Lan, Jiahe; Zhang, Rui; Yan, Zheng; Wang, Jie; Chen, Yu; Hou, Ronghui

Adversarial attacks and defenses in Speaker Recognition Systems

Published in:
Journal of Systems Architecture

DOI:
10.1016/j.sysarc.2022.102526

Published: 01/06/2022

Document Version
Publisher's PDF, also known as Version of record

Journal of Systems Architecture 127 (2022) 102526

Contents lists available at ScienceDirect

Journal of Systems Architecture

journal homepage: www.elsevier.com/locate/sysarc

Adversarial attacks and defenses in Speaker Recognition Systems: A survey

Jiahe Lan [a], Rui Zhang [a], Zheng Yan [a,b,*], Jie Wang [a], Yu Chen [c], Ronghui Hou [a]

[a] State Key Laboratory on Integrated Services Networks, School of Cyber Engineering, Xidian University, China
[b] Department of Communications and Networking, Aalto University, Finland
[c] School of Information Systems and Technology, Lucas College and Graduate School of Business, San Jose State University, USA

ARTICLE INFO

Keywords:
Speaker recognition system
Adversarial attacks
Adversarial examples

ABSTRACT

Speaker recognition has become very popular in many application scenarios, such as smart homes and smart assistants, due to ease of use for remote control and economic-friendly features. The rapid development of SRSs is inseparable from the advancement of machine learning, especially neural networks. However, previous work has shown that machine learning models are vulnerable to adversarial attacks in the image domain, which inspired researchers to explore adversarial attacks and defenses in Speaker Recognition Systems (SRS). Unfortunately, existing literature lacks a thorough review of this topic. In this paper, we fill this gap by performing a comprehensive survey on adversarial attacks and defenses in SRSs. We first introduce the basics of SRSs and concepts related to adversarial attacks. Then, we propose two sets of criteria to evaluate the performance of attack methods and defense methods in SRSs, respectively. After that, we provide taxonomies of existing attack methods and defense methods, and further review them by employing our proposed criteria. Finally, based on our review, we find some open issues and further specify a number of future directions to motivate the research of SRSs security.

1. Introduction

Biometrics such as fingerprint, face, and voiceprint is widely used for user identification and authentication [1,2]. Speaker recognition, as a technology that recognizes a speaker's identity through his/her voiceprint [3,4], has attracted special attention from both academia and industry due to its ease of use for remote control and economic-friendly features. The last decade has seen a dramatic improvement in Speaker Recognition Systems (SRSs), which can be divided into Speaker Identification Systems (SISs) and Speaker Verification Systems (SVSs) according to different tasks. The former is responsible for identifying which enrolled speaker utters an input, and the identification result is an enrolled speaker. The latter aims to verify whether an input is uttered by a claimed speaker, and the verification result is yes or no. SRSs have been deployed in both classical and emerging Internet-of-Things (IoT) devices [5], such as smartphones, laptops, smart speakers, and smart homes.

The rapid development of SRSs is inseparable from the advancement of Neural Networks (NNs), especially Deep Neural Networks (DNNs). While SRSs based on traditional methods, such as i-vector [6] and Gaussian Mixture Model (GMM) [7], have prospered for decades, they are being replaced by NN-based methods due to the strong ability of NNs. However, previous work has demonstrated that NNs are susceptible to adversarial attacks [8]. Adversarial attacks mean that an adversary utilizes adversarial examples, which are generated by adding small perturbations, i.e., adversarial perturbations, into clean samples, to make a machine learning model misbehave. Adversarial attacks were first conducted in the image field. Szegedy et al. [8] successfully fooled an image classification model using adversarial examples. After that, adversarial attacks have gained widespread attention in the image field and many effective attack methods, such as Fast Gradient Sign Method (FGSM) [9] and Basic Iterative Method (BIM) [10], have been proposed. Defense methods have also been extensively studied, such as feature squeezing [11] and adversarial training [9]. Akhtar et al. [12] and Yuan et al. [13] comprehensively reviewed existing adversarial attack methods and defense methods in the image field, respectively.

Inspired by the advancement of adversarial attacks in the image field, an increasing number of researchers pay their attention to adversarial attacks in the audio field. As the most widely used voice processing system, the speech recognition system was successfully deceived by adversarial examples in 2015 [14]. Three years later, Kreuk et al. [15] first successfully attacked an SRS through FGSM, which proves the effectiveness of adversarial attacks in SRSs. Since then, adversarial attacks and defenses in SRSs started to draw special attention.

Several researchers have surveyed the security of SRSs from different perspectives. Wu et al. [16] presented a study on spoofing attacks,

* Corresponding author at: State Key Laboratory on Integrated Services Networks, School of Cyber Engineering, Xidian University, China.
E-mail address: zyan@xidian.edu.cn (Z. Yan).

https://doi.org/10.1016/j.sysarc.2022.102526
Received 31 August 2021; Received in revised form 27 February 2022; Accepted 21 April 2022
Available online 30 April 2022

including impersonation, relay, speech synthesis, and voice conversion, and countermeasures in SVSs. However, adversarial attacks and defenses were not mentioned in their work. SISs were also missed. Das et al. [17] demonstrated the security vulnerabilities of SVSs from the perspective of attackers by considering both spoofing attacks and adversarial attacks. Nevertheless, they ignored SISs and did not consider how to evaluate the performance of different attack methods or defense methods. Abdullah et al. [18] researched adversarial attack methods on both speaker recognition and speech recognition comprehensively. However, they paid little attention to defense methods. Overall, a comprehensive review on the recent advance of adversarial attacks and defenses in SRSs shall be done.

In this paper, we fill the gap by providing a thorough survey on adversarial attacks and defenses in SRSs. We first illustrate the basics of SRSs. At the same time, we introduce several basic concepts and threat models related to adversarial attacks and defenses. After that, we propose taxonomies of attack methods and defense methods respectively. The former includes optimization-based attacks and signal processing-based attacks, and the latter involves proactive defenses and passive defenses. Based on classification and evaluation criteria, we comprehensively review attack methods and defense methods in SRSs and analyze their pros and cons. Finally, we figure out several unsolved issues and suggest future research directions. Specifically, the contributions of this paper are as follows:

- This paper proposes two sets of evaluation criteria for adversarial attacks and defenses in SRSs, respectively.
- This paper provides taxonomies of existing adversarial attacks and defense methods.
- Based on the taxonomies, this paper comprehensively reviews the existing attacks and defenses in SRSs, and evaluates them by employing our proposed criteria.
- This paper points out several open issues and suggests future research directions based on systematic reviews.

The rest of this paper is organized as follows. Section 2 gives a brief overview of SRSs and introduces the basic concepts related to adversarial attacks and defenses. In Section 3, we propose two sets of criteria for evaluating adversarial attacks and defenses, respectively. Section 4 provides the taxonomy of adversarial attacks in SRSs and conducts a thorough review of existing attacks, followed by the taxonomy of countermeasures and a thorough review on them in Section 5. In Section 6, we highlight open issues and propose future research directions. Finally, we conclude this paper in the last section.

2. Preliminaries

In this section, we first introduce the basics of SRSs followed by the definitions of the three important concepts in this paper, i.e., adversarial examples, adversarial perturbations and adversarial attacks. Finally, we list possible threat models in adversarial attacks and defenses.

2.1. Speaker recognition systems overview

Fig. 1 shows an overview of a typical SRS. The SRS includes three modules, i.e., preprocessing module, feature extraction module, and model inference module. Meanwhile, the lifecycle of an SRS involves three stages, i.e., training stage, enrollment stage, and recognition stage.

When audio is input, the preprocessing module first filters out background noise and high-frequency signals beyond the frequency range of human voices. Feature extraction algorithms are then used to generate a feature vector that reduces dimensions of the audio by capturing its most important features and characteristics. Various feature extraction algorithms have been proposed, such as Mel-Frequency Cepstral Coefficients (MFCC) [19], Spectral Subband Centroid (SSC) [20], and Perceptual Linear Predictive (PLP) [21]. Among them, MFCC is the most popular one in SRSs due to its ability to expose important acoustic features, similar to human ears. After that, the feature vector is passed to a model for either training or inferencing.

In the training stage, corpora are used to train the SRS and adjust system parameters to obtain a capable SRS. After that, multiple speakers enroll in the SRS. All enrolled speakers form a speaker group. The SRS calculates and stores a feature vector for every enrolled speaker, which is used in the recognition stage. In the recognition stage, the SRS is responsible for recognizing the identity of unknown input audio.

According to the difference of recognition tasks, SRSs can be divided into two categories: Speaker Verification and Speaker Identification. For an arbitrary input audio x, a Speaker Identification System (SIS) determines which enrolled speaker utters x. Regarding a Speaker Verification System (SVS), an unknown speaker not only needs to input his/her audio y but also needs to claim his/her identity. The SVS determines whether y is uttered by the claimed speaker, and the verification result is "yes" or "no".

2.2. Adversarial examples generation

Adversarial attack is a kind of attack method that an adversary generates adversarial examples to make a machine learning model misbehave. An *adversarial example* refers to specifically crafted input designed to look normal to humans but causes misbehaviors of a machine learning model [8]. Given an input x with its corresponding label y, and a well-trained machine learning model $f(\cdot)$, an adversarial example x' can be constructed as:

$$x' = x + \delta \wedge f(x', \theta) \neq y \wedge \|\delta\| < \epsilon$$

Here, δ is called *adversarial perturbation*, which is the noise that is added to a clean sample to make it an adversarial example [8]. θ is the parameter of the machine learning model $f(\cdot)$. The hyperparameter ϵ is used to control the maximum perturbation generated. Suppose $L(\cdot)$ is the loss function and y' is the target label, the adversarial perturbation δ can be calculated by

$$\begin{aligned} &\min L(f(x+\delta, \theta), y') \\ &s.t. \ \|\delta\| < \epsilon \end{aligned}$$

To solve the above formula, many attack methods have been proposed. We first introduce a classic and well-known method, FGSM [9]. The adversarial perturbation can be calculated by the following formula.

$$\delta = \eta sign(\nabla_x f(x, \theta))$$

Here, η is the magnitude of the perturbation. The adversarial example x' is calculated as: $x' = x + \delta$. Kurakin et al. [10] proposed BIM which iterates FGSM for multiple rounds. The adversarial example is generated in multiple iterations.

$$x'_0 = x$$

$$x'_{n+1} = Clip(x'_n + \eta sign(\nabla_x f(x'_n, \theta)))$$

Here, $Clip(\cdot)$ is a function which limits the change of the generated adversarial example in each iterations.

2.3. Threat models

In this subsection, we introduce threat models of attack methods and defense methods, respectively.

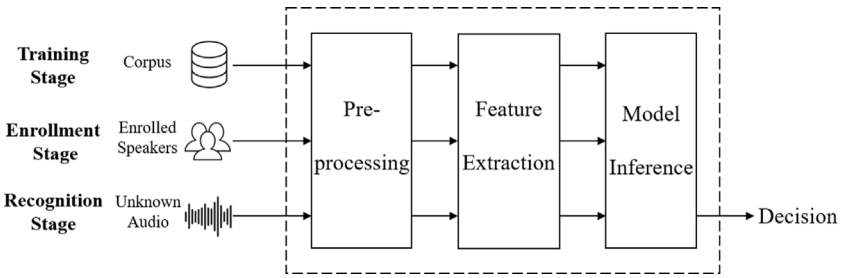

Fig. 1. An overview of a typical SRS.

2.3.1. Threat models of attack methods

The threat model of an attack method is related to the adversary's goal, which is also called adversarial specificity, and the adversary's knowledge about target attack models. Based on the adversary's goal, adversarial attacks can be categorized into targeted attacks and untargeted attacks. Based on the adversary's knowledge about target attack models, attacks can be divided into white-box attacks, gray-box attacks, and black-box attacks. The specific definitions are as follows.

Targeted Attacks: since an SVS only has two possible decisions (i.e., yes or no), we regard all adversarial attacks on SVSs as targeted attacks. In terms of an SIS, the targeted attack means a speaker is maliciously identified as a specific speaker.

Non-targeted Attacks: it means a speaker fails to be identified or is identified as any other speaker in an SIS.

White-box Attacks: the adversary has full knowledge about a target model, such as model architecture, parameters, gradients, layer outputs, input and output pairs, etc.

Gray-box Attacks: the adversary only knows part of knowledge about a target model.

Black-box Attacks: the adversary can only get the input and output pairs of a target model.

2.3.2. Threat models of defense methods

The threat model of a defense method is related to the adversary's knowledge about defense methods. Based on the adversary's knowledge about defense methods, attacks can be divided into adaptive attacks and non-adaptive attacks. Their specific definitions are as follows.

Adaptive Attacks: the adversary has full knowledge about a deployed defense method. Using this knowledge, the adversary can adjust its attack methods to overcome the deployed defense method. But a strong defense method can still counter this type of attack.

Non-adaptive Attacks: the adversary does not have any knowledge about a deployed defense method. Resisting non-adaptive attacks is the minimum requirement that a feasible defense method needs to meet.

3. Evaluation criteria

In this section, we propose two sets of evaluation criteria (As shown in Fig. 2). One is used to evaluate attack methods against SRSs, thus analyzing the performance of each attack method. The other is used for evaluating defense methods, thus analyzing the performance of each defense method.

3.1. Evaluation criteria for attack methods

In this subsection, we put forward evaluation criteria for attack methods from three aspects: practicability, imperceptibility, and effectiveness.

3.1.1. Practicability

Practicability refers to the ability of an attack method to be used in the real world. We introduce the following five metrics to evaluate the practicability of an attack method.

Transferability: it is the ability of an adversarial example to continue to make an impact on SRSs other than the one created it. In the real world, SRSs are usually black-box to adversaries. Thus, investigating transferable attack methods is meaningful since they can be used to attack different SRSs. According to the difference between two SRSs before and after transferring, transferability can be classified into cross-feature, cross-dataset, cross-model, etc. Cross-feature, cross-dataset, and cross-model indicate that two SRSs differ from feature extraction technologies, training datasets, and model frameworks, respectively.

Universality: it is the ability of an adversarial perturbation to fool a given model on any clean samples with high probability. If an adversary can generate a universal adversarial perturbation, he can obtain adversarial examples by adding the perturbation to any clean samples effortlessly, which helps to achieve real-time attacks.

Attack Media: adversarial examples can be fed into an SRS via different medium, each of which introduces different challenges such as background noise and distortion. There are three common attack media including over-line, over-air, and over-telephone-network. Over-line attacks indicate that a Waveform Audio file is directly fed to an SRS. They are the easiest to execute since over-line ensures lossless transmission. Over-air attacks indicate that an audio file played by a speaker is fed to an SRS. Although over-air attacks are more difficult to implement than over-line attacks since the quality of adversarial examples will decrease due to the background noise, attenuation, and multi-path effects during transmission, they are close to the real world. Over-telephone-network attacks are more difficult to implement than over-air attacks since adversarial examples not only pass through the air but also the telephone network. Serious signal processing operations, such as jitter and compression, will further reduce the quality of adversarial examples. To summarize, the overall difficulty of achieving adversarial attacks in the above media is over-line $<$ over-air $<$ over-telephone-network.

Distance: the farther the adversarial example travels in loss medium, the worse the audio quality. Therefore, distance is used to measure the farthest distance an adversarial example can spread without losing its ability to attack.

Commercial SRSs: it is used to evaluate whether an attack method can successfully attack commercial SRSs, such as Azure Verification API [22] and Azure Attestation API [23], which are more complex and have higher security requirements. Once a commercial SRS is attacked successfully, the security and privacy of its users and the reputation of the vendor will be severely damaged.

3.1.2. Imperceptibility

Imperceptibility indicates that adversarial perturbations impair utterances very slightly for human perception. Since if ordinary people can distinguish an adversary example and its corresponding clean

Fig. 2. Evaluation criteria of attack and defense methods.

sample, the adversarial example is likely to be discarded before entering SRSs. We introduce the following four metrics to evaluate the imperceptibility of adversarial perturbations.

Types of Adversarial Audio: depending on the attack type and scenario, adversaries can produce different types of adversarial audio. Adversarial audio can be categorized into three classes, including noisy, clean, and inaudible audio. Noisy audio sounds like noise to humans but is considered legitimate audio by SRSs. Clean audio is perturbed at such low intensity that human listeners cannot perceive these perturbations at all, even though there is a hidden command embedded in it. However, SRSs can detect and execute these embedded commands. Inaudible audio is generated by an adversary exploiting the characteristics of human auditory system. On the one hand, the human auditory system can only perceive frequency that ranges from 20 Hz to 20 kHz. Therefore, audio whose frequency is beyond 20 kHz cannot be heard by humans but can be recognized by SRSs. On the other hand, frequency masking which refers to the phenomenon that one faint but audible sound becomes inaudible in the presence of another louder audible sound can also be used to generate inaudible adversarial audio. To summarize, the imperceptibility in the above audio types is noisy audio<clean audio<inaudible audio.

Perturbation Norm: it indicates the restricted l_p-norm of the perturbations to make them imperceptible. l_2 and l_∞ are two commonly used metrics. l_2 measures the Euclidean distance between the clean sample and the corresponding adversarial example, and l denotes the maximum change direction between the clean sample and the corresponding adversarial example.

Human Perception: it is used to evaluate whether human perception has been considered, in other words, whether user studies (e.g., ABX test) have been implemented in a paper. Human perception is the most direct metric to measure the imperceptibility of adversarial examples.

Signal-to-noise Ratio (SNR): it is defined as the ratio of signal power to the noise power and is often expressed in decibels (dB). The noise is the adversarial perturbations. A larger SNR value indicates a smaller perturbation.

3.1.3. Effectiveness

Effectiveness is used to quantitatively measure the performance of an attack method. We introduce the following six metrics to evaluate the effectiveness of attack methods. Generation time and attack success rate are direct metrics to measure the effectiveness of an attack method. The remaining four metrics are indirect metrics since they are used to evaluate the performance of an SRS. The difference between these four metrics before and after the SRS is attacked can be used to measure the effectiveness of an attack method.

Generation Time: it refers to the time required to generate an adversarial example. The less the generation time, the more efficient the attack method.

Attack Success Rate (ASR): for targeted attacks, ASR refers to a proportion of adversarial examples that are recognized as the targeted speaker. For non-targeted attacks, ASR indicates a proportion of adversarial examples that are recognized as other speakers rather than the original speaker.

Recognition Accuracy (RA): it is a proportion of utterances correctly recognized by an SRS.

False-positive Rate (FPR): it is a proportion of utterances of non-original speakers that are recognized as the original speaker.

False-negative Rate (FNR): it is a proportion of utterances of original speakers that are recognized as non-original speakers.

Table 1
Comparison of adversarial attacks against SRSs.

Ref	Type	Task	AK-M	AS	Practicability					Imperceptibility				Effectiveness				Attack Method
					Tr	Un	AM	Dis (m)	CS	ToA	PN	HP	SNR (dB)	GT	BA	AA	Met(%)	
[15]	O	SVS	W	Ta	?	×	L	-	×	C	l_∞	✓	?	?	87.5/4.88	25.75/94.63	RA/FPR	FGSM
			B		✓										81.55/12	58.93/46		
[24]	O	SVS	W	Ta	?	×	L	-	×	C	l_∞	✓	?	?	7.2/7.2	96.87/97.64	FPR/EER	FGSM
			B		✓										6.62	74.32	EER	
[25]	O	SIS	G	Uta	?	✓	L	-	×	C	l_2	?	44.13	Real-time	-	97.5	ASR	Generative Network
				Ta									48.53			97.2		
[26]	O	SIS	W	Ta	?	✓	L, A	?	×	C	?	?	9.42	real-time	–	90.19	ASR	Optimization + RIR
[27]	O	SIS	W	Ta	?	✓	L, A	1.6–3	✓	C	l_2	?	8.3	real-time	–	96.9	ASR	AdvPulse
[28]	O	SIS, SVS	B	Ta	✓	×	L, A	1	✓	C	l_∞	✓	30.2	minutes	–	99	ASR	FakeBob
[29]	SP	SIS, SVS	B	Ta	✓	×	L, A	0.3	✓	N	–	?	?	seconds	–	100	ASR	TDI, RPG, HFA, TS
[30]	SP	SIS	B	Uta	✓	×	L, TN	–	✓	C	L_2	✓	?	seconds	–	10–20	ASR	Filtering Out Low-intensity Components
[31]	SP	SIS	W	Ta	?	×	L	–	✓	In	L_∞	✓	34.111	?	–	93.8	ASR	Frequency Masking

AK-M: Adversary's Knowledge about Models; AS: Adversarial Specificity; Tr: Transferability; Un: Universality; AM: Attack Media; Dis: Distance; CS: Commercial SRSs; ToA: Types of Adversarial Audio; PN: Perturbation Norm; HP: Human Perception; SNR: Signal-to-noise Ratio; GT: Generation Time; BA: Before Attack; AA: After Attack; Met: Metrics; O: Optimization; SP: Signal Processing; W: White-box; G: Gray-box; B: Black-box; Ta: Targeted; Uta: Untargeted; L: Line; A: Air; TN: Telephone Network; C: Clean; N: Noisy; In: Inaudible; ✓: satisfied; ×: not satisfied; ?: not discussed; -: not available.

Equal Error Rate (EER): it refers to the rate when FPR is equal to FNR. The lower the EER, the greater the performance of an SRS.

3.2. Evaluation criteria for defenses

From the perspective of practicability and effectiveness, we propose a set of criteria to evaluate the performance of defense methods in SRSs.

3.2.1. Practicability

Practicability refers to the ability of a defense method to be used in the real world. We introduce the following three metrics to evaluate the practicability of a defense method.

Generality: it is the ability of a defense method to resist different attack methods. In the real world, designers of SRSs usually cannot know which attack methods will be used by adversaries in advance. Therefore, the more attack methods a defense method can resist, the more practical it is. Based on the number of attack methods that can be resisted, we divide defense methods into three levels: high-, medium-, and low-generality. Firstly, highly general defense methods can resist any attack methods. Because attack–agnostic defense methods do not rely on any attack methods, they can be used to resist any attack methods theoretically. Therefore, attack–agnostic defense methods are high-generality. Secondly, in addition to attack–agnostic defense methods, there are some defense methods that rely on attack methods. For example, a defense method needs adversarial examples, which can be generated by FGSM or other attack methods, to adjust parameters. In terms of them, the transferable defense method is medium-generality since it can not only detect adversarial examples generated by the attack method it depends on, but also detect adversarial examples generated by other attack methods. At last, the non-transferable defense method, which can only resist adversarial examples generated by the attack method it depends on, is low-generality.

Defense Media: similar to attack media, there are three common defense media including over-line, over-air, and over-telephone-network. In Section 3.1.1, we analyze that the overall difficulty of achieving adversarial attacks in the above media is over-line < over-air < over-telephone-network. The overall difficulty of achieving defenses in the above media is over-line > over-air > over-telephone-network.

Defendable Attacks: for a defense method, its corresponding defendable attacks refer to the attacks that can be resisted, as proved by experiments. The more attacks that can be resisted, the more effective the defense method is.

3.2.2. Effectiveness

Effectiveness is used to quantitatively measure the performance of a defense method. Six metrics are introduced to evaluate the effectiveness of a defense method. Defense time and detection accuracy are direct metrics, while, the same as Section 3.1.3, the remaining four metrics, including RA, FPR, FNR, and EER, are indirect metrics. We can evaluate the ability of a defense method through the difference between the four metrics before and after deploying the defense method. We do not repeat the definition of RA, FPR, FNR, and EER in this subsection.

Defense Time: it refers to the time required for a defense method to detect or purity an input utterance.

Detection Accuracy (DA): it refers to the ratio of correct discrimination on adversarial examples. The higher the DA is, the more effective the defense method is.

4. Adversarial attacks against SRSs

In this section, we first propose a taxonomy of existing adversarial attack methods against SRSs. Then, we review them based on the taxonomy. Furthermore, we evaluate and compare these attack methods (as shown in Table 1) with our proposed criteria in Section 3.1.

4.1. Taxonomy of adversarial attacks against SRSs

Adversarial attacks against SRSs can be categorized as optimization-based attacks and signal processing-based attacks. The optimization-based attacks generate adversarial examples by solving an optimization problem that is obtained by formalizing the purpose of adversarial attacks. In recent years, several optimization-based adversarial example generation algorithms have been proposed, such as FGSM [9] and BIM [10]. The signal processing-based attacks use signal processing techniques to generate adversarial examples. Although these attacks do not directly target machine learning models embedded in SRSs, they can still force machine learning models to misbehave.

4.2. Optimization-based attacks

Inspired by the success of adversarial attacks in computer vision and speech recognition, Kreuk et al. [15] tried to fool a DNN-based SVS by adversarial examples for the first time. Adversarial examples were

generated by FGSM, which is l_∞ perturbation norm. They verified the effectiveness of adversarial attacks in SVSs by a white-box attack. Then they deployed two black-box attacks to explore the transferability of adversarial examples they generated. Experimental results show that the adversarial examples are transferable in both cross-feature and cross-dataset settings. Finally, through human perception experiments, they found that human listeners cannot distinguish between clean samples and adversarial examples, which reflected the imperceptibility of adversarial perturbations generated by FGSM. However, since the purpose of this work was to verify that SVSs are susceptible to adversarial examples, the authors did not make efforts to pursue a better attack effect and did not attack commercial SRSs. In addition, they ignored generation time and over-air attacks, which are more practical than over-line attacks.

Li et al. [24] found that [15] has two limitations. On the one hand, [15] studied the impact of adversarial attacks on DNN-based SVSs, but ignored GMM-based SVSs. On the other hand, cross-model transferability was not discussed. To overcome the two limitations, they deployed a white-box attack to verify that GMM-based SVSs are also subject to adversarial examples. They then implemented three black-box experiments to study the transferability of adversarial examples generated by FGSM. The results show that adversarial examples are transferable in cross-feature, cross-dataset, and cross-model settings. Finally, they demonstrated the imperceptibility of adversarial examples they generated to human listeners by human perception experiments. Although [24] solved two limitations of [15], the attack performance of [24] was not good enough since it generated adversarial examples by simple FGSM. In addition, they did not discuss generation time and over-air attacks.

The aforementioned work cannot achieve real-time attacks since they generated different adversarial perturbations for different audio, which usually costs a lot of time. Real-time attacks are inseparable from universal adversarial perturbations since universal adversarial perturbations can be added to any samples to generate effective adversarial examples. If an adversary obtains universal adversarial perturbations in advance, it only needs one addition operation to generate an adversarial example, which is time-efficient. Li et al. [25] tried to generate universal adversarial perturbations for SincNet [32], a state-of-the-art speaker identification model, with a generative network, which can learn the mapping from a low-dimensional normal distribution to a universal adversarial perturbation subspace. They conducted both untargeted attacks and target attacks against SincNet [32] in gray-box settings. The results show the existence of universal adversarial perturbations. In addition, they studied the imperceptibility of universal adversarial perturbations by SNR. However, they did not deploy over-air attack experiments and did not try to attack commercial SRSs.

Xie et al. [26] almost simultaneously conducted similar work. While Li et al. [25] generated universal perturbations with a generative network, Xie et al. [26] generated universal perturbations with a conventional optimization-based approach. They attacked an SIS in white-box settings successfully. After that, they tried to attack the SIS in the real world, i.e., over-air attacks. However, they failed because of the attenuation and multi-path effects of sound in the propagation process. They then assumed that the adversary has the knowledge of the room's layout and took Room Impulse Response (RIR) into consideration to enhance the robustness of adversarial examples. The results show that considering RIR can greatly increase the ASR in over-air attacks. However, the high attack success rate came at the cost of low SNR. That means they need to make a trade-off between the attack effect and the imperceptibility of the adversarial perturbations.

Although [25,26] generated universal adversarial perturbations successfully, Li et al. [27] found that they only focused on static-speech attack scenarios, but ignored streaming-speech attack scenarios. In static-speech attack scenarios, an adversary should obtain universal adversarial perturbations and clean samples to generate adversarial examples before feeding adversarial examples into an SRS. On the contrary, in streaming-speech scenarios, an SRS takes streaming audio inputs (e.g., live human speech) and an adversary can fool the SRS by playing universal adversarial perturbations through a nearby loudspeaker. There is no doubt that streaming-speech scenarios are closer to the real world than static-speech scenarios. Therefore, Li et al. [27] designed AdvPulse, a method to generate a subsecond-level adversarial perturbation that can be added at any point of a streaming audio input to launch targeted adversarial attacks. In other words, they generated universal adversarial perturbations against SRSs in streaming-speech attack scenarios. In addition, when a loudspeaker and an SRS are less than three meters apart, adversarial examples they generated can successfully deceive the SRS. This work is advanced. Unfortunately, they only considered white-box settings but ignored black-box settings.

Chen et al. [28] proposed an attack method named Fakebob, which formalizes the generation of adversarial examples into an optimization problem. In this optimization problem, a score threshold and the strength of adversarial perturbations are considered. To solve the optimization problem, they proposed an approach that utilizes a novel algorithm to estimate the threshold, a natural evolution strategy (NES) to estimate gradient, and finally the BIM method is applied to generate adversarial examples. This work is currently the most comprehensive one. It has four main contributions. Firstly, they effectively implemented targeted attacks on SRSs in black-box settings and the ASR can reach 99%. Secondly, they considered all possible SRSs based on different tasks, including SVSs and SISs. Thirdly, they did many experiments, including over-line and over-air experiments, on both open source systems and commercial systems, which prove the transferability and practicability of Fakebob. Fourthly, they tried four defense methods to defend against Fakebob. Experimental results show that Fakebob still affects the performance of victim systems, which demonstrates the robustness of Fakebob. Meanwhile, they employed human perception experiments to explore the imperceptibility of adversarial perturbations. However, FakeBob took several minutes to generate an adversarial example, which limits its wide use in the real world.

4.3. Signal processing-based attacks

Abdullah et al. [29] noticed that different audio samples may have the same feature vector when being transformed by acoustic feature extraction algorithms (e.g., MFCC) and nearly all SRSs appear to rely on several feature extraction algorithms. Based on this knowledge, they successfully attacked a commercial SVS and a commercial SIS in black-box settings. They first obtained desired audio (e.g., *OK, Google* uttered by Alice). Then they designed four methods to obfuscate the desired audio as much as possible to generate obfuscated audio, i.e., adversarial examples. The four methods include Time Domain Inversion (TDI), Random Phase Generation (RPG), High Frequency Addition (HFA), and Time Scaling (TS). Although this work achieved efficient attacks since an adversarial example could be generated in a few seconds, adversarial examples it generated were noise in human perception which were easy to be noticed by human listeners.

Abdullah et al. [30] designed an attack method to circumvent surveillance in telephone networks. They assumed that SISs rely on the components of audio that are non-essential for human comprehension. In order to find out the non-essential components, they first split audio into phonemes. Then, signal decomposition algorithms were used to decompose the phoneme into individual components and corresponding intensities. Since the low-intensity components are less perceptible to humans, the low-intensity components are likely what is looked for. Therefore, adversarial examples are generated by filtering out low-intensity components of every phoneme. They employed an untargeted attack in black-box settings to prove the effectiveness of the attack method they proposed. The results show that when only one phoneme in an utterance is perturbed, the ASR of the utterance can reach 10%–20% for most phonemes. An adversary can perturb multiple phonemes in an utterance to achieve a high ASR. They further demonstrated that

the attack method is transferable in a cross-model setting. Although they achieved untargeted attacks and over-telephone-network attacks, most adversaries cannot obtain the right to monitor the telephone network. Targeted attacks and over-air attacks against intelligent voice assistants embedded in smart devices are the mainstream of adversarial attacks against SRSs now and even in the future.

While [30] successfully achieved untargeted attacks by removing low-intensity components of clean samples, Wang et al. [31] tried to achieve target attacks by adding additional components to clean samples. They also aimed to generate inaudible adversarial perturbations, instead of maintaining a slight noise to the clean sample. Inspired by previous work [33,34] on speech recognition systems, they first obtained the desired audio (e.g., *OK, Google*) as adversarial perturbations. Then, they leveraged frequency masking, which refers to the phenomenon that one faint but audible sound becomes inaudible in the presence of another louder audible sound, to hide adversarial perturbations in normal audio, such as birdsong and white noise. They successfully attacked a DNN-based SIS by inaudible adversarial audio in human perception. Unfortunately, they deployed experiments in white-box settings rather than more practical black-box settings. In addition, they did not explore the transferability of adversarial examples they generated and did not explore over-air attacks.

4.4. Comparison and discussion

In Section 4, we comprehensively review adversarial attack methods in SRSs and compare them in Table 1. Based on Table 1, we conclude our review as below.

Among all the works we reviewed in this section [15,24–31], the methods based on optimization [15,24–28] account for two-third, while the methods based on signal processing [29–31] account for one-third. The first type of method can be transferred from computer vision directly since they generate adversarial examples from digital vectors rather than raw audio or images. For example, Goodfellow et al. [9] proposed FGSM to deceive an image classification model, while Kreuk et al. [15] and Li et al. [24] also used FGSM to fool SRSs successfully. The methods based on signal processing cannot be transferred in different domains directly since they leverage acoustic signal processing techniques to generate adversarial examples. For example, Abdullah et al. [30] generated adversarial examples by filtering out low-intensity components in raw audio. Obviously, this type of method cannot be used to deceive image processing tasks.

We observe that researchers not only focus on simple white-box attacks but also pay attention to black-box attacks. More than half of reviewed works [15,24,28–30] explored the effectiveness of attack methods in black-box settings. In addition, all reviewed works except [30] achieved targeted attacks, which are more difficult than untargeted attacks.

Three of all reviewed works [25–27] generated universal adversarial perturbations, which help adversaries to achieve practical real-time attacks. More than half of reviewed methods [15,24,28–30] proved that adversarial examples they generated are transferable. Transferable adversarial examples can not only deceive the SRS that generates them, but also deceive other SRSs. In other words, transferable adversarial examples can deceive multiple SRSs, while non-transferable adversarial examples can only deceive the SRS that generates them. Therefore, transferable adversarial examples are more practical in the real world. In addition, four of all reviewed studies [26–29] considered over-air attacks and deployed experiments to explore over-air attacks. Among them, the effective attack distance of adversarial examples generated by [29] is only 0.3 m, while the effective attack distance of adversarial examples generated by [27] can reach 3 m. More than half of reviewed attack methods [27–31] can successfully attack commercial SRSs.

Adversarial examples generated by the optimization-based attack methods [15,24–28] sound clean to humans. Adversarial examples generated by the three attack methods based on signal processing [29–31] are clean, noisy, and inaudible to humans, respectively. Additionally, more than half of reviewed works [15,24,28,30,31] deployed human perception experiments to explore the imperceptibility of adversarial perturbations. At the same time, four reviewed works [25–28,31] measured SNR of adversarial examples to quantitatively represent the relationship between audio and perturbations in adversarial examples.

Three of all reviewed works [25–27] achieve real-time attacks since they generate universal adversarial perturbations. The optimization-based attack method proposed in [28] takes several minutes to generate an adversarial example, while the signal processing-based attack methods proposed in [29,30] only take several seconds. This indicates that the methods based on signal processing are more time-efficient than the optimization-based attack methods.

5. Defenses against SRSs

In this section, we first propose a taxonomy of existing defenses against SRSs. After that, we review and evaluate some defense methods against adversarial attacks in SRSs (as shown in Table 2) by applying the criteria proposed in Section 3.2.

5.1. Taxonomy of defenses against SRSs

There are two types of defense methods against adversarial attacks: (1) proactive defenses, (2) passive defenses. Proactive defense methods employ adversarial data augmentation to retrain original models such that they can be robust to adversarial examples. Passive defense methods defend against adversarial attacks by adding new components rather than modifying original models. According to the function of new components, passive defense methods can be divided into detection methods and purification methods. When an adversarial example is identified, a detection method aims to refuse it to enter systems, while a purification method aims to feed it to systems after removing adversarial perturbations.

5.2. Proactive defenses

Wang et al. [35] proposed adversarial regularization based on adversarial examples to defend against adversarial attacks. Adversarial regularization aims to seek the worst sample around an input sample and then use the worst sample to optimize an SRS. They used adversarial examples generated by FGSM and virtual adversarial training based on local distributional smoothness (LDS) to attack a DNN-based SVS. It is worth noting that virtual adversarial training can calculate adversarial perturbations for unlabeled samples. To the best of our knowledge, this is the first work to apply virtual adversarial training into SVSs. After that, they leveraged FGSM and LDS to regularize the SVS, respectively. They showed that adversarial regularization is a medium-general defense method through several experiments. However, the performance of adversarial regularization is not good enough, since the EER of the regularized SVS only decreased slightly. In addition, this work did not consider adaptive attacks and did not mention defense time.

Many previous works have explored defense methods to resist spoofing attacks, including synthesis, convert and replay attacks, for SVSs. However, spoofing countermeasure models are still vulnerable to adversarial attacks. To address this issue, Wu et al. [36] proposed two defense methods, one is proactive, i.e., adversarial training, and the other is passive, i.e., spatial smoothing, to improve the robustness of SVS spoofing countermeasure models. We will introduce spatial smoothing in Section 5.3. Adversarial training refers to a defense method that uses adversarial data augmentation to retrain the model to enhance the robustness of the model. They retrained a DNN-based SVS by adversarial data augmentation which was generated by Projected Gradient Descent (PGD) method. They proved that the performance of adversarial training is better than spatial smoothing through several experiments, which is because adversarial training is model-specific. However, adversarial training is low-generality. In addition, defense time and adaptive attacks were not mentioned in the paper.

Table 2
Comparison of defense methods against SRSs.

Ref	Type	Task	AK-DM	Practicability			Effectiveness					Defense method
				Ge	DM	DfA	DT	Ori	AA	AD	Met (%)	
[35]	Pro	SVS	Nad	Me	L	FGSM, LDS	?	4.87	11.89	8.31	EER	Adversarial regularization
[36]	Pro	SVS	Nad	Lo	L	PGD	?	99.99	37.06	98.60	DA	Adversarial training
[37]	P-D	SVS	Nad	Me	L	BIM, JSMA	?	5.97/ -	39.87/ -	0.18/ 99.83	EER/ DA	Separate detection network
[38]	P-D	SVS	Nad	Hi	L	BIM	?	2.24/ 2.56	71.83/ 74.92	10.66/ 24.68	FPR/ FNR	Voting for the right answer
			Ad							13.29/ 27.75		
[39]	P-D	SVS	Nad	Hi	L	BIM	?	2.88/ -	99.33/ -	-/ 99.76	EER/ DA	Neural vocoders
[36]	P-P	SVS	Nad	Hi	L	PGD	?	99.99	48.32	93.95	DA	Spatial smoothing
[40]	P-P	SVS, SIS	Nad	Me	L	FGSM, MT, PGD, DDN	?	0.89	13.81	3.62	EER	Adversarial Separate network
[41]	P-P	SVS	Nad	Hi	L	BIM	?	8.87	66.02	22.94	EER	Cascaded self-supervised learning models
			Ad							40.69		

AK-DM: Adversary's Knowledge about Defense Methods; Ge: Generality; DM: Defense Media; DfA: Defendable Attacks; DT: Defense Time; Ori: Original; AA: After Attack; AD: After Defense; Met: Metrics; Pro: Proactive; P-D: Passive-Detection; P-P: Passive-Purification; Ad: Adaptive; Nad: Non-adaptive; Hi: High; Me: Medium; Lo: Low; L: Line; ?: not discussed; -: not available.

5.3. Passive defenses

Passive defense methods utilize new components to detect or purify adversarial examples. In this subsection, we first review detection methods, followed by a review of purification methods.

5.3.1. Detection methods

Although adversarial training is effective, it is difficult to obtain adversarial data augmentation since we need to label every adversarial example. Inspired by [42,43], Li et al. [37] made the first attempt to defend SVSs against adversarial attacks with a separate detection network which is a VGG-like network structure. The separate detection network not only avoids retraining well-developed SVSs but also can combine with countermeasures against spoofing attacks to obtain a powerful defense method. They first respectively adjusted parameters of two separate detection networks using adversarial examples generated by BIM and Jacobian-based Saliency Map Attack (JSMA) to obtain two different separate detection networks, i.e., a BIM-based separate detection network and a JSMA-based separate detection network. Then, they proved that the BIM-based separate detection network can not only detect adversarial examples generated by BIM but also detect adversarial examples generated by JSMA to a certain extent. Similarly, the JSMA-based separate detection network can also detect adversarial examples generated by BIM to a certain extent. In other words, the defense method they proposed, i.e. the separate detection network, is effective and medium-general. However, they did not consider adaptive attacks, which are more challenging, and also did not mention defense time.

The above defense methods [35–37] require knowledge of the attack methods used by adversaries. However, it is impractical for SRSs designers to know which attack methods will be implemented by adversaries in advance. Therefore, Wu et al. [38] proposed a highly general defense method called voting for the right answer. It means that whether an input utterance is accepted by the SVS is determined by the similarity between the input utterance and the enrollment utterance and the similarities between the enrollment utterance and neighbors of the input utterance which are some samples randomly selected around the input utterance. As its name suggests, this method means that the input utterance and its neighbors are voting on whether to accept the input. They used adversarial examples generated by BIM to attack a DNN-based SVS in white-box settings and considered both adaptive attacks and non-adaptive attacks. Although the proposed defense method is simple and effective, there are still two issues: (1) the performance of the defense method is related to some parameters that are difficult to select; (2) during defense, the SVS needs to run many times for every sample since we need to calculate the similarities between its neighbors and the enrollment utterance, which is time-consuming; (3) they did not discuss defense time quantitatively.

Wu et al. [39] also proposed a highly general defense method. They leveraged Parallel WaveGan, a neural vocoder, to re-synthesize the input utterance, and then used the difference between the SVS scores, i.e., the similarity with the enrollment utterance, for the input and re-synthesized utterance to determine whether the input utterance is an adversarial example. Since neural vocoders can purify adversarial perturbations, the large difference in SVS scores indicates the input utterance is an adversarial example. This is the first work to adopt neural vocoders as shields to detect adversarial examples for SVSs, and it showed neural vocoders are effective to detect adversarial examples by several experiments. Meanwhile, this work also clarified by experiments that the defense method slightly affected clean samples. However, it did not analyze adaptive attacks and defense time.

5.3.2. Purification methods

Inspired by [44], Wu et al. [36] proposed a highly general defense method based on spatial smoothing. The reason is that implementing smoothing does not need extra training efforts. In image processing, spatial smoothing uses nearby pixels to smooth the central pixel. According to different weighting mechanisms of nearby pixels, spatial smoothing can be divided into many categories, such as median filter, mean filter, and Gaussian filter. The authors leveraged these filters for SVSs, and employed experiments to prove the effectiveness of spatial smoothing. They further explored the combination of spatial smoothing and adversarial training, which achieves better performance. However, they ignored adaptive attacks and did not analyze defense time.

Zhang et al. [40] designed an adversarial separation network (AS-Net) to defend against adversarial attacks in SRSs. AS-Net aims to eliminate adversarial perturbations and restore natural clean utterances. Two optimized components, including compression structure and speaker quality loss, are introduced. The former is responsible for reconstructing adversarial perturbations, and the latter supervises whether the restored utterances generated by AS-Net are correctly

labeled by the target SRS. They deployed a lot of experiments to defend against FGSM, PGD, Decoupled Direction and Norm (DDN) and Momentum attack (MT), which show the effectiveness and medium-generality of AS-Net in DNN-based SRSs. In addition, they compared the performance of different countermeasures, including adversarial training, feature-squeezing, and AS-Net. The results show that AS-Net significantly outperformed other countermeasures. However, they did not consider defense time and ignored adaptive attacks.

Wu et al. [41] proposed a highly general defense method based on cascaded self-supervised learning models, which possesses the ability to mitigate superficial perturbations in the input utterance after pre-training. Transformer encoder representations from alteration (TERA) as an advanced self-supervised learning model was used to construct the defense method. They generated adversarial examples by BIM to attack a DNN-based SVS. It is worth noting that they considered both adaptive attacks and non-adaptive attacks. The results show the defense method is effective on both of them to some extent and it is more difficult to defend against adaptive attacks. However, the experimental results also show that the defense method they proposed has a negative impact on clean samples. In addition, defense time was neglected.

5.4. Comparison and discussion

In Section 5, we comprehensively review the existing works about defense methods [35–41] for adversarial attacks in SRSs. Meanwhile, we compare all the defense methods reviewed in this section in Table 2. Based on Table 2, we summarize our review as below.

Among all the studies reviewed in this section, passive defense methods [36–41] are distinctly overwhelming with three-quarters of all reviewed papers. Passive defense methods are so popular since they can be deployed in any SRSs to defend against adversarial attacks. However, proactive defense methods [35,36] cannot be transferred between different SRSs since they are model-specific, which makes them receive less attention than passive defense methods.

We observe that all reviewed works deployed experiments to defend against non-adaptive attacks, while only two works [38,41] try to defend against adaptive attacks that are more threatening than non-adaptive attacks.

Half of reviewed defense methods [36,38,39,41] are highly general. Highly general defense methods are favored by researchers since they can defend against any attack methods theoretically. In addition, all reviewed works considered and deployed over-line attacks. Over-line ensures lossless transmission of adversarial examples. Therefore, defense methods that can defend against over-line attacks can also defend against over-air and over-telephone-network attacks.

All reviewed works employ experiments to show the effectiveness of defense against conventional attack methods, such as FGSM, BIM, and PGD. However, they do not defend against some advanced attacks, such as AdvPluse [27] and FakeBob [28]. At last, defense time, an important evaluation criterion of practicability, is missed in discussion in all reviewed works.

6. Open issues and future directions

6.1. Open issues

By reviewing and comparing the above literature with our proposed criteria, we figure out several open issues for adversarial attacks and defenses in SRSs.

First, it is difficult to enhance the robustness of SRSs. Adversarial training, which utilizes adversarial data augmentation to retrain SRSs, is an effective way to enhance the robustness of SRSs. However, obtaining adversarial data augmentation, i.e., adversarial example and its true label pairs, is time-consuming since we need to manually label each adversarial example. Therefore, how to enhance the robustness of SRSs efficiently is still an open and tough issue.

Second, it is not convenient to directly compare the performance of attack methods or defense methods proposed in different works. This is caused by the differences in experimental settings and evaluation metrics applied in different works. For example, the defense methods proposed in [38,39] were used to defend against adversarial examples generated by BIM. However, we cannot directly compare the performance of them since [38,39] used different evaluation metrics, i.e., [38] used FPR and FAR, and [39] used EER and DA. Similarly, we cannot compare the performance of defense methods proposed in [40,41] since they were used to defend against adversarial examples generated by different attack methods. All in all, the differences in experimental settings and evaluation metrics among different works hinder researchers from comparing the performance of attacks methods and defense methods in a direct way. Uniform evaluation metrics or criteria should be defined and adopted.

Third, signal processing-based attacks receive little attention. On one hand, although signal processing-based attacks are more efficient than optimization-based attacks as discussed in Section 5.4, to the best of our knowledge, only three articles [29–31] raise signal processing-based attacks in SRSs, which are far less than optimization-based attacks. On the other hand, all defense methods are proposed to defend against adversarial examples generated by optimization-based attacks, such as FGSM and BIM, while ignoring signal processing-based attacks. We cannot judge if current defense methods can defend against signal processing-based attacks. In short, signal processing-based attacks have not been paid sufficient attention in the current literature.

Fourth, the literature still lacks research on poisoning attacks against SRSs. The poisoning attacks refer to adding malicious data into training data, resulting in a biased model. Previous work has shown that poisoning attacks can seriously threaten the security and privacy of ML models in the image field [45–47]. However, as the structure of SRSs is more complex than image processing systems, little work pays attention to poisoning attacks against SRSs.

6.2. Future directions

We suggest several future research directions motivated by the above open issues as below.

Firstly, applying virtual adversarial training [48] to enhance the robustness of SRSs is worthy of deep-insight research. Virtual adversarial training can relieve the pressure of labeling adversarial examples since it retrains SRSs in semi-supervised settings. In addition, virtual adversarial training has low computational costs and a small number of hyperparameters. Therefore, virtual adversarial training in SRSs may be an interesting attempt to enhance the robustness of SRSs.

Secondly, unified adversarial attack and defense evaluation frameworks should be established. Specifically, the attack evaluation framework should clarify target SRSs, which include both open source systems and commercial systems, and attack evaluation metrics as shown in Section 3.1.3. The defense evaluation framework should include defendable attacks, which are used to generate adversarial examples for evaluating the performance of defense methods, and defense evaluation metrics as shown in Section 3.2.2. To comprehensively evaluate the performance of defense methods, both optimization-based attacks and signal processing-based attacks should be included into defendable attacks. In short, establishing unified frameworks for both attack and defense is an effective way to help researchers compare the performance of different methods of adversarial attack and defense to stimulate mutual development. Thirdly, the research on the security of preprocessing module and feature extraction module should be strengthened. As shown in Fig. 1, the SRS includes three modules, i.e., preprocessing module, feature extraction module and model inference module. The structure of SRSs is more complex than image recognition systems due to the addition of preprocessing and feature extraction. Each module introduces a surface of attacks, causing

exploitable vulnerabilities from the perspective of an adversary. Researchers in the audio field proposed some signal processing-based attacks that utilize vulnerabilities of the preprocessing module or the feature extraction module, such as those mentioned in [29–31] for SRSs and in [14,49] for speech recognition systems. However, there are still many unknown vulnerabilities. Therefore, we recommend strengthening the vulnerability mining of preprocessing module and feature extraction module. Correspondingly, defense methods should also be studied to enhance the robustness of these two modules. We believe that such an arms race will promote the security of SRSs.

Finally, we suggest studying poisoning attacks against SRSs and corresponding defense methods. On one hand, state-of-the-art SRSs require a huge amount of training data and it is common to collect these data from potentially untrustworthy sources (e.g., edge devices). Therefore, it is easy to poison the training dataset of SRSs. On the other hand, federated learning is a popular method to train ML models in a somehow privacy-preserving way, including ML models used in SRSs. However, it is difficult to guarantee that each party participating in federated learning is honest and trustworthy. A malicious party may deliberately use poisoned data for model training resulting in security and privacy threats. Therefore, it is interesting and promising to research poisoning attacks against SRSs and corresponding defense methods, especially in the context of federated learning, as well as other learning models.

7. Conclusion

In this paper, we overviewed the adversarial attacks and attack countermeasures in SRSs. We proposed two sets of criteria to evaluate the performance of adversarial attacks and defense methods. Based on our proposed taxonomies of existing adversarial attacks and defense methods, we reviewed existing adversarial attacks and defense methods by employing our proposed criteria, respectively. Through thorough review and analysis, we figured out several open research issues and highlighted future research directions to motivate the research of SRSs security.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China under Grant 62072351; in part by the Academy of Finland under Grant 308087, Grant 335262, Grant 345072 and Grant 350464; in part by the Open research project of ZheJiang Lab under grant 2021PD0AB01; in part by the Shaanxi Innovation Team Project under Grant 2018TD-007; and in part by the 111 Project, China under Grant B16037.

References

[1] R. Zhang, Z. Yan, A survey on biometric authentication: Toward secure and privacy-preserving identification, IEEE Access 7 (2019) 5994–6009.

[2] Z. Yan, S. Zhao, A usable authentication system based on personal voice challenge, in: International Conference on Advanced Cloud and Big Data, CBD, Chengdu, China, August 13-16, 2016.

[3] T.F. Zheng, L. Li, Robustness-Related Issues in Speaker Recognition, Springer, 2017.

[4] X. Wang, Z. Yan, R. Zhang, P. Zhang, Attacks and defenses in user authentication systems: A survey, J. Netw. Comput. Appl. 188 (2021) 103080.

[5] R. Zhang, Z. Yan, X. Wang, R. Deng, VOLERE: Leakage resilient user authentication based on personal voice challenges, IEEE Trans. Dependable Secure Comput. (2022).

[6] N. Dehak, R. Dehak, P. Kenny, N. Brümmer, P. Ouellet, P. Dumouchel, Support vector machines versus fast scoring in the low-dimensional total variability space for speaker verification, in: 10th Annual Conference of the International Speech Communication Association, INTERSPEECH, Brighton, United Kingdom, September 6-10, 2009.

[7] C.E. Rasmussen, The infinite Gaussian mixture model, in: Advances in Neural Information Processing Systems, NIPS, Denver, Colorado, USA, November 29 - December 4, 1999.

[8] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I.J. Goodfellow, R. Fergus, Intriguing properties of neural networks, in: 2nd International Conference on Learning Representations, ICLR, Banff, AB, Canada, April 14-16, 2014.

[9] I.J. Goodfellow, J. Shlens, C. Szegedy, Explaining and harnessing adversarial examples, in: 3rd International Conference on Learning Representations, ICLR, San Diego, CA, USA, May 7-9, 2015.

[10] A. Kurakin, I.J. Goodfellow, S. Bengio, Adversarial examples in the physical world, in: 5th International Conference on Learning Representations, ICLR, Toulon, France, April 24-26, 2017.

[11] W. Xu, D. Evans, Y. Qi, Feature squeezing: Detecting adversarial examples in deep neural networks, in: 25th Annual Network and Distributed System Security Symposium, NDSS, San Diego, California, USA, February 18-21, 2018.

[12] N. Akhtar, A.S. Mian, Threat of adversarial attacks on deep learning in computer vision: A survey, IEEE Access 6 (2018) 14410–14430.

[13] X. Yuan, P. He, Q. Zhu, X. Li, Adversarial examples: Attacks and defenses for deep learning, IEEE Trans. Neural Netw. Learn. Syst. 30 (9) (2019) 2805–2824.

[14] T. Vaidya, Y. Zhang, M. Sherr, C. Shields, Cocaine noodles: Exploiting the gap between human and machine speech recognition, in: 9th USENIX Workshop on Offensive Technologies, WOOT '15, Washington, DC, USA, August 10-11, 2015.

[15] F. Kreuk, Y. Adi, M. Cissé, J. Keshet, Fooling end-to-end speaker verification with adversarial examples, in: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, Calgary, AB, Canada, April 15-20, 2018.

[16] Z. Wu, N.W.D. Evans, T. Kinnunen, J. Yamagishi, F. Alegre, H. Li, Spoofing and countermeasures for speaker verification: A survey, Speech Commun. 66 (2015) 130–153.

[17] R.K. Das, X. Tian, T. Kinnunen, H. Li, The attacker's perspective on automatic speaker verification: An overview, in: 21st Annual Conference of the International Speech Communication Association, INTERSPEECH, Virtual Event, Shanghai, China, October 25-29, 2020.

[18] H. Abdullah, K. Warren, V. Bindschaedler, N. Papernot, P. Traynor, SoK: The faults in our ASRs: An overview of attacks against automatic speech recognition and speaker identification systems, in: 42nd IEEE Symposium on Security and Privacy, SP, San Francisco, CA, USA, May 24-27, 2021.

[19] V. Tiwari, MFCC and its applications in speaker recognition, Int. J. Emerg. Technol. 1 (1) (2010) 19–22.

[20] K.K. Paliwal, Spectral subband centroid features for speech recognition, in: 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, Seattle, Washington, USA, May 12-15, 1998.

[21] H. Hermansky, Perceptual linear predictive (PLP) analysis of speech, J. Acoust. Soc. Am. 87 (4) (1990) 1738–1752.

[22] Azure verification api, https://github.com/Microsoft/Cognitive-SpeakerRecognition-Python/tree/master/Verifcation.

[23] Azure identification api, https://github.com/Microsoft/Cognitive-SpeakerRecognition-Python/tree/master/Identification.

[24] X. Li, J. Zhong, X. Wu, J. Yu, X. Liu, H. Meng, Adversarial attacks on GMM I-Vector based speaker verification systems, in: 2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, Barcelona, Spain, May 4-8, 2020.

[25] J. Li, X. Zhang, C. Jia, J. Xu, L. Zhang, Y. Wang, S. Ma, W. Gao, Universal adversarial perturbations generative network for speaker recognition, in: IEEE International Conference on Multimedia and Expo, ICME, London, UK, July 6-10, 2020.

[26] Y. Xie, C. Shi, Z. Li, J. Liu, Y. Chen, B. Yuan, Real-time, universal, and robust adversarial attacks against speaker recognition systems, in: 2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, Barcelona, Spain, May 4-8, 2020.

[27] Z. Li, Y. Wu, J. Liu, Y. Chen, B. Yuan, AdvPulse: Universal, synchronization-free, and targeted audio adversarial attacks via subsecond perturbations, in: 2020 ACM SIGSAC Conference on Computer and Communications Security, CCS, Virtual Event, USA, November 9-13, 2020.

[28] G. Chen, S. Chen, L. Fan, X. Du, Z. Zhao, F. Song, Y. Liu, Who is real Bob? Adversarial attacks on speaker recognition systems, in: 42nd IEEE Symposium on Security and Privacy, SP, San Francisco, CA, USA, May 24-27, 2021.

[29] H. Abdullah, W. Garcia, C. Peeters, P. Traynor, K.R.B. Butler, J. Wilson, Practical hidden voice attacks against speech and speaker recognition systems, in: 26th Annual Network and Distributed System Security Symposium, NDSS, San Diego, California, USA, February 24-27, 2019.
[30] H. Abdullah, M.S. Rahman, W. Garcia, K. Warren, A.S. Yadav, T. Shrimpton, P. Traynor, Hear "No Evil, See "Kenansville"*: Efficient and transferable black-box attacks on speech recognition and voice identification systems, in: 2021 42nd IEEE Symposium on Security and Privacy, SP, San Francisco, CA, USA, May 24-27, 2021.
[31] Q. Wang, P. Guo, L. Xie, Inaudible adversarial perturbations for targeted attack in speaker recognition, in: 21st Annual Conference of the International Speech Communication Association, INTERSPEECH, Virtual Event, Shanghai, China, October 25-29, 2020.
[32] M. Ravanelli, Y. Bengio, Speaker recognition from raw waveform with SincNet, in: IEEE Spoken Language Technology Workshop, SLT, Athens, Greece, December 18-21, 2018.
[33] L. Schönherr, K. Kohls, S. Zeiler, T. Holz, D. Kolossa, Adversarial attacks against automatic speech recognition systems via psychoacoustic hiding, in: 26th Annual Network and Distributed System Security Symposium, NDSS, San Diego, California, USA, February 24-27, 2019.
[34] Y. Qin, N. Carlini, G.W. Cottrell, I.J. Goodfellow, C. Raffel, Imperceptible, robust, and targeted adversarial examples for automatic speech recognition, in: Proceedings of the 36th International Conference on Machine Learning, ICML, Long Beach, California, USA, June 9-15, 2019.
[35] Q. Wang, P. Guo, S. Sun, L. Xie, J.H.L. Hansen, Adversarial regularization for end-to-end robust speaker verification, in: 20th Annual Conference of the International Speech Communication Association, INTERSPEECH, Graz, Austria, September 15-19, 2019.
[36] H. Wu, S. Liu, H. Meng, H. Lee, Defense against adversarial attacks on spoofing countermeasures of ASV, in: 2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, Barcelona, Spain, May 4-8, 2020.
[37] X. Li, N. Li, J. Zhong, X. Wu, X. Liu, D. Su, D. Yu, H. Meng, Investigating robustness of adversarial samples detection for automatic speaker verification, in: 21st Annual Conference of the International Speech Communication Association, INTERSPEECH, Virtual Event, Shanghai, China, October 25-29, 2020.
[38] H. Wu, Y. Zhang, Z. Wu, D. Wang, H.-y. Lee, Voting for the right answer: Adversarial defense for speaker verification, in: 22st Annual Conference of the International Speech Communication Association, INTERSPEECH, Brno, Czechia, August 30 - September 3, 2021.
[39] H. Wu, P.-c. Hsu, J. Gao, S. Zhang, S. Huang, J. Kang, Z. Wu, H. Meng, H.-y. Lee, Spotting adversarial samples for speaker verification by neural vocoders, 2021, arXiv preprint arXiv:2107.00309.
[40] H. Zhang, L. Wang, Y. Zhang, M. Liu, K.A. Lee, J. Wei, Adversarial separation network for speaker recognition, in: 21st Annual Conference of the International Speech Communication Association, INTERSPEECH, Virtual Event, Shanghai, China, October 25-29, 2020.
[41] H. Wu, X. Li, A.T. Liu, Z. Wu, H. Meng, H. Lee, Adversarial defense for automatic speaker verification by cascaded self-supervised learning models, in: 2021 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, Toronto, on, Canada, June 6-11, 2021.
[42] Z. Gong, W. Wang, W.-S. Ku, Adversarial and clean data are not twins, 2017, arXiv preprint arXiv:1704.04960.
[43] S. Samizade, Z. Tan, C. Shen, X. Guan, Adversarial example detection by classification for deep speech recognition, in: 2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, Barcelona, Spain, May 4-8, 2020.
[44] Y. Dong, Y. Yao, Secure mmwave-radar-based speaker verification for IoT smart home, IEEE Internet Things J. 8 (5) (2021) 3500–3511.
[45] H. Aghakhani, D. Meng, Y.-X. Wang, C. Kruegel, G. Vigna, Bullseye polytope: A scalable clean-label poisoning attack with improved transferability, in: 2021 IEEE European Symposium on Security and Privacy, EuroS&P.
[46] A. Shafahi, W.R. Huang, M. Najibi, O. Suciu, C. Studer, T. Dumitras, T. Goldstein, Poison frogs! targeted clean-label poisoning attacks on neural networks, Adv. Neural Inf. Process. Syst. 31 (2018).
[47] C. Zhu, W.R. Huang, H. Li, G. Taylor, C. Studer, T. Goldstein, Transferable clean-label poisoning attacks on deep neural nets, in: International Conference on Machine Learning, PMLR, 2019.
[48] T. Miyato, S. Maeda, M. Koyama, S. Ishii, Virtual adversarial training: A regularization method for supervised and semi-supervised learning, IEEE Trans. Pattern Anal. Mach. Intell. 41 (8) 1979–1993.
[49] N. Carlini, P. Mishra, T. Vaidya, Y. Zhang, M. Sherr, C. Shields, D.A. Wagner, W. Zhou, Hidden voice commands, in: 25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016.

Jiahe Lan received the B.S. degree in Network Engineering from Xidian University, Xi'an, China, in 2020, where she is currently pursuing the M.S. degree with the School of Cyber Engineering. Her research interests include adversarial machine learning, voice processing systems, and human–computer interaction.

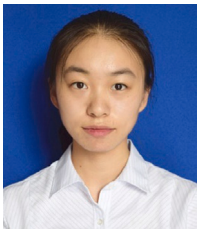

Rui Zhang received the B.S. degree in Computer Science and Technology from CUMT, Xuzhou, China, in 2016. Now she is major in Cyber Security and studying for a Ph.D. in Xidian University in Xi'an, China. Her research interests are in information security, authentication, and privacy preserving in social network.

Zheng Yan is a full professor at the School of Cyber Engineering, Xidian University, China and a Finnish Academy Fellow at the Aalto University, Finland. She received the Doctor of Science in Technology in 2007 from the Helsinki University of Technology (i.e., current Aalto University), Finland. Before joining academia in 2011, she was a senior researcher at the Nokia Research Center, Helsinki, Finland, since 2000. Her research interests are in trust, security, privacy, and security-related data analytics, focusing on autonomous and privacy-enhanced trust management for securing networking and communications. Her research leads to 310+ peer reviewed scientific articles, two solely authored books and nine conference proceedings. She is an inventor of 110+ patents, 87 of them have been adopted in industry. Some of her patented techniques have been referred by international standards and widely used in practice. She is an area or associate editor of Information Fusion, IEEE Network Magazine, IEEE Internet of Things Journal, Information Sciences, and JNCA, etc. She served as a general chair or program chair for over 30 international conferences including IEEE TrustCom2015, NSS/ICA3PP/IEEE CIT 2017, IEEE Blockchain2018, and IFIP Networking 2021. She is a founding steering committee co-chair of IEEE International Conference on Blockchain. She received many awards and honors in recent years, including 2021 N²Women Star in Computer Networking and Communications, Distinguished Inventor Award of Nokia (2020) for her significant technology contributions, Aalto ELEC Impact Award (2020) for patent contributions to Finnish society, Elsevier 2020 highly cited Chinese researcher, the 2017 Best Journal Paper Award issued by IEEE Communication Society Technical Committee on Big Data, the Best Paper Award of SpaCCS2019, more than ten IEEE Outstanding/Distinguished Leadership/Service Awards (2010–2019), and the Outstanding Associate Editor of 2017 and 2018 for IEEE Access Journal. She has offered more than 20 keynotes and invited talks in international conferences. She is a Fellow of IET and a senior member of IEEE.

Jie Wang received the B.S. degree in Network Engineering from Xidian University in 2020, where he is currently pursuing the Ph.D. degree with the School of Cyber Engineering. His research interests include trust management, machine learning and blockchain.

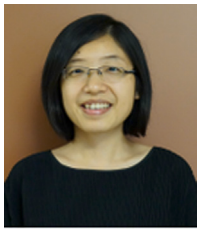

Yu Chen is an assistant professor in the School of Information Systems and Technology in the Lucas College of Business at San Jose State University. Prior to that, she was a postdoctoral researcher at University of California, Irvine. Dr. Chen received her Ph.D. from Swiss Federal Institute of Technology at Lausanne, Master's degrees from both Aalto University and Northwestern University of Science and Technology, and Bachelor's degree from Huazhong University of Science and Technology.

Ronghui Hou received the B.Eng., M.Eng., and Ph.D. degrees in communication engineering from Northwestern Polytechnical University in 2002, 2005, and 2007, respectively. She was a Post-Doctoral Fellow with the Department of Electrical and Electronic Engineering, The University of Hong Kong, from 2007 to 2009. Since 2009, she has been with Xidian University, China, where she is currently a Professor with the Department of Telecommunication Engineering. Her research interests include network quality of service issues, routing algorithm design, and wireless networks.