

# DeAR: A Deep-Learning-Based Audio Re-recording Resilient Watermarking

Chang Liu<sup>1</sup>, Jie Zhang<sup>1,2</sup>, Han Fang<sup>\*3</sup>, Zehua Ma<sup>1</sup>,  
Weiming Zhang<sup>\*1</sup>, Nenghai Yu<sup>1</sup>

<sup>1</sup>University of Science and Technology of China,

<sup>2</sup>University of Waterloo,

<sup>3</sup>National University of Singapore

hichangliu@mail.ustc.edu.cn, jiezhangsp@gmail.com, fanghan@nus.edu.sg,

mzh045@mail.ustc.edu.cn, zhangwm@ustc.edu.cn, ynh@ustc.edu.cn

## Abstract

Audio watermarking is widely used for leaking source tracing. The robustness of the watermark determines the traceability of the algorithm. With the development of digital technology, audio re-recording (AR) has become an efficient and covert means to steal secrets. AR process could drastically destroy the watermark signal while preserving the original information. This puts forward a new requirement for audio watermarking at this stage, that is, to be robust to AR distortions. Unfortunately, none of the existing algorithms can effectively resist AR attacks due to the complexity of the AR process. To address this limitation, this paper proposes DeAR, a deep-learning-based audio re-recording resistant watermarking. Inspired by DNN-based image watermarking, we pioneer a deep learning framework for audio carriers, based on which the watermark signal can be effectively embedded and extracted. Meanwhile, in order to resist the AR attack, we delicately analyze the distortions that occurred in the AR process and design the corresponding distortion layer to cooperate with the proposed watermarking framework. Extensive experiments show that the proposed algorithm can resist not only common electronic channel distortions but also AR distortions. Under the premise of high-quality embedding (SNR=25.86 dB), in the case of a common re-recording distance (20 cm), the algorithm can effectively achieve an average bit recovery accuracy of 98.55%.

## Introduction

As an efficient method to trace the source of leakage, digital watermarking technology has been widely studied. There have been many excellent works in image (Fang et al. 2018), audio (Erfani, Pichevar, and Rouat 2016), and video (Asikuzzaman and Pickering 2017) watermarking. The two most important properties that audio watermarking should satisfy are fidelity and robustness. Fidelity ensures the normal use of the watermarked audio. Robustness guarantees that even if the audio is distorted (MPEG encoding, noise addition, audio re-recording, etc.), the embedded watermark can still be extracted losslessly (Wu, Su, and Kuo 1999).

Most traditional audio watermarking methods are concerned with the robustness of digital distortions in the electronic channel because most audio copying occurs in the dig-




Figure 1: The re-recording operation preserves the content information while destroying the watermark information in the audio to hide the source of the leakage.

ital channel. However, with the miniaturization of recording devices, audio re-recording (AR) has become a more convenient and converted way to copy audios. Meanwhile, AR could effectively preserve the audio content and significantly damage the embedded watermark signal; attackers can easily and stealthily steal the audio content without any watermark evidence, as shown in Fig. 1. Therefore, ensuring robustness to AR becomes the urgent property of audio watermarking at this stage.

Currently, the research field of audio watermarking is still dominated by traditional mathematical algorithms, which try to search for an invariant feature before and after distortion to conduct watermark embedding. Most used features are in the transform domain, such as discrete cosine transform (DCT), discrete wavelet transform (DWT), and fast Fourier transform (FFT) (Su et al. 2018; Xiang et al. 2017; Natgunanathan et al. 2012; Liu, Huang, and Huang 2018). However, due to the complexity of the AR process itself, quantitatively and delicately analyzing the distortions and finding the robust feature in this process is a nontrivial task. Therefore, none of the existing algorithms can resist AR distortion very well. This drives the demand to seek a learnable feature that is adapted to AR distortions.

In recent years, deep-learning-based watermarking has achieved tremendous success in the field of image and video watermarking (Zhu et al. 2018; Tancik, Mildenhall, and Ng 2020; Zhang et al. 2019; Luo et al. 2021). The main backbone of such a framework is an end-to-end auto-encoder-like architecture that contains an encoder, a distortion layer, and a decoder. The central part to ensure robustness is the distortion layer which tries to generate the distorted image for

\*Corresponding authors.

training (Zhu et al. 2018), and it is essential to make the distortion layer differentiable. However, the AR process is a complex and non-differential process. Although some algorithms currently focus on deep learning-based information hiding studies in the audio field (Kreuk et al. 2020; Jiang et al. 2020), they cannot handle the robustness requirement against distortions, including AR distortion.

To address the robustness of AR distortion, in this paper, we propose DeAR, a deep-learning-based audio re-recording resistant watermarking method, which can effectively resist AR distortion at different distances in the real world. Specifically, we adopt the classic deep-learning-based watermarking framework, where an encoder and a decoder are jointly trained for watermark embedding and extraction. To achieve robustness against re-recording, we first analyze the re-recording process from the effects of sound propagation in the air and the processing of microphones and speakers. According to the analysis, we delicately model the AR distortion with several differential operations (environment reverberation, band-pass filtering, and Gaussian noise) and serve such operations as the distortion layer to cooperate with the proposed framework. In addition, motivated by traditional audio watermarking, we change the input of the time domain signal into the low-frequency coefficients of the audio for better robustness. Specifically, we introduce the differential time-frequency transform and its corresponding reverse transform into the end-to-end training process, which can automatically search for the optimal embedding frequency rather than be determined by a pre-defined rule such as Singular Value Decomposition (SVD). Experimental results verify that the proposed method can achieve satisfying robustness against audio re-recording at different distances and is also resilient to other common distortions.

The primary contributions of our work are concluded as follows:

- We are the first to propose a deep-learning-based audio watermarking against audio re-recording (AR), DeAR, which can achieve watermark embedding and extraction in an end-to-end manner rather than based on handcrafted rules. According to the audio cover, we flexibly adapt audio frequency transform and one-dimensional convolution operations for better implementation.
- To achieve robustness against AR, we first analyze the distortions induced by AR and model it as a distortion pipeline composed of environment reverberation, band-pass filtering, and Gaussian noise. The whole pipeline is further absorbed into the training of DeAR.
- Extensive experiments demonstrate that the proposed method can achieve robustness against audio re-recording and common electronic channel distortions while guaranteeing the requirement of fidelity. In addition, some ablation studies further verify our design and prove the flexibility of DeAR.

## Related Work

### Traditional Audio Watermarking

Traditional audio watermarking (AW) mainly embeds watermark information in the time domain and the transform do-

main. Time-domain based AW (Cvejic and Seppanen 2004; Natgunanathan and Xiang 2010) is simple and efficient but not robust enough. In comparison, transform-domain based AW (Bansal et al. 2015; Su et al. 2018) can achieve better robustness. For example, Su *et al.* (Su et al. 2018) conducted DWT transform and DCT transform on the audio and appended the watermark information in the low-frequency band to achieve robustness. However, its watermark extraction is non-blind, requiring the access to the raw audio. Besides, they did not consider robustness against re-recording process. Based on the above method, Liu *et al.* (Liu, Huang, and Huang 2018) first considered the re-recording robustness in traditional framework. Specifically, they simplify the practical process as an idealized noise appending process, which only influences the amplitude of the target audio. However, Liu’s method does not perform well in a practical scenario. We explained that re-recording is composed of not only the additive noise but also other different distortions such as the convolutional noise (namely, environment reverberation) and noise induced by microphones and speakers, which is demonstrated in many existing works (Peddinti et al. 2015; Yakura and Sakuma 2018).

### Real-world Air Channel Distortion

In this paper, we regard the re-recording process as an air channel distortion. Eliminating the side effect of the air channel distortion has always been an inspiring topic in automatic speech recognition (ASR) tasks. To address this issue, massive perturbed data are collected for training noise adaptive acoustic models, which is not applicable for the audio watermarking task. Except that, the most similar research field is robust adversarial audio of speech-to-text models, which guarantees robustness against air channel distortion. In detail, the corresponding studies (Qin et al. 2019; Yakura and Sakuma 2018) modeled these distortions by expectation over transformation (EOT) operations, which is incorporated in the optimization procedure of the adversarial audio. Motivated by this, we first analyze the re-recording process from the effects of sound propagation in the air and the processing of microphones and speakers. Then, we model the re-recording as a distortion pipeline, which is composed of environment reverberation, band-pass filtering, and Gaussian noise. This distortion pipeline can easily cooperate with the deep-learning-based audio watermarking framework.

### Deep-learning-based Image Watermarking

In recent years, some deep-learning (DL)-based image watermarking schemes have been proposed. For example, Zhu *et al.* (Zhu et al. 2018) proposed the first end-to-end image watermarking framework. As shown in Fig. 2, the popular framework contains three parts: an encoder responsible for watermark embedding, a differential distortion layer to simulate the transmission process, and a decoder structure responsible for watermark extraction. The main challenge is making the non-differential distortion differential, which is necessary for end-to-end training.

In (Zhu et al. 2018), Zhu *et al.* leveraged a differential approximations for non-differential JPEG compression, which




Figure 2: Illustration of the deep-learning-based image watermarking framework.

was further improved by subsequent works (Ahmadi et al. 2020; Ying et al. 2021). Based on (Zhu et al. 2018), Tancik *et al.* (Tancik, Mildenhall, and Ng 2020) further considered the distortion caused by printing and camera-shooting, Jia (Jia et al. 2020) realized resilience to 3D rendering, and Zhang *et al.* (Zhang et al. 2020) even achieved robustness against model extraction.

Based on DL-based image watermarking, there have been a few attempts at DL-based audio steganography. For example, Jiang *et al.* (Jiang et al. 2020) converted raw audio to 2D mel-spectrum images by short-time Fourier transform (STFT) to satisfy the above image watermarking framework. However, STFT and its inverse transform can intrinsically induce the loss of embedded information. In addition, this method lacks the consideration of any robustness during transmission, including audio re-recording (AR).

## Method

As shown in Fig. 3, DeAR is mainly composed of three parts, namely, an encoder for watermark embedding, a decoder for watermark extraction, and a distortion layer to enhance robustness against audio re-recording. All of the above parts are jointly trained, and we introduce each part in detail.

### Watermark Embedding

We use  $A$  to represent a single-channel raw audio with length  $N$ . Rather than directly feeding the raw audio  $A$  into the encoder, we first transfer it into the frequency domain by a differential DWT and obtain the corresponding approximate coefficients  $A_{ac}$  and detail coefficients  $A_{dc}$ , *i.e.*,

$$A_{ac}, A_{dc} = \text{DWT}(A), \quad (1)$$

where the length of  $A_{ac}$  and  $A_{dc}$  is half the original audio signal, namely,  $N/2$ . Motivated by traditional audio watermarking, we propose embedding the watermark into the low frequency of the raw audio, namely, leveraging  $A_{ac}$  as the audio cover. Meanwhile,  $A_{dc}$  is deposited for subsequent audio reconstruction. The goal of the encoder is to embed the watermark information  $W$  into  $A_{ac}$ . As shown in Fig. 3, we want the encoder **En** to generate a stealthy residual  $\mathcal{R}$  and further stamp it onto the original  $A_{ac}$  to generate the watermarked approximate coefficients  $A_{ac}^W$ , *i.e.*,

$$A_{ac}^W = \mathbf{En}(A_{ac}, W) * S + A_{ac}, \quad (2)$$

where  $S$  is the strength factor and set as 1 by default. To satisfy the fidelity requirement, we constrain the watermarked approximate coefficients  $A_{ac}^W$  acoustically consistent with

the original one  $A_{ac}$ . To achieve this, we introduce a basic loss  $\mathcal{L}_e$  for the encoder’s training, that is, we adopt the widely-used mean square error MSE as  $\mathcal{L}_e$ , *i.e.*,

$$\mathcal{L}_e = \text{MSE}(A_{ac}, A_{ac}^W) = \frac{2}{N} \sum_{i=1}^{N/2} (A_{ac}(i) - A_{ac}^W(i))^2. \quad (3)$$

To further improve the fidelity and minimize the domain gap between  $A_{ac}^W$  and  $A_{ac}$ , we introduce an extra discriminator **D** for adversarial training with the encoder **En**, and adversarial loss  $\mathcal{L}_d$  will let **En** embed watermarks better so that **D** cannot distinguish  $A_{ac}^W$  from the watermark-free  $A_{ac}$ , *i.e.*,

$$\mathcal{L}_d = \log(1 - \mathbf{D}(A_{ac}^W)). \quad (4)$$

Meanwhile, for **D**,  $\mathcal{L}_D = \log(1 - \mathbf{D}(A_{ac})) + \log(\mathbf{D}(A_{ac}^W))$ .

### Watermark Extraction

Given the watermarked approximate coefficients  $A_{ac}^W$ , the decoder **De** needs to recover watermark  $W'$  as consistent as the original watermark  $W$ . To achieve this, we introduce the watermark loss  $\mathcal{L}_w$ , namely, the MSE between the original watermark  $W$  and the extracted watermark  $W'$ , *i.e.*,

$$\mathcal{L}_w = \text{MSE}(W, W'). \quad (5)$$

It should be emphasized that we adopt the binary watermark  $W \in \{-1, 1\}^L$  instead of  $\{0, 1\}^L$ , which is better for forensics. Meanwhile, this helps the MSE-based constraint work better.

### Audio Re-recording Modeling

To enhance the robustness against the audio re-recording (AR) process, we further insert a distortion layer between the encoder and the decoder. As mentioned above, it is essential to make the distortion layer differentiable, which can prevent gradient interruption during end-to-end learning. However, the AR process is a complex and non-differential process. To overcome this challenge, we learn from the studies on real-world air channel distortion (Qin et al. 2019; Yakura and Sakuma 2018), and introduce a differential audio re-recording operation **DAR**, which consists of three components: environment reverberation, band-pass filtering, and Gaussian noise. Because **DAR** is a processing pipeline operated on the time domain, it cannot be directly applied to  $A_{ac}^W$ . Therefore, after generating  $A_{ac}^W$ , we adopt inverse DWT (IDWT) to transform the watermarked approximate coefficients  $A_{ac}^W$  back to the watermarked audio  $A^W$ , with the corresponding deposited  $A_{dc}$ , *i.e.*,

$$A^W = \text{IDWT}(A_{ac}^W, A_{dc}). \quad (6)$$

**Environment Reverberation.** The impulse Response (IR) is the environment’s reaction when presented with a brief input signal. It describes the acoustic characteristics of an environment, particularly of interest being the behavior of reverberation in the space. IR can reproduce the reverberation in the captured environment by convolution. From different microphones, room environments, and speakers, we can collect diverse base IR  $p$  to form a set  $P$ . Given a target

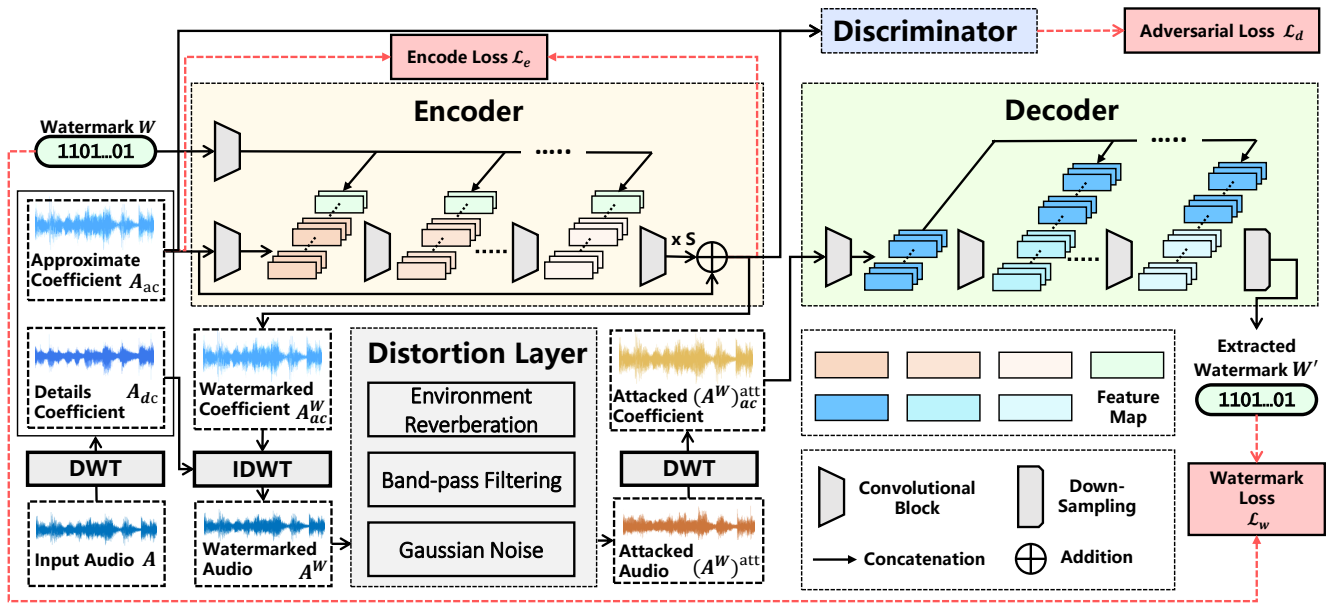


Figure 3: The overall framework of the proposed DeAR.

audio  $A$ , we randomly select a base IR  $p$  from the dataset  $P$ , and operate convolution  $Conv(\cdot)$  on  $A$  by  $p$  to simulate the Environment Reverberation (ER). *i.e.*,

$$ER(A) = Conv(A, p), \quad \text{where } p \in P. \quad (7)$$

Here, we follow the previous work (Palomäki, Brown, and Barker 2004) and leverage the acoustic impulse responses dataset<sup>1</sup> that was collected using four mics and various degrees of reverberation in the varechoic chamber at Bell Labs.

**Band-pass Filtering.** Since the frequency band of human hearing is limited, the widely-used normal range is from 500 Hz to 2000 Hz. Based on this, commonly-used speakers will not play audio with a too high or too low-frequency band. Meanwhile, the microphone will also process the playing audio, usually cutting off the frequency band outside the normal range to reduce noise, that is, a basic denoising process. Therefore, to simulate the distortions caused by the inherent characteristics of devices such as the speaker and the microphone, we apply a frequency band-pass filtering  $BF(\cdot)$  operation to the watermarked audio. Given a target audio  $A$ , we conduct  $BF(\cdot)$  as follows:

$$BF(A) = LF[HF[A, \alpha], \beta], \quad (8)$$

where  $LH[\cdot]$  and  $HF[\cdot]$  represent the low-pass filtering and the high-pass filtering, respectively. And  $\alpha$  and  $\beta$  denote the corresponding thresholds of  $HF[\cdot]$  and  $LF[\cdot]$ .

**Gaussian Noise.** In addition to the above two components, we introduce the popular Gaussian noise to simulate the random noise induced by indeterminate factors during AR process. Gaussian noise is a kind of additive noise, that is widely

used in current automatic speech recognition (ASR) systems (Yin et al. 2015) to enhance the robustness against random environmental noise. Specifically, we operate  $GN(\cdot)$  on audio  $A$  by directly superimposing noise  $\omega \sim \mathcal{N}(0, \sigma^2)$ , *i.e.*,

$$GN(A) = A + \omega, \quad \text{where } \omega \sim \mathcal{N}(0, \sigma^2). \quad (9)$$

It shall be mentioned that the  $\sigma$  is audio-aware and determined by the pre-defined signal-noise-ratio, which is randomly sampled from 20 dB to 25 dB.

**The pipeline of the distortion layer.** DAR is set as the distortion layer to enhance robustness against AR process. The pipeline of DAR is as follows:

$$DAR(\cdot) = GN(BF(ER(\cdot))). \quad (10)$$

Given a watermarked audio  $A^W$ , we can finally obtain the attacked audio  $(A^W)^{att}$ , *i.e.*,

$$(A^W)^{att} = DAR(A^W). \quad (11)$$

Afterward, we leverage DWT to acquire its corresponding approximate coefficients  $(A^W)^{att}_{ac}$  and feed it into the decoder  $De$  for watermark extraction, *i.e.*,

$$(A^W)^{att}_{ac}, (A^W)^{att}_{dc} = DWT((A^W)^{att}), \quad (12)$$

$$W' = De((A^W)^{att}_{ac}).$$

## More Details of DeAR

**Network Structures.** For the encoder  $En$ , we adopt a fully convolutional network, which keeps the size of the feature maps in each layer unchanged. Similar to  $En$ , the decoder  $De$  and the discriminator  $D$  also utilize a fully convolutional network but further append convolutional downsampling module and mean operation for final extraction and binary classification, respectively. To remedy the information loss

<sup>1</sup><https://www1.icsi.berkeley.edu/Speech/papers/gelbart-m/pointers/>



during forward propagation, we leverage the skip connection to stack the initial information with the feature map in each layer for both **En** and **De**. More importantly, we leverage 1D convolution rather than 2D convolution in basic networks, which is effective for the 1D audio waveform.

**Loss functions.** During the training stage, we jointly train the encoder and the decoder, and the whole loss function  $\mathcal{L}$  can be formulated as follows:

$$\mathcal{L} = \lambda_e \mathcal{L}_e + \lambda_d \mathcal{L}_d + \lambda_w \mathcal{L}_w, \quad (13)$$

where  $\lambda_e$ ,  $\lambda_d$ , and  $\lambda_w$  are used to balance the three terms.

## Experiments

### Experiment Settings

**Dataset.** We conduct our experiments on FMA (Defferrard et al. 2017), a famous music analysis dataset in which 12000 audios are utilized for the training of the proposed DeAR, and 200 randomly selected audios are adopted as testing audios. The sampling frequencies are all 44.1kHz and are processed as wav audio with 500k sample points in length.

**Metrics.** To measure the fidelity of the watermarked audio, the signal-to-noise ratio (SNR) is adopted with the following definition:

$$\text{SNR} = 10 \cdot \log \left( \frac{\sum_{i=1}^N A(i)^2}{\sum_{i=1}^N [A^W(i) - A(i)]^2} \right).$$

Besides, we take the average bit recovery accuracy  $\overline{\text{ACC}}$ , calculated from 200 testing audios, to evaluate the robustness of watermarking schemes.

**Implementation Details.** In the training process of DeAR, we set  $\lambda_e = 150$ ,  $\lambda_w = 1$  and  $\lambda_d = 0.01$ , and utilize Adam(Kingma and Ba 2014) with a learning rate of  $10^{-4}$  for optimization by default. And we empirically set the threshold of the high-pass filtering ( $HF[\cdot]$ ) and that of the low-pass filtering ( $LF[\cdot]$ ), namely,  $\alpha$  and  $\beta$  as 1 kHz and 4 kHz for training the model robust to re-recording.

In the testing process, the same watermark bit sequence of 100 bits is embedded for all testing audios. We adopt methods the most relevant to the proposed DeAR as the baseline methods, *i.e.*, Liu’s method (Liu, Huang, and Huang 2018) and Su’s method (Su et al. 2018). For a fair comparison, we slightly modify the parameter of Su’s method for better robustness<sup>2</sup>.

For re-recording experiments, we use one consumer-grade speaker, SENNHEISER Sp10, as the playback device and a consumer-grade microphone, ATR2100-USB, to re-record the played audio. The corresponding experimental scenario is shown in Fig. 4, where 5 cm is set as the default distance between the speaker and the microphone. Considering the de-synchronization introduced by digital-to-analog conversion, we design a synchronization strategy and apply it to DeAR and all baseline methods for a fair comparison.

<sup>2</sup>The upper bound of the interval of the search space of intensity factor is set to 2.




Figure 4: A practical experimental environment to evaluate robustness against audio re-recording.

Metrics	DeAR	Liu	Su
SNR (dB)	<b>25.86</b>	25.81	24.94
$\overline{\text{ACC}}$ (%)	<b>99.18</b>	77.09	56.00

Table 1: Quantitative comparison with the baseline methods.

Specifically, we shift the target re-recorded audio within a pre-defined range (3/44.1 s, 8/44.1 s) and calculate the corresponding bit recovery accuracy ( $\overline{\text{ACC}}$ ), in which the highest one is regarded as the final  $\overline{\text{ACC}}$ .

### Comparison Results

**Fidelity.** We first compare the fidelity of the proposed DeAR with the baseline methods. As shown in Table 1, DeAR achieves a 25.86 SNR, which outperforms the baseline methods (SNR is 25.81 and 24.94 for Liu’s method (Liu, Huang, and Huang 2018) and Su’s methods (Su et al. 2018), respectively). Furthermore, we provide one visual example of the watermarked audio for qualitative comparison in Fig. 5. We can observe that the proposed DeAR and Liu’s method tend to modify the original audio adaptively, while Su’s method prefers a relatively large and uniform noise pattern. Compared with Liu’s method, DeAR induces more slight modification to guarantee fidelity.

**Robustness Against Audio Re-recording.** We compare robustness against audio re-recording in this experiment and provide quantitative results in Table 1. With comparable fidelity, DeAR outperforms the baseline methods by a large margin (above 20% and 40%, respectively). Su’s method is fragile to the audio re-recording (AR) process because it only considers the common distortions during the digital transmission but lacks the consideration of AR. For Liu’s method, we believe that its limited performance (77.09%) is due to its AR model being an idealized additive noise, which is inconsistent with the practical scenario. In Fig. 6, we provide a visual example of the noise induced by the practical re-recording process, which is a more complex noise pattern with an even larger amplitude than the original audio. In addition to the default distance (5 cm), we further conduct a controlled comparison with Liu’s method at different




Figure 5: Qualitative comparison of fidelity. The top row shows the original audio and watermarked audios of DeAR and the baseline methods, wherein the black component represents the  $10\times$  residuals. To better illustrate their difference, the enlarged views of the same regions are shown in the middle and bottom rows.

Distance (cm)	5	20	50	100
DeAR	<b>99.18</b>	<b>98.55</b>	<b>93.40</b>	<b>92.68</b>
Liu	77.09	82.64	74.76	66.02

Table 2: Comparison of robustness against re-recording at different distances.




Figure 6: Visual example of noise induced by audio re-recording.

distances. As shown in Table 2, our method performs better over a wide range of distances. As the distance increases, the robustness against AR suffers a corresponding degradation but is still acceptable (all above 90%).

**Robustness Against Other Common Distortions.** To compare robustness more comprehensively, we further evaluate it under other common distortions during digital transmission, namely, Gaussian noise under different signal-to-noise ratios (10 dB, 15 dB, 20 dB, 25 dB), MP3 compression (64 kbps, 128 kbps), Band-pass (1 kHz high-pass, 4 kHz low-pass), Re-sampling (the watermarked audios are resampled to 90% of the original sampling frequency and

Distortions		DeAR	Su	Liu
Gaussian Noise	10 dB	87.13 / 98.67	<b>100.0</b>	60.09
	15 dB	93.77 / 99.89	<b>100.0</b>	62.41
	20 dB	96.10 / 99.99	<b>100.0</b>	65.76
	25 dB	96.96 / 99.99	<b>100.0</b>	71.61
MP3	64 kbps	95.22 / <b>99.94</b>	97.98	88.54
	128 kbps	97.00 / <b>99.97</b>	98.00	98.53
Band-pass	1 kHz	98.17 / <b>100.0</b>	<u>57.02</u>	99.14
	4 kHz	92.12 / 99.06	<b>99.99</b>	<u>50.57</u>
Re-sampling		97.04 / <b>100.0</b>	<b>100.0</b>	99.37
Dropout		97.14 / <b>99.99</b>	99.08	69.63
Amplitude Modification		97.14 / <b>99.99</b>	98.00	99.84
Re-quantization		97.11 / <b>100.0</b>	<b>100.0</b>	94.22
Median Filtering		96.16 / <b>100.0</b>	<b>100.0</b>	90.80

Table 3: Robustness against other common distortions. We provide the ACC of the default / the enhanced DeAR.

then resampled back), Dropout (drop (zero) a sample every 100 sample points), Amplitude Modification (90% of the original audios), Re-quantization (requantize down to 8 bits/sample), and Median Filtering (window size of 3). As shown in Table 3, Su’s method achieves excellent robustness in most cases except under the 1 kHz high-pass (underlined data), which may be the reason for its fragility to audio re-recording. Liu’s method cannot extract the watermark well, facing 4 kHz low-pass. In contrast, DeAR is robust under all types of distortions.

Furthermore, we propose the enhanced DeAR by adding additional distortions to the distortion layer in addition to the AR distortion. And one distortion is randomly selected for each mini-batch during training. The additional distortions contains the above common distortions except Gaussian noise and band-pass filtering. For example, it only involves Gaussian noise under 10 dB for the enhanced training. As

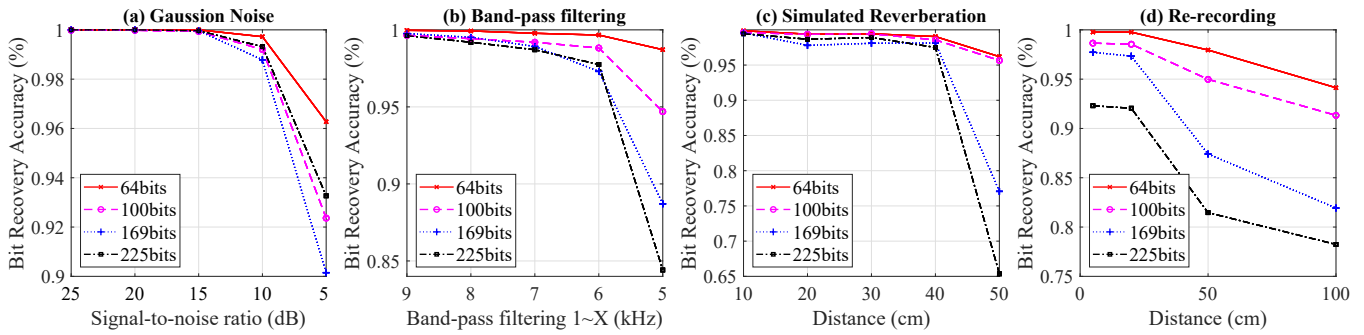


Figure 7: The influence of different embedding bits on robustness.

Distance (cm)	5	20	50	100
Default	<b>99.18</b>	<b>98.55</b>	93.40	<b>92.68</b>
Enhanced	98.65	98.54	<b>94.98</b>	91.33

Table 4: Comparison of robustness against AR between the default and the enhanced DeAR.

shown in Table 3 and Table 4, the enhanced DeAR can achieve better robustness against other common distortions while still preserving robustness against audio re-recording.

## Ablation Study

**The Influence of Different Embedding Bits.** We further explore the influence of different embedding bits on robustness. We take 64 bits, 100 bits, 169 bits, and 225 bits as examples to evaluate the robustness against Gaussian noise, band-pass, simulated reverberation, and audio re-recording. We evaluate the robustness of the model with different embedding bits under the exact SNR requirement ( $SNR = 26 \pm 0.2$ ). For each distortion, we also test on different strengths. We shall note that the simulated reverberation is implemented by convolving watermarked audios and impulse responses. The impulse response here is selected in a conference room at different distances between the speaker and the microphone. As shown in Fig. 7, in most cases, we shall sacrifice the robustness ability to some extent if we want to embed more information into the target audio.

## The Importance of Each Component of the Re-recording Modeling.

To fully verify the importance of the proposed re-recording modeling, we re-train the model with the modified distortion layer three times. In each training, one component of three distortions in the distortion layer is removed. Besides the default consumer-grade microphone, we also utilize the widely-used smartphone (*e.g.* Apple iPhone12 pro) to re-record the watermarked audio. As shown in Table 5, all designed distortion components, *i.e.*, environment reverberation, band-pass filtering, and Gaussian noise could improve the robustness against audio re-recording process. Among them, the reverberation is much necessary because it induces an accuracy improvement of more than 25%.

Device	Default	w/o ER	w/o BF	w/o GN
Microphone	<b>99.18</b>	73.63	98.52	75.80
Smartphone	<b>98.57</b>	66.17	92.64	75.25

Table 5: The robustness performance ( $\overline{ACC}$ ) of different configurations under different devices.

Strength Factor	0.2	0.5	0.8	1	1.2
SNR (dB)	39.84	31.88	27.79	25.86	24.27
$\overline{ACC}$ (%)	79.58	94.81	98.48	99.18	99.49

Table 6: Performance of DeAR with different strength factors.

**Flexibility with Strength Factor.** Given an audio coefficient and watermark as input, the well-trained encoder outputs a watermark residual, which is further superimposed on the input. In a practical scenario, we are able to flexibly adjust the strength of the residual to balance the trade-off between fidelity and robustness. In Table 6, we provide the quantitative results of fidelity and robustness against audio re-recording under different strength factors, which demonstrate the flexibility of the proposed DeAR.

## Conclusion

To the best of our knowledge, we are the first to propose using deep neural networks for audio watermarking and achieve robustness against the audio re-recording (AR) process. To achieve this, we jointly train an encoder and decoder for watermark embedding and extraction, between which a distortion layer that simulates the re-recording process is further inserted to enhance robustness. Extensive experiments demonstrate that our method outperforms the baseline methods in terms of both fidelity and resilience to the AR process. Furthermore, some ablation studies are also conducted to verify the importance of our design and the flexibility in practical scenarios.

## Acknowledgments

This work was supported in part by the Natural Science Foundation of China under Grant 62072421, 62002334, 62102386, 62121002 and U20B2047 and Students' Innovation and Entrepreneurship Foundation of USTC under Grant XY2022X01CY.

## References

- Ahmadi, M.; Norouzi, A.; Karimi, N.; Samavi, S.; and Emami, A. 2020. ReDMark: Framework for residual diffusion watermarking based on deep networks. *Expert Systems with Applications*, 146: 113157.
- Asikuzzaman, M.; and Pickering, M. R. 2017. An overview of digital video watermarking. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(9): 2131–2153.
- Bansal, N.; Deolia, V. K.; Bansal, A.; and Pathak, P. 2015. Comparative analysis of LSB, DCT and DWT for Digital Watermarking. In *2015 2nd International Conference on Computing for Sustainable Global Development (INDIA-Com)*, 40–45. IEEE.
- Cvejic, N.; and Seppanen, T. 2004. Increasing robustness of LSB audio steganography using a novel embedding method. In *International Conference on Information Technology: Coding and Computing, 2004. Proceedings. ITCC 2004.*, volume 2, 533–537. IEEE.
- Defferrard, M.; Benzi, K.; Vandergheynst, P.; and Bresson, X. 2017. FMA: A Dataset for Music Analysis. In *18th International Society for Music Information Retrieval Conference (ISMIR)*.
- Erfani, Y.; Pichevar, R.; and Rouat, J. 2016. Audio watermarking using spikegram and a two-dictionary approach. *IEEE transactions on information forensics and security*, 12(4): 840–852.
- Fang, H.; Zhang, W.; Zhou, H.; Cui, H.; and Yu, N. 2018. Screen-shooting resilient watermarking. *IEEE Transactions on Information Forensics and Security*, 14(6): 1403–1418.
- Jia, J.; Gao, Z.; Chen, K.; Hu, M.; Min, X.; Zhai, G.; and Yang, X. 2020. RIHOOP: robust invisible hyperlinks in offline and online photographs. *IEEE Transactions on Cybernetics*.
- Jiang, S.; Ye, D.; Huang, J.; Shang, Y.; and Zheng, Z. 2020. SmartSteganography: Light-weight generative audio steganography model for smart embedding application. *Journal of Network and Computer Applications*, 165: 102689.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kreuk, F.; Adi, Y.; Raj, B.; Singh, R.; and Keshet, J. 2020. Hide and Speak: Towards Deep Neural Networks for Speech Steganography. *Proc. Interspeech 2020*, 4656–4660.
- Liu, Z.; Huang, Y.; and Huang, J. 2018. Patchwork-based audio watermarking robust against de-synchronization and recapturing attacks. *IEEE transactions on information forensics and security*, 14(5): 1171–1180.
- Luo, X.; Li, Y.; Chang, H.; Liu, C.; Milanfar, P.; and Yang, F. 2021. Dvmark: A deep multiscale framework for video watermarking. *arXiv preprint arXiv:2104.12734*.
- Natgunanathan, I.; and Xiang, Y. 2010. A novel bipolar time-spread echo hiding based watermarking method for stereo audio signals. In *2010 4th International Conference on Signal Processing and Communication Systems*, 1–7. IEEE.
- Natgunanathan, I.; Xiang, Y.; Rong, Y.; Zhou, W.; and Guo, S. 2012. Robust patchwork-based embedding and decoding scheme for digital audio watermarking. *IEEE Transactions on Audio, Speech, and Language Processing*, 20(8): 2232–2239.
- Palomäki, K. J.; Brown, G. J.; and Barker, J. P. 2004. Techniques for handling convolutional distortion with missing data' automatic speech recognition. *Speech Communication*, 43(1-2): 123–142.
- Peddinti, V.; Chen, G.; Povey, D.; and Khudanpur, S. 2015. Reverberation robust acoustic modeling using i-vectors with time delay neural networks. In *Sixteenth Annual Conference of the International Speech Communication Association*.
- Qin, Y.; Carlini, N.; Cottrell, G.; Goodfellow, I.; and Raffel, C. 2019. Imperceptible, robust, and targeted adversarial examples for automatic speech recognition. In *International conference on machine learning*, 5231–5240. PMLR.
- Su, Z.; Zhang, G.; Yue, F.; Chang, L.; Jiang, J.; and Yao, X. 2018. SNR-constrained heuristics for optimizing the scaling parameter of robust audio watermarking. *IEEE Transactions on Multimedia*, 20(10): 2631–2644.
- Tancik, M.; Mildenhall, B.; and Ng, R. 2020. Stegastamp: Invisible hyperlinks in physical photographs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2117–2126.
- Wu, C.-P.; Su, P.-C.; and Kuo, C.-C. J. 1999. Robust audio watermarking for copyright protection. In *Advanced Signal Processing Algorithms, Architectures, and Implementations IX*, volume 3807, 387–397. SPIE.
- Xiang, Y.; Natgunanathan, I.; Peng, D.; Hua, G.; and Liu, B. 2017. Spread spectrum audio watermarking using multiple orthogonal PN sequences and variable embedding strengths and polarities. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(3): 529–539.
- Yakura, H.; and Sakuma, J. 2018. Robust audio adversarial example for a physical attack. *arXiv preprint arXiv:1810.11793*.
- Yin, S.; Liu, C.; Zhang, Z.; Lin, Y.; Wang, D.; Tejedor, J.; Zheng, T. F.; and Li, Y. 2015. Noisy training for deep neural networks in speech recognition. *EURASIP Journal on Audio, Speech, and Music Processing*, 2015(1): 1–14.
- Ying, Q.; Zhou, H.; Zeng, X.; Xu, H.; Qian, Z.; and Zhang, X. 2021. Hiding Images into Images with Real-world Robustness. *arXiv preprint arXiv:2110.05689*.
- Zhang, J.; Chen, D.; Liao, J.; Fang, H.; Zhang, W.; Zhou, W.; Cui, H.; and Yu, N. 2020. Model watermarking for image processing networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 12805–12812.



Zhang, K. A.; Xu, L.; Cuesta-Infante, A.; and Veeramachaneni, K. 2019. Robust invisible video watermarking with attention. *arXiv preprint arXiv:1909.01285*.

Zhu, J.; Kaplan, R.; Johnson, J.; and Fei-Fei, L. 2018. Hidden: Hiding data with deep networks. In *Proceedings of the European conference on computer vision (ECCV)*, 657–672.