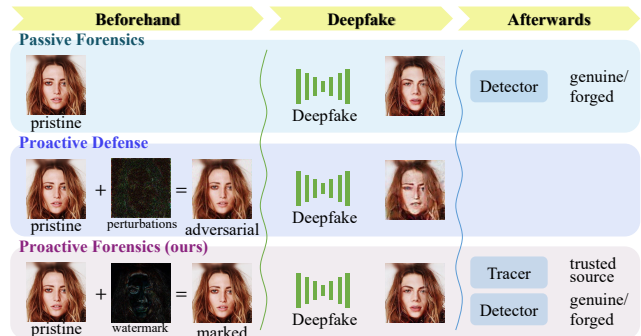
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

MM '23, October 29–November 3, 2023, Ottawa, ON, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0108-5/23/10...\$15.00

<https://doi.org/10.1145/3581783.3612471>



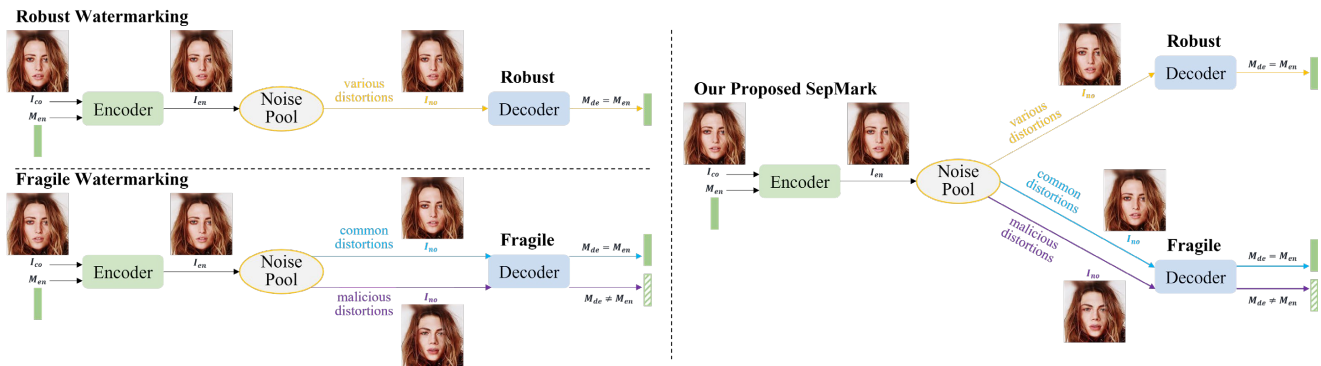
**Figure 1: Illustration of different types of Deepfake countermeasures: passive forensics pays attention to the task of Deepfake detection while the pristine face is non-protected; proactive defense concentrates on the task of Deepfake destruction; and our proactive forensics focuses on both the tasks of source tracing and Deepfake detection.**

of the 31st ACM International Conference on Multimedia (MM '23), October 29–November 3, 2023, Ottawa, ON, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3581783.3612471>

## 1 INTRODUCTION

Until now, AI-generated content has been fully fueled in the deep learning community, such as Generative Pre-trained Transformer [61, 62] and Stable Diffusion [63]. In this era, everyone can be a master manipulator who creates authentic-seeming content with just one click [1]. However, the breakthrough progress in multimedia generation has once again brought Deepfake, which mainly refers to the technology of faking face using deep learning [68, 76, 77], to the forefront. Generally speaking, the development of the forgery generation is usually ahead of its forensic counterpart in time [13, 36]. Despite its significance in face de-identification [4, 40] and privacy protection [53, 82], some individuals intentionally produce and spread Deepfakes with high visual deception to distort the truths and manipulate public opinions. In the context of “seeing is not always believing”, uncontrollable Deepfakes will eventually lead to unprecedented trust crisis and moral panic. Therefore, the forensics of Deepfakes is becoming increasingly important and urgent.

Passive forensics has been introduced since the emergence of Deepfakes [88]. Hand-craft features such as eye blink [43], head pose [78], color inconsistency [41], etc. work well for detecting them, at an early stage. However, more clever creators using adversarial training can readily fix these shallow manipulation artifacts



**Figure 2: The pipelines of robust watermarking (left top), fragile watermarking (left bottom), and our proposed SepMark (right). For comparison and clarity, the additional adversary discriminator is omitted here. In the prior encoder-decoder structure, an encoder and a decoder are jointly trained by sampling various distortions from the noise pool, where the robust decoder can resist against all of the distortions and extract the watermark in a high level of robustness; alternatively, the fragile one is selectively sensitive to malicious distortions and the watermark can only be extracted in a lower level of robustness. SepMark innovatively integrates a single encoder with two separable decoders, enabling the embedded watermark to be extracted separately at different levels of robustness.**

[12, 29]. Vanilla deep-learned features were demonstrated to be effective [64], while suffering from generalization when encountering unseen forgery methods. Therefore, some data-driven models aim to learn common artifacts, such as blending boundary [42], spectral anomaly [15], temporal inconsistency [23], etc., while avoiding overfitting specific forgery method. Nevertheless, their performances when the inquiry face undergoes several post-processing operations or adversarial perturbations are unclear [5, 24, 30, 45, 55], rendering that the in-the-wild forensics remains a challenge due to limited robustness [22, 28, 91]. Despite a handful of proactive defense methods using adversarial attacks have been dedicated to disrupting or nullifying the creation of Deepfakes [1, 20, 26, 27, 65, 71, 79, 80], they may contribute to the detection task [9, 73], but not to the source tracing task [47, 59, 72] which is another focus of our proactive forensics. We highlight the differences between these countermeasures in Fig. 1, and the application of proactive forensics will be detailed later in Section 3.1.

To our knowledge, invisible watermarking is the art of imperceptibly concealing additional messages within valuable media without influencing their nature-looking [11]. According to whether the original image is available during watermark extraction, invisible watermarking can be further categorized as blind and non-blind. Since the extraction needs only the marked image, blind watermarking has increasingly broader applications in modern multimedia ecology [6, 18, 81], including copyright protection, source tracing, content authentication, etc. Therefore, we leverage blind watermarking to combat Deepfakes in this paper. Thanks to the efforts made by deep learning in blind watermarking [89], encoder-decoder-based deep watermarking has attracted more interest owing to its effectiveness and flexibility [34, 56], compared to traditional watermarking [57, 75]. Still, in accordance with Fig. 2, let us conduct an in-depth analysis of why current deep watermarking with single-function cannot meet the proactive forensics well.

Robust watermarking places a high priority on robustness, emphasizing that the embedded watermark can still be extracted after

severe distortions [89]. This is crucial to protect the copyright and trace the source, even though the marked image has major portions occluded or has been captured by a screenshot [33, 70]. However, the overly high level of robustness may also allow the watermark to survive both before and after Deepfake, making it difficult to distinguish between genuine and forged images. To make matters worse, the original source identification is preserved even after the image has been forged. Therefore, relying solely on robust watermarks cannot detect Deepfakes while tracing their trusted sources.

Fragile watermarking focuses on full or selective fragility, which is opposite to robust watermarking, and emphasizes that the embedded watermark will change accordingly with the marked image under any or specific distortions [2]. Indeed, it mainly indicates whether the content has been altered by malicious distortions, but cannot provide information about the source of original trustworthy image due to the disappeared watermark [56]. More importantly, it is hard to tell apart the forged image from pristine images without prior knowledge, since all of them have no correct watermark [69]. Therefore, solely relying on fragile watermarks cannot always trace the trusted source or detect Deepfakes.

To accomplish both the tasks of source tracing and Deepfake detection, the most intuitive manner is to embed one robust watermark followed by another fragile watermark [66], while the two watermarks may significantly affect each other and concurrently degrade the visual quality. Another intuition is to embed these independent watermarks into separate regions of interest, foreground and background of the face image, for example, while we argue that the division of the regions may require the participation of extra sophisticated face parsing models. More importantly, the watermarks are more likely to be perceived and attacked by virtue of the knowledge of the interest regions. To remedy these limitations of multi-embedding, a learning-based, end-to-end SepMark is proposed as a whole, corresponding to the new paradigm where one elegant embedding has two distinct extraction manners. To conclude, we have the following remarkable contributions:

1) We propose for the first time the deep separable watermarking, SepMark, which brings a new paradigm to current deep watermarking, where a single encoder embeds one watermark elegantly, while two decoders can extract the watermark separately at different levels of robustness. On top of that, our proposed SepMark also naturally provides a novel proactive forensics solution for unified source tracing and Deepfake detection.

2) This is the very first end-to-end learning architecture for deep separable watermarking, which consists of a single encoder, a discriminator, and two separable decoders trained by sampling different types of distortions from a random forward noise pool (RFNP). The robust decoder resists against all the common and malicious distortions while the semi-robust one is selectively sensitive to malicious distortions.

3) Extensive experiments conducted on face images demonstrate the high level of robustness of robust decoder under numerous distortions, and the selective fragility of semi-robust decoder under malicious Deepfake distortions. Therefore, the separable decoders can reliably trace the trusted source of the marked face and detect whether it has been altered since being marked.

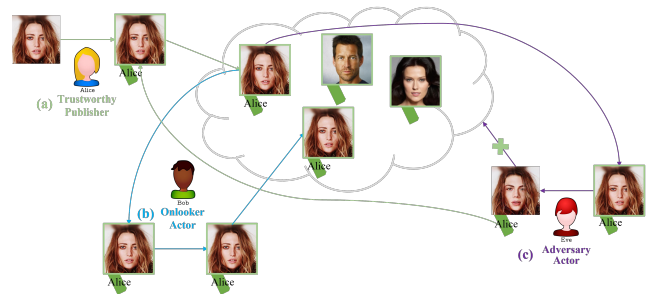
## 2 RELATED WORK

### 2.1 Deepfake Passive Forensics

Passive forensics, especially self-supervised Deepfake detection, has recently gained much attention [85]. The fundamental principle among them is the synthesis of self-craft Deepfakes driven by positive examples. In general, their detection performances depend on the prior knowledge introduced during the production of synthetic data. FWA [44] targets the artifacts of affine face warping due to inconsistent resolutions between inner and outer faces. Face X-ray [42] focuses on the artifacts of blending boundary that reveal whether the inquiry image can be decomposed into the blending of two faces from different sources. Similarly, I2G [86] aims to generate the blended faces with inconsistent source features through a blending pipeline. SBIs [67] applies two random transformations to the single image, one as the source and the other as the target of the blending process, to reproduce more challenging negative samples. Moreover, ICT [14] considers the identity consistency in the inquiry image, such that the inner face and outer face of different identities are swapped to form synthetic data. Following these additional cues [7, 21, 87], more general modeling of Deepfake distortions can take a step forward towards generalizing to unseen forgeries.

### 2.2 Deepfake Proactive Defense

Similar to proactive forensics, proactive defense requires the addition of imperceptible signals to the pristine image in advance. Deepfake destruction can be achieved either by disrupting or nullifying the Deepfake model through the injection of adversarial perturbations. The former can be regarded as a non-targeted attack, where the attacked image after Deepfake should be visually abnormally and as unnatural as possible compared to the original forged one [26, 27, 65, 71, 80]. Conversely, the latter is a targeted attack, where the attacked image after Deepfake should be as similar as possible to the original pristine one [20, 79]. Inspired by universal adversarial perturbations [54], CMUA [26] generates image-universal cross-model perturbations to attack multiple Deepfake models rather

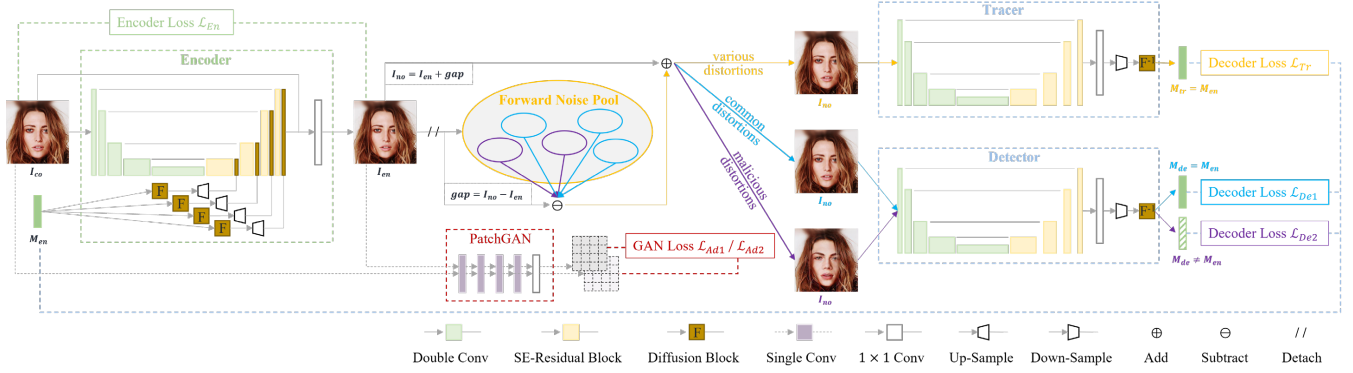


**Figure 3: Illustrative example of proactive forensics. (a) Trustworthy publisher Alice embeds her identification into the published face. (b) Onlooker actor Bob reposts this face without affecting the content semantics. (c) Adversary actor Eve manipulates this face to distort the truth.**

than the specific one. Anti-Forgery [71] further enhances the robustness of perceptual-aware perturbations against various distortions. However, the visually perceptible artifacts might still escape the detection since passive detectors do not share a similar perspective as human eyes [73]. Moreover, counterparts like MagDR [9] can detect and remove the adversarial perturbations; even the Deepfakes can be further reconstructed normally.

### 2.3 Deep Watermarking

Motivated by the success of deep learning in computer vision tasks, DNN-based models have been investigated in blind watermarking for stronger robustness. HiDDeN [89] is the first end-to-end learning architecture for robust watermarking, which consists of an encoder, a discriminator, a noise layer and a decoder. The key to guaranteeing robustness is the data augmentation of the marked image with the differential noise layer. To resist against non-differential JPEG distortion, MBRS [34] randomly switches real JPEG and simulated differential JPEG for each mini-batch training. CIN [52] incorporates an invertible neural network to learn a joint representation between watermark embedding and extraction, where a non-invertible attention-based module is additionally devised against non-additive quantization noise. In another distortion-agnostic line [51], TSDL [48] proposes a two-stage learning architecture, where both the encoder and decoder are first trained in a noise-free manner, following that only the decoder is fine-tuned to adapt to black-box distortions. ASL [83] models the forward noise as pseudo-noise that interacts with the marked image, allowing the gradient to be propagated to the encoder without passing through the noise layer during backward propagation. FIN [17] handles the black-box distortion through an invertible noise layer similar to invertible denoising [49]; it is designed to simulate the distortion, serves as a noise layer during training, and is also applied as a preprocessing step before extraction to eliminate the distortion. Moreover, LFM [74], StegaStamp [70], RIHOOP [32], and PIMoG [16] learn to resist against physical distortions like printer-scanner, printer-camera, and screen-camera that are cross-media. To our knowledge, although deep robust watermarking has been developed a lot, deep fragile watermarking remains largely unexplored.



**Figure 4: Overall architecture of SepMark.** The encoder  $En$  receives the cover image  $I_{co}$  and encoded message  $M_{en}$  as input and produces the encoded image  $I_{en}$ . The random forward noise pool  $RFNP$  samples different types of pseudo-noise  $gap$  in a detached forward propagation and interacts with  $I_{en}$  in both standard forward and backward propagation, resulting in the common or maliciously distorted image  $I_{no}$ . The tracer  $Tr$  extracts the message  $M_{tr}$  which is identical to  $M_{en}$  from the arbitrary distorted image. The detector  $De$  extracts the correct message from the common distorted image but the wrong message from the maliciously distorted image. Additionally, a PatchGAN  $Ad$  is used to distinguish each pristine and encoded image patch.

### 3 PROACTIVE FORENSICS: SEPMARK

#### 3.1 Problem Statement

Instead of detecting which are forged afterwards, SepMark is dedicated to ensuring beforehand which are genuine and tracing their trusted sources. This proactive forensics should work on the pristine face in advance, where “pristine” means the face has not been processed by SepMark in this paper. In the real world, a trustworthy content publisher uploads his/her pristine face to social networks and makes it public to anyone. Unfortunately, some adversary individuals may download the face and manipulate it to spread misinformation, which causes severe impacts on social trust and public interest. Moreover, relying on ex-post forensics, the spread of Deepfakes cannot be immediately blocked before it is confirmed.

Here is a more illustrative example regarding the suggested proactive forensics shown in Fig. 3. To provide preventive measures and more reliable evidence for the pristine face, SepMark embeds one watermark into the face in advance (Fig. 3-a), such as the publisher identification and time stamp, which can identify the trusted source of the marked face and at the same time indicate it has been processed by SepMark. Subsequently, the embedded watermark can be extracted separately at different levels of robustness. Concretely, both the robust and semi-robust watermarks can be extracted correctly under common distortions that have little effect on image semantics (Fig. 3-b), such as JPEG compression, resizing, contrast adjustments, etc. In contrast, under malicious distortions such as Deepfake, which is mainly addressed here (Fig. 3-c), the robust watermark can still be extracted correctly, while the semi-robust one is extracted incorrectly, implying that the face image has been altered since it was marked. In this manner, any individuals or social networks can verify whether the original source of the face comes from a trustworthy publisher. Furthermore, even without the knowledge of original watermark, it is possible to compare the two extracted watermarks to determine whether the marked face has been manipulated after publication, and block the further spreading of the suspicious image with mismatched watermarks.

In summary, SepMark accomplishes both provenance tracking and blind detection of Deepfake without knowing the original image.

Following Fig. 4, let us formulate SepMark in detail. Given an arbitrary pristine face which also refers to as the cover image  $I_{co}$ , the encoder  $En$  embeds encoded message  $M_{en}$  into the image  $I_{co}$ , thus obtaining an encoded image  $I_{en}$ . After that, the separable Tracer  $Tr$  and Detector  $De$  can extract the messages  $M_{tr}$  and  $M_{de}$  from the common or maliciously distorted image  $I_{no}$ , respectively. Consequently, depending on the separable decoders  $Tr$  and  $De$ , we can divide the proactive forensics into the following cases.

Case 1: If the robust decoder  $Tr$  is available, the extracted  $M_{tr}$  is expected to always be identical to the embedded  $M_{en}$ , whether the marked image is noised by common or malicious distortions, to reliably trace the trusted source.

Case 2: If the semi-robust decoder  $De$  is available, it can extract the correct message from the common distorted image but the wrong message from the image noised by malicious distortions, to distinguish the type of the distortion.

Case 3: Only when both the decoders  $Tr$  and  $De$  are available, the tracing of the trusted source and the detection of malicious distortions can be realized simultaneously. More specifically, if  $M_{de} \approx M_{en}$  or  $M_{de} \approx M_{tr}$ , SepMark will indicate that the face is genuine and comes from the trusted source  $M_{tr}$ . Otherwise, if  $M_{de} \not\approx M_{en}$  or  $M_{de} \not\approx M_{tr}$ , the face has been forged since it was marked, and the original trustworthy one can be further acquired from the source  $M_{tr}$  if necessary.

#### 3.2 Model Architecture

To achieve end-to-end training of deep separable watermarking, our proposed SepMark consists of a single encoder  $En$ , a random forward noise pool  $RFNP$ , two separable decoders  $Tr$  and  $De$ , and an additional adversary discriminator  $Ad$ , as can be seen in Fig. 4. In a nutshell, 1)  $En$  receives a batch of  $I_{co} \in \mathbb{R}^{B \times 3 \times H \times W}$  and a batch of  $M_{en} \in \{-\alpha, \alpha\}^{B \times L}$  to produce  $I_{en}$ ; 2)  $RFNP$  randomly samples different distortions as pseudo-noise and interacts with  $I_{en}$ ,

resulting in  $I_{no}$ ; 3)  $Tr$  and  $De$  separately extract  $M_{tr}$  and  $M_{de}$  from  $I_{no}$  with different levels of robustness; and 4)  $Ad$  classifies pristine or encoded for each image patch of  $I_{en}$  relative to  $I_{co}$ .

**3.2.1 Single Encoder.** In the main line,  $I_{co}$  is fed into a UNet-like structure. Firstly, in the downsampling stage, “conv-in-relu” with stride 2 halves the spatial dimensions, and “Double Conv (2× conv-in-relu)” doubles the number of feature channels. In the parallel lines,  $M_{en}$  is fed into several “Diffusion Block (FC layer-Conv)” to inject redundancy, where its length  $L$  is expanded to  $L \times L$ . Secondly, in the upsampling stage, nearest-neighbour interpolation doubles the spatial dimensions and “conv-in-relu” halves the number of feature channels, followed by the concatenation with both the corresponding feature maps from the downsampling stage and the redundant message after interpolation operation at each resolution. Moreover, “SE-Residual Block” inserted with Squeeze-and-Excitation [25] replaces batch normalization in the original bottleneck residual block with instance normalization, and is applied to the concatenated features to enhance the information interaction between different channels. Finally, we further concatenate the aforementioned feature maps with the cover image via a skip connection, and fed the output into a  $1 \times 1$  convolution to produce the encoded image  $I_{en}$ .

**3.2.2 Random Forward Noise Pool.** Considering that numerous distortions and a specific noise layer cannot adapt to all of them, it is proposed to introduce various distortions to our noise pool for the combined training. Moreover, given that the simulated distortion struggles to approach the real complex distortion, we directly get arbitrary real distortion  $gap = I_{no} - I_{en}$  in a detached forward propagation. In this way, both the standard forward and backward propagation can be realized through differentiable  $I_{no} = I_{en} + gap$ . It is noteworthy that the difference between  $RFNP$  and ASL [83] lies in the random sampling of different distortions in each mini-batch while not that of only one distortion. Empirically, this change benefits our stable training and explains the instance normalization in our encoder design. To be specific, random sampling from common and malicious distortions, common distortions only, and malicious distortions only results in three batches of  $I_{no}$ .

**3.2.3 Separable Decoders.** We simply borrow the UNet-like structure from our encoder for the two decoders. Note that after the concatenations of corresponding feature maps at each resolution, we fed the final output into a  $1 \times 1$  convolution to produce a single-channel residual image. After that, the residual image is down-sampled to  $L \times L$ , and an “Inverse Diffusion Block (FC layer)” is used to extract the decoded message of length  $L$ . More importantly, due to different inputs and optimization objectives during training, the decoders have the same structural design but do not share model parameters; see our carefully designed loss functions later. To elaborate further, the robust  $Tr$  receives the image  $I_{no}$  noised by various distortions, while the semi-robust  $De$  receives the common and maliciously distorted ones, respectively, in each mini-batch training. Moreover,  $Tr$  is optimized by the objective  $M_{tr} = M_{en}$  under arbitrary distortions. In comparison,  $De$  is optimized by the objective  $M_{de} = M_{en}$  under common distortions, and the objective  $M_{de} \neq M_{en}$  under malicious distortions.

**3.2.4 Adversary Discriminator.** To supervise the visual quality of encoded image  $I_{en}$ , we also resort to adversarial training, where an

additional discriminator  $Ad$  has trained alternately with the main encoder-decoder. In particular, vanilla GAN is replaced with Patch-GAN [31, 90] to classify each image patch as pristine or encoded. Specifically, the discriminator is a fully-convolutional structure consisting of several “conv-in-leakyrelu” blocks, and finally, a  $1 \times 1$  convolution outputs the ultimate prediction map. However, the GAN training process was found usually unstable and even non-convergence when using the standard GAN loss [19]. This behavior also appears in our training process, and we conjecture that it might also be related to the task of message extraction, in addition to adversarial generation and discrimination. To improve the stability of the training, we use the GAN loss of RaLSGAN [35] instead, and we find that this adjustment makes our training more stable. This might be mainly attributed to the transformation from absolute true or false into relative true or false.

### 3.3 Loss Functions

We found simple  $L_2$  constraint is sufficient for the end-to-end training of SepMark. Note that  $\theta$  indicates trainable parameters, and  $\theta_{co}$  and  $\theta_{ma}$  only denote common and malicious distortions sampled by  $RFNP$ , respectively. We have the following loss functions.

For the encoder, its encoded image  $I_{en}$  should be visually similar to the input image  $I_{co}$ :

$$\mathcal{L}_{En} = L_2(I_{co}, En(\theta, I_{co}, M_{en})) = L_2(I_{co}, I_{en}) \quad (1)$$

For the decoders, the robust Tracer  $Tr$  should be able to extract  $M_{tr}$ , which is identical to  $M_{en}$ , from the image  $I_{no}$  noised by arbitrary common and malicious distortions:

$$\mathcal{L}_{Tr} = L_2(M_{en}, Tr(\theta, RFNP(\theta_{co} + \theta_{ma}, I_{en}))) = L_2(M_{en}, M_{tr}) \quad (2)$$

where

$$I_{en} = En(\theta, I_{co}, M_{en}) \quad (3)$$

The semi-robust Detector  $De$  is expected to be fault-tolerant towards the common distorted image  $I_{no}$ :

$$\mathcal{L}_{De1} = L_2(M_{en}, De(\theta, RFNP(\theta_{co}, I_{en}))) = L_2(M_{en}, M_{de}) \quad (4)$$

In contrast,  $De$  should be selectively sensitive to malicious distortions  $\theta_{ma}$ :

$$\mathcal{L}_{De2} = L_2(0, De(\theta, RFNP(\theta_{ma}, I_{en}))) = L_2(0, M_{de}) \quad (5)$$

where the  $L_2$  distance between 0 and  $M_{de}$  is optimized. It should be emphasized that the binary message  $\{0, 1\}$  had been transferred to the zero-mean  $\{-\alpha, \alpha\}$ . Therefore, if the extracted message is greater than 0, it is recovered as binary 1, otherwise as binary 0. This loss term makes the extracted  $M_{de}$  close to random guesses.

For the discriminator, its parameters are updated by:


$$\mathcal{L}_{Ad1} = L_2(Ad(\theta, I_{co}) - \overline{Ad}(\theta, I_{en}), 1) + L_2(Ad(\theta, I_{en}) - \overline{Ad}(\theta, I_{co}), -1) \quad (6)$$

where

$$\overline{Ad}(I) = \frac{1}{B} \sum_{i=1}^B Ad(I^{i \times 3 \times H \times W}) \quad (7)$$

$$Ad(\theta, I_{en}) = Ad(\theta, En(I_{co}, M_{en})) \quad (8)$$

The relativistic average discriminator here estimates the probability that the cover image  $I_{co}$  is more pristine than the encoded image  $I_{en}$



**Figure 5: Subjective visual quality under typical distortions. From top to bottom are the cover image  $I_{co}$ , the encoded image  $I_{en}$ , the noised image  $I_{no}$ , the residual signal of  $\mathcal{N}(|I_{co} - I_{en}|)$ , and  $\mathcal{N}(|I_{no} - I_{en}|)$ , where  $\mathcal{N}(I) = (I - \min(I))/(\max(I) - \min(I))$ . Each column corresponds to one distortion. Image size:  $128 \times 128$ . For more malicious distortions, please see Appendix A.1.**

on average; please refer to the literature [35] for more details. Note that the discriminator also participates in updating the encoder  $En$ :

$$\mathcal{L}_{Ad2} = L_2(Ad(I_{co}) - \overline{Ad}(I_{en}), -1) + L_2(Ad(I_{en}) - \overline{Ad}(I_{co}), 1) \quad (9)$$

where

$$Ad(I_{en}) = Ad(En(\theta, I_{co}, M_{en})) \quad (10)$$

To summarize, the total loss for the main encoder-decoder can be formulated by:

$$\mathcal{L}_{Total} = \lambda_1 \mathcal{L}_{Ad2} + \lambda_2 \mathcal{L}_{En} + \lambda_3 \mathcal{L}_{Tr} + \lambda_4 \mathcal{L}_{De1} + \lambda_5 \mathcal{L}_{De2} \quad (11)$$

where  $\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5$  are the weights for respective loss terms.

## 4 EXPERIMENTS

### 4.1 Implementation Details

The experiments mainly conduct on CelebA-HQ [37, 39], where 24183/2993/2824 face images are used for training, valuation, and testing, referencing the official splits of CelebA [50]. Additionally, the testing set of CelebA, which includes 19962 face images, and the valuation set of COCO [46], which contains 5000 images of common objects, are employed to evaluate the generalizability. Unless stated otherwise, all the images are resized to two resolutions of  $128 \times 128$  and  $256 \times 256$ , due to the limit of computation resource. For the malicious distortions, SimSwap [8], GANimation [60], and StarGAN [10] are selected as they are typical Deepfakes, which belong to representative face swapping, expression reenactment, and attribute editing, respectively. In detail, the target face used for swapping is randomly selected from the valuation set of CelebA, the target expression used for reenactment is randomly selected from the driving images “eric\_andre”<sup>1</sup>, the attribute set we edit is {Male, Young, BlackHair, BlondHair, BrownHair}, and all of them are implemented based on their released pre-trained model for resulting convincing Deepfakes. Note that the face images are further background-removed by OpenFace [3] only when testing GANimation, to match the requirements of the input. Without loss of generality, the set of common distortions is {Identity, JpegTest, Resize,

<sup>1</sup>[https://github.com/vipermu/ganimation/tree/master/animations/eric\\_andre/attribute\\_images](https://github.com/vipermu/ganimation/tree/master/animations/eric_andre/attribute_images)

**Table 1: Objective visual quality of the encoded image  $I_{en}$ .**

Model	Image Size	Message Length	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
MBRS [34]	$128 \times 128$	30	33.0456	0.8106	0.0141
CIN [52]	$128 \times 128$	30	42.4135	0.9628	0.0006
PIMoG [16]	$128 \times 128$	30	37.7271	0.9470	0.0086
SepMark	$128 \times 128$	30	38.5112	0.9588	0.0028
FaceSigns [56]	$256 \times 256$	128	32.3319	0.9211	0.0260
SepMark	$256 \times 256$	128	38.5646	0.9328	0.0080

GaussianBlur, MedianBlur, Brightness, Contrast, Saturation, Hue, Dropout, SaltPepper, GaussianNoise}. The detailed parameters of the above distortions can be seen at the bottom of Fig. 5. Note that cropping is not included in the set since it is difficult to understand as a common distortion for the complete face.

Our SepMark is implemented by PyTorch [58] and executed on NVIDIA RTX 4090. Since we address numerous distortions rather than a specific one, SepMark uses the combined training fashion where the noise pool consists of all the common distortions and the malicious distortions, namely SimSwap, GANimation, and StarGAN (Male). The whole training lasted for 100 epochs with a batch size of 16, and as a rule of thumb, we adjusted the Adam optimizer [38] where  $lr = 0.0002$ ,  $\beta_1 = 0.5$  for stable adversarial training. Besides,  $\lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5$  in Eq. (11) are set to 0.1, 1, 10, 10, 10 in our training, respectively. Intuitively, the message range  $\alpha$  in our SepMark is set to 0.1, and we show that  $\alpha$  is essential for the trade-off between visual quality and robustness in the ablation study.

**Baselines.** As far as we know, this is the first work on deep separable watermarking, we therefore have to adopt robust watermarking methods such as MBRS [34], CIN [52], PIMoG [16], and fragile watermarking method FaceSigns [56] as our baselines. Two independent models, one trained on images of size  $128 \times 128$  and messages of length 30, the other trained on images of size  $256 \times 256$  and messages of length 128, for comparison with the pre-trained models of baselines to demonstrate the effectiveness of SepMark.

**Table 2: Robustness test for the common distorted image  $I_{no}$ , where “Tracer” means the BER between  $M_{tr}$  and  $M_{en}$ , “Detector” means the BER between  $M_{de}$  and  $M_{en}$ , and “Both” means the BER between  $M_{tr}$  and  $M_{de}$ .**

Distortion	128 × 128						256 × 256			
	MBRS [34]	CIN [52]	PIMoG [16]	SepMark			FaceSigns [56]	SepMark		
				Tracer	Detector	Both		Tracer	Detector	Both
Identity	0.0000%	0.0000%	0.0366%	0.0000%	0.0000%	0.0000%	0.0136%	0.0000%	0.0000%	0.0000%
JpegTest	0.2597%	2.7514%	19.5562%	0.2136%	0.2172%	0.2656%	0.8258%	0.0075%	0.0069%	0.0133%
Resize	0.0000%	0.0000%	0.0767%	0.0059%	0.0212%	0.0201%	1.0726%	0.0000%	0.0000%	0.0000%
GaussianBlur	0.0000%	22.7786%	0.1169%	0.0024%	0.0035%	0.0012%	0.1671%	0.0000%	0.0274%	0.0274%
MedianBlur	0.0000%	0.0307%	0.0992%	0.0012%	0.0012%	0.0000%	0.0977%	0.0000%	0.0000%	0.0000%
Brightness	0.0000%	0.0000%	1.3443%	0.0059%	0.0106%	0.0142%	10.8196%	0.0017%	0.0030%	0.0047%
Contrast	0.0000%	0.0000%	0.8121%	0.0012%	0.0024%	0.0012%	0.0334%	0.0000%	0.0000%	0.0000%
Saturation	0.0000%	0.0000%	0.0803%	0.0000%	0.0000%	0.0000%	0.7113%	0.0000%	0.0000%	0.0000%
Hue	0.0000%	0.0000%	0.1523%	0.0000%	0.0012%	0.0012%	8.3780%	0.0000%	0.0000%	0.0000%
Dropout	0.0000%	0.0000%	0.4828%	0.0000%	0.0000%	0.0000%	17.5615%	0.0058%	0.0000%	0.0058%
SaltPepper	0.0000%	0.0378%	2.3667%	0.0413%	0.0106%	0.0472%	12.3238%	0.0008%	0.0003%	0.0011%
GaussianNoise	0.0000%	0.0000%	12.7396%	0.7460%	0.8735%	0.8994%	7.0697%	0.0578%	0.0622%	0.0930%
Average	0.0216%	2.1332%	3.1553%	0.0848%	0.0951%	0.1042%	4.9228%	0.0061%	0.0083%	0.0121%

**Table 3: Robustness test for the maliciously distorted image  $I_{no}$ .**

Distortion	128 × 128						256 × 256			
	MBRS [34]	CIN [52]	PIMoG [16]	SepMark			FaceSigns [56]	SepMark		
				Tracer	Detector	Both		Tracer	Detector	Both
SimSwap	19.3744%	48.5068%	8.6745%	13.8255%	50.8829%	50.8558%	49.9463%	7.9068%	45.9117%	45.6713%
GANimation	0.0000%	0.0000%	0.4802%	0.0000%	36.7938%	36.7938%	45.4172%	0.0020%	43.8524%	43.8527%
StarGAN (Male)	18.3133%	67.2568%	9.2044%	0.1157%	52.6003%	52.6027%	50.2617%	0.0033%	45.4624%	45.4641%
StarGAN (Young)	17.0562%	69.0805%	8.7465%	0.1074%	52.3678%	52.3737%	50.3649%	0.0030%	45.5319%	45.5333%
StarGAN (BlackHair)	19.2233%	58.7913%	10.5312%	0.1416%	49.2434%	49.2599%	50.2576%	0.0019%	45.1299%	45.1296%
StarGAN (BlondHair)	18.4478%	72.9733%	10.3506%	0.1712%	48.6084%	48.5930%	50.9829%	0.0022%	44.6953%	44.6953%
StarGAN (BrownHair)	17.6381%	79.5857%	9.0675%	0.0980%	52.3820%	52.3855%	50.7884%	0.0017%	45.9950%	45.9939%
Average	15.7219%	56.5992%	8.1507%	2.0656%	48.9827%	48.9806%	49.7170%	1.1316%	45.2255%	45.1915%

**Evaluation Metrics.** For the evaluation of objective visual quality, we report the average PSNR, SSIM, and LPIPS [84] of the encoded images throughout the testing set. For the robustness test, the average bit error ratio (BER) is utilized as the default evaluation metric. In our SepMark, the BER of semi-robust Detector under malicious distortions should approach 50%; in other cases, the BER of both Tracer and Detector should approach 0%. Suppose that the embedded message  $M_E$  and the extracted message  $M_D$ , formally,

$$BER(M_E, M_D) = \frac{1}{B} \times \frac{1}{L} \times \sum_{i=1}^B \sum_{j=1}^L (\mathcal{B}(M_E^{i \times j}) - \mathcal{B}(M_D^{i \times j})) \times 100\% \quad (12)$$

where

$$\mathcal{B}(M) = \begin{cases} 1 & M > 0, \\ 0 & M \leq 0. \end{cases} \quad (13)$$

Note that we have strictly adjusted  $\mathcal{B}(\cdot)$  for the baselines to the same as their original setting, avoiding affecting their performances.

## 4.2 Visual Quality

We evaluate the visual quality of encoded images from the objective and subjective perspectives. As we can see in Table 1, SepMark has slightly better performance than PIMoG and is superior to MBRS and FaceSigns by a large margin, in terms of PSNR, SSIM, and LPIPS. Moreover, no clear artifacts can be observed by the naked eye from the encoded images, as visualized in Fig. 5. We speculate that the

satisfied performance mainly derives from our stable adversarial training. Although CIN yields impressive objective visual quality, its invertible neural network is not flexible to be compatible with our architecture of deep separable watermarking.

## 4.3 Robustness Test

Primarily, the robust Tracer should have a low BER under arbitrary distortions to trace the trusted source, while the semi-robust Detector should have a low BER under common distortions but a high BER under malicious ones to distinguish the type of the distortion.

Under the common distortions presented in Table 2, the BER of Tracer approaches 0% on average, reaching nearly optimal robustness compared to the baselines in addition to competitive MBRS. Moreover, the BER of Detector also gains reduction compared to CIN, PIMoG, and FaceSigns. By observation from the “Both” column, it can be drawn that the separable decoders extract nearly the same message when encountering common distortions. We do not intend to demonstrate that SepMark is definitely superior to baselines in terms of robustness, but rather that its sufficient robustness against common distortions.

Under the malicious distortions presented in Table 3, the Tracer also has ideal robustness compared to the baselines, and its BER reaches the lowest on average. As expected, the BER of Detector approaches 50% on average, close to random guesses. Our Detector successfully holds selective fragility to malicious distortions, while

**Table 4: Ablation experiments. Range of message  $M_{en}$  where  $\alpha \in \{0.15, 0.05\}$ . Image size:  $128 \times 128$ . Gain: “↓” means better.**

Distortion	$\alpha = 0.15$		$\alpha = 0.05$	
	PSNR = 38.1806		PSNR = 40.8058	
	Tracer	Detector	Tracer	Detector
Identity	0.0000%	0.0000%	0.0000%	0.0000%
JpegTest	0.1062%	0.1204%	0.6941%	0.7495%
Resize	0.0012%	0.0000%	0.0614%	0.0696%
GaussianBlur	0.0000%	0.0000%	0.1464%	0.5076%
MedianBlur	0.0000%	0.0000%	0.0141%	0.0106%
Brightness	0.0012%	0.0012%	0.0035%	0.0083%
Contrast	0.0012%	0.0000%	0.0224%	0.0224%
Saturation	0.0000%	0.0000%	0.0000%	0.0000%
Hue	0.0000%	0.0000%	0.0000%	0.0000%
Dropout	0.0012%	0.0012%	0.0024%	0.0047%
SaltPepper	0.0059%	0.0083%	0.4308%	0.3778%
GaussianNoise	0.4273%	0.4403%	2.0515%	2.2238%
SimSwap	15.8015%	54.1714%	20.0802%	48.3239%
GANimation	0.0037%	53.9691%	0.0012%	29.5494%
StarGAN (Male)	0.0331%	58.0914%	0.3175%	45.4544%
StarGAN (Young)	0.0378%	57.7679%	0.3671%	45.2325%
StarGAN (BlackHair)	0.0425%	60.0248%	0.4355%	45.3978%
StarGAN (BlondHair)	0.0354%	56.9795%	0.4202%	45.3872%
StarGAN (BrownHair)	0.0189%	59.2694%	0.2927%	45.0342%
Average Common	0.0454%	0.0476%	0.2856%	0.3312%
Average Malicious	↓ <b>0.0394%</b>	↓ <b>0.0475%</b>	↑ <b>0.2008%</b>	↑ <b>0.2961%</b>
	2.2818%	57.1819%	3.1306%	43.4828%
	↑ <b>0.2162%</b>	—	↑ <b>1.0650%</b>	—

the overly robust MBRS, PIMoG, and Tracer are not. It can also be observed that the separable decoders extract mismatched messages when encountering malicious distortions.

Intriguingly, the above even holds true for the pristine image that is without the embedded message, and we show in Table 6 in the Appendix that the two extracted messages are mismatched under malicious distortions, while they are matched to some extent under common distortions. Also in Table 7, we investigate the observation that the model trained on images of size  $128 \times 128$  can work on input images of arbitrary size. Please refer to Table 8 for the cross-dataset setting on CelebA; it has minor fluctuation in the BER of each distortion compared to that on CelebA-HQ, as well as in the PSNR of encoded images, thereby validating the excellent generalizability. We also show in Table 9 that the model trained on face images can surprisingly work on the natural images in COCO.

#### 4.4 Ablation Study

For the ablation in an effective setup, we first examine the range of embedded message, i.e., the hyper-parameter  $\alpha$ . Accordingly, as depicted in Table 4, a smaller  $\alpha$  can achieve higher visual quality but with reduced robustness, while a larger  $\alpha$  can achieve stronger robustness but with decreased visual quality. Recall that  $\alpha$  bridges the trade-off between visual quality and robustness. By the way, the bold values in the table mean the performance gains compared to the original counterparts (see Tables 2 and 3). The omitted term indicates the BER comparison under malicious distortion, as they are close to 50%, but higher is not necessarily better.

It is worth mentioning that we also try the two-stage learning, where one decoder is jointly trained with the left components, followed by training with only the other decoder for an additional 100

**Table 5: Ablation experiments. One decoder is trained followed by the other. Image size:  $128 \times 128$ .**

Distortion	Tracer→Detector		Detector→Tracer	
	PSNR = 39.3359		PSNR = 42.6461	
	Tracer	Detector	Tracer	Detector
Identity	0.0000%	0.0000%	0.0012%	0.0000%
JpegTest	0.6645%	0.5430%	2.0267%	2.4906%
Resize	0.0295%	0.0212%	0.1039%	0.0720%
GaussianBlur	0.0012%	0.0000%	0.0260%	0.0189%
MedianBlur	0.0012%	0.0012%	0.0059%	0.0201%
Brightness	0.0000%	0.0012%	0.0035%	0.0047%
Contrast	0.0035%	0.0035%	0.0106%	0.0071%
Saturation	0.0000%	0.0000%	0.0000%	0.0000%
Hue	0.0000%	0.0000%	0.0000%	0.0000%
Dropout	0.0035%	0.0059%	0.0047%	0.0024%
SaltPepper	0.1333%	0.0024%	0.0342%	0.2998%
GaussianNoise	1.5557%	1.5321%	3.5057%	4.1218%
SimSwap	17.9438%	52.1966%	29.1678%	48.5045%
GANimation	0.0062%	49.2074%	0.0198%	47.3840%
StarGAN (Male)	0.2491%	54.3366%	0.4001%	43.1256%
StarGAN (Young)	0.2656%	53.9731%	0.4603%	43.0961%
StarGAN (BlackHair)	0.3376%	53.2153%	0.5453%	43.0146%
StarGAN (BlondHair)	0.4084%	54.2292%	0.5382%	42.6039%
StarGAN (BrownHair)	0.2077%	53.5694%	0.4344%	42.9792%
Average Common	0.1994%	0.1759%	0.4769%	0.5865%
Average Malicious	↑ <b>0.1146%</b>	↑ <b>0.0808%</b>	↑ <b>0.3921%</b>	↑ <b>0.4914%</b>
	2.7741%	52.9611%	4.5094%	44.3868%
	↑ <b>0.7085%</b>	—	↑ <b>2.4438%</b>	—

epochs. From Table 5 we can see that training the robust decoder first and then training the semi-robust decoder, or conversely, is inferior to the end-to-end training of SepMark in terms of robustness (see Tables 2 and 3). One hypothesis is that this is likely due to the unbalanced trade-off with visual quality. Furthermore, we conduct the toy example where MBRS first embeds a robust watermark into the image and then uses FaceSigns to embed a second fragile watermark and vice versa, causing the failure of robust decoding and the deteriorated visual quality, as seen in Table 10. All of these again demonstrate the effectiveness of the proposed SepMark.

## 5 CONCLUSION

In this paper, we propose the first end-to-end learning architecture for deep separable watermarking, dubbed SepMark, and naturally apply it to proactive Deepfake forensics, realizing unified source tracing and Deepfake detection. Thanks to the separable decoders that extract the embedded watermark with different levels of robustness, the robust decoder resists against numerous distortions while the semi-robust one is selectively sensitive to malicious distortions. Experimental results show that SepMark has made an encouraging start on typical Deepfakes, to provide preventive measures and more reliable evidence for trustworthy faces. Further work involves the more general modeling of Deepfake distortions that can generalize the watermarking to unseen Deepfakes.

## ACKNOWLEDGMENTS

This work is supported by National Natural Science Foundation of China (Grant Nos. U22A2030, 61972142), National Key R&D Program of China (Grant No. 2022YFB3103500).

## REFERENCES

- [1] Shivangi Aneja, Lev Markhasin, and Matthias Nießner. 2022. TAFIM: Targeted Adversarial Attacks against Facial Image Manipulations. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XIV*. Springer, 58–75.
- [2] Vishal Asnani, Xi Yin, Tal Hassner, Sijia Liu, and Xiaoming Liu. 2022. Proactive image manipulation detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15386–15395.
- [3] Tadas Baltrusaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. 2018. Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*. IEEE, 59–66.
- [4] Jingyi Cao, Bo Liu, Yunqian Wen, Rong Xie, and Li Song. 2021. Personalized and invertible face de-identification by disentangled identity information manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3334–3342.
- [5] Nicholas Carlini and Hany Farid. 2020. Evading deepfake-image detectors with white-and black-box attacks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*. 658–659.
- [6] Qinwei Chang, Leichao Huang, Shaoteng Liu, Hualuo Liu, Tianshu Yang, and Yexin Wang. 2022. Blind Robust Video Watermarking Based on Adaptive Region Selection and Channel Reference. In *Proceedings of the 30th ACM International Conference on Multimedia*. 2344–2350.
- [7] Liang Chen, Yong Zhang, Yibing Song, Lingqiao Liu, and Jue Wang. 2022. Self-supervised learning of adversarial example: Towards good generalizations for deepfake detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 18710–18719.
- [8] Renwang Chen, Xuanhong Chen, Bingbing Ni, and Yanhao Ge. 2020. Simswap: An efficient framework for high fidelity face swapping. In *Proceedings of the 28th ACM International Conference on Multimedia*. 2003–2011.
- [9] Zhikai Chen, Lingxi Xie, Shanmin Pang, Yong He, and Bo Zhang. 2021. Magdr: Mask-guided detection and reconstruction for defending deepfakes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9014–9023.
- [10] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8789–8797.
- [11] Ingemar Cox, Matthew Miller, Jeffrey Bloom, Jessica Fridrich, and Ton Kalker. 2007. *Digital watermarking and steganography*. Morgan kaufmann.
- [12] Feng Ding, Guopu Zhu, Yingcan Li, Xinpeng Zhang, Pradeep K Atrey, and Siwei Lyu. 2021. Anti-forensics for face swapping videos via adversarial training. *IEEE Transactions on Multimedia* 24 (2021), 3429–3441.
- [13] Brian Dolhansky, Joanna Bittton, Ben Pflaum, Jikuo Lu, Russ Howes, Menglin Wang, and Cristian Canton Ferrer. 2020. The deepfake detection challenge (dfdc) dataset. *arXiv preprint arXiv:2006.07397* (2020).
- [14] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Ting Zhang, Weiming Zhang, Nenghai Yu, Dong Chen, Fang Wen, and Baining Guo. 2022. Protecting celebrities from deepfake with identity consistency transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9468–9478.
- [15] Ricard Durall, Margret Keuper, and Janis Keuper. 2020. Watch your up-convolution: Cnn based generative deep neural networks are failing to reproduce spectral distributions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 7890–7899.
- [16] Han Fang, Zhaoyang Jia, Zehua Ma, Ee-Chien Chang, and Weiming Zhang. 2022. PIMoG: An Effective Screen-shooting Noise-Layer Simulation for Deep-Learning-Based Watermarking Network. In *Proceedings of the 30th ACM International Conference on Multimedia*. 2267–2275.
- [17] Han Fang, Yupeng Qiu, Kejiang Chen, Jiye Zhang, Weiming Zhang, and Ee-Chien Chang. 2023. Flow-based robust watermarking with invertible noise layer for black-box distortions. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37.
- [18] Jiayun Fu, Bin B Zhu, Haidong Zhang, Yayi Zou, Song Ge, Weiwei Cui, Yun Wang, Dongmei Zhang, Xiaojing Ma, and Hai Jin. 2022. ChartStamp: Robust Chart Embedding for Real-World Applications. In *Proceedings of the 30th ACM International Conference on Multimedia*. 2786–2795.
- [19] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- [20] Ziweng He, Wei Wang, Weinan Guan, Jing Dong, and Tieniu Tan. 2022. Defeating DeepFakes via Adversarial Visual Reconstruction. In *Proceedings of the 30th ACM International Conference on Multimedia*. 2464–2472.
- [21] Kai Hong and Xiaoyu Du. 2021. Self-Supervised Deepfake Detection by Discovering Artifact Discrepancies. (2021).
- [22] Yang Hou, Qing Guo, Yihao Huang, Xiaofei Xie, Lei Ma, and Jianjun Zhao. 2023. Evading DeepFake Detectors via Adversarial Statistical Consistency. *arXiv preprint arXiv:2304.11670* (2023).
- [23] Juan Hu, Xin Liao, Jinwen Liang, Wenbo Zhou, and Zheng Qin. 2022. Finfer: Frame inference-based deepfake detection for high-visual-quality videos. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 36. 951–959.
- [24] Juan Hu, Xin Liao, Wei Wang, and Zheng Qin. 2021. Detecting compressed deepfake videos in social networks using frame-temporality two-stream convolutional network. *IEEE Transactions on Circuits and Systems for Video Technology* 32, 3 (2021), 1089–1102.
- [25] Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7132–7141.
- [26] Hao Huang, Yongtao Wang, Zhaoyu Chen, Yuze Zhang, Yuheng Li, Zhi Tang, Wei Chu, Jingdong Chen, Weisi Lin, and Kai-Kuang Ma. 2022. Cmu-watermark: A cross-model universal adversarial watermark for combating deepfakes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 989–997.
- [27] Qidong Huang, Jie Zhang, Wenbo Zhou, Weiming Zhang, and Nenghai Yu. 2021. Initiative defense against facial manipulation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 1619–1627.
- [28] Yihao Huang, Felix Juefei-Xu, Qing Guo, Lei Ma, Xiaofei Xie, Weikai Miao, Yang Liu, and Geguang Pu. 2021. Dodging DeepFake detection via implicit spatial-domain notch filtering. *arXiv preprint arXiv:2009.09213* (2021).
- [29] Yihao Huang, Felix Juefei-Xu, Run Wang, Qing Guo, Lei Ma, Xiaofei Xie, Jianwen Li, Weikai Miao, Yang Liu, and Geguang Pu. 2020. Fakepolisher: Making deepfakes more detection-evasive by shallow reconstruction. In *Proceedings of the 28th ACM international conference on multimedia*. 1217–1226.
- [30] Shehzeen Hussain, Paarth Neekhara, Malhar Jere, Farinaz Koushanfar, and Julian McAuley. 2021. Adversarial deepfakes: Evaluating vulnerability of deepfake detectors to adversarial examples. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 3348–3357.
- [31] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [32] Jun Jia, Zhongpai Gao, Kang Chen, Menghan Hu, Xiongkuo Min, Guangtao Zhai, and Xiaokang Yang. 2020. RIHOOP: robust invisible hyperlinks in offline and online photographs. *IEEE Transactions on Cybernetics* 52, 7 (2020), 7094–7106.
- [33] Jun Jia, Zhongpai Gao, Dandan Zhu, Xiongkuo Min, Guangtao Zhai, and Xiaokang Yang. 2022. Learning Invisible Markers for Hidden Codes in Offline-to-online Photography. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2273–2282.
- [34] Zhaoyang Jia, Han Fang, and Weiming Zhang. 2021. Mbrs: Enhancing robustness of dnn-based watermarking by mini-batch of real and simulated jpeg compression. In *Proceedings of the 29th ACM international conference on multimedia*. 41–49.
- [35] Alexia Jolicoeur-Martineau. 2018. The relativistic discriminator: a key element missing from standard GAN. *arXiv preprint arXiv:1807.00734* (2018).
- [36] Felix Juefei-Xu, Run Wang, Yihao Huang, Qing Guo, Lei Ma, and Yang Liu. 2022. Countering malicious deepfakes: Survey, battleground, and horizon. *International Journal of Computer Vision* 130, 7 (2022), 1678–1734.
- [37] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196* (2017).
- [38] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [39] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. 2020. Maskgan: Towards diverse and interactive facial image manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5549–5558.
- [40] Dongze Li, Wei Wang, Kang Zhao, Jing Dong, and Tieniu Tan. 2023. RIDDLE: Reversible and Diversified De-identification with Latent Encryptor. *arXiv preprint arXiv:2303.05171* (2023).
- [41] Haodong Li, Bin Li, Shunquan Tan, and Jiwu Huang. 2020. Identification of deep network generated images using disparities in color components. *Signal Processing* 174 (2020), 107616.
- [42] Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Baining Guo. 2020. Face x-ray for more general face forgery detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 5001–5010.
- [43] Yuezun Li, Ming-Ching Chang, and Siwei Lyu. 2018. In ictu oculi: Exposing ai created fake videos by detecting eye blinking. In *2018 IEEE international workshop on information forensics and security (WIFS)*. IEEE, 1–7.
- [44] Yuezun Li and Siwei Lyu. 2018. Exposing deepfake videos by detecting face warping artifacts. *arXiv preprint arXiv:1811.00656* (2018).
- [45] Xin Liao, Yumei Wang, Tianyi Wang, Juan Hu, and Xiaoshuai Wu. 2023. FAMM: Facial Muscle Motions for Detecting Compressed Deepfake Videos over Social Networks. *IEEE Transactions on Circuits and Systems for Video Technology* (2023).
- [46] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V* 13. Springer, 740–755.
- [47] Yuzhen Lin, Han Chen, Emanuele Maiorana, Patrizio Campisi, and Bin Li. 2022. Source-ID-Tracker: Source Face Identity Protection in Face Swapping. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 01–06.

- [48] Yang Liu, Mengxi Guo, Jian Zhang, Yuesheng Zhu, and Xiaodong Xie. 2019. A novel two-stage separable deep learning framework for practical blind watermarking. In *Proceedings of the 27th ACM International conference on multimedia*. 1509–1517.
- [49] Yang Liu, Zhenyue Qin, Saeed Anwar, Pan Ji, Dongwoo Kim, Sabrina Caldwell, and Tom Gedeon. 2021. Invertible denoising network: A light solution for real noise removal. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 13365–13374.
- [50] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*. 3730–3738.
- [51] Xiyang Luo, Ruohan Zhan, Huiwen Chang, Feng Yang, and Peyman Milanfar. 2020. Distortion agnostic deep watermarking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 13548–13557.
- [52] Rui Ma, Mengxi Guo, Yi Hou, Fan Yang, Yuan Li, Huizhu Jia, and Xiaodong Xie. 2022. Towards Blind Watermarking: Combining Invertible and Non-invertible Mechanisms. In *Proceedings of the 30th ACM International Conference on Multimedia*. 1532–1542.
- [53] Vahid Mirjalili, Sebastian Raschka, and Arun Ross. 2020. PrivacyNet: Semi-adversarial networks for multi-attribute face privacy. *IEEE Transactions on Image Processing* 29 (2020), 9400–9412.
- [54] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, Omar Fawzi, and Pascal Frossard. 2017. Universal adversarial perturbations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1765–1773.
- [55] Paarth Neekhara, Brian Dolhansky, Joanna Bitton, and Cristian Canton Ferrer. 2021. Adversarial threats to deepfake detection: A practical perspective. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 923–932.
- [56] Paarth Neekhara, Shehzeen Hussain, Xinqiao Zhang, Ke Huang, Julian McAuley, and Farinaz Koushanfar. 2022. FaceSigns: semi-fragile neural watermarks for media authentication and countering deepfakes. *arXiv preprint arXiv:2204.01960* (2022).
- [57] Zhicheng Ni, Yun-Qing Shi, Nirwan Ansari, and Wei Su. 2006. Reversible data hiding. *IEEE Transactions on circuits and systems for video technology* 16, 3 (2006), 354–362.
- [58] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [59] Pengfei Pei, Xianfeng Zhao, Yun Cao, Jinchuan Li, and Xuyuan Lai. 2021. Vision Transformer Based Video Hashing Retrieval for Tracing the Source of Fake Videos. *arXiv preprint arXiv:2112.08117* (2021).
- [60] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. 2018. Ganimation: Anatomically-aware facial animation from a single image. In *Proceedings of the European conference on computer vision (ECCV)*. 818–833.
- [61] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [62] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
- [63] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10684–10695.
- [64] Andreas Rossler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. 2019. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE/CVF international conference on computer vision*. 1–11.
- [65] Nataniel Ruiz, Sarah Adel Bargal, and Stan Sclaroff. 2020. Disrupting deepfakes: Adversarial attacks against conditional image translation networks and facial manipulation systems. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*. Springer, 236–251.
- [66] Hong Shen and Bo Chen. 2012. From single watermark to dual watermark: a new approach for image watermarking. *Computers & Electrical Engineering* 38, 5 (2012), 1310–1324.
- [67] Kaede Shiohara and Toshihiko Yamasaki. 2022. Detecting deepfakes with self-blended images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 18720–18729.
- [68] Jingxiang Sun, Xuan Wang, Yong Zhang, Xiaoyu Li, Qi Zhang, Yebin Liu, and Jue Wang. 2022. Fenerf: Face editing in neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7672–7682.
- [69] Pu Sun, Yuezun Li, Honggang Qi, and Siwei Lyu. 2022. Faketracer: Exposing Deepfakes with Training Data Contamination. In *2022 IEEE International Conference on Image Processing (ICIP)*. IEEE, 1161–1165.
- [70] Matthew Tancik, Ben Mildenhall, and Ren Ng. 2020. Stegastamp: Invisible hyperlinks in physical photographs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2117–2126.
- [71] Run Wang, Ziheng Huang, Zhikai Chen, Li Liu, Jing Chen, and Lina Wang. 2022. Anti-Forgery: Towards a Stealthy and Robust DeepFake Disruption Attack via Adversarial Perceptual-aware Perturbations. *arXiv preprint arXiv:2206.00477* (2022).
- [72] Run Wang, Felix Juefei-Xu, Meng Luo, Yang Liu, and Lina Wang. 2021. Faketagger: Robust safeguards against deepfake dissemination via provenance tracking. In *Proceedings of the 29th ACM International Conference on Multimedia*. 3546–3555.
- [73] Xueyu Wang, Jiajun Huang, Siqi Ma, Surya Nepal, and Chang Xu. 2022. Deepfake disrupter: The detector of deepfake is my friend. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14920–14929.
- [74] Eric Wengrowski and Kristin Dana. 2019. Light field messaging with deep photographic steganography. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1515–1524.
- [75] Xiaoshuai Wu, Tong Qiao, Yanli Chen, Ming Xu, Ning Zheng, and Xiangyang Luo. 2022. Sign steganography revisited with robust domain selection. *Signal Processing* 196 (2022), 108522.
- [76] Zhiliang Xu, Zhibin Hong, Changxing Ding, Zhen Zhu, Junyu Han, Jingtuo Liu, and Errui Ding. 2022. Mobilefaceswap: A lightweight framework for video face swapping. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 2973–2981.
- [77] Zhiliang Xu, Hang Zhou, Zhibin Hong, Ziwei Liu, Jiaming Liu, Zhizhi Guo, Junyu Han, Jingtuo Liu, Errui Ding, and Jingdong Wang. 2022. Styleswap: Style-based generator empowers robust face swapping. In *European Conference on Computer Vision*. Springer, 661–677.
- [78] Xin Yang, Yuezun Li, and Siwei Lyu. 2019. Exposing deep fakes using inconsistent head poses. In *ICASSP 2019–2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 8261–8265.
- [79] Chin-Yuan Yeh, Hsi-Wen Chen, Hong-Han Shuai, De-Nian Yang, and Ming-Syan Chen. 2021. Attack as the best defense: Nullifying image-to-image translation gans via limit-aware adversarial attack. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 16188–16197.
- [80] Chin-Yuan Yeh, Hsi-Wen Chen, Shang-Lun Tsai, and Sheng-De Wang. 2020. Disrupting image-translation-based deepfake algorithms with adversarial attacks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision Workshops*. 53–62.
- [81] Ning Yu, Vladislav Skripniuk, Sahar Abdelnabi, and Mario Fritz. 2021. Artificial fingerprinting for generative models: Rooting deepfake attribution in training data. In *Proceedings of the IEEE/CVF International conference on computer vision*. 14448–14457.
- [82] Zhuowen Yuan, Zhengxin You, Sheng Li, Zhenxing Qian, Xinpeng Zhang, and Alex Kot. 2022. On Generating Identifiable Virtual Faces. In *Proceedings of the 30th ACM International Conference on Multimedia*. 1465–1473.
- [83] Chaoning Zhang, Adil Karjauv, Philipp Benz, and In So Kweon. 2021. Towards robust deep hiding under non-differentiable distortions for practical blind watermarking. In *Proceedings of the 29th ACM international conference on multimedia*. 5158–5166.
- [84] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 586–595.
- [85] Hanqing Zhao, Wenbo Zhou, Dongdong Chen, Weiming Zhang, and Nenghai Yu. 2022. Self-supervised transformer for deepfake detection. *arXiv preprint arXiv:2203.01265* (2022).
- [86] Tianchen Zhao, Xiang Xu, Mingze Xu, Hui Ding, Yuanjun Xiong, and Wei Xia. 2021. Learning self-consistency for deepfake detection. In *Proceedings of the IEEE/CVF international conference on computer vision*. 15023–15033.
- [87] Jiaran Zhou and Yuezun Li. 2022. Detection-by-Simulation: Exposing DeepFake via Simulating Forgery using Face Reconstruction. In *2022 IEEE 5th International Conference on Multimedia Information Processing and Retrieval (MIPR)*. IEEE, 210–215.
- [88] Peng Zhou, Xintong Han, Vlad I Morariu, and Larry S Davis. 2017. Two-stream neural networks for tampered face detection. In *2017 IEEE conference on computer vision and pattern recognition workshops (CVPRW)*. IEEE, 1831–1839.
- [89] Jiren Zhu, Russell Kaplan, Justin Johnson, and Li Fei-Fei. 2018. Hidden: Hiding data with deep networks. In *Proceedings of the European conference on computer vision (ECCV)*. 657–672.
- [90] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*. 2223–2232.
- [91] Bojia Zi, Minghao Chang, Jingjing Chen, Xingjun Ma, and Yu-Gang Jiang. 2020. Wildeepfake: A challenging real-world dataset for deepfake detection. In *Proceedings of the 28th ACM international conference on multimedia*. 2382–2390.

## A APPENDIX

### A.1 Additional Visual Examples

(a) Image size:  $128 \times 128$ (b) Image size:  $256 \times 256$ 

Figure 6: More malicious distortions regarding Deepfakes.

### A.2 Additional Experimental Results

Table 6: More investigations on the pristine image  $I_{Co}$ . “Both” means the BER between  $M_{Tr}$  and  $M_{de}$ .

Distortion	$128 \times 128$		$256 \times 256$	
	PSNR = +∞	Both	PSNR = +∞	Both
Identity	18.3534%		15.8034%	
JpegTest	14.3272%		16.7296%	
Resize	17.1341%		15.4153%	
GaussianBlur	18.0217%		17.9925%	
MedianBlur	17.6121%		17.0536%	
Brightness	18.7901%		16.7889%	
Contrast	18.6202%		16.8386%	
Saturation	18.3392%		15.9481%	
Hue	18.2649%		15.8679%	
Dropout	18.3534%		15.8034%	
SaltPepper	27.8199%		23.0745%	
GaussianNoise	16.0966%		18.7992%	
SimSwap	50.5406%		47.4623%	
GANimation	48.6370%		49.8244%	
StarGAN (Male)	50.4049%		49.0832%	
StarGAN (Young)	50.6209%		49.2378%	
StarGAN (BlackHair)	49.5609%		48.9125%	
StarGAN (BlondHair)	48.7028%		48.9916%	
StarGAN (BrownHair)	50.1771%		49.0907%	
Average Common	18.4777%		17.1763%	
Average Malicious	49.8063%		48.9432%	

Table 7: More investigations on  $128 \times 128$  trained model across arbitrary input size of image  $I_{Co}$ .

Distortion	$256 \times 256$		$512 \times 512$	
	PSNR = 35.4984	PSNR = 35.2925	PSNR = 35.2925	PSNR = 35.2925
	Tracer	Detector	Tracer	Detector
Identity	0.0000%	5.2467%	0.0000%	25.1287%
JpegTest	0.1735%	0.9679%	0.1062%	0.7177%
Resize	0.0508%	2.9592%	0.0165%	18.6532%
GaussianBlur	0.0106%	11.0127%	0.0035%	33.1704%
MedianBlur	0.0071%	9.0061%	0.0035%	30.9301%
Brightness	0.0153%	5.8735%	0.0165%	25.0118%
Contrast	0.0177%	6.2028%	0.0106%	25.3234%
Saturation	0.0000%	5.4958%	0.0000%	25.3966%
Hue	0.0000%	5.5902%	0.0000%	26.1036%
Dropout	0.0165%	7.6546%	0.0224%	26.1674%
SaltPepper	0.1558%	2.0491%	0.2325%	5.9832%
GaussianNoise	0.5323%	3.9341%	0.3895%	8.1161%
SimSwap	2.5484%	49.9115%	1.9677%	49.1407%
GANimation	0.0679%	46.8494%	0.1136%	44.9160%
StarGAN (Male)	0.2668%	50.2101%	0.0873%	49.8430%
StarGAN (Young)	0.2821%	50.0390%	0.0814%	49.4653%
StarGAN (BlackHair)	0.2986%	50.6232%	0.0755%	50.8263%
StarGAN (BlondHair)	0.3553%	49.5574%	0.0921%	48.2177%
StarGAN (BrownHair)	0.2644%	50.3057%	0.0767%	50.3836%
Average Common	0.0816%	5.4994%	0.0668%	20.8919%
Average Malicious	↓0.0032%	↑5.4043%	↓0.0180%	↑20.7968%
	0.5834%	49.6423%	0.3563%	48.9704%
	↓1.4822%	—	↓1.7093%	—

Table 8: Robustness test for the image  $I_{No}$  under common and malicious distortions on cross-dataset CelebA.

Distortion	$128 \times 128$		$256 \times 256$	
	PSNR = 37.8865	PSNR = 38.7051	Tracer	Detector
	Tracer	Detector	Tracer	Detector
Identity	0.0000%	0.0417%	0.0063%	0.0002%
JpegTest	0.2233%	0.2934%	0.0247%	0.0194%
Resize	0.0134%	0.0272%	0.0001%	0.0002%
GaussianBlur	0.0058%	0.0416%	0.0004%	0.0022%
MedianBlur	0.0037%	0.0730%	0.0003%	0.0002%
Brightness	0.0094%	0.0579%	0.0160%	0.0016%
Contrast	0.0068%	0.0698%	0.0034%	0.0002%
Saturation	0.0000%	0.0519%	0.0056%	0.0002%
Hue	0.0000%	0.0401%	0.0070%	0.0002%
Dropout	0.0020%	0.0625%	0.4647%	0.0004%
SaltPepper	0.0406%	0.0242%	0.0017%	0.0005%
GaussianNoise	0.7461%	0.8952%	0.1594%	0.1708%
SimSwap	12.2428%	50.7506%	10.2216%	34.2429%
GANimation	0.0007%	39.6265%	0.0001%	45.5291%
StarGAN (Male)	0.1209%	50.8483%	0.0019%	45.5910%
StarGAN (Young)	0.1358%	50.6294%	0.0025%	45.6882%
StarGAN (BlackHair)	0.1479%	48.0443%	0.0021%	46.1099%
StarGAN (BlondHair)	0.1840%	46.7939%	0.0020%	44.6107%
StarGAN (BrownHair)	0.1196%	50.8680%	0.0014%	45.9944%
Average Common	0.0876%	0.1399%	0.0575%	0.0163%
Average Malicious	↑0.0028%	↑0.0448%	↑0.0514%	↑0.0080%
	1.8502%	48.2230%	1.4617%	43.9666%
	↓0.2154%	—	↑0.3301%	—

Table 9: Robustness test for the common distorted image  $I_{No}$  on cross-dataset COCO.

Distortion	$128 \times 128$		$256 \times 256$	
	PSNR = 35.1500	PSNR = 35.5506	Tracer	Detector
	Tracer	Detector	Tracer	Detector
Identity	0.2100%	1.4440%	0.0420%	0.0475%
JpegTest	1.6027%	3.6627%	0.1691%	0.1439%
Resize	3.9873%	2.8900%	0.2661%	0.2502%
GaussianBlur	1.1447%	4.0467%	0.0698%	0.1478%
MedianBlur	0.7460%	3.4340%	0.1005%	0.0878%
Brightness	0.4560%	1.9213%	0.1745%	0.1500%
Contrast	0.3860%	1.7367%	0.0756%	0.0639%
Saturation	0.2333%	1.5240%	0.0438%	0.0567%
Hue	0.2293%	1.6453%	0.0428%	0.0495%
Dropout	0.8800%	2.8593%	0.2595%	0.1906%
SaltPepper	1.8013%	1.9353%	0.1944%	0.1355%
GaussianNoise	1.0473%	2.2393%	0.2186%	0.2402%
Average Common	1.0603%	2.4449%	0.1381%	0.1303%
	↑0.9755%	↑2.3498%	↑0.1320%	↑0.1220%

Table 10: Toy example. One watermark is embedded followed by the other. Image size:  $256 \times 256$ .

Distortion	MBRS→FaceSigns		FaceSigns→MBRS	
	PSNR = 32.1240	PSNR = 32.0076	MBRS	FaceSigns
	MBRS	FaceSigns	MBRS	FaceSigns
Identity	0.3378%	0.0133%	0.0000%	0.0340%
JpegTest	5.0135%	0.8648%	0.1546%	1.5589%
Resize	5.7166%	1.1027%	1.6975%	1.5783%
GaussianBlur	6.9225%	0.1328%	3.6559%	0.1676%
MedianBlur	5.1384%	0.1151%	2.3408%	1.5160%
Brightness	1.7552%	0.0764%	1.0235%	0.0943%
Contrast	1.7989%	0.0318%	0.9658%	0.0515%
Saturation	0.3329%	0.7660%	0.0000%	0.8219%
Hue	0.9683%	8.0228%	0.0296%	9.3728%
Dropout	4.5914%	17.6177%	38.8410%	5.9916%
SaltPepper	40.2750%	12.4776%	32.2704%	13.3133%
GaussianNoise	25.7147%	0.8963%	12.5065%	1.2598%
SimSwap	46.5361%	49.8661%	45.8514%	49.8559%
GANimation	2.7499%	45.2734%	0.2568%	45.5127%
StarGAN (Male)	11.8858%	50.1740%	7.7480%	50.2683%
StarGAN (Young)	11.7510%	50.3737%	7.8313%	50.4302%
StarGAN (BlackHair)	13.0742%	50.1859%	8.9936%	50.2606%
StarGAN (BlondHair)	10.3177%	50.9066%	6.8103%	51.1755%
StarGAN (BrownHair)	11.0078%	50.7384%	7.1055%	50.9381%
Average Common	8.2138%	3.5098%	7.7905%	2.9800%
Average Malicious	↑8.2077%	↑3.5015%	↑7.7844%	↑2.9717%
	15.3318%	49.6454%	12.0853%	49.7773%
	↑14.2002%	—	↑10.9537%	—

For each figure in the Appendix, from top to bottom are the cover image  $I_{co}$ , the encoded image  $I_{en}$ , the noised image  $I_{no}$ , the residual signal of  $\mathcal{N}(|I_{co} - I_{en}|)$ , and  $\mathcal{N}(|I_{no} - I_{en}|)$ .



Figure 7: SimSwap. Dataset: CelebA-HQ. Image size:  $256 \times 256$ .



Figure 8: GANimation: the face images are processed by OpenFace only when testing, to match the requirements of the input. Dataset: CelebA-HQ. Image size:  $128 \times 128$ .



Figure 9: StarGAN (Male). Dataset: CelebA-HQ. Image size:  $256 \times 256$ .



Figure 10: StarGAN (Young). Dataset: CelebA-HQ. Image size:  $256 \times 256$ .



Figure 11: StarGAN (BlackHair). Dataset: CelebA-HQ. Image size:  $128 \times 128$ .



Figure 12: StarGAN (BlondHair). Dataset: CelebA-HQ. Image size:  $128 \times 128$ .




Figure 13: StarGAN (BrownHair). Dataset: CelebA-HQ. Image size:  $128 \times 128$ .




Figure 14: JpegTest. Dataset: COCO. Image size:  $256 \times 256$ .




Figure 15: JpegTest. Dataset: COCO. Image size:  $128 \times 128$ .