WaveFuzz: A Clean-Label Poisoning Attack to Protect Your Voice

Yunjie Ge[1] , **Qian Wang**[1*] , **Jingfeng Zhang**[2] , **Juntao Zhou**[1] , **Yunzhu Zhang**[1] , **Chao Shen**[3] ,

[1]School of Cyber Science and Engineering, Wuhan University, Wuhan, Hubei, China
[2]RIKEN Center for Advanced Intelligence Project, Tokyo, Japan
[3]School of Cyber Science and Engineering, Xi'an Jiaotong University, Xi'an, Shaanxi,China

{yunjiege, qianwang} @whu.edu.cn, jingfeng.zhang@riken.jp,
{ yunzhuzhang, juntaozhou} @whu.edu.cn, chaoshen@mail.xjtu.edu.cn

Abstract

People are not always receptive to their voice data being collected and misused. Training the audio intelligence systems needs these data to build useful features, but the cost for getting permissions or purchasing data is very high, which inevitably encourages hackers to collect these voice data without people's awareness. To discourage the hackers from proactively collecting people's voice data, we are the first to propose a clean-label poisoning attack, called WaveFuzz, which can prevent intelligence audio models from building useful features from protected (poisoned) voice data but still preserve the semantic information to the humans. Specifically, WaveFuzz perturbs the voice data to cause Mel Frequency Cepstral Coefficients (MFCC) (typical representations of audio signals) to generate the poisoned frequency features. These poisoned features are then fed to audio prediction models, which degrades the performance of audio intelligence systems. Empirically, we show the efficacy of WaveFuzz through attacking two representative types of intelligent audio systems, *i.e.*, speaker recognition system (SR) and speech command recognition system (SCR). For example, the accuracies of models are declined by 19.78% when only 10% of the poisoned voice data is to fine-tune models, and the accuracies of models declined by 6.07% when only 10% of the training voice data is poisoned. Consequently, WaveFuzz is an effective technique that enables people to fight back to protect their own voice data, which sheds new light on ameliorating privacy issues.

1 Introduction

Due to the prevalence of social networks (*e.g.*, Facebook, Twitter, and YouTube), individuals prefer to share their daily lives on the web. For example, individuals interact with their friends and acquaintances with videos or audios in these nets. Therefore, such rich data at never-before-seen scales is contained on the web. Unfortunately, social networking sites al-
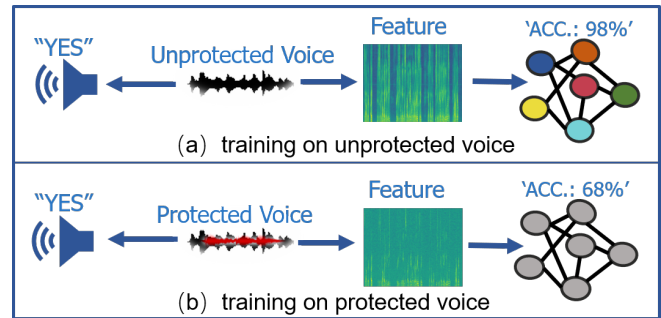

Figure 1: The protected voice example sounds like normal voice example but will produce the wrong features to the audio model, which degrades the accuracy of malicious models.

ways lack data protection, meaning that hackers can easily collect and misuse users' voice data as shown in Fig. 1 (a). Nowadays, training high-performance audio intelligence systems requires massive voice data to build valuable features (*e.g.*, voiceprint, pronunciation, and phoneme). Hackers prefer to collect these voice data without users' awareness from the web, *e.g.*, using the crawler, owing to the expensive cost of getting permissions or purchasing voice data. However, people are not always receptive to their voice data being collected and used since voice data contains various information, including voiceprint and speech content. Collecting and using people's voice data will bring privacy leakage concerns for the individual. Despite the great success in protecting user privacy (*e.g.*, face) [Liu *et al.*, 2017], it is still an open challenge to mitigate privacy leaks about people's voice data.

To discourage the hacker from proactively collecting people's voice data, we propose a clean-label poison attack, called WaveFuzz, which achieves the ability to generate indistinguishable protected voice data. Specifically, WaveFuzz elaborately perturbs voice data by maximizing the feature distance between unprotected (clean) voice data and protected (poisoned) voice data and minimizing their input distance, where features are extracted by MFCC (the de facto standard used by most intelligent audio systems). Then, the poisoned features extracted from poisoned voice data are fed into audio prediction models, where the performance of the models will inevitably be degraded. Besides, we utilize a penalty term to limit the difference between benign and perturbed voice data

---

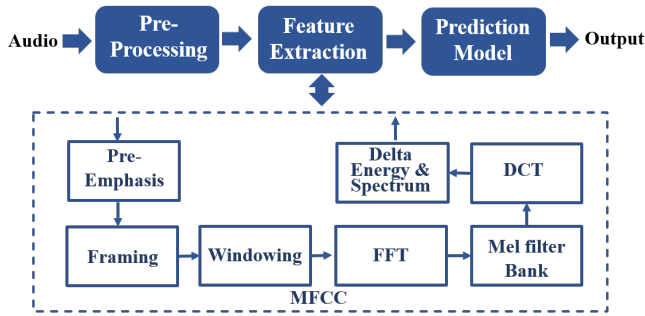*Qian Wang is the corresponding author.

Figure 2: The typical architecture of audio recognition systems and the overview of MFCC.

to preserve the semantic meaning. Therefore, the protected voice data can transfer the correct information to humans but wrong knowledge to the machine. Moreover, the hacker is hard to use the poisoned samples to training modes and suffers from a punishment as shown in Fig. 1 (b), *i.e.*, performance decline .

We conduct experiments on two representative types of intelligent audio systems: SR and SCR. Moreover, we assume that the hacker may use users' voice data to fine-tune models or train models from scratch. For fine-tuning models, our WaveFuzz can reduce the average of model accuracies by 19.78%, 14.98%, and 10.74% when 10%, 5%, and 0.5% of the poisoned voice data is to fine-tune models. For training model from scratch, our WaveFuzz can reduce the average of model accuracies by 6.07%, 4.87%, and 3.86% when 10%, 5%, and 0.5% of voice data is poisoned. Besides, WaveFuzz outperforms the baseline method by 9.42%, 6.81%, and 4.79% of accuracy decline, when 10%, 5%, and 0.5% of the training voice data is poisoned. We also make a human study to validate that semantic meaning is preserved, where about 90% volunteers can correctly understand the semantic meaning of poisoned voice data.

2 Background and Related Work

In this section, we introduce the components of intelligence audio recognition systems and then review related studies.

2.1 Background

The typical intelligence audio recognition system consists of three major components (as shown in Fig. 2): pre-processing, feature extraction, and the prediction model. Given a raw audio input, pre-processing will retain the frequencies in the range of human hearing and the segments higher than a specific energy threshold. Then, feature extraction will obtain the features from the processed audio by a feature extraction algorithm, such as Mel-frequency Cepstral Coefficients (MFCC) [Muda *et al.*, 2010], Linear Predictive Coefficient (LPC) [Itakura, 1975], and Perceptual Linear Predictive (PLP) [Hermansky, 1990]. It is noted that MFCC is the de facto standard used in the audio field. Finally, the prediction model takes the extracted features as input and outputs the decisions.

MFCC [Muda *et al.*, 2010] can extract the parametric representation from an acoustic signal, which is in line with the

acoustic characteristics of the human ear. As shown in Fig. 2, MFCC consists of the following steps: 1) pre-emphasis increases the energy of signal at a higher frequency; 2) framing segments the speech samples into a small frame; 3) hamming windowing considers the next block and integrates all the closest frequency lines; 4) Fast Fourier transform (FFT) converting each frame from the time domain into the frequency domain; 5) The energy spectrum is fed to a set of Mel-scale filter banks; 6)Discrete Cosine Transform (DCT) is used to convert the log Mel spectrum into the time domain; 7) Delta energy & delta spectrum are used to add features related to the change in spectral features over time.

2.2 Related Work

In this section, we simply review the related studies, including privacy protection methods and poisoning attacks.

Privacy Protection: To protect the users' privacy in the social networks, a bulk of studies [Devaux *et al.*, 2009; Wilber *et al.*, 2016] utilized various obfuscation methods to blur human faces in a photo which can evade face detection. Moreover, several studies [Szegedy *et al.*, 2014; Oh *et al.*, 2017] used adversarial examples to protect photo privacy. The defender generates the adversarial example by adding invisible perturbation to the benign image, where the adversarial example can mislead the recognition models.

Poisoning attack: The poisoning attacks can be divided into two classes [Jagielski *et al.*, 2018] based on the attack behaviors, including availability attacks and targeted attacks. Availability attacks [Nelson *et al.*, 2008; Xiao *et al.*, 2012] introduce failures in the final classifier by tampering with the training dataset, which have been the focus of many research efforts over time. Targeted attacks [Gu *et al.*, 2017; Xiao *et al.*, 2015; Muñoz-González *et al.*, 2017] affect specific data points for the victim model. Recently, clean label poisoning attacks [Aghakhani *et al.*, 2021; Geiping *et al.*, 2020; Huang *et al.*, 2020; Zhu *et al.*, 2019] have been proposed. The poisoned image example needs to be similar to the legal example but can perturb the model. Recent, VENOMAVE[Aghakhani *et al.*, 2020] is proposed to induce a targeted misclassification against automatic speech recognition (ASR) model using data poisoning. For example, the poisoned ASR model classifies the target frame x_τ as state Z, but performs well on the clean input. However, our paper aims to prevent unauthorized exploitation of voice data, a new aspect of data privacy and different from VENOMAVE.

Despite great success in protecting the privacy of users' images, there remains an urgent problem of preserving individual voices from being misused.

3 Method

In the following, we introduce our WaveFuzz (see Fig. 3), the first clean label poisoning attack to discourage hackers from collecting users' voice data.

Insight. Given a voice example, the feature extraction component will extract the features as the input fed into the prediction model. Therefore, if the current features generated by the feature extraction algorithm are wrong, the wrong features will damage the prediction model inevitably. Moreover,
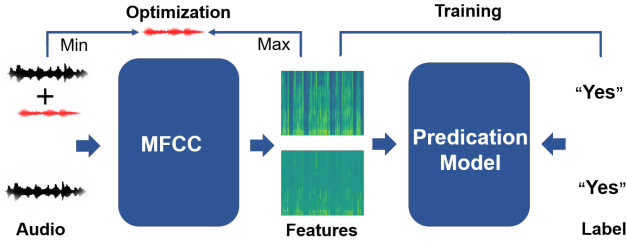
Figure 3: The overview of WaveFuzz: The defender generates corrupted points by maximizing the distance of the MFCC features between normal and corrupted voice data, and minimizing their input distance. Then, poisoned features are fed into the prediction model will cause intelligence systems' performance to decline.

Algorithm 1 WaveFuzz

Input: Normal example x; MFCC threshold ε;
Output: poisoned example x^*
1: initialization: $\delta^{(0)}$
2: $i = 0$
3: **while** $\left\| \mathcal{MF}(x + \delta^{(i)}) - \mathcal{MF}(x) \right\|_p < \varepsilon$ **do**
4:     $\delta^{(i)} \leftarrow E.q.\ 3$
5:     $i = i + 1$
6: **end while**
7: $x^* = x + \delta^{(i)}$
8: **return** x^*

---

the feature spectrum is not intuitive, making it difficult for humans to judge whether the spectrum is correct.

Threat Model. We discourage the hacker from proactively collecting users' voice data by generating protected voice data. These protected voice examples are sound normal but can degrade the performance of the systems that want to steal the information from users' voice data. We assume that the malicious systems use the MFCC to extract the features since MFCC is the most popular feature extraction algorithm in the audio field. Besides, we assume that the defender does not know the information of the prediction model.

3.1 WaveFuzz

Considering the popularity of MFCC, we use it as the feature extraction algorithm and denote it as $\mathcal{MF}(\cdot)$. MFCC will output features $\mathcal{MF}(x)$ for a voice example x, and the features are fed into the prediction model.

Given a voice input x and the original label y, we want to craft a corrupted example x^* that sounds like x, but $\mathcal{MF}(x^*)$ has a huge gap (distance) with $\mathcal{MF}(x)$. To achieve that, the overall solution can be divided into two steps. In the first step, we need to push the current example away from the original feature boundary. In other words, we need to find such example x^* which meets the following constraint:

$$\left\| \mathcal{MF}(x^*) - \mathcal{MF}(x) \right\|_p > \varepsilon, \tag{1}$$

where ε is a pre-defined threshold, and $\left\| \cdot \right\|_p$ calculates the distance. We think that if the threshold ε is large enough, the features $\mathcal{MF}(x^*)$ are away from the feature boundary of the current class .

Task	SR		SCR	
Model	VggVox	Resnet-18	Attention	CNN
Dataset	Voxceleb	Voxceleb	Speech Command V1	Speech Command V2
Feature Extraction	FFT	MFCC	MFCC	FFT
Accuracy	82.20%	79.07%	93.78%	91.00%

Table 1: The information about the malicious models, the accuracies, and the datasets.

The current poisoned samples can compromise malicious models, where the models are hard to learn from the samples. However, the protected samples remain a weakness. That is, the individual may not understand the semantic content. To solve the problem, we make the poisoned example x^* sound like the normal example x in the second step. Therefore, the poisoned example x^* should still satisfy the second constraint:

$$\left\| x^* - x \right\|_p < \epsilon, \tag{2}$$

where ϵ is a threshold to limit the difference between the two examples.

It is difficult to find such a example x^* by solving the restricted problem. Therefore, we turn the restricted problem into the following optimization function:

$$\arg \min_{\delta} \ - \left\| \mathcal{MF}(x + \delta) - \mathcal{MF}(x) \right\|_p + \alpha \cdot \left\| \delta \right\|_p, \tag{3}$$

where δ is the elaborate noise added to the original example, and α is a parameter to balance the attack performance and stealthiness. Besides, we use the L_2-norm to calculate the distance between the two feature vectors.

The overall pseudo-code is shown in Algorithm 1, which is easy to implement. We perform the optimization function (E.q. 3) until the distance between $\mathcal{MF}(x)$ and $\mathcal{MF}(x^*)$ is longer than the preselected threshold.

4 Experiment

In this section, we consider that the hackers may use users' voice data to fine-tune pre-trained models (*e.g.*, model updating) or train models from scratch.

4.1 Experiment Setup

The information about the malicious models and the datasets is summarized in Table 1. We introduce the datasets and the relevant models used in our experiments in the following.

- Voxceleb[1]: It contains about 100,000 voices from 1,251 celebrities, extracted from interview videos uploaded to YouTube. VggVox is a speaker recognition model proposed by [Chung *et al.*, 2018].

- SpeechCommand[2]: It contains 64,727 utterances from 1,881 speakers. Version 1 (V1) includes ten words. Version 2 (V2) adds four more command words. The detail

---

[1]https://www.robots.ox.ac.uk/ vgg/data/voxceleb/
[2]https://tensorflow.google.cn/datasets/catalog/speech_commands

Model	PR(%)	DAcc.(%)	SNR	ϵ	α
		TASK:SR			
VggVox Model	10	30.34	11.13	400	0.1
	5	24.55	11.15		
	1	22.48	11.19		
	0.5	20.31	11.25		
		TASK:SCR			
Attention Model	10	14.63	4.14	300	0.1
	5	10.64	4.92		
	1	7.42	5.40		
	0.50	6.18	5.12		
CNN Model	10	14.37	7.89	1700	0.1
	5	9.75	7.80		
	1	7.37	7.73		
	0.50	5.75	7.69		

Table 2: WaveFuzz against models that are fine-tuned by the poisoned voice data.

Model	PR(%)	DAcc.(%)	SNR	ϵ	α
		TASK:SR			
VggVox Model	10	7.28	10.14	1000	0.1
	5	5.23	10.28		
	1	4.00	10.04		
	0.5	2.15	10.27		
Resnet-18 Model	10	2.48	9.96	1000	0.1
	5	1.84	10.14		
	1	1.54	10.23		
	0.5	3.39	10.05		
		TASK:SCR			
Attention Model	10	7.89	3.88	1700	0.1
	5	6.97	3.90		
	1	6.31	3.83		
	0.50	4.68	3.91		
CNN Model	10	6.62	6.89	1700	0.1
	5	5.45	7.45		
	1	5.35	7.26		
	0.50	5.23	7.76		

Table 3: WaveFuzz against models that are trained from scratch by the poisoned voice data.

information of the Attention-based model and the convolutional neural networks (CNN) is in [de Andrade *et al.*, 2018] and [Arik *et al.*, 2017], respectively, where the two models are used to solve the speech command recognition task.

We quantify the following evaluation goals.

- DAcc.: DAcc. is defined as: $DAcc. = Acc.(\mathcal{F}) - Acc.(\mathcal{F}^*)$, where $Acc.(\mathcal{F})$ is the accuracy of the model trained on clean dataset, and $Acc.(\mathcal{F})$ is the accuracy of the model trained on poisoned dataset. Besides, $Acc.(\mathcal{F})$ of all malicious models is reported in Table 1.

- PR: Poison rate (PR) represents the rate of poisoned data in training data. PR is defined as: $PR = m/n$, where m denotes the number of the poisoned examples, n denotes the number of the benign examples.

- SNR: We use SNR to describe the perturbation on the poisoned example x^*. $SNR = 10 \lg(\frac{x^*-x}{x})$, where x denotes the original example.

4.2 Implementation and Training Details

We generate malicious voice files following the formats of the datasets we use, *e.g.*, sampling rate and bit number. For fine-tuning, we directly use the poisoned voice data to fine-tune a pre-trained model. For training from scratch, we perturb a PR rate training samples. We will re-write the current samples using the corresponding poisoned samples. Given an voice sample, we use WaveFuzz to generate a compared poisoned sample and Adam optimizer [Kingma and Ba, 2014] to optimize the sample. Besides, We utilize a substitute model to calculate a proper ϵ where the output is changed when the feature distance between the pure and poisoned samples is longer than the ϵ.

4.3 Attack Performance

To fully evaluate the performance of our WavFuzz, we target two classical intelligence audio systems (*i.e.*, SR and SCR).

Here, the hackers use the poisoned voice data to fine-tune their pre-trained models or train models from scratch. The attack results are reported in Table 2 and Table 3. We make the following key observations.

Performance Degradation. Based on the results, we find that all the models experienced some degradation in accuracy, implying that the poisoned voice data reduces the expected performance of the malicious models. The performance of the models is severely compromised, especially when the hackers only use the poisoned voice data to fine-tune their models. For example, we only use the 5% poison data points to fine-tune the models, and the precisions of the malicious models are decreased by 24.55%, 10.64%, and 9.75%, respectively. Moreover, for training the models from scratch, our WaveFuzz is still effective to hinder the models from learning useful information from the protected voice data. Interestingly, the average model accuracy drops by 3.87% even when only 0.5% of the samples were poisoned. We also notice the abnormal results from Resnet-18, where only an DAcc. of 2.48 on an PR of 10%. The effect we attribute to the poor ability of the model where the model accuracy is only 79.07%.

DAcc. vs PR. To validate the effectiveness of our WaveFuzz, we conduct the poisoning attack under different PRs (*e.g.*, 10%, 5%, 1%, and 0.5%). A high PR value means that more protected voice samples are collected and used them to fine-tune or train the models. Based on these experimental results, we find a consistent conclusion that the more poisoned samples used for training, the more severely the performance of the model suffers.

Transferability. The malicious models we use adopt two feature extraction algorithms (*i.e.*, MFCC and FFT). Our WaveFuzz still mitigates privacy leakage since all the model

TASK	Poisoned audio				benign audio			
	Normal(%)	Noise A(%)	Noise B(%)	Lable Acc.(%)	Normal(%)	Noise A(%)	Noise B(%)	Label Acc.(%)
SCR	59.50	33.00	7.50	90.50	94.50	5.50	0.00	96.00
SR	69.00	27.00	4.00	88.50	90.75	8.00	1.25	93.50

Table 4: Results of the human perception evaluation on poisoned audio generated by WaveFuzz and the corresponding benign audio.
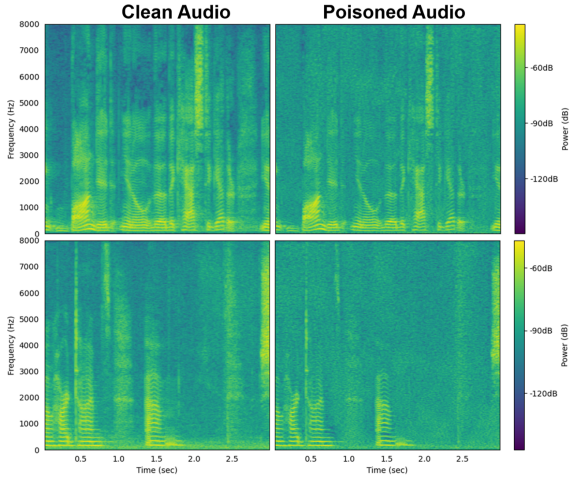

Figure 4: Spectrograms of the benign and poisoned voice examples generated by our WaveFuzz.

Figure 5: Waves of the benign and poisoned voice examples generated by our WaveFuzz.

performances are damaged. There is no doubt that the performance of WaveFuzz will be decreased when it encounters some feature extraction algorithms different from MFCC. To further improve the generalization ability of WaveFuzz, we plan to integrate various feature extraction algorithms to disturb the benign audio example as future work.

Visualization. Fig. 4 shows the spectrograms of the benign voice and poisoned voice examples. First, the two spectrograms are different, proving that the poisoned voice features are confused. Second, though there are apparent differences between the two spectrograms, we cannot visually judge which is wrong since the spectrogram is not intuitionistic. Fig. 5 shows the waveforms of the benign voice and poisoned voice examples. The two waveforms are close and indistinguishable, indicating that the poisoned voice data can successfully transfer the information to others, and they do not feel abnormal.

4.4 Human Study

We conduct a human study to evaluate the imperceptibility of the corrupted samples. In this experiment, the volunteers are asked to do two studies, where we survey 20 volunteers aged from 18 to 26, including nine males and eleven females. In the first study, given an voice example, the volunteers need to label the example, where they give the content of example for SCR task or judge who is the speaker from three candidates for SR task. We use Label Acc. to represent the rate of the poisoned samples understood by the users. If the label given by the user is the same as the original label, we believe that the current poisoned example can be successfully understood by the user. In the second study, the volunteers listen to voice
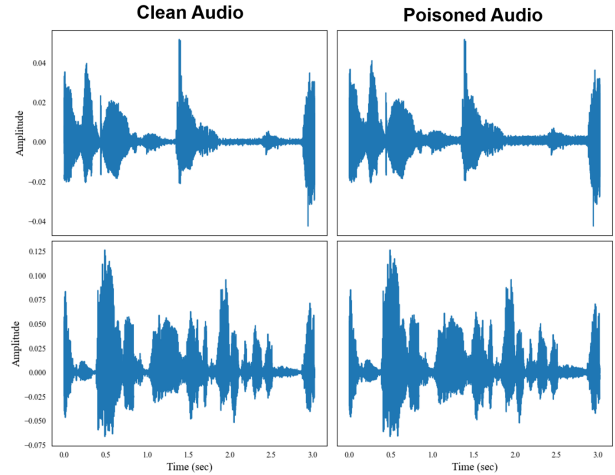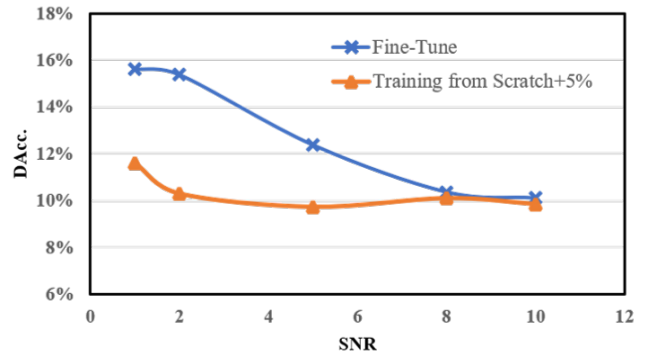

Figure 6: The DAcc. under different SNR on CNN, where we manual add the DAcc. from training from scratch by 5% to draw.

examples and give their views, *i.e.*, "normal", "noise but acceptable" (NoiseA), and "noise and unacceptable" (NoiseB). Specifically, we pick ten pairs of clean and poisoned voice examples for each audio task, where we use the clean voice examples as the reference. Table 4 presents the results of the experiments on human study, showing that our poisoned voice can successfully transfer the semantic meaning to humans.

4.5 Compare with Baseline and Trade-off

Here, we discuss that our WaveFuzz can damage the model performance, not caused by adding some noises. We report the results only adding random noise (gaussian noise) as the reference. Table 5 shows the DAcc. caused by WaveFuzz and random noise under the same level of PR and SNR . The DAcc. shows that our WaveFuzz will cause more de-

PR	WaveFuzz		Random noise	
	DAcc.(%)	SNR	DAcc.(%)	SNR
Hacker's Goal: Fine-tuning CNN				
10%	14.37	7.89	3.62	7.91
5%	9.75	7.90	2.75	7.68
1%	7.37	7.73	2.12	7.64
0.50%	5.75	7.69	1.87	7.73
Hacker's Goal: Training CNN from Scratch				
10%	6.62	6.89	2.26	7.02
5%	5.45	7.45	0.62	7.24
1%	5.35	7.26	0.12	7.51
0.50%	5.23	7.76	0.77	7.69
Hacker's Goal: Training VggVox from Scratch				
10%	7.28	10.14	1.23	10
5%	5.23	10.28	-2.77	10
1%	4.00	10.23	0.01	10
0.50%	3.39	10.05	0.31	10

Table 5: The comparison of accuracy decline caused by our Wave-Fuzz and random noise. We fine-tune CNN, and train CNN and VggVox from scratch.

cline in the model performance, where the average values of the differences are 6.49%, 4.77%, and 5.28% for fine-tuning CNN, and training CNN and VggVox from scratch respectively. Besides, the difference between DAcc.s from WaveFuzz and random noise increases with the poisoning rate, *e.g.*, 3.88% (PR=0.5%) to 10.75% (PR=10%) for fine-tuning, 3.77% (PR=0.5%) to 5.21% (PR=10%) for training from scratch.

Besides, we also conduct an experiment to analyze the trade-off between the DAcc. and the SNR (see Fig. 6), where the targeted model is the CNN. It is clear that the DAcc. decreases as the SNR increases, whereas a low SNR indicates more noise added in the clean example. The high SNR or slight noise limits us from fully interfering with the sample, so the feature distance between the original example and the poisoned example is not huge enough to confuse the malicious model. Hence, there is a trade-off between the attack performance and stealthiness.

4.6 WaveFuzz against Pre-Trained Models

We evaluate whether our poisoned voice data can confuse the pre-trained automatic recognition systems. ECAPA-TDNN [Desplanques *et al.*, 2020], an SR model, achieves the state-of-the-art performance. DeepSpeech [Hannun *et al.*, 2014] is a popular automatic speech recognition (ASR) model, and Alibaba[3] is an ASR online service. Malicious entities may extract private information from the voice to analyze users' behaviors, which will lead to severe problems of privacy leakage. We generate the poisoned voice by our WaveFuzz and use it to "attack" the pre-trained models. If the models returns the output that is different from the label of the corresponding benign voice, we think the malicious model is hard to extract the information contained in the voice. The results are shown in Table 6, where we calculate accuracies

---

[3]https://ai.aliyun.com/nls/asr

Task	Model	Normal Acc.(%)	Poisoned Acc.(%)	SNR	ϵ
SR	VggVox	82.29	66.15	10.27	1000
	ECAPA-TDNN	96.62	54.75	4.03	1700
SCR	Attention	93.28	60.12	3.88	1400
	CNN	91.28	58.24	7.89	1600
Task	Model	Normal WER(%)	Poisoned WER(%)	SNR	ϵ
ASR	DeepSpeech	23.24	121.12	3.66	3500
	Alibaba	12.2	72.21	3.54	3500

Table 6: Classification accuracies of pre-trained models on the clean and poisoned voice data, respectively.

on the benign and poisoned voice data. Particularly, we use word error rate (WER) for ASR task. Based on these results, we observe that the protected voice could not be correctly recognized by these recognition systems, meaning that the poisoned voice data is hard to by analyzed by the pre-trained models.

5 Discussion

Previously, we evaluated our WaveFuzz on SR models with the input speech containing short sentences (∼5 words) and SCR models for single-word command. In this section, we discuss the possibility of preventing more sophisticated automatic speech recognition (ASR), *e.g.*, Kaldi [4], from stealing useful information from long sentences. Here, we use the Kaldi, a classic ASR model, to discuss the possibility. When PR=0.5%, the word error rate (WER) is decreased by 0.29%, where the model generalization ability is improved. WER increases as PR further increases. For example, the word error rate (WER) increases by 0.42% when PR=10%. The results show that our WaveFuzz is not effective enough against Kaldi. The reason is that the prediction model of Kaldi is robust to poisoned features and complicated. We plan to design more effective methods to avoid various information in voice data extracted by various audio models as future work.

6 Conclusion

In this paper, we proposed a clean label poisoning attack, WaveFuzz, which can mitigate the risk of users' voice data being misused and maintain its messaging functionality. Specifically, we generate the protected voice data by maximizing the feature distance between the original voice data and poisoned voice data. We test two classic intelligence audio systems. The models can not steal information from the users' voice data and suffer from the punishment (*i.e.*, performance decline). Moreover, protected voice data can avoid being analyzed by pre-trained recognition systems, fooling these malicious systems. We make a human study to show that poisoned voice data can be understood by other users, indicating that its messaging function is normal. Consequently, WaveFuzz is a simple and effective baseline for mitigating privacy leakage of the user's voice. We think our work sheds new light on ameliorating privacy issues.

---

[4]http://www.kaldi-asr.org/

References

[Aghakhani *et al.*, 2020] Hojjat Aghakhani, Lea Schönherr, Thorsten Eisenhofer, Dorothea Kolossa, Thorsten Holz, Christopher Kruegel, and Giovanni Vigna. Venomave: Targeted poisoning against speech recognition. *arXiv preprint arXiv:2010.10682*, 2020.

[Aghakhani *et al.*, 2021] Hojjat Aghakhani, Dongyu Meng, Yu-Xiang Wang, Christopher Kruegel, and Giovanni Vigna. Bullseye polytope: A scalable clean-label poisoning attack with improved transferability. In *Proc. of IEE EuroS&P*, pages 159–178, 2021.

[Arik *et al.*, 2017] Sercan Ömer Arik, Markus Kliegl, Rewon Child, Joel Hestness, Andrew Gibiansky, Christopher Fougner, Ryan Prenger, and Adam Coates. Convolutional recurrent neural networks for small-footprint keyword spotting. In *Proc. of Interspeech*, pages 1606–1610, 2017.

[Chung *et al.*, 2018] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. Voxceleb2: Deep speaker recognition. *CoRR*, 2018.

[de Andrade *et al.*, 2018] Douglas Coimbra de Andrade, Sabato Leo, Martin Loesener Da Silva Viana, and Christoph Bernkopf. A neural attention model for speech command recognition, 2018.

[Desplanques *et al.*, 2020] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck. ECAPA-TDNN: emphasized channel attention, propagation and aggregation in TDNN based speaker verification. In *Proc. of Interspeech*, pages 3830–3834, 2020.

[Devaux *et al.*, 2009] Alexandre Devaux, Nicolas Paparoditis, Frédéric Precioso, and Bertrand Cannelle. Face blurring for privacy in street-level geoviewers combining face, body and skin detectors. In *Proc. of MVA*, pages 86–89, 2009.

[Geiping *et al.*, 2020] Jonas Geiping, Liam H. Fowl, W. Ronny Huang, Wojciech Czaja, Gavin Taylor, Michael Moeller, and Tom Goldstein. Witches' brew: Industrial scale data poisoning via gradient matching. In *Proc. of ICLR*, 2020.

[Gu *et al.*, 2017] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017.

[Hannun *et al.*, 2014] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al. Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567*, 2014.

[Hermansky, 1990] Hynek Hermansky. Perceptual linear predictive (plp) analysis of speech. *the Journal of the Acoustical Society of America*, pages 1738–1752, 1990.

[Huang *et al.*, 2020] W Ronny Huang, Jonas Geiping, Liam Fowl, Gavin Taylor, and Tom Goldstein. Metapoison: Practical general-purpose clean-label data poisoning. *arXiv preprint arXiv:2004.00225*, 2020.

[Itakura, 1975] Fumitada Itakura. Line spectrum representation of linear predictor coefficients of speech signals. *The Journal of the Acoustical Society of America*, pages S35–S35, 1975.

[Jagielski *et al.*, 2018] Matthew Jagielski, Alina Oprea, Battista Biggio, Chang Liu, Cristina Nita-Rotaru, and Bo Li. Manipulating machine learning: Poisoning attacks and countermeasures for regression learning. In *Proc. of IEEE S&P*, pages 19–35, 2018.

[Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[Liu *et al.*, 2017] Yujia Liu, Weiming Zhang, and Nenghai Yu. Protecting privacy in shared photos via adversarial examples based stealth. *Secur. Commun. Networks*, 2017:1897438:1–1897438:15, 2017.

[Muda *et al.*, 2010] Lindasalwa Muda, Mumtaj Begam, and Irraivan Elamvazuthi. Voice recognition algorithms using mel frequency cepstral coefficient (mfcc) and dynamic time warping (dtw) techniques. *arXiv preprint arXiv:1003.4083*, 2010.

[Muñoz-González *et al.*, 2017] Luis Muñoz-González, Battista Biggio, Ambra Demontis, Andrea Paudice, Vasin Wongrassamee, Emil C Lupu, and Fabio Roli. Towards poisoning of deep learning algorithms with back-gradient optimization. In *Proc. of AISec*, pages 27–38, 2017.

[Nelson *et al.*, 2008] Blaine Nelson, Marco Barreno, Fuching Jack Chi, Anthony D Joseph, Benjamin IP Rubinstein, Udam Saini, Charles Sutton, J Doug Tygar, and Kai Xia. Exploiting machine learning to subvert your spam filter. *LEET*, pages 1–9, 2008.

[Oh *et al.*, 2017] Seong Joon Oh, Mario Fritz, and Bernt Schiele. Adversarial image perturbation for privacy protection A game theory perspective. In *Proc. of IEEE ICCV*, pages 1491–1500, 2017.

[Szegedy *et al.*, 2014] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *Proc. of ICLR*, 2014.

[Wilber *et al.*, 2016] Michael J. Wilber, Vitaly Shmatikov, and Serge J. Belongie. Can we still avoid automatic face detection? In *Proc. of IEEE WACV*, pages 1–9, 2016.

[Xiao *et al.*, 2012] Han Xiao, Huang Xiao, and Claudia Eckert. Adversarial label flips attack on support vector machines. In *Proc. of ECAI*, pages 870–875. 2012.

[Xiao *et al.*, 2015] Huang Xiao, Battista Biggio, Gavin Brown, Giorgio Fumera, Claudia Eckert, and Fabio Roli. Is feature selection secure against training data poisoning? In *Proc. of ICML*, pages 1689–1698, 2015.

[Zhu *et al.*, 2019] Chen Zhu, W Ronny Huang, Hengduo Li, Gavin Taylor, Christoph Studer, and Tom Goldstein. Transferable clean-label poisoning attacks on deep neural nets. In *International Conference on Machine Learning*, pages 7614–7623, 2019.