Black-box Adversarial Attacks on Commercial Speech Platforms with Minimal Information

Baolin Zheng^{1,*}, Peipei Jiang^{1,*}, Qian Wang^{1,†}, Qi Li², Chao Shen³, Cong Wang⁴, Yunjie Ge¹, Qingyang Teng¹, and Shenyi Zhang¹

¹ School of Cyber Science and Engineering, Wuhan University

² Institute of Network Sciences and Cyberspace, Tsinghua University; BNRist

³ School of Cyber Science and Engineering, Xi'an Jiaotong University

⁴ Department of Computer Science, City University of Hong Kong

{baolinzheng, ppjiang, qianwang}@whu.edu.cn, qli01@tsinghua.edu.cn, chaoshen@xjtu.edu.cn, congwang@cityu.edu.hk, {yunjiege, qingyangteng, shenyizhang}@whu.edu.cn

ABSTRACT

Adversarial attacks against commercial black-box speech platforms, including cloud speech APIs and voice control devices, have received little attention until recent years. Constructing such attacks is difficult mainly due to the unique characteristics of time-domain speech signals and the much more complex architecture of acoustic systems. The current “black-box” attacks all heavily rely on the knowledge of prediction/confidence scores or other probability information to craft effective adversarial examples (AEs), which can be intuitively defended by service providers without returning these messages. In this paper, we take one more step forward and propose two novel adversarial attacks in more practical and rigorous scenarios. For commercial cloud speech APIs, we propose Occam, a decision-only black-box adversarial attack, where only final decisions are available to the adversary. In Occam, we formulate the decision-only AE generation as a discontinuous large-scale global optimization problem, and solve it by adaptively decomposing this complicated problem into a set of sub-problems and cooperatively optimizing each one. Our Occam is a one-size-fits-all approach, which achieves 100% success rates of attacks (SRoA) with an average SNR of 14.23dB, on a wide range of popular speech and speaker recognition APIs, including Google, Alibaba, Microsoft, Tencent, iFlytek, and Jingdong, outperforming the state-of-the-art black-box attacks. For commercial voice control devices, we propose NI-Occam, the first non-interactive physical adversarial attack, where the adversary does not need to query the oracle and has no access to its internal information and training data. We, for the first time, combine adversarial attacks with model inversion attacks, and thus generate the physically-effective audio AEs with high

transferability without any interaction with target devices. Our experimental results show that NI-Occam can successfully fool Apple Siri, Microsoft Cortana, Google Assistant, iFlytek and Amazon Echo with an average SRoA of 52% and SNR of 9.65dB, shedding light on non-interactive physical attacks against voice control devices.

CCS CONCEPTS

• Security and privacy → Software and application security.

KEYWORDS

Speech recognition; speaker recognition; adversarial attacks; black-box attacks

ACM Reference Format:

Baolin Zheng^{1,*}, Peipei Jiang^{1,*}, Qian Wang^{1,†}, Qi Li², Chao Shen³, Cong Wang⁴, Yunjie Ge¹, Qingyang Teng¹, and Shenyi Zhang¹. 2021. Black-box Adversarial Attacks on Commercial Speech Platforms with Minimal Information. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21)*, November 15–19, 2021, Virtual Event, Republic of Korea. ACM, New York, NY, USA, 22 pages. <https://doi.org/10.1145/3460120.3485383>

1 INTRODUCTION

Nowadays, with the advance of speech and speaker recognition technologies, they are reshaping the way we interact with ubiquitous smart devices. More specifically, automatic speech recognition (ASR) technologies [41] allow machines to understand human voices, while speaker recognition (SR) technologies [37, 71] enable machines to identify a person from the characteristics of his/her voices¹. As a result, ASR and SR have become universal in our daily lives, ranging from personal voice assistants (PVAs) [14, 46] to biometric authentication [20, 79] on various smart devices. The popularity of such speech services allows people to greatly enjoy the convenience of integrating speech as a new input for smart devices to perform daily and even complicated tasks. For example, Amazon has released Alexa [1] and Auto SDK [4] that allow users to make credit card payments and control vehicles with voice interaction, respectively.

Despite their wide applications, the excessive use of voice commands in safety-critical systems, like autonomous driving [38, 59] and biometric identification, also poses potential safety hazards.

¹To facilitate differentiation, in the following discussions we abbreviate automatic speech recognition and speaker recognition to ASR and SR, respectively.

* The first two authors contributed equally to this work.

† Qian Wang is the corresponding author.

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

CCS '21, November 15–19, 2021, Virtual Event, Republic of Korea

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8454-4/21/11...\$15.00

<https://doi.org/10.1145/3460120.3485383>

A line of recent researches [12, 23, 92] have extensively demonstrated the vulnerability of acoustic systems to numerous types of abnormal audios, such as noises and inaudible ultrasounds. These attacks, however, can be easily detected and/or defended by differentiating and analyzing the nature (*i.e.*, legitimate or malicious) of the received audio signals. Inspired by the resounding success of adversarial attacks against image recognition systems [24, 40, 62, 78], more recent researches have begun to investigate the feasibility of adversarial examples in the audio domain [47], as shown in Table 1.

The very first attempts made by Carlini *et al.* [25] and Yuan *et al.* [90] have shown that ASR systems are also inherently vulnerable to audio AEs in the white-box scenarios, where the attackers can make full use of a prior knowledge of the structure and parameters inside the system. When it comes to black-box settings, however, the success of adversarial attacks in the image domain is hard to be ported to the audio domain, mainly owing to the multiple non-trivial challenges presented by the unique characteristics of time-domain speech signals and the much more complex architecture of acoustic systems.

As evidenced by [80], Taori *et al.* combined genetic algorithms [84] and gradient estimation [28], which have been proved effective in the image domain [13, 19], to carry out a black-box attack against open-source DeepSpeech [42] but with only a success rate of 35%. Besides, considering that the attacker requires query access to the last layer (*i.e.*, logits) of DNNs inside the DeepSpeech, such attacks become unrealistic when applied to the closed-source ASR systems. More recently, Chen *et al.* [32] showed that when the confidence scores are publicly known, several commercial ASR systems are vulnerable to adversarial inputs, but with only a very limited number of target commands. Almost instantaneously, Du *et al.* [35] and Chen *et al.* [26] presented the first black-box adversarial attacks against SR systems that both heavily relied on prediction scores. However, such kind of scores defined by the attackers or provided by speech service APIs may be useless to legitimate users, so service providers can easily hide these scores to defend against the above black-box attacks. Besides, there also exist a significant number of speech service APIs (*e.g.*, Alibaba Short Speech Recognition API [11], iFlytek Speech-to-Text API [3]) which do not return any intermediate information (*e.g.*, confidence/prediction scores or other probabilities), except for the final decision results, *e.g.*, final transcriptions in ASR systems and user-ids in SR systems.

From the practical perspective, the prior efforts are important but not satisfactory enough with respect to the minimum information as required by the adversary to launch a successful black-box attack. We may ask: “*is it possible to launch effective and practical audio adversarial attacks against commercial black-box speech platforms with the minimum information?*” We are facing this opportunity already, more facing a big challenge mainly due to the extreme lack of information about the target model. Specifically, the acoustic system, which involves non-linear feature extraction steps to cope with the intricate frequency feature changes in the time dimension, is much more complicated than the image processing system. Furthermore, a speech vector usually contains nearly one hundred thousand variables, far exceeding the hundreds or thousands of pixels in images, *i.e.*, MNIST and CIFAR-10 are 28×28 and 32×32 respectively. As reported in [88], the explicit interdependencies

among the massive number of variables significantly hinder the successful construction of audio AEs.

1.1 Our Works

Generally speaking, there are mainly two types of black-box speech platforms. One is commercial Cloud Speech APIs that provide audio services to users, and the other is commercial Voice Control Devices, such as Apple Siri, which perform the speech-to-text task in the physical world. In this paper, we present two attack schemes, Occam², a decision-only attack on cloud speech APIs, and NI-Occam, a non-interactive physical attack on voice control devices.

Occam. In our first design, we take one more step forward and focus on real-world threat scenarios where the adversary has access to an oracle (target model) which returns only its final decision. We propose Occam, a decision-based black-box adversarial attack against cloud speech APIs. We demonstrate that various commercial speech API services, such as Google Cloud Speech-to-Text, Alibaba Cloud Speech-to-Text, and Microsoft Azure Speech Service, are inherently vulnerable to audio AEs generated by our Occam, even if no internal information is exposed to the adversary.

Our key idea of Occam is to formulate the decision-based black-box attack against smart acoustic systems as a discontinuous large-scale global optimization problem, on the basis of the final discrete decision (the attacker can only obtain) and a large number of optimization variables incurred by the speech. Inspired by this observation, we develop a novel technique called CC-CMA-ES, which applies a cooperative co-evolution (CC) framework to the powerful covariance matrix adaptation evolution strategy (CMA-ES), to solve the large and complex problem in the strictly black-box setting. More specifically, CC-CMA-ES first decomposes the complicated problem into a set of smaller and simpler sub-problems, and then uses CMA-ES to cooperatively optimize each one by modeling their local geometries. To improve the attack efficiency, we further propose an adaptive counterpart, which allows the sub-problem size and the decomposition strategy to self-adapt to the environmentally changeable evolution process.

We conduct extensive experiments to evaluate our attack capabilities on both speech and speaker recognition tasks, and also compare it with five decision-based black-box methods to demonstrate the superiority of Occam. We first craft audio adversarial examples against the local DeepSpeech model in the strictly black-box setting, achieving perfect success rates in both targeted and untargeted attacks. Then, we launch black-box adversarial attacks on a wide range of commercial speech-to-text API services, including Google, Microsoft, Alibaba, Tencent, and iFlytek, with success rates of 100% and an average SNR of 14.37dB. Furthermore, we verify the attack effectiveness against commercial SR systems including Microsoft and Jingdong. It still achieves success rates of 100% and an average SNR of 14.07dB.

NI-Occam. In our second design, we further probe the possibility of launching more rigorous and practical attacks on voice control devices, where the adversary still has no access to internal information and training data of the oracle, and does not even need to make

²Occam comes from the famous Occam’s razor that plurality should not be posited without necessity, indicating that our attacks against speech platforms can be performed with minimal information.

Table 1: An overview of the state-of-the-art adversarial audio attacks against ASR and SR systems.

Method	Threat Scenario	Attack Type [‡]	Task	Knowledge	Commercial [§]	Queries	SRoA [†]
Carlini <i>et al.</i> [25]	White-box	Digital	ASR	Gradient	✗	~1000	100%
CommanderSong [90]	White-box	Digital	ASR	Gradient	✗	~100	100%
Taori <i>et al.</i> [80]	Black-box	Digital	ASR	Prediction score	✗	~300,000	<40%
SGEA [82]	Black-box	Digital	ASR	Prediction score	✗	~300,000	100%
Devil's Whisper [32]	Black-box	Digital	ASR	Confidence score	✓	~1500	~50%
		Physical					<100%
SirenAttack [35]	Black-box	Digital	ASR	Gradient	✗	~1000	100%
			SR	Prediction score	✗	~7500	100%
FakeBob [26]	Black-box	Digital	SR	Prediction score	✗	~5000	100%
Ours	Black-box	Digital	ASR	Final decision	✓	~30,000	100%
			SR			~10,000	
		Physical	ASR	None [#]	✓	0	~50%

Note that, (i) ‡: “Digital” means that audio AEs are injected into target systems, while “Physical” means that audio AEs are played over-the-air. (ii) §: “✓” means that the target model is the commercial platforms, otherwise “✗”. (iii) †: SRoA denotes the success rate of attack. It calculates the proportion of adversarial examples that can successfully attack target systems. (iv) #: Our physical attack focuses on the non-interactive setting where the adversary has no access to the oracle (the target model).

queries to probe it. We, for the first time, propose a non-interactive physical attack, named NI-Occam, which successfully attacks many commercial voice control devices without any interaction. We show that NI-Occam works well in real-world attack scenarios, where audio AEs are played over-the-air.

Our key idea of NI-Occam is to combine adversarial attacks with model inversion attacks [29, 36, 94]. More specifically, we make the attempt to recover the key parts of natural commands audio that are critical for speech recognition on the original example via the gradient information. Since these two audios are well blended together in model inversion process, it is difficult to be separated by human ears, which significantly hinders people from recognizing the malicious audio. Finally, our proposed NI-Occam can successfully fool Apple Siri, Microsoft Cortana, Google Assistant, iFlytek, and Amazon Echo with an average SRoA of 52%. Human perception experiments further show that after being heard once, only 6.4% audio AEs can be recognized as target commands by volunteers.

We emphasize that our attacks have the following highlights:

- 1) *Practicality*. They are able to attack commercial black-box platforms in the real-world scenarios without any prior knowledge;
- 2) *Generality*. They are able to attack a wide range of commercial cloud speech APIs and voice control devices;
- 3) *Effectiveness*. They are able to automatically and easily generate audio AEs with high success rates of attack.

Contribution. Our major contributions are summarized as follows.

- *Generic black-box attacks with the minimum required information.* We present a novel decision-only audio adversarial attack, named Occam, under the strictly black-box scenario where the attackers can rely solely on the final decisions available in any application cases, and this is quite different from the state-of-the-art black-box adversarial attacks against commercial Cloud Speech APIs. In this sense, our attack strategy is the first one that can fool both commercial ASR and SR services, to our best knowledge.
- *Effective attacks with a perfect success rate of attack.* We thoroughly evaluate our attack on a wide range of popular open-source and commercial (A)SR systems, including Google, Alibaba, Microsoft, Tencent, iFlytek, Jingdong, and DeepSpeech systems. Extensive experiments demonstrate that our attack is highly effective with a

success rate of 100% and an average SNR of 14.23dB on commercial speech services, outperforming the state-of-the-art black-box attacks on commercial cloud speech APIs.

- *Practical physical attacks without any interaction.* We explore the possibility of generating audio AEs and playing them against the commercial voice control devices over-the-air. We thus for the first time propose a new non-interactive physical attack, named NI-Occam, which can successfully fool various voice control devices, including Apple Siri, Microsoft Cortana, Google Assistant, iFlytek and Amazon Echo, without any feedback information from the target devices. The experimental results show that our over-the-air attack can achieve an average success rate of 52% with SNR of 9.65dB. This observation is shedding light on non-interactive physical attacks against voice control devices.

2 BACKGROUND

2.1 Speech Recognition

Automatic speech recognition (ASR) systems allow machines to automatically convert speeches into texts, and have found tremendous applications. Typically, an ASR system consists of four main components: pre-processing, feature extraction, acoustic model, and language model, as in Figure 1. Pre-processing plays an important role to filter out the frequencies beyond the range of human hearing and the segments below a specific energy threshold in the raw audio. Then, features are extracted via a feature extraction algorithm, such as Mel-frequency Cepstral Coefficients (MFCC) [63], Linear Predictive Coefficient (LPC) [49], Perceptual Linear Predictive (PLP) [44], *etc.* Different from image classification models [56] which take pixels as input, the acoustic model takes the extracted features as input, and outputs the probability of phonemes. In early ASR systems, Hidden Markov Model is one of the preferred techniques [16]. With the stupendous advance of deep learning, deep neural networks (DNNs) [45] and especially recurrent neural networks (RNNs) [41, 70] have become the dominant choices today. An ASR system will finally produce the correct transcriptions conforming to grammatical rules via the language model.

Figure 1: The architecture of a typical ASR system.

Figure 2: The architecture of a typical SR system.

2.2 Speaker Recognition

Due to the tremendous advance of DNN, speaker recognition (SR) systems are becoming increasingly popular in biometric authentication [30, 58, 91]. The systems, which allow machines to correctly identify a person from his/her unique characteristics of voices, have various applications such as bank services and forensic tests. The key step then is to extract users' voice features from a series of utterances which essentially consist of the underlying text information and the features of the speaker.

Figure 2 provides the overview of a typical SR system [54], which consists of two phases: enrollment and evaluation. Generally, in the enrollment phase, a set of background speakers are required to train the background model in the offline phase such that a speaker can provide a few utterances online to create a new specific speaker model. The technologies for generating such models can be generally summarized as three types: i-vector-PLDA [65, 66], GMM-UBM [72, 73], and DNN [77, 91]. During the evaluation phase, the unknown speaker's voice is taken as input and scored by the speaker models in the library. Based on the resulting scores, the decision model will generate the final recognition result.

There are two important sub-tasks in speaker recognition: speaker verification (SV) [72] and speaker identification (SI) [52]. The former is to validate whether the current user is legitimate, *i.e.*, output either accept or reject. The latter aims to figure out the identity of the speaker among a set of enrolled ones. According to its text dependence, SR systems can also be divided into text-dependent and text-independent [71]. The difference is that the text-dependent approach requires all speakers to utter pre-defined sentences. Despite higher accuracy, it is only used for the SV sub-task.

2.3 Adversarial Examples

Despite their great success, the vulnerabilities of neural networks to adversarial examples have recently been extensively studied [24, 40, 62, 78]. An adversary can slightly revise the legitimate inputs to generate adversarial examples for fooling neural networks [78]. Generally speaking, there are two types of adversarial examples: untargeted and targeted ones, also known as dodging attacks and

impersonation attacks, respectively. Let $f(x) : x \in \mathcal{X} \rightarrow y \in \mathcal{Y}$ denote the recognition model that maps the input x into the corresponding output prediction y . Given the original input x with the prediction y such that $f(x) = y$, an untargeted adversarial example x^* in the dodging attacks can be represented as:

$$x^*, \quad \text{s.t.} \quad f(x^*) \neq y, D(x, x^*) = \|x - x^*\|_p \leq \epsilon, \quad (1)$$

where $D(x, x^*)$ is the distance between an original input x and an untargeted adversarial example x^* , ϵ is a parameter used to limit this distance, and p is typically 0, 2, or ∞ . Similarly, in the impersonation attacks, for an original input x and a specific y^* such that $f(x) \neq y^*$, a targeted adversarial example x^* can be represented as:

$$x^*, \quad \text{s.t.} \quad f(x^*) = y^*, D(x, x^*) \leq \epsilon. \quad (2)$$

3 OCCAM: A DECISION-ONLY DIGITAL ATTACK AGAINST CLOUD SPEECH APIS

3.1 Threat Model

Nowadays, many commercial cloud speech platforms offer both ASR and SR API services, *e.g.*, Microsoft Azure Speech and Speaker Recognition Service API, and thus third-party developers can access commercial APIs if they have enrolled or paid for the services. The service providers, on the other hand, will not expose any parameters or the architecture of the target model because the internal information is commercially sensitive. Actually, a number of API services, *e.g.*, iFlytek, Alibaba, Tencent, and Jingdong, provide only the final decision results without exposing any other information. Therefore, it is important to explore a generic attack against both ASR and SR commercial APIs in this decision-based scenario. Note that although ASR tasks are different from SR tasks, we can treat them as the same problem in this design because the construction of audio AEs for ASR and SR APIs can be formulated as the same optimization problem.

In this section, our target is commercial cloud speech services that open their APIs to the public, and we assume that the adversary intends to generate AEs against both ASR and SR services without any internal knowledge of the target model. More specifically, the adversary can only query the target model and obtain its final decision, which is a strict but more practical assumption in real-world applications. Considering the adversary's knowledge of the original audio, we make two different assumptions for the ASR and SR tasks respectively. For ASR systems, we assume that the adversary knows nothing about the dataset, and thus we utilize Text-to-Speech Service API to generate the audio of the target text. For SR systems, the adversary only needs to collect the victim's one voice sample, which is readily available owing to the serious leakage of personal information, *e.g.*, public videos in social media.

3.2 Problem Formulation

For an acoustic system, given a voice $x \in \mathbb{R}^n$ and a specific y^* such that $f(x) \neq y^*$, a targeted AE x^* can be described by

$$x^*, \quad \text{s.t.} \quad f(x^*) = y^*, D(x, x^*) \leq \epsilon. \quad (3)$$

In the white-box setting, the problem can be formulated as

$$\min_{x^*} \mathcal{L}(x^*) = \mathcal{D}(x^*, x) + c \cdot \mathcal{J}(x^*, y^*), \quad (4)$$

Figure 3: An illustration of continuous (left) and discontinuous (right) optimization problem.

Figure 4: The architecture of our Occam.

where $\mathcal{L}(\cdot)$ is the objective function, $\mathcal{J}(\cdot, \cdot)$ the loss function to check how well x^* meets the adversarial requirement, and c the adjustment parameter [25]. In the black-box setting without internal knowledge, we can reformulate it as an optimization problem as

$$\min_{x^*} \mathcal{L}(x^*) = \begin{cases} \mathcal{D}(x^*, x), & \text{if } f(x^*) = y^*, \\ +\infty, & \text{otherwise.} \end{cases} \quad (5)$$

Note that the $\mathcal{L}(x^*)$ is equal to $+\infty$ when x^* is not adversarial. Thus, we try to find the adversarial region and minimize its distance from the original audio. Similarly, the construction of untargeted audio AEs in our decision-based black-box attack can be reformulated as

$$\min_{x^*} \mathcal{L}(x^*) = \begin{cases} \mathcal{D}(x^*, x), & \text{if } f(x^*) \neq f(x), \\ +\infty, & \text{otherwise.} \end{cases} \quad (6)$$

3.3 Technical Challenges

As mentioned above, compared to the white-box and score-based settings, it is quite challenging to craft decision-based black-box adversarial examples against acoustic systems, since the attacker has no internal knowledge about the target model, except for the final decision corresponding to the query. Thus, we are facing several design challenges. First, depending on whether the output decision matches the adversarial target or not, the optimization space can be divided into a large non-adversarial region and a very small adversarial region. Hence, different from almost all previous adversarial voice attacks which feature a continuous optimization problem, the construction of audio AEs in the strictly black-box scenario poses a difficult discontinuous optimization problem due to the extreme lack of information, as shown in Figure 3. Second, the acoustic systems are extremely complicated, and we have to deal with intricate feature changes of the audios in the time dimension. Finally, because the audio sampling rate is very high (e.g., 16kHz),

Algorithm 1 Occam

Input: The original audio x , the initial adversarial audio x' , the input space dimension n , the attack objective function $\mathcal{L}(\cdot)$, the binary search time b and the total number of queries T .

Output: The adversarial audio sample x^* .

1: Initialize the following parameters:

- Covariance matrix $C = I_n$, $x^* = x'$, $t = 0$;
- $\delta, \mu, c_c, c_{cov} \in \mathbb{R}_+$, $m, \lambda \in \mathbb{Z}_+$ (See Section 3.4.5).

2: Query b times to find the new adversarial point x^* close to the decision boundary via binary search algorithm;

3: $t = t + b$;

4: Choose a grouping strategy and a group number m according to the adaptive scheme (See Section 3.4.4);

5: Decompose the audio vector into m disjoint parts with each part having $s = n/m$ dimensions;

6: Set $sub = 1$ to start an optimization cycle;

7: Extract two $s \times 1$ vectors x_{sub}^* and x_{sub} and an $s \times s$ covariance matrix C_{sub} from x^* , x and C , respectively;

8: **for** $i = 0$ to λ **do**

9: Sample $z \sim \mathcal{N}(0, \sigma^2 \cdot C_{sub})$;

10: Generate one complete solution $solu$ using $x_{sub}^* + \mu(x_{sub} - x_{sub}^*) + z$ and collaborative information from other subspaces;

11: **if** $\mathcal{L}(solu) < \mathcal{L}(x^*)$ **then**

12: $x_{sub}^* = x_{sub}^* + \mu(x_{sub} - x_{sub}^*) + z$;

13: Update C_{sub} according to Eqs. (8) and (9);

14: **end if**

15: **end for**

16: $t = t + \lambda$;

17: Use x_{sub}^* and C_{sub} to update x^* and C , respectively;

18: **if** $t \geq T$ **then**

19: **return** x^* .

20: **else if** $sub < m$ **then**

21: $sub++$ and go to step 5;

22: **else**

23: Go to step 2;

24: **end if**

the large number of optimization variables in the speech vector presents another challenge, *i.e.*, the curse of dimensionality, especially when there is a clear interdependence among the variables in audio AEs [88]. The reason behind it is that as the number of optimization variables increases, the complexity of the problem grows exponentially, and the nature of the problem may also change.

3.4 Our Method

As noted in our threat model, the lack of internal knowledge (e.g., structures, parameters, gradients, and scores) about the target model further exacerbates the difficulty of crafting AEs. All the adversary can do is to send a limited number of queries to probe the system as far as possible, and obtain the corresponding final decisions. More specifically, the initiation of our decision-based black-box adversarial attack only needs the final decision, e.g., final transcription, and

this kind of audio AE generation can be formulated as a discontinuous and large-scale global optimization problem. To solve it, we resort to the large-scale black-box optimization approach.

Design Overview. Note that we are going to fool both commercial ASR and SR services with the minimum required information from the target model. To address the challenges, we propose a new class of cooperative co-evolution methods to generate effective audio AEs. Our method is mainly developed from CC-CMA-ES [60]. However, we cannot directly apply CC-CMA-ES to constructing audio AEs. In the audio domain, the correlations between variables will change in the dynamic evolution process, which is not considered in [60]. To solve this problem, we devise an adaptive scheme to make our strategy self-adapt to the environmentally changeable evolution process, the core of our cooperative co-evolution framework.

In the literature, the CMA-ES [43], known as an efficient evolutionary algorithm, has already demonstrated its good performance on many problems. But it will lose its effectiveness when applied to large-scale global optimization problems due to “the curse of dimensionality”. A dimensionality reduction strategy, which uses the bilinear interpolation method to project the original space ($112 \times 112 \times 3$) to a lower-dimensional search space (e.g., $45 \times 45 \times 3$), has been carefully devised to effectively create adversarial images against face recognition models [34]. However, this method is also not applicable to the audio domain. This is because bilinear interpolation works in two directions on images, while the inputs to the commercial Cloud Speech APIs are one-dimensional vectors from speeches. Thus, we for the first time introduce the general cooperative co-evolution (CC) framework to construct audio AEs in the strictly black-box setting as shown in Figure 4. Our CC framework can scale up CMA-ES to deal with large-scale optimization problems we are facing, by decomposing these challenging problems into a set of simpler and smaller sub-problems and cooperatively optimizing each of them. Considering that the group size and the decomposition strategy play a crucial role in CC, we further propose an adaptive scheme to improve the attack efficiency by letting the size of sub-problems and the decomposition strategy self-adapt to the environmentally changeable evolution process.

Our Occam is presented in Alg. 1. In each optimization cycle, the original problem is decomposed into a set of smaller and simpler subproblems according to the selected grouping strategy. Then, a new offspring is generated from the current solution in each subproblem by a subspace CMA-ES whose parameters are extracted from a global CMA-ES, and the complete candidate solution can be further obtained by using the collaborative information from other subproblems and the generated offspring. The objective function is further used to evaluate the two solutions and choose the better one, based on which we update the covariance matrix accordingly. Next, we will describe each step of the algorithm in detail.

3.4.1 Initialization. As shown in Eq. (5), the optimization routine should start from an adversarial point, because the value of the objective function is equal to $+\infty$ when the input is not adversarial. We first initialize x^* with a natural adversarial sample distant from the original audio. More specifically, we utilize the text-to-speech API service to synthesize the desired speech as an initial adversarial audio sample against ASR systems. For SR systems, we initialize x^* using an audio of the target speaker, which can be obtained

Figure 5: The diagram of CMA-ES.

from the speaker’s self-recorded songs and videos posted on public social networks. Since the initial input audio is adversarial and the original one is not adversarial, we can first utilize the binary search algorithm to effectively approach the decision boundary.

3.4.2 Covariance Matrix Adaptation Evolution Strategy. CMA-ES, known as an efficient derivative-free method, generates offsprings by sampling a multivariate normal distribution with covariance C , i.e., $\mathcal{N}(0, C)$. To facilitate understanding, we provide a brief summary of CMA-ES, as shown in Figure 5. Since the covariance is the measure of the relationship between two random variables, it can use the selected sample distribution to estimate the covariance for learning dependencies between variables [43]. Due to the unreliability of this estimation for small samples, the covariance matrix adaptation using the history information and the evolution path P , a sequence of successive and normalized steps, thus has been introduced. By adaptively updating the estimated covariance, CMA-ES is able to find better search directions, and achieve a powerful local search by modeling local geometries.

Our framework uses a simple yet effective variant of CMA-ES, i.e., (1+1)-CMA-ES [48]. This version generates one candidate solution from current solution by sampling a random noise, and selects a better solution according to its objective function. Since the direction of optimization in our attack is to find a new adversarial audio closer to the original audio according to the objective function $\mathcal{L}(\cdot)$, we adopt a modified (1+1)-CMA-ES [34] to improve its efficiency by adding a bias term $\mu(x - x^*)$ to current solution x^* as

$$x_{i+1}^* \sim \mathcal{N}(x_i^* + \mu(x - x_i^*), \sigma^2 \cdot C), \quad (7)$$

where σ is the global step size, C is the covariance matrix that determines the shape of the distribution, and μ is a parameter that controls the degree of proximity towards the original audio. Furthermore, we can directly remove candidate solutions that are farther from the original audio, regardless of whether it is adversarial.

Note that the covariance matrix C plays a very important role in CMA-ES since it models the local geometries to improve the local search efficiency. However, the adaptation of covariance matrix with the complexity of $O(n^3)$ may be infeasible when the input dimension n is huge. To speed up the computation, the covariance matrix C can be simplified as a diagonal matrix [34] and updated adaptively by the evolution path P as

$$P = (1 - c_c)P + \sqrt{c_c(2 - c_c)} \frac{z}{\sigma}, \quad (8)$$

$$C = (1 - c_{cov})C + c_{cov}P(P)^T, \quad (9)$$

where c_c and c_{cov} are the parameters controlling the adaptation of P and C , respectively. The update enlarges the variance along the past successful directions for future search.

3.4.3 Cooperative Co-evolution. It has been proven that the performance of evolution algorithms may drop significantly [60, 89], as the dimensionality of the problem increases, because the complexity of the problem grows exponentially and the property of the problem may also change. To scale up CMA-ES to the high-dimensional optimization problem, we use cooperative co-evolution (CC) to conduct the large-scale black-box optimization in a divide-and-conquer manner, by decomposing the large-scale problem into several smaller sub-problems and optimizing each sub-problem alternately and iteratively. Considering that each subproblem is only a part of the original problem, the collaborative information from other sub-problems is required to evaluate individuals in the current sub-problem. Generally speaking, the best solutions of each sub-problem in the current cycle are used as the collaborative information. However, considering the query limitation in our attack, we propose a greedy strategy to produce and update the collaborative information in our attack design. More specifically, we do not optimize these sub-problems concurrently. Instead, we optimize them alternatively and iteratively. Therefore, when optimizing a subproblem, we can evolve its values of the variables, which are used to replace those related to the current subproblem in the best solution, and generate a complete solution. The solution is further evaluated by the objective function $\mathcal{L}(\cdot)$ to locate a better solution. After the optimization of each subspace, the collaborative information will be updated correspondingly. Since CC is a general framework based on the divide-and-conquer strategy for solving large-scale black-box optimization problems, it is generalizable to other black-box methods that are also trapped in these problems, and CC can also improve their effectiveness.

3.4.4 Adaptive Decomposition. The grouping strategy, which determines how to assign variables to different groups, plays a crucial role in CC. However, there is insufficient knowledge about the correlations between variables, making manually devising or choosing the most suitable grouping strategy extremely hard when applying CC. Therefore, we propose an adaptive approach, which puts several popular grouping strategies into a candidate pool, and adaptively selects a proper decomposition strategy from it. The candidate pool includes four grouping strategies: *Static grouping* (SG), *Random grouping* (RG), *Min-variance grouping* (MiVG) and *Max-variance grouping* (MaVG). We adopt SG to preserve the information in the time domain, and RG [89] can help randomly allocate variables to subspaces for improving the probability of placing two interacting variables in the same subspace. Considering that the covariance matrix is used to model the local geometries of the search directions, thus we adopt MiVG and MaVG, which are devised for CC-CMA-ES. Actually, they can minimize or maximize the diversity of the diagonal values of the variables in the same subspace. At the beginning of each optimization cycle, we randomly select a decomposition strategy. Based on its performance, we calculate the selection probabilities of each grouping strategy for the next cycle.

We observe that the levels of interdependency among variables in the audio vector will change significantly from the natural audio to the audio AE during the optimization process [88]. Therefore, we propose to adaptively adjust the group size to capture different interdependency levels in the dynamic evolution process. Then, we can make a good trade-off between the effectiveness and the

efficiency of the optimization. The details of the adaptive decomposition algorithm can be seen in Appendix A.

3.4.5 Parameter Adjustment. Our algorithm contains many hyper-parameters, such as δ , μ , λ , c_c , c_{cov} and b . Following [34], we set $\delta = 0.001 \cdot D(x^*, x)$, $c_c = 0.01$, $c_{cov} = 0.001$ and $b = 15$. We further set $\lambda = 30$ and $\mu = 0.08$, because μ has an important impact on the search process, we need to carefully tune μ . Finally, we adopt the 1/5th success rule [15] to update μ as

$$\mu = \begin{cases} 1.5\mu, & \text{if a better solution is obtained,} \\ 1.5^{-1/4}\mu, & \text{otherwise.} \end{cases} \quad (10)$$

4 NI-OCCAM: A NON-INTERACTIVE PHYSICAL ATTACK AGAINST VOICE CONTROL DEVICES

4.1 Threat Model

In this section, our target is commercial voice control devices. We consider the most rigorous and practical assumption, called non-interactive physical setting, where the adversary makes no query to the oracle. Compared to prior physical attacks, the key advantage of non-interactive physical attacks is that we do not need to query the target devices for effective audio AE generation, thus saving the potentially large query cost. In this sense, this attack is the most practical one in the real world.

4.2 Technical Challenges

Compared to our decision-only adversarial attacks, non-interactive black-box setting is much more challenging since it further breaks the dependence on the final decision in the decision-only black-box scenario. That is to say, the target model is completely unknown to the adversary. Moreover, voice control devices also present additional challenges that the constructed audio AEs should remain robust even if they are played in the physical world. By physical attacks against voice control devices, we mean that audio AEs are played by a speaker and recorded by the device. Since the effectiveness of audio AEs is greatly affected by the reverberation of the environment, and perturbations from the speaker and the microphone [75, 86], it is really difficult to launch physical attacks. These two obstacles pose severe challenges to craft effective audio AEs.

4.3 Our Method

As described in our threat model, the adversary will not issue any queries to probe the target model and obtain no feedback in the non-interactive black-box setting. Thus, our Occam cannot be applied in this case. Intuitively, it is almost impossible to directly construct an audio adversarial example against the target model by solving the optimization problem without any interaction. In fact, we are facing the problem of generating AEs with no information during the whole attacking process. In the image domain, there have been works that demonstrated AEs crafted for the target model are able to attack other unknown models, which is called the transferability of AEs. However, in the audio domain, the poor transferability of audio AEs [35, 90] among different ASR systems indicates that we cannot directly leverage the transferability property especially when there are no interactions with the target model.

Figure 6: The main architecture of the Kaldi model.

Observing that the ultimate goal of speech recognition systems is to perform the task of converting natural speech into text, we believe that the inclusion of the characteristics of natural command audios in the constructed audio AEs may help improve their transferability. Inspired by model inversion attacks [29, 36, 94] that aim to recover input data or its sensitive attributes via the model output, we propose NI-Occam to craft audio AEs, where the command voice is recreated and implicitly embedded in the original music via the gradient update. The main reason behind is that it is hard for people to perform speech separation on our constructed audio AEs and further recognize the malicious speech commands.

Finally, audio AEs we constructed, just like natural command audios, will remain robust in the over-the-air attack and naturally effective in the physical world. Besides, compared to cloud speech APIs, voice control devices are more vulnerable to audio AEs containing command audios since they are more sensitive to voice commands than speech APIs, which has been demonstrated in Devil’s Whisper [32]. Thus, we propose NI-Occam, which realizes non-interactive physical attacks against voice control devices. Notably, this is the first physical attack that can effectively create audio AEs without any feedback information. Previous physical attacks, *e.g.*, Devil’s Whisper and FakeBob, rely much on returned scores to generate effective audio AEs, while our attack is a non-interactive one, *i.e.*, requiring no access to the target devices.

Our proposed NI-Occam is presented in Alg. 2. To be specific, we choose the open-source Kaldi model (ASpIRE Chain Model) as the substitute model and the inversion model due to its simple structure of the neural network and the excellent performance. In Figure 6, it can be seen that the Kaldi model obtains MFCC feature through feature extraction and takes these features as the input of DNN, while the output of DNN is the probability density function (pdf). According to the idea of “pdf-id sequence matching” proposed in CommanderSong [90], we can recover the command audios from their pdf-id sequences via the gradient inversion. Different from model inversion attacks that start from the Gaussian noise $z \sim \mathcal{N}(0, \sigma^2)$, adversarial attacks start from the original example x . Therefore, we add Gaussian noise z into the original example x to address this problem. The audio AEs can be described as

$$x^* = \arg \min_{x^*} \mathcal{J}(z + x^*, y) \quad s.t. \quad \|x - x^*\|_{\infty} < \epsilon, \quad (11)$$

where y is the probability value of the target pdf-id sequence, $\mathcal{J}(\cdot, \cdot)$ is the loss function [90]. Note that, the Gaussian noise z we added to the original example x is very large. This is to alleviate the impact of the original example on the model inversion process. To facilitate convergence, we gradually attenuate the size of Gaussian noise z in the iterative process. Besides, very recent works [83, 87] have shown that the AdaBelief method [96] is beneficial to improve the transferability of adversarial examples, so we use the AdaBelief optimizer to solve Eq. (11). The details of implementation are: the learning rate α is set to 0.003, the standard deviation σ is set to 0.25, and the size of perturbation ϵ is set to 0.3.

Algorithm 2 NI-Occam

Input: The original example x , the command audio x' , the loss function $\mathcal{J}(\cdot, \cdot)$, the Kaldi model f , the learning rate α , the standard deviation σ and the size of perturbation ϵ .

Output: The produced audio adversarial examples.

```

1: Obtain the target pdf-id sequence  $y$  through the Kaldi model  $f$ 
   and the command audio  $x'$ ;
2: Initialize  $x^*$  with  $x$ ,  $X^*$  with  $\emptyset$ , and the learning rate of the
   AdaBelief Optimizer with  $\alpha$ ;
3: while not converged yet do
4:   Sample  $z \sim \mathcal{N}(0, \sigma^2)$ ;
5:   Use the AdaBelief Optimizer to minimize  $\mathcal{J}(z + x^*, y)$  and
     update  $x^*$ ;
6:   Clip  $x^*$  into the  $\epsilon$  vicinity of  $x$ ;
7:    $\sigma = 0.998 \times \sigma$ ;
8:    $X^* = X^* \cup x^*$ ;
9: end while
10: return  $X^*$ .
  
```

5 IMPLEMENTATION AND EVALUATIONS

5.1 Experiment Settings

Experiment Design. To evaluate the performance of Occam, *i.e.*, our decision-only digital attacks on cloud speech APIs, we design four sets of experiments: attacks on open-source ASR systems, attacks on ASR services, attacks on SR services, and human perception of the audio AEs³. For the attacks on ASR services, we choose ten frequently-used voice commands⁴ that are expected to be recognized by the target systems. The original audios are selected from three datasets including Common Voice [8], Song [90], and LibriSpeech [67]. We select ten test samples in each experiment to evaluate the performance. To generate audio AEs, we need to query the target commercial speech-to-text cloud services⁵ in Table 2 and get the final transcription in return according to our proposed algorithms. To access the commercial cloud APIs, we first register on the platforms and then use audio AEs to query the speech-to-text services according to API keys provided by the platforms. For the attacks against SR services, we conduct our targeted attack on three systems (also in Table 2). For Microsoft SI and Jingdong SV, we choose ten people from Voxceleb dataset [64] and enroll four utterances for each person. Since Microsoft SV is text-dependent, we collect five volunteers’ voice data, and each volunteer is required to enroll 10 fixed sentences. To generate AEs, we query the APIs for 10,000 times for each target person.

To evaluate the performance of NI-Occam, *i.e.*, our non-interactive physical attacks on voice control devices, we design two sets of experiments: attacks on popular commercial voice assistants and human perception of the audio AEs. We evaluate NI-Occam on Google Assistant, Apple Siri, Microsoft Cortana, iFlytek, and Amazon Echo, as shown in Table 3. We generate 10 sets of audio AEs locally, whose target phrases of AEs are the same as those in the

³The evaluations of attacks on open-source ASR systems and human perception are presented in Appendixes D.1 and D.3, respectively

⁴The target transcripts include *call my wife, make it warmer, navigate to my home, open the door, open the website, play music, send a text, take a picture, turn off the light, and turn on the airplane mode.*

⁵Details on the datasets, target systems, and hardware are given in Appendix C.

Table 2: Details of the commercial speech and speaker recognition services.

Task	Commercial Services	Return results
Speech Recognition	Google STT [‡] [5]	D ^b +S [‡]
	Microsoft ASR [6]	D+S
	Alibaba SSR [§] [11]	D
	Tencent SSR [10]	D
	iFlytek [3]	D
Speaker Recognition	Microsoft SI [‡] [6]	D+S
	Microsoft SV [†] [6]	D+S
	Jingdong SV [9]	D

Note that, (i) [‡]: “STT” means “Speech-to-Text”. We use the command_and_search model in Google STT API. (ii) [§]: “SSR” means “Short Speech Recognition”. (iii) [‡]: “SI” means “Speaker Identification”. (iv) [†]: “SV” means “Speaker Verification”. (v) ^b: “D” means the model returns the final decision. (vi) [‡]: “S” means the model returns the confidence score or confidence level.

Table 3: Details of the commercial voice control devices.

Voice Assistant	Version	Device	Audio Source	Speaker
Apple Siri	13.6.1	iPhone 11	Default media player, ThinkPad X1 Carbon	JBL Clip3
iFlytek	10.0.8	iPhone 11		
Microsoft Cortana	3.3.3	Samsung C9		
Google Assistant	2.5.1	Nokia 7plus		
Amazon Echo	631499520	Echo 1st gen		

evaluation of Occam. We then play them using a JBL Clip 3 portable speaker near the target devices in a quite laboratory. The distance between the speaker and the target devices is around 15cm.

Methods for Comparison. We compare Occam with two state-of-the-art black-box adversarial attacks against ASR and SR services, *i.e.*, Devil’s Whisper [32] and FakeBob [26]. Since Devil’s Whisper originally utilized confidence scores to filter the synthetic audio data, in the follow-up evaluations, we omit this step to adapt Devil’s Whisper to the decision-based attack. To better evaluate the effectiveness of Occam, we also select five decision-based attacks for comparison: the boundary attack [21], the HopSkipJump attack (HSJA) [27], the opt-attack [33], the evolutionary attack⁶ [34], and the Differential Evolution attack (DEA)⁷ [68].

We compare NI-Occam with a straightforward non-interactive attack, *i.e.*, the superposition attack. In this attack, we just directly superimpose the original example and the command audio. Since this procedure does not require any knowledge of the target systems, the superposition attack is non-interactive.

Evaluation Metrics. We use the success rate of attack (SRoA) to evaluate the effectiveness of AEs. SRoA calculates the proportion of AEs that can successfully attack ASR or SR services. Besides, we use SNR⁸ to describe the perturbation on audio AEs and the number of queries on the target model to indicate the efficiency of the attacks. It is worth noting that the success rate or the SNR is not the only metric that determines whether an AE can successfully

attack the target system. An effective AE should fool both the model and the human. Hence, we should combine both the SRoA and SNR when evaluating the effectiveness of AEs. Note that, to illustrate this statement, we conduct a user study to analyze how SNRs affect the human perception. The results show that a lower SNR makes it easier for the users to notice or even recognize the AEs. Due to the space limit, we present the detailed results in Appendix E.

5.2 Evaluation on Cloud Speech APIs

Effectiveness of Occam on ASR services. Table 4 shows the performance of the targeted attacks on different commercial speech services after 30,000 queries (10,000 queries on Google). With regard to SNR, Occam performs the best among the seven attacks, reaching the best SNR of 17.84dB. Notably, DEA only obtains an average SNR of 6.37dB, which means that the perturbations of the audio AEs are very large. The results demonstrate that as the gradient-free optimization method in Occam’s cooperative co-evolution framework, CMA-ES is a better choice than the differential evolution. The reason behind this may be that CMA-ES is more suitable than DEA to solve the non-separable optimization, since CMA can well learn the dependencies between variables. However, the average SNR of AEs generated by the evolutionary attack (which is based solely on CMA-ES) can only achieve 7.11dB, while Occam has an average SNR of 14.37dB. This indicates that although CMA-ES can solve the non-separable optimization problem, when regarding to complex speech data, CMA-ES becomes ineffective to solve the discontinuous large-scale global optimization problem. Hence, CMA-ES alone cannot deal with complex speech data well. The above results demonstrate that Occam can effectively manage high-dimensional audio data.

Compared to Devil’s Whisper, Occam can achieve 100% SRoAs on all API services, while Devil’s Whisper can only achieve an average SRoA of 54%. This indicates that Devil’s Whisper is less effective in fooling the target speech recognition model than Occam. Figure 10 in Appendix D.2 shows the waveforms and spectrograms of the original audio and the adversarial audios generated from Occam and Devil’s Whisper. We can see from the waveforms that the audio AE generated from Occam is almost the same as the original audio. However, the differences in Devil’s Whisper are more noticeable and thus more likely to be perceived by humans. Besides, although we only choose ten commands in the experiments, Occam can generate AEs of arbitrary phrases⁹, while Devil’s Whisper can only generate a limited number of target phrases with a trained model. These observations suggest our decision-based Occam is more effective, powerful, and practical. We also evaluate the untargeted attacks. Due to space limitation, related results are given in Appendix D.2. **Effectiveness of Occam on SR services.** We evaluate the SNRs and SRoAs of the AEs against SR services, as in Table 5. Among the attacks against SR services, Occam still achieves 100% SRoAs, indicating that the audio AEs can be successfully recognized as the target person (or mislead the SR system). FakeBob can only achieve a 1% SRoA. Note that we tested FakeBob with 200 instead of 10 AEs. Since the success rate of FakeBob is too low, we enlarge the AE set to create an effective AE. We find that the results given by FakeBob are higher than those in Table 5. This is probably because

⁶In our experiments, the bilinear interpolation in the evolutionary attack is removed in the audio domain as it is not applicable to the time-series speech data.

⁷We give a detailed description on DEA in Appendix B.

⁸SNR calculates the ratio of the signal power of the original audio to the noise power. A higher SNR indicates a lower noise level.

⁹To illustrate this, we further evaluate Occam on a large group of target phrases. Related results are presented in Appendix G.

Table 4: Experimental results on targeted attacks on commercial cloud speech-to-text APIs.

Cloud service	Boundary Attack		Opt-Attack		Evolutionary Attack		HSJA		DEA		Devil's Whisper		Occam	
	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR
Google STT	10/10	6.80	10/10	5.52	10/10	6.11	10/10	5.38	10/10	6.52	5/10	9.13	10/10	15.33
Microsoft ASR	10/10	8.71	10/10	6.45	10/10	6.60	10/10	4.59	10/10	7.35	8/10	9.91	10/10	10.74
Alibaba SSR	10/10	9.32	10/10	5.70	10/10	7.20	10/10	8.11	10/10	5.44	3/10	8.21	10/10	17.84
Tencent SSR	10/10	11.03	10/10	6.43	10/10	9.14	10/10	8.49	10/10	6.06	4/10	9.64	10/10	16.70
iFlytek	10/10	6.78	10/10	6.73	10/10	6.50	10/10	5.08	10/10	6.47	7/10	9.27	10/10	11.22
Average	10/10	8.53	10/10	6.17	10/10	7.11	10/10	6.33	10/10	6.37	5.4/10	9.23	10/10	14.37

Note that, Devil's Whisper originally utilized confidence scores to filter the synthetic audio data. We omit this step to adapt Devil's Whisper to the decision-based attack.

Table 5: Experimental results on targeted and untargeted attacks on cloud speaker recognition APIs.

Cloud service	Boundary Attack		Opt-Attack		Evolutionary Attack		HSJA		DEA		FakeBob		Occam	
	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR
Microsoft SI	10/10	7.13	10/10	7.01	10/10	7.25	10/10	4.87	10/10	2.84	0/200	N/A	10/10	14.31
Microsoft SV	10/10	8.48	10/10	5.85	10/10	9.93	10/10	4.27	10/10	3.05	0/200	N/A	10/10	13.25
Jingdong SV	10/10	6.26	10/10	6.42	10/10	5.97	10/10	5.89	10/10	5.67	2/200	35.30	10/10	13.78
Microsoft SI (untargeted)	10/10	9.61	10/10	9.95	10/10	9.04	10/10	9.43	10/10	6.53	2/200	40.21	10/10	14.92
Average	10/10	7.87	10/10	7.31	10/10	8.05	10/10	6.12	10/10	4.52	1/200	37.76	10/10	14.07

Note that, (i) we only evaluate untargeted attacks against Microsoft Azure SI. For speaker verification systems, the goal of untargeted attacks is the same as that of targeted attacks. (ii) N/A denotes "not available". Since there is no effective AE against Microsoft SI and SV in FakeBob, the SNR is not available.

Table 6: Experimental results of the physical attacks.

Devices		Siri	iFlytek	Cortana	Google	Echo
NI-Occam	SRoA	6/10	6/10	6/10	4/10	4/10
	SNR (dB)	9.81	9.09	9.58	10.36	9.42
Super-position attack	SRoA	5/10	1/10	1/10	1/10	2/10
	SNR (dB)	7.00	7.00	7.00	7.00	7.00

the upgrading of Microsoft speaker recognition systems makes the transferability of the AEs in FakeBob no longer effective. As for SNR, the AEs from Occam can achieve pretty good SNRs, outperforming other decision-based attacks. Although effective AEs from FakeBob can achieve SNR as high as 35.3dB, the SRoA is too low to effectively mislead the SR services and thus far from being practical.

5.3 Evaluation of NI-Occam against Voice Control Devices

Effectiveness. We test the effectiveness of NI-Occam on 5 voice control devices, and the results are given in Table 6¹⁰. If the audio AE can be correctly recognized by the devices as the target command within 3 attempts (play the AEs within three times), we consider this AE successful. Overall, NI-Occam achieves an average SRoA of 52% and SNR of 9.65dB. Note that NI-Occam is a non-interactive attack that requires no access to the target devices, which is very practical since some devices like Apple Siri do not provide a programmable API. For example, Devil's Whisper [32] fails to generate effective AEs when confronted with devices that do not return confidence scores. We also find that NI-Occam performs well on Apple Siri, while Devil's Whisper failed in attacking updated versions of Siri.

¹⁰Detailed results on individual commands can be found in Appendix F (see Table 12).

This indicates that NI-Occam is more effective in leveraging the useful transferability of the AEs and can generate successful AEs with minimal information from the target devices. We also evaluate a simple non-interactive attack as the baseline, *i.e.*, the superposition attack. Since the AEs from the superposition attack are constructed by superimposing two audios, the SNR is adjustable. We set the SNRs as 7.00dB, smaller than those of NI-Occam. A smaller SNR means the audio of the target command is more obvious in the audio AE, making it easier for the devices to recognize. However, the superposition attack can only achieve an average SRoA of 20% with a 7.00dB SNR. More importantly, the superposition attack is easily perceived by human, which we will discuss shortly.

We also evaluate the impact of the number of attempts on SRoA with a larger group of target phrases. We observe that increasing the query attempts can help increase SRoA to 70%, and our NI-Occam can also perform well on large sets of 60 commands with an SRoA of 71.7%. Detailed results are given in Appendixes F and G. **Human Perception.** Although SNR describes the proportion of noise, it cannot fully reflect the imperceptibility of the audio. For example, if the noise can well fit the background, even though the signal has a low SNR, the users cannot perceive it. On the contrary, if the noises all appear in a small piece of audio, although the overall SNR is high, the users may readily perceive the command.

To evaluate the performance of NI-Occam on human perception, we surveyed 37 volunteers aged from 19 to 24 (who are sensitive to sound), including 21 males and 16 females. In the experiment, we first show some examples of "noise", "normal", "talking", and "recognized", and then ask the volunteers to listen to the audio AEs and tell their views about them. Specifically, the audio AEs are generated from NI-Occam and the superposition attack and can successfully attack the devices. Each volunteer ranks 6, 6, 6, 4, and

Table 7: Evaluation results on human perception of the non-interactive physical attacks.

Devices	Method	Normal (%)	Noise (%)	Talking (%)	Once-recognize (%)	Twice-recognize (%)
Apple Siri	NI-Occam	12.4	61.7	25.9	4.1	5.9
	Superposition attack	0.0	0.0	100	88.1	91.9
iFlytek	NI-Occam	9.6	58.3	32.1	5.1	6.9
	Superposition attack	0.0	0.0	100	86.5	94.6
Microsoft Cortana	NI-Occam	14.2	56.7	29.1	6.1	8.3
	Superposition attack	0.0	0.0	100	83.8	86.5
Google Assistant	NI-Occam	15.0	55.7	29.4	4.8	8.2
	Superposition attack	0.0	0.0	100	83.8	86.5
Amazon Echo	NI-Occam	2.7	68.2	29.1	12.2	14.2
	Superposition attack	0.0	0.0	100	86.5	94.6

Note that, (i) “Normal” means that the volunteer regards the audio as a normal audio. (ii) “Noise” means that the volunteer can feel some noises. (iii) “Talking” means that the volunteer can hear talking in the audio. If the volunteer thinks there is talking in the audio, he/she is then asked to recognize the content of the talking. (iv) The audio will be labeled as “once-recognize” or “twice-recognize” if the volunteer recognizes over half of the content after listening to the audio once or twice, respectively.

4 successful AEs from NI-Occam on Apple Siri, iFlytek, Microsoft Cortana, Google Assistant, and Amazon Echo, respectively¹¹.

Table 7 presents the results of the experiments on human perception. Overall, NI-Occam performs much better than the superposition attack. More than 67% volunteers think the audio generated from NI-Occam is normal or just noisy, while 100% volunteers can recognize the AEs from the superposition attack. This is because the audios crafted in the superimposing manner can be more easily noticed due to the human ear’s excellent ability of speech separation. We also assume that a proper noise level is acceptable because the background environment may be noisy, and the equipment may emit some noise due to a temporary fault. The results show that NI-Occam is stealthy enough and cannot be easily perceived.

6 RELATED WORK

6.1 Audio Adversarial Examples

Adversarial Examples Against ASR Systems. Despite the great success of adversarial examples in the image domain, the transcription of spontaneous speech poses a more significant challenge for crafting audio AEs. The early results have clearly indicated that ASR systems are inherently vulnerable to AEs in white-box settings. Among others, Carlini *et al.* [25] was the first to implement an iterative optimization-based attack with a success rate of 100% on the end-to-end Mozilla DeepSpeech model. However, it takes approximately one hour for their attack to produce one adversarial example on a single NVIDIA 1080Ti GPU, and the crafted adversarial sample fails when being played over the air. Concurrently, CommanderSong [90], which embedded malicious commands into popular songs, was reported to have successfully attacked against Kaldi [2]. It further launched a very limited over-the-air attack, which is heavily dependent on the recording devices, speakers and room settings. Efforts were also made by [86] to build robust over-the-air AEs by utilizing impulse responses to simulate the reverberation with a success rate of around 60%. Moreover, Chen *et al.* [31] achieved a 90% success rate of over-the-air attacks over the attack distance of up to 6m by further removing device- and environment-specific features. Besides, imperceptible audio AEs

were produced in [74] via the psychoacoustic model. Both imperceptible and robust audio AEs were constructed in [69] against the Lingvo ASR system, with a success rate of 50%.

Compared to the above attacks, our attacks have the following superiorities: 1) Our attacks are more practical since these attacks have access to the internal information of target model; 2) Instead of only targeting one open-source system in the white-box setting, our attacks evaluate the robustness of many representative audio processing systems in real-world scenarios.

While adversarial example based white-box attacks have obtained excellent results against open-source ASR systems, its black-box counterpart hasn’t made big progress until recently. By leveraging the last layer (*i.e.*, logits) of DNNs inside the DeepSpeech, Taori *et al.* [80] obtained the fitness score of adversarial inputs, and combined genetic algorithms with the gradient estimation to attack DeepSpeech. In addition to the rather low success rates even after 300,000 queries, this type of black-box attack is not applicable to commercial systems. Following this work, selective gradient estimation attack [82] was proposed to achieve a success rate of 98% in this setting. Moreover, multi-objective genetic algorithms were also introduced to start a black-box adversarial attack against ASR systems [53]. However, due to the relatively large word error rate (WER) after many evolutions in the attack, it becomes ineffective for generating audio AEs. A very recent work, Devil’s Whisper [32], utilized confidence scores exposed by commercial ASR systems to launch the black-box attack. However, it is worth noting that there are many commercial ASR systems that do not return any score information, and service providers are also apt to hide these scores to reduce the risk of adversarial attacks, considering that these information almost has no effect on the user experience and may be mainly exploited by malicious attackers.

There are two key differences: 1) Our Occam, as a generic attack, requires only the final decisions to generate audio AEs, effective to both commercial ASR and SR services; 2) Our NI-Occam requires no access to the targeted devices, but still can generate effective audio AEs. Our attacks are more practical in real world scenarios.

Adversarial Examples Against SR Systems. When it comes to adversarial example generation against SR systems, relatively little work has been done in both white-box and black-box cases. Kreuk *et al.* [55] presented a white-box adversarial attack against the DNN-based speaker verification system. Gong *et al.* [39] demonstrated

¹¹The AEs are the ones that can fool the devices, as shown in Table 6.

the vulnerability of speaker identification system to audio AEs in the white-box scenario. Obviously, these attacks require access to internal structures and parameters of the target systems, and they are impractical when facing commercial SR systems.

In a recent work, SirenAttack [35] was presented to launch a black-box attack against a number of classification-oriented acoustic systems, including the SR system via the predicted probabilities/scores. More recently, FakeBob [26] also utilized the predicted probabilities/scores to conduct a black-box adversarial attacks against Talentedsoft API [7] with a success rate of 100%. The main limitation of SirenAttack and FakeBob is that they lose the effectiveness when applied to commercial SR systems, which usually hide the predicted scores. For example, the Microsoft Azure SR API service only provides the decision (*i.e.*, the predicted speaker) along with three confidence levels (*i.e.*, low, normal, or high) to users. Our attack shows its superiority to prior attacks by achieving an attack success rate of 100% against commercial SR systems even if the service provider conceals the prediction score information.

6.2 Other Types of Attacks

In addition to audio AEs, researchers have also discovered that intelligent voice systems are vulnerable to other types of attacks, including misinterpretation attack and hidden voice attack.

Misinterpretation Attacks. Kumar *et al.* [57] conducted an empirical analysis of misinterpretations and investigated security implications on Amazon Alexa, based on which they introduced a new attack called skill squatting to surreptitiously route users to malicious third-party public services. Along this direction, Zhang *et al.* [93] reported similar attacks against ASR systems, where a malicious skill with the similarly pronounced name or paraphrased name was exploited to impersonate a benign skill. Also targeting at ASR systems, Zhang *et al.* [95] designed a linguistic-model guided fuzzing tool called LipFuzzer to systematically discover misinterpretations leading to malicious attacks.

Hidden Voice Attacks. Besides, by either exploiting knowledge of the feature extraction algorithm or hardware vulnerabilities in microphone circuits, the adversary can embed hidden commands into an audio carrier, in the form of noises, thereby compromising the intelligent voice systems. To achieve this goal, hidden voice command [23] utilized inverse MFCC to craft obfuscated commands against ASR systems in a “black-box” manner with considerable human effort for obtaining feedbacks. Four different perturbations were introduced to generate noise-like adversarial audios against ASR and SR systems [12]. Moreover, DolphinAttack [92] was devised to modulate voice commands on inaudible ultrasounds, which can be interpreted by the ASR system, by exploiting the non-linearity of the microphone circuits. However, compared to adversarial example based attacks, these attacks could be easily defended or perceived.

Our initial idea stems from the image-based adversarial attacks. We have overcome particular challenges for generating audio AEs. Compared to the state-of-the-art black-box attacks on acoustic systems [26, 32, 35, 80], ours is more generic and practical.

7 DISCUSSIONS

We discuss four possible countermeasures to defend against our Occam and NI-Occam below, with detailed performance results of countermeasures in Appendix H (Tables 14 and 15).

Local Smoothing. Because audio AEs are carefully constructed by adding small perturbations, they can be mitigated by local smoothing. We can apply a sliding window with the median filter to the adversarial audio signals. Given a data point x_i , we replace it with the average of k samples before and after it, *i.e.*, $[x_{i-h}, \dots, x_i, \dots, x_{i+h}]$. Since the adversarial perturbation is carefully constructed in our attacks, the audio AEs may become less effective after local smoothing transformation. For example, when $h = 1$, the SRoA of Occam drops from 100% to 20%, while the SRoA of NI-Occam drops from 60% to 40%. When $h = 3$, all AEs of Occam fail, and NI-Occam remains an SRoA of 40%. We can find that NI-Occam is more robust to local smoothing. The reason is that the recognized key parts of natural command audios are recovered from audio AEs in NI-Occam, thus making them more robust.

Downsampling. Based on the sampling theory, the high-frequency information in the audios will be lost after the downsampling process, which may disrupt perturbations in the audio AEs. Thus, audio AEs crafted in our Occam will fail to work after the downsampling and upsampling process, *e.g.*, audio AEs are first downsampled to 12kHz and then upsampled to 16kHz as indicated in our experiments. Nonetheless, NI-Occam remains a better SRoA, *i.e.*, 30% when the downsampling rate (DSR) is 12kHz, indicating that the AEs of NI-Occam have a greater chance of resisting downsampling. While downsampling can help to mitigate our attacks, if the downsampling/upsampling rates are known to the attacker, this countermeasure will be invalid. This is because the adversary can directly compensate the generated example for the information lost in the downsampling and upsampling process.

Temporal Dependency Based Approach. The inherent temporal dependency in audio data was recently utilized in [88] to detect audio AEs due to the disruption of the temporal information. Namely, the first k ($0 \sim 1$) portions of the whole audio were selected and recognized as S_k , and the k portions of the transcription of the whole audio $S_{\{whole,k\}}$ is obtained and compared with S_k . By checking the consistency between $S_{\{whole,k\}}$ and S_k , audio AEs crafted by Occam can be easily detected. This approach has a strong discriminative ability in identifying AEs targeting at ASR systems, and theoretically it can identify almost all audio AEs against ASR systems. Even so, Occam is still effective in attacking SR API services since AEs constructed for SR systems preserve the temporal information like natural audios. For instance, Occam also achieves a success rate of 80% against Microsoft Azure speaker identification API when our audio AEs were randomly split into two parts, *i.e.*, k and $1 - k$ portions of the whole audio.

The defense [88] is better suited for ASR tasks since the temporal dependency is stronger. To evaluate the effectiveness of the temporal dependency based approach against (NI-)Occam for ASR, we build a dataset of 40 audio AEs generated using (NI-)Occam with 40 natural command audios, and calculate the detection rates. For each audio sample, we randomly select k in the range of $[0.2, 0.8]$ and split the audio into two pieces. We then calculate the consistency of the split audios and the whole audio by the word error rate (WER). With a varying WER threshold, we can obtain the ROC curve and finally calculate the area under curve (AUC), where a higher AUC indicates better detection performance. The AUC is 100% in Occam, showing that the defense can successfully detect all AEs generated by Occam. However, the AUC drops to 68% when classifying AEs

from NI-Occam, which means that NI-Occam is robust to the temporal dependency based approach. This is because the adversarial perturbation generated by NI-Occam usually only occurs in a small piece of audio, and splitting the audio does not disrupt the temporal dependency. Therefore, it is hard for the classifier to detect the AEs generated using NI-Occam by analyzing the temporal information.

Adversarial Training. Adversarial training [40, 61] is one of the most effective defenses against adversarial attacks in the image domain [17]. The basic idea of adversarial training is to train models on AEs to make the models robust to AEs. It can be formulated as a mini-max optimization problem:

$$\min_{\theta} \mathbb{E}_{(x,y) \sim D} \left[\max_{\|\delta\|_p < \epsilon} \mathcal{J}(\theta, x + \delta, y) \right], \quad (12)$$

where θ denotes the deep learning model, (x, y) denotes the original data point, δ is the adversarial perturbation, and $\|\cdot\|_p$ is the p-norm. Here, the worst-case samples for the given model are found in the inner maximization problem to train a more robust model via the outer minimization operation.

So far, adversarial training has been extensively studied on image classification tasks [18, 22, 76, 81, 85]. However, there is relatively little research on this defense for speech recognition tasks, which are more complex and challenging. Besides, generating strong audio AEs incurs high computation costs. For example, [25] reports that it takes approximately one hour to construct an audio AE in a single NVIDIA 1080Ti GPU, far exceeding the time of generating an adversarial image. In order to study the effect of adversarial training on ASR systems, we evaluate the performance of adversarial training on the open-source Kaldi. For solving the inner maximization problem of Eq. (12) that, our targeted NI-Occam is not suitable as it performs the minimization operation, and thus we adopt the untargeted projected gradient descent (PGD) attack [61], which is usually used in adversarial training, as follows:

$$x^{i+1} = x + \mathcal{P}_{\epsilon,p}(x^i - x + \alpha \cdot \text{sign}(\nabla_{x^i} \mathcal{J}(\theta, x^i, y))), \quad (13)$$

where $\mathcal{P}_{\epsilon,p}$ is a projection operator on L_p ball, and α the step size.

According to our experimental results on Kaldi (Mini LibriSpeech model¹²) (see Table 15 in Appendix H), adversarial training is indeed an effective defense against our targeted NI-Occam attack, e.g., the SRoA of the AEs drops to 30% when $\epsilon = 0.002$. But, meanwhile, the WER of the model is increasing from 10.69 to 19.82, which means that the accuracy of the speech recognition drops about 10%. Moreover, with the increase of ϵ , both the SRoA and model accuracy will be significantly reduced. Considering the extreme case when almost all audio AEs fail, the model accuracy drops about 20%. And if ϵ was further increased to eliminate our attack, e.g., approaching $\epsilon = 0.006$, adversarial training is not even able to converge. The reason may be that the audio vector contains many values very close to zero, and they are changed a lot with a high ϵ .

Overall, adversarial training on Kaldi is effective against our attacks. However, while the SRoA can be reduced by 70%~90%, it also inevitably brings a significant performance degradation, i.e., a 10%~20% accuracy loss, and such loss is unacceptable for commercial ASR systems. Moreover, adversarial training will significantly increase the training time and costs [51, 76], particularly

for ASR systems. For example, for the Mini LibriSpeech dataset containing only 5 hours of audio data, it took about 10 days for us to adversarially train Kaldi model on six NVIDIA 2080Ti GPUs, while commonly-used voice datasets, like LibriSpeech and Common Voice, have around 1000 hours of voice data), making it almost impractical on large-scale models and datasets. Thus, even with adversarial training our attack cannot be prevented. Therefore, service providers may not be willing to adopt the defense for “black-box” commercial models due to the issues above.

Remarks. We point out several potential research directions to obtain more practical attacks. Since the human perception experiment has shown that some audio AEs can be recognized and regarded as abnormal audios, it is necessary to further improve the stealthiness of constructed audio AEs. Moreover, the physical attack against voice control devices in this work will fail when the distance is long or in a very noisy environment. For example, when we play the AEs at a distance of 50cm from the devices in a noisy cafe, the SRoAs of NI-Occam against Amazon Echo and Apple Siri decrease to 20%. Thus, it is still challenging to launch physical and black-box attacks on ASR systems at long distances or in noisy environments, which requires more efforts in the future. Besides, robust physical adversarial attacks against SR systems in the decision-only and non-interactive settings should also be put on the agenda. Finally, it is also an interesting and meaningful job to achieve black-box adversarial attacks against both ASR and SR systems on one audio AE, because some smart speakers, such as Apple HomePod, will first perform speaker identification on the input audio.

8 CONCLUSION

In this paper, we proposed two novel black-box adversarial attacks against commercial speech platforms, Occam and NI-Occam. Occam constructs audio AEs against cloud speech APIs in the decision-based black-box setting. It is effective under the strictly black-box scenario where the attackers can rely solely on the final decisions. Extensive experiments on targeted and untargeted attacks against a wide range of popular open-source and commercial ASR and SR systems demonstrated the effectiveness of Occam. Occam achieves an average SNR of 14.37dB and 100% SRoA on commercial ASR systems, outperforming the state-of-the-art black-box audio adversarial attacks. NI-Occam is the first non-interactive physical attack, simple but effective, which can successfully fool commercial devices without needing any feedback information from the target devices. Extensive experiments showed the effectiveness of the AEs from NI-Occam in attacking Apple Siri, Microsoft Cortana, Google Assistant, iFlytek, and Amazon Echo with an average SRoA of 52% and SNR of 9.65dB.

ACKNOWLEDGMENTS

This work was partially supported by the National Key R&D Program of China (2020AAA0107701), NSFC under Grants U20B2049, 61822207, 61822309, 61773310, 62132011, and U1736205, BNRist under Grant BNR2020RC01013, RGC of Hong Kong under Grants CityU 11217819, CityU 11217620, and R6021-20F, and Laboratory for AI-Powered Financial Technologies.

¹²It is a TDNN based chain Kaldi model trained on Mini-Librispeech dataset. The URL is https://github.com/kaldi-asr/kaldi/tree/master/egs/mini_librispeech/s5.

REFERENCES

- [1] 2014. Alexa. <https://developer.amazon.com/en-US/alexa>.
- [2] 2014. Kaldi. <http://kaldi-asr.org>.
- [3] 2015. iFlytek Speech-to-Text. <https://www.xfyun.cn/services/voicedictation>.
- [4] 2018. Alexa Auto SDK. <https://developer.amazon.com/en-US/alexa/alexa-auto/sdk>.
- [5] 2018. Google Cloud Speech-to-Text. <https://cloud.google.com/speech-to-text>.
- [6] 2018. Microsoft Azure Speech Service. <https://azure.microsoft.com/en-us/services/cognitive-services/speech-services/>.
- [7] 2018. Talentedsoft. <http://www.talentedsoft.com/>.
- [8] 2020. Common Voice Dataset. <https://voice.mozilla.org/en/datasets>.
- [9] 2020. Jingdong Speaker Recognition. <https://www.jdcloud.com/en/products/voiceprint-recognition>.
- [10] 2020. Tencent Short Speech Recognition. <https://cloud.tencent.com/document/product/1093>.
- [11] 2021. Alibaba Short Speech Recognition. <https://www.alibabacloud.com/zh/product/intelligent-speech-interaction>.
- [12] Hadi Abdullah, Washington Garcia, Christian Peeters, Patrick Traynor, Kevin R. B. Butler, and Joseph Wilson. 2019. Practical Hidden Voice Attacks Against Speech and Speaker Recognition Systems. In *Proc. of NDSS*.
- [13] Moustafa Alzantot, Yash Sharma, Supriyo Chakraborty, Huan Zhang, Cho-Jui Hsieh, and Mani B. Srivastava. 2019. Genattack: Practical Black-box Attacks with Gradient-free Optimization. In *Proc. of GECCO*.
- [14] Tawfiq Ammari, Jofish Kaye, Janice Y. Tsai, and Frank Bentley. 2019. Music, Search, and IoT: How People (Really) Use Voice Assistants. *ACM Transactions on Computer-Human Interaction* 26, 3 (2019), 17:1–17:28.
- [15] Anne Auger. 2009. Benchmarking the (1+1) Evolution Strategy with One-fifth Success Rule on the BBOB-2009 Function Testbed. In *Proc. of GECCO*.
- [16] Steve Austin, Chris Barry, Yen-Lu Chow, Man Derr, Owen Kimball, Francis Kubala, John Makhoul, Paul Placeway, William Russell, Richard M. Schwartz, and George Yu. 1989. Improved HMM Models for High Performance Speech Recognition. In *Proc. of Human Language Technology (HLT)*.
- [17] Tao Bai, Jinqi Luo, Jun Zhao, Bihan Wen, and Qian Wang. 2021. Recent Advances in Adversarial Training for Adversarial Robustness. In *Proc. of IJCAI*.
- [18] Yogesh Balaji, Tom Goldstein, and Judy Hoffman. 2019. Instance Adaptive Adversarial Training: Improved Accuracy Tradeoffs in Neural Nets. *CoRR abs/1910.08051* (2019).
- [19] Arjun Nitin Bhagoji, Warren He, Bo Li, and Dawn Song. 2018. Practical Black-box Attacks on Deep Neural Networks Using Efficient Query Mechanisms. In *Proc. of ECCV*.
- [20] Debnath Bhattacharyya, Rahul Ranjan, Farkhod Alisherov, Minkyu Choi, et al. 2009. Biometric Authentication: A Review. *International Journal of u-and e-Service, Science and Technology* 2, 3 (2009), 13–28.
- [21] Wieland Brendel, Jonas Rauber, and Matthias Bethge. 2018. Decision-based Adversarial Attacks: Reliable Attacks Against Black-box Machine Learning Models. In *Proc. of ICLR*.
- [22] Qi-Zhi Cai, Chang Liu, and Dawn Song. 2018. Curriculum Adversarial Training. In *Proc. of IJCAI*.
- [23] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David A. Wagner, and Wenchao Zhou. 2016. Hidden Voice Commands. In *Proc. of USENIX Security*.
- [24] Nicholas Carlini and David Wagner. 2017. Towards Evaluating the Robustness of Neural Networks. In *Proc. of IEEE S&P*.
- [25] Nicholas Carlini and David Wagner. 2018. Audio Adversarial Examples: Targeted Attacks on Speech-to-text. In *Proc. of IEEE SPW*.
- [26] Guangke Chen, Sen Chen, Lingling Fan, Xiaoning Du, Zhe Zhao, Fu Song, and Yang Liu. 2021. Who is Real Bob? Adversarial Attacks on Speaker Recognition Systems. In *Proc. of IEEE S&P*.
- [27] Jianbo Chen, Michael I. Jordan, and Martin J. Wainwright. 2020. HopSkipJumpAttack: A Query-efficient Decision-based Attack. In *Proc. of IEEE S&P*.
- [28] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. 2017. Zeroth Order Optimization Based Black-box Attacks to Deep Neural Networks without Training Substitute Models. In *Proc. of ACM AISec*.
- [29] Si Chen, Ruoxi Jia, and Guo-Jun Qi. 2020. Improved Techniques for Model Inversion Attacks. *CoRR abs/2010.04092* (2020).
- [30] Si Chen, Kui Ren, Sixu Piao, Cong Wang, Qian Wang, Jian Weng, Lu Su, and Aziz Mohaisen. 2017. You Can Hear But You Cannot Steal: Defending Against Voice Impersonation Attacks on Smartphones. In *Proc. of IEEE ICDCS*.
- [31] Tao Chen, Longfei Shanguan, Zhenjiang Li, and Kyle Jamieson. 2020. Metamorph: Injecting Inaudible Commands into Over-the-air Voice Controlled Systems. In *Proc. of NDSS*.
- [32] Yuxuan Chen, Xuejing Yuan, Jiangshan Zhang, Yue Zhao, Shengzhi Zhang, Kai Chen, and XiaoFeng Wang. 2020. Devil's Whisper: A General Approach for Physical Adversarial Attacks Against Commercial Black-box Speech Recognition Devices. In *Proc. of USENIX Security*.
- [33] Minhao Cheng, Thong Le, Pin-Yu Chen, Huan Zhang, Jinfeng Yi, and Cho-Jui Hsieh. 2019. Query-efficient Hard-label Black-box Attack: An Optimization-based Approach. In *Proc. of ICLR*.
- [34] Yinpeng Dong, Hang Su, Baoyuan Wu, Zhifeng Li, Wei Liu, Tong Zhang, and Jun Zhu. 2019. Efficient Decision-based Black-box Adversarial Attacks on Face Recognition. In *Proc. of IEEE/CVF CVPR*.
- [35] Tianyu Du, Shouling Ji, Jinfeng Li, Qinchun Gu, Ting Wang, and Raheem Beyah. 2020. SirenAttack: Generating Adversarial Audio for End-to-end Acoustic Systems. In *Proc. of ACM AsiaCCS*.
- [36] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. In *Proc. of ACM SIGSAC*.
- [37] Sadaoki Furui. 1997. Recent Advances in Speaker Recognition. *Pattern Recognition Letters* 18, 9 (1997), 859–872.
- [38] Andreas Geiger, Philip Lenz, and Raquel Urtasun. 2012. Are We Ready for Autonomous Driving? The Kitti Vision Benchmark Suite. In *Proc. of IEEE/CVF CVPR*.
- [39] Yuan Gong and Christian Poellabauer. 2017. Crafting Adversarial Examples for Speech Paralinguistics Applications. *CoRR abs/1711.03280* (2017).
- [40] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. Explaining and Harnessing Adversarial Examples. In *Proc. of ICLR*.
- [41] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech Recognition with Deep Recurrent Neural Networks. In *Proc. of IEEE ICASSP*.
- [42] Awni Y. Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng. 2014. Deep Speech: Scaling up End-to-end Speech Recognition. *CoRR abs/1412.5567* (2014).
- [43] Nikolaus Hansen. 2016. The CMA Evolution Strategy: A Tutorial. *CoRR abs/1604.00772* (2016).
- [44] Hynek Hermansky. 1990. Perceptual Linear Predictive (PLP) Analysis of Speech. *The Journal of the Acoustical Society of America* 87, 4 (1990), 1738–1752.
- [45] Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. 2012. Deep Neural Networks for Acoustic Modeling in Speech Recognition: The Shared Views of Four Research Groups. *IEEE Signal Processing Magazine* 29, 6 (2012), 82–97.
- [46] Matthew B. Hoy. 2018. Alexa, Siri, Cortana, and More: An Introduction to Voice Assistants. *Medical Reference Services Quarterly* 37, 1 (2018), 81–88.
- [47] Shengshan Hu, Xingcan Shang, Zhan Qin, Minghui Li, Qian Wang, and Cong Wang. 2019. Adversarial Examples for Automatic Speech Recognition: Attacks and Countermeasures. *IEEE Communications Magazine* 57, 10 (2019), 120–126.
- [48] Christian Igel, Thorsten Suttrop, and Nikolaus Hansen. 2006. A Computational Efficient Covariance Matrix Update and a (1+1)-CMA for Evolution Strategies. In *Proc. of ACM GECCO*.
- [49] Fumitada Itakura. 1975. Line Spectrum Representation of Linear Predictor Coefficients of Speech Signals. *The Journal of the Acoustical Society of America* 57, S1 (1975), S35–S35.
- [50] Arindam Jati, Chin-Cheng Hsu, Monisankha Pal, Raghuveer Peri, Wael AbdAlmageed, and Shrikanth Narayanan. 2021. Adversarial Attack and Defense Strategies for Deep Speaker Recognition Systems. *Computer Speech & Language* 68 (2021), 101199.
- [51] Sonal Joshi, Jesús Villalba, Piotr Żelasko, Laureano Moro-Velázquez, and Najim Dehak. 2021. Study of Pre-processing Defenses Against Adversarial Attacks on State-of-the-art Speaker Recognition Systems. *CoRR abs/2101.08909* (2021).
- [52] William S. Mohn Jr. 1971. Two Statistical Feature Evaluation Techniques Applied to Speaker Identification. *IEEE Trans. Comput.* 20, 9 (1971), 979–987.
- [53] Shreya Khare, Rahul Aralikatte, and Senthil Mani. 2018. Adversarial Black-box Attacks for Automatic Speech Recognition Systems Using Multi-objective Genetic Optimization. *CoRR abs/1811.01312* (2018).
- [54] Tomi Kinnunen and Haizhou Li. 2010. An Overview of Text-independent Speaker Recognition: From Features to Supervectors. *Speech Communication* 52, 1 (2010), 12–40.
- [55] Felix Kreuk, Yossi Adi, Moustapha Cissé, and Joseph Keshet. 2018. Fooling End-to-end Speaker Verification with Adversarial Examples. In *Proc. of IEEE ICASSP*.
- [56] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proc. of NeurIPS*.
- [57] Deepak Kumar, Riccardo Paccagnella, Paul Murley, Eric Hennenfent, Joshua Mason, Adam Bates, and Michael Bailey. 2018. Skill Squatting Attacks on Amazon Alexa. In *Proc. of USENIX Security*.
- [58] Yun Lei, Nicolas Scheffer, Luciana Ferrer, and Mitchell McLaren. 2014. A Novel Scheme for Speaker Recognition Using a Phonetically-aware Deep Neural Network. In *Proc. of IEEE ICASSP*.
- [59] Jesse Levinson, Jake Askeland, Jan Becker, Jennifer Dolson, David Held, Soeren Kammel, J Zico Kolter, Dirk Langer, Oliver Pink, Vaughan Pratt, et al. 2011. Towards Fully Autonomous Driving: Systems and Algorithms. In *Proc. of IEEE Intelligent Vehicles Symposium*.
- [60] Jinpeng Liu and Ke Tang. 2013. Scaling up Covariance Matrix Adaptation Evolution Strategy Using Cooperative Coevolution. In *Proc. of Intelligent Data Engineering and Automated Learning (IDEAL)*.

- [61] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *Proc. of ICLR*.
- [62] Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. 2016. Deepfool: A Simple and Accurate Method to Fool Deep Neural Networks. In *Proc. of IEEE/CVF CVPR*.
- [63] LindsaIwa Muda, Mumtaj Begam, and I. Elamvazuthi. 2010. Voice Recognition Algorithms Using Mel Frequency Cepstral Coefficient (MFCC) and Dynamic Time Warping (DTW) Techniques. *CoRR* abs/1003.4083 (2010).
- [64] Arsha Nagrani, Joon Son Chung, and Andrew Senior. 2017. VoxCeleb: A Large-scale Speaker Identification Dataset. In *Proc. of INTERSPEECH*.
- [65] Mahesh Kumar Nandwana, Luciana Ferrer, Mitchell McLaren, Diego Castán, and Aaron Lawson. 2019. Analysis of Critical Metadata Factors for the Calibration of Speaker Recognition Systems. In *Proc. of INTERSPEECH*.
- [66] Phani Sankar Nidadavolu, Vicente Iglesias, Jesús Villalba, and Najim Dehak. 2019. Investigation on Neural Bandwidth Extension of Telephone Speech for Improved Speaker Recognition. In *Proc. of IEEE ICASSP*.
- [67] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. LibriSpeech: An ASR Corpus Based on Public Domain Audio Books. In *Proc. of IEEE ICASSP*.
- [68] Kenneth V. Price. 2013. Differential Evolution. In *Handbook of Optimization*. 187–214.
- [69] Yao Qin, Nicholas Carlini, Garrison W. Cottrell, Ian J. Goodfellow, and Colin Raffel. 2019. Imperceptible, Robust, and Targeted Adversarial Examples for Automatic Speech Recognition. In *Proc. of ACM ICML*.
- [70] Kanishka Rao, Hasim Sak, and Rohit Prabhavalkar. 2017. Exploring Architectures, Data and Units for Streaming End-to-end Speech Recognition with RNN-Transducer. In *Proc. of IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*.
- [71] Douglas A Reynolds. 2002. An Overview of Automatic Speaker Recognition Technology. In *Proc. of IEEE ICASSP*.
- [72] Douglas A. Reynolds, Thomas F. Quatieri, and Robert B. Dunn. 2000. Speaker Verification Using Adapted Gaussian Mixture Models. *Digital Signal Processing* 10, 1-3 (2000), 19–41.
- [73] Douglas A. Reynolds and Richard C. Rose. 1995. Robust Text-independent Speaker Identification Using Gaussian Mixture Speaker Models. *IEEE Transactions on Speech and Audio Processing* 3, 1 (1995), 72–83.
- [74] Lea Schönherr, Katharina Kohls, Steffen Zeiler, Thorsten Holz, and Dorothea Kolossa. 2019. Adversarial Attacks Against Automatic Speech Recognition Systems via Psychoacoustic Hiding. In *Proc. of NDSS*.
- [75] Lea Schönherr, Steffen Zeiler, Thorsten Holz, and Dorothea Kolossa. 2019. Robust Over-the-air Adversarial Examples Against Automatic Speech Recognition Systems. *CoRR* abs/1908.01551 (2019).
- [76] Ali Shafahi, Mahyar Najibi, Amin Ghiasi, Zheng Xu, John P. Dickerson, Christoph Studer, Larry S. Davis, Gavin Taylor, and Tom Goldstein. 2019. Adversarial Training for Free!. In *Proc. of NeurIPS*.
- [77] David Snyder, Daniel Garcia-Romero, Daniel Povey, and Sanjeev Khudanpur. 2017. Deep Neural Network Embeddings for Text-independent Speaker Verification. In *Proc. of INTERSPEECH*.
- [78] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. 2014. Intriguing Properties of Neural Networks. In *Proc. of ICLR*.
- [79] Qian Tao and Raymond Veldhuis. 2010. Biometric Authentication System on Mobile Personal Devices. *IEEE Transactions on Instrumentation and Measurement* 59, 4 (2010), 763–773.
- [80] Rohan Taori, Amog Kamsetty, Brenton Chu, and Nikita Vemuri. 2019. Targeted Adversarial Examples for Black Box Audio Systems. In *Proc. of IEEE SPW*.
- [81] Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian J. Goodfellow, Dan Boneh, and Patrick D. McDaniel. 2018. Ensemble Adversarial Training: Attacks and Defenses. In *Proc. of ICLR*.
- [82] Qian Wang, Baolin Zheng, Qi Li, Chao Shen, and Zhongjie Ba. 2021. Towards Query-efficient Adversarial Attacks Against Automatic Speech Recognition Systems. *IEEE Transactions on Information Forensics and Security* 16 (2021), 896–908.
- [83] Yixiang Wang, Jiqiang Liu, and Xiaolin Chang. 2021. Generalizing Adversarial Examples by AdaBelief Optimizer. *CoRR* abs/2101.09930 (2021).
- [84] Darrell Whitley. 1994. A Genetic Algorithm Tutorial. *Statistics and Computing* 4, 2 (1994), 65–85.
- [85] Eric Wong, Leslie Rice, and J. Zico Kolter. 2020. Fast Is Better Than Free: Revisiting Adversarial Training. In *Proc. of ICLR*.
- [86] Hiromu Yakura and Jun Sakuma. 2019. Robust Audio Adversarial Example for a Physical Attack. In *Proc. of IJCAI*.
- [87] Bo Yang, Hengwei Zhang, Yuchen Zhang, Kaiyong Xu, and Jindong Wang. 2021. Adversarial Example Generation with AdaBelief Optimizer and Crop Invariance. *CoRR* abs/2102.03726 (2021).
- [88] Zhuolin Yang, Bo Li, Pin-Yu Chen, and Dawn Song. 2019. Characterizing Audio Adversarial Examples Using Temporal Dependency. In *Proc. of ICLR*.
- [89] Zhenyu Yang, Ke Tang, and Xin Yao. 2008. Multilevel Cooperative Coevolution for Large Scale Optimization. In *Proc. of IEEE CEC*.
- [90] Xuejing Yuan, Yuxuan Chen, Yue Zhao, Yunhui Long, Xiaokang Liu, Kai Chen, Shengzhi Zhang, Heqing Huang, Xiaofeng Wang, and Carl A. Gunter. 2018. CommanderSong: A Systematic Approach for Practical Adversarial Voice Recognition. In *Proc. of USENIX Security*.
- [91] Chunlei Zhang and Kazuhito Koishida. 2017. End-to-end Text-independent Speaker Verification with Triplet Loss on Short Utterances. In *Proc. of INTERSPEECH*.
- [92] Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyan Xu. 2017. DolphinAttack: Inaudible Voice Commands. In *Proc. of ACM CCS*.
- [93] Nan Zhang, Xianghang Mi, Xuan Feng, Xiaofeng Wang, Yuan Tian, and Feng Qian. 2019. Dangerous Skills: Understanding and Mitigating Security Risks of Voice-controlled Third-party Functions on Virtual Personal Assistant Systems. In *Proc. of IEEE S&P*.
- [94] Yuheng Zhang, Ruoxi Jia, Hengzhi Pei, Wenxiao Wang, Bo Li, and Dawn Song. 2020. The Secret Revealer: Generative Model-inversion Attacks Against Deep Neural Networks. In *Proc. of IEEE/CVF CVPR*.
- [95] Yangyong Zhang, Lei Xu, Abner Mendoza, Guangliang Yang, Phakpoom Chinpruthiwong, and Guofei Gu. 2019. Life after Speech Recognition: Fuzzing Semantic Misinterpretation for Voice Assistant Applications. In *Proc. of NDSS*.
- [96] Juntang Zhuang, Tommy Tang, Yifan Ding, Sekhar C. Tatikonda, Nicha C. Dvornek, Xenophon Papademetris, and James S. Duncan. 2020. AdaBelief Optimizer: Adapting Stepsizes by the Belief in Observed Gradients. In *Proc. of NeurIPS*.

A ADAPTIVE DECOMPOSITION ALGORITHM

Here we present the detailed adaptive decomposition algorithm in Occam. The cooperative co-evolution in Occam is illustrated in Figure 7. Intuitively, one may want to probe all possible group sizes and choose the one with the best performance. However, it is very impractical because of the prohibitively high computational overheads. Instead, we propose a dynamic and self-adapting grouping strategy. In our design, we let the group size vary in different stages according to its performance, which is evaluated in a “pilot” manner. More specifically, we first divide the variables into subgroups of candidate group sizes and optimize one of the subgroups. The candidate group size is selected in a gradual manner, *i.e.*, the number of subgroups is twice or half as large as that in the previous stage. We then determine the group size in the next stage according to the performance of pilot test. A candidate size will be adopted in the next stage if the performance of the pilot test is better. In this way, the grouping strategy greatly reduces computation costs since the pilot test only involves optimizing a small group of variables.

The details of the grouping algorithm are described as follows. Since the degree of interdependence between the variables from the natural audio and the audio AE gradually increases, we first initialize the number of groups as $m = 1$.

- 1) Create two lists $R = \{r_1, r_2, r_3, r_4\}$ and $\Delta = \{\delta_{m/2}, \delta_m, \delta_{2m}\}$ to record the impact of the four grouping strategies in the candidate pool and different group sizes, respectively. Initialize $R = 0$, $\Delta = 0$ and $m = 1$.
- 2) Compute the probabilities of selecting each grouping strategy in the candidate pool as

$$prob_j = \frac{e^{r_j}}{\sum_{j=1}^4 e^{r_j}}, j = 1, 2, 3, 4. \quad (14)$$

- 3) Select a grouping strategy according to probabilities $\{prob_j\}_{j=1}^4$.
- 4) Group the variables into m subgroups $\mathcal{G}_m = \{G_1, \dots, G_m\}$ according to the selected grouping strategy.

Figure 7: Illustration of cooperative co-evolution in Occam.

- 5) Optimize the subspaces on $\mathcal{G}_m$ until the current stage ends. Update r_i and δ_m as

$$r_j = \delta_m = \left| \frac{(v - v')}{v \cdot m} \right|, \quad (15)$$

where v is the best fitness of the last cycle, and v' is the best fitness of the current cycle.

- 6) Run pilot test.
- i. If $\delta_{m/2} = 0$ and $m/2 \geq 1$, group the variables into $m/2$ subgroups $\mathcal{G}_{m/2} = \{G_1, \dots, G_{m/2}\}$ according to the selected grouping strategy, and optimize G_1 to obtain the performance record

$$\delta_{m/2} = \left| \frac{(v - v')}{v} \right|. \quad (16)$$

- ii. If $\delta_{2m} = 0$ and $2m \leq n$, group the variables into $2m$ subgroups $\mathcal{G}_{2m} = \{G_1, \dots, G_{2m}\}$ according to the selected grouping strategy, and optimize G_1 to obtain the performance record

$$\delta_{2m} = \left| \frac{(v - v')}{v} \right|. \quad (17)$$

- 7) If $\delta_{m/2}$ is the maximum among Δ , then set $m = m/2$, $\delta_{2m} = \delta_m$, $\delta_m = \delta_{m/2}$, and $\delta_{m/2} = 0$. Else if δ_{2m} is the maximum, then set $m = 2m$, $\delta_{m/2} = \delta_m$, $\delta_m = \delta_{2m}$, and $\delta_{2m} = 0$.
- 8) Go to step 3.

B DIFFERENTIAL EVOLUTION

Differential evolution is a gradient-free optimization method like CMA-ES. To sufficiently justify the design choice of CMA-ES, we also include the evaluations on the differential evolution attack (DEA). Our experimental results (see Tables 8 and 4) show that DEA cannot perform well on audio data. Here we present the algorithm of differential evolution, as shown in Alg. 3. To more easily generate offspring that is closer to the original audio, a bias term is also added into the mutation step, and the crossover step is removed.

Algorithm 3 Differential Evolution Attack

Input: The original audio x , the initial adversarial audio x' , the input space dimension n , the Population NP , and the total number of iterations G .

Output: The adversarial audio sample x^* .

- 1: Initialize the following parameters:

- $NP = 10$, $G = 3000$, $F = 0.5$,
- $g = 0$, $\sigma = 0.003$, $\mu = 0.2$.

- 2: **for** $i = 0$ to NP **do**

- 3: Sample $z \sim \mathcal{N}(0, \sigma^2)$;

- 4: $X_i^0 = x' + z$;

- 5: **end for**

- 6: **while** $g < G$ **do**

- 7: ▷ Mutation

- 8: **for** $i = 0$ to NP **do**

- 9: Generate three random integers r_1, r_2 and $r_3 \in \{0, 1, \dots, NP - 1\}$, with $r_1 \neq r_2 \neq r_3$;

- 10: $Xm_i = X_{r_1}^g + \text{rand}(0, \mu)(x - X_{r_1}^g) + F \times (X_{r_2}^g - X_{r_3}^g)$;

- 11: **end for**

- 12: ▷ Selection

- 13: **for** $i = 0$ to NP **do**

- 14: **if** $\mathcal{L}(Xm_i) < \mathcal{L}(X_i^g)$ **then**

- 15: $X_i^{g+1} = Xm_i$;

- 16: **end if**

- 17: **end for**

- 18: $g \leftarrow g + 1$;

- 19: **end while**

- 20: **return** x^* .

C DETAILED EXPERIMENT SETTING

C.1 Datasets

Common Voice. Common Voice [8] is the largest open-source human voice dataset launched by Mozilla, including 28 different languages and nearly 1,800 validated hours of voice data. The Common Voice dataset is used to train ASR systems and test the effectiveness of audio AEs.

Song. Song is released by CommanderSong [90] and also used in Devil's Whisper [32]. It contains 20 songs divided into four categories on average, including the soft, popular, rock, and rap. Similar to Devil's Whisper [32], we select a total of 10 songs in the soft and popular categories, which are less noisy and easier to disguise, to evaluate the performance of Occam against speech recognition APIs.

LibriSpeech. LibriSpeech [67] is an English speech corpus that consists of 1,000 hour reading audio sampled at 16kHz.

VoxCeleb. VoxCeleb [64] is an open-source large-scale audio dataset that contains more than 100,000 utterances from more than 7,000 celebrities. It consists of short human voice clips extracted from interview videos uploaded to YouTube. In our experiments on the SR systems, we use part of VoxCeleb as the enrollment data.

C.2 Target Systems

DeepSpeech. DeepSpeech [42] is a state-of-the-art ASR system that performs speech-to-text tasks. It is usually used as the target

model in adversarial attacks [53, 80]. Note that although DeepSpeech is open-source, Occam does not require any internal knowledge about the target model.

Commercial Cloud Speech APIs. Speech-to-text API services are widely used in many applications. To evaluate Occam on commercial services, we choose six popular API services, including Google Cloud Speech-to-Text [5], Microsoft Azure Speech Service [6], Alibaba Short Speech Recognition [11], Tencent Short Speech Recognition [10], and iFlytek voice dictation [3]. For the Google’s Speech-to-Text API, we select the “command_and_search model” as the target models. The characteristics of these APIs are listed in Table 2. Note that although some APIs provide confidence scores, our attack does not require such information.

Commercial Speaker Recognition Systems. We test Occam on the Microsoft Azure’s speaker recognition system and the Jingdong speaker recognition system [9]. Microsoft Azure’s API [6] can perform speaker identification (who is the speaker) and speaker verification (whether the speaker is legal). It only returns the decision (*i.e.*, the predicted speaker) along with three confidence levels (*i.e.*, low, normal, or high). Jingdong’s API can perform speaker verification and only returns the final result, *i.e.*, accept or reject.

Commercial Voice Control Devices. We test NI-Occam on five commercial voice control devices, *i.e.*, Apple Siri, iFlytek, Microsoft Cortana, Google Assistant, Amazon Echo. In our experiments, Apple Siri, iFlytek, Microsoft Cortana, Google Assistant are applications installed in on-the-shelf smartphones, and Amazon Echo is an intelligent voice-controlled speaker.

C.3 Hardware

In our experiments, we conduct the attacks against DeepSpeech on a server equipped with four Nvidia 2080Ti GPUs, a six-core Intel(R) Xeon(R) W-2133 CPU 3.60GHz, 62 Gigabyte RAM, and a Hard Drive with 1.37 Terabytes. For the experiment on speech API services, we adopt several laptops, including a Lenovo ThinkPad X1 Carbon 4th with Core Intel i5-6200U CPU 2.30GHz and 8 Gigabyte RAM, a Microsoft Corporation Surface Pro 6 with Core Intel i7-8650U CPU 1.90GHz and 8 Gigabyte RAM, and a Lenovo ThinkPad X1 Carbon 5th with Core Intel i7-7500U CPU 2.70GHz and 8 Gigabytes RAM. Besides, we attack the voice assistant including Apple Siri with version 13.6.1 and iFlytek with version 10.0.8 on an iPhone 11, Cortana with version 3.3.3 on a Samsung C9000 with version 3.3.3, and Google Assistant with version 2.5.1 on a Nokia 7 plus. We play the audio AEs using a JBL Clip 3 portable speaker.

D SUPPLEMENTARY EVALUATION OF OCCAM

D.1 Evaluation on Open-source ASR Systems

We evaluate the effectiveness of the audio AEs generated by the attacks in Table 8. These AEs are generated after 200,000 queries (100,000 times on Song¹³) on DeepSpeech. Since Genetic Algorithm-based Attack (GAA) [80] and Selective Gradient Estimation Attack (SGEA) [82] need to leverage the prediction scores of the model, the AEs achieve higher SNRs than other decision-based attacks.

¹³The original data from Song requires fewer queries because the signal power of the songs is relatively strong.

Figure 8: SNRs of the audio AEs after querying DeepSpeech for different times over (a) Common Voice and (b) Song.

Figure 9: SNRs of the AEs from targeted attacks against the commercial services.

However, neither of them can successfully attack DeepSpeech with a 100% SRoA-ASR. In the other six decision-based attacks, the initial example is originally adversarial and then trained to approach the target example. Therefore, the SRoA of these attacks is always 100%. However, it is worth noting that SRoA is not the only metric that determines whether an AE can successfully attack the target system. An effective AE needs to fool both the model and the human, where SNR must be used as another important factor to measure the effectiveness of AEs. Although the six decision-based attacks can all achieve a 100% SRoA-ASR, only the AEs generated by our attack obtain high SNRs (with the best SNR of 13.80dB).

To evaluate the efficiency of Occam, we tested the SNRs of the AEs generated after different numbers of queries on DeepSpeech. As shown in Figure 8, Occam achieves a high SNR within a small number of queries. Note that the dataset Song requires fewer queries because the audio in Song is less noisy and more powerful, making it easier to generate an audio AE with a high SNR. We also find that the growth rate of SNR significantly decreases with the increase of SNR. For example, the SNR of our AEs can quickly converge to 12.86dB, while the evolutionary attack may require hundreds of thousands of queries to increase SNR from 1.20dB to 12.86dB. Hence, the results show that the co-evolution algorithm has a higher bound of SNR and a faster convergence rate than the evolution algorithm on audio data, validating the effectiveness of the CC framework in constructing audio AEs.

Table 8: Experimental results on targeted attacks against DeepSpeech.

Dataset	GAA		SGEA		Boundary Attack		Opt-Attack		Evolutionary Attack		HSJA		DEA		Occam	
	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR	SRoA	SNR
Songs	1/10	23.78	4/10	21.62	10/10	8.03	10/10	2.88	10/10	1.20	10/10	5.88	10/10	-0.26	10/10	12.86
LibriSpeech	1/10	14.72	7/10	14.60	10/10	4.80	10/10	2.93	10/10	1.09	10/10	1.94	10/10	-1.41	10/10	11.30
Common Voice	1/10	17.69	8/10	19.14	10/10	3.41	10/10	2.01	10/10	0.25	10/10	1.78	10/10	-0.52	10/10	13.80
Average	1/10	18.73	6.3/10	18.45	10/10	5.41	10/10	2.61	10/10	0.85	10/10	3.2	10/10	-0.73	10/10	12.65

Table 9: Experimental results of untargeted attacks on commercial cloud speech-to-text APIs.

Model	Boundary Attack	Opt-Attack	Evolutionary Attack	HSJA	DEA	Occam
Google STT	26.43	19.63	31.20	30.04	16.31	31.43
Microsoft ASR	15.43	11.89	27.37	19.60	11.08	28.68
Alibaba SSR	17.84	15.05	27.85	20.66	17.49	29.57
Tencent SSR	20.38	17.02	30.16	21.51	18.03	31.68
iFlytek	42.66	29.07	38.59	39.82	26.61	40.46

Note that, (i) this table shows the SNRs of the audio AEs generated after 200 queries. (ii) This table does not show the success rate because all of the six attacks achieve a success rate of 100%. (iii) The best and the second best results on each API are marked in bold with underling and bold font, respectively.

D.2 Evaluation on Cloud Speech APIs

For completeness, we give some additional experimental results as a complement to Section 5.2.

We present Figure 9 to show the SNRs with confidence intervals of 68.2% (-std, +std). The results show that Occam has the highest SNRs among all the decision-based attacks. Figure 10 shows the waveforms and spectrograms of the original audio and the adversarial audios generated by the seven attacks. It can be seen that the waveforms of the original audio and the audio AEs generated by Occam are almost the same. However, the differences in other attacks are more noticeable and thus more likely to be perceived by humans.

As for the untargeted attacks, the results in Table 9 show that the SNRs of AEs generated from untargeted attacks are much higher than those from targeted attacks. In this experiment, the performance of the evolutionary attack is comparable to Occam. Since the implementation of untargeted attacks is much simpler than the targeted attack (e.g., 200 queries are enough), the optimization problem on untargeted attacks can be easily solved without being decomposed into a set of simpler sub-problems. Thus, approaches without a cooperative co-evolution framework can also work well. However, concerning the large-scale and complex problems of generating targeted audio AEs, these approaches become ineffective. Besides, Devil’s Whisper does not explicitly support untargeted attacks. To launch the untargeted attack by using the methodology of Devil’s Whisper, the attacker needs to intentionally set a wrong phrase as the target phrase. This step incurs a large amount of queries. In contrast, our methodology can successfully generate effective audio AEs within 200 queries.

D.3 Evaluation on Human Perception

Similar to the human perception experiments on NI-Occam, we also evaluate the performance of Occam about the human perception. The experiment settings are the same as those in Section 5.3.

For each commercial service, the volunteers need to listen to 10 pieces of audio AEs generated from the boundary attack, the opt-attack, the evolutionary attack, HSJA, DEA, Devil’s Whisper, and Occam. The audio AEs are generated by querying the target systems for 200,000 times with original audios from the dataset Song. In all, we get 37×10 samples from each attack on each commercial service (a total of 37×10×6 samples for one attack). Note that, since Devil’s Whisper does not achieve a 100% SRoA, the volunteers only rate the successful AEs. Each volunteer ranks 3, 5, 7, 8, and 4 successful AEs from Devil’s Whisper on Alibaba, Google, iFlytek, Microsoft, and Tencent, respectively.

Table 10 presents the results of the experiments on human perception. Overall, Occam has the highest rate of “normal”. More than 70% volunteers think the audio generated from Occam is “normal” or “noise”, which is comparable to Devil’s Whisper. However, about 50% AEs from Devil’s Whisper fail to fool the ASR systems in the first place. The results show that Occam is more effective in fooling both the model and the human than other possible attacks.

E A HUMAN STUDY ON AUDIO AES WITH DIFFERENT SNRS

Recall that we use two important metrics to evaluate the effectiveness of the audio AEs, *i.e.*, the SRoA (Success Rate of Attack) and SNR. Both metrics are crucial to determine whether an AE can successfully attack the target system, since an effective AE needs to fool both the model and the human simultaneously. To demonstrate the necessity of SNR as an evaluation criterion, we conduct a small human study to help understand how the metric of SNR affects the effectiveness and quality of an audio AE.

In this experiment, we surveyed 30 volunteers who are sensitive to sound, including 17 males and 13 females. Similar to the human perception experiment (See Section 5.3), we first show some examples of “noise”, “normal”, “talking”, and “recognized”, and then ask the volunteers to listen to the audio AEs and tell their views about them. The audio AEs are generated from Occam against Alibaba

Table 10: Evaluation results on human perception of different attacks in the digital world.

Commercial Service	Method	Normal (%)	Noise (%)	Talking (%)	Once-recognize (%)	Twice-recognize (%)
Alibaba SSR	Boundary Attack	26.76	23.24	50.00	21.35	28.92
	DEA	6.76	23.24	70.00	45.14	52.70
	Evolutionary Attack	2.70	32.16	65.14	43.24	46.76
	HSJA	8.65	37.30	54.05	31.62	36.22
	Opt-Attack	4.59	29.73	65.68	38.65	43.24
	Devil's Whisper	8.11	51.35	40.54	14.41	18.02
	Occam	50.54	35.41	14.05	3.24	6.22
Google STT	Boundary Attack	4.59	21.89	73.51	35.68	39.73
	DEA	0.54	21.35	78.11	34.59	42.70
	Evolutionary Attack	3.51	17.30	79.19	33.24	42.16
	HSJA	5.14	15.68	79.19	37.03	45.68
	Opt-Attack	6.76	12.70	80.54	40.00	46.22
	Devil's Whisper	18.38	60.54	21.08	7.03	8.11
	Occam	40.00	38.38	21.62	6.22	14.59
iFlytek	Boundary Attack	3.51	23.78	72.70	33.24	43.51
	DEA	0.00	23.78	76.22	34.59	41.08
	Evolutionary Attack	5.14	23.24	71.62	33.24	39.73
	HSJA	1.62	20.27	78.11	35.68	40.70
	Opt-Attack	3.51	19.73	76.76	37.03	43.24
	Devil's Whisper	17.76	48.26	33.98	8.11	13.13
	Occam	22.16	41.08	36.76	14.05	21.08
Microsoft ASR	Boundary Attack	13.51	25.95	60.54	27.57	29.19
	DEA	8.11	23.24	68.65	31.08	37.57
	Evolutionary Attack	1.08	23.78	75.14	36.22	44.59
	HSJA	1.08	17.84	81.08	39.19	47.03
	Opt-Attack	4.05	19.73	76.22	36.76	43.51
	Devil's Whisper	29.05	45.27	25.68	8.78	13.85
	Occam	31.08	28.92	40.00	16.22	24.05
Tencent SSR	Boundary Attack	29.73	28.11	42.16	21.62	29.19
	DEA	1.08	20.27	78.65	34.59	37.57
	Evolutionary Attack	14.59	33.24	52.16	28.65	31.08
	HSJA	18.65	30.81	50.54	29.19	32.16
	Opt-Attack	11.62	19.73	68.65	35.68	41.62
	Devil's Whisper	12.84	37.84	49.32	10.14	15.54
	Occam	45.41	34.05	20.54	5.41	7.84
Average	Boundary Attack	15.62	24.59	59.78	27.89	34.11
	DEA	3.30	22.38	74.32	36.00	42.32
	Evolutionary Attack	5.41	25.95	68.65	34.92	40.86
	HSJA	7.03	24.38	68.59	34.54	40.76
	Opt-Attack	6.11	20.32	73.57	37.62	43.57
	Devil's Whisper	17.23	48.65	34.12	9.69	13.73
	Occam	37.84	35.57	26.59	9.03	14.76

Note that, (i) "Normal" means that the volunteer regards the audio as a normal audio. (ii) "Noise" means that the volunteer can feel some noises. (iii) "Talking" means that the volunteer can hear talking in the audio. If the volunteer thinks there is talking in the audio, he/she is then asked to recognize the content of the talking. (iv) The audio will be labeled as "once-recognize" or "twice-recognize" if the volunteer recognizes over half of the content after listening to the audio once or twice, respectively.

Figure 10: Waveforms and spectrograms of the original audio and adversarial audios generated by the targeted attacks against Alibaba SSR. The audio AEs can be recognized as “call my wife”.

and Tencent, with SNRs ranging from 6dB to 16dB. Specifically, we generate 10 AEs for each SNR group. As shown in Table 11, when the SNR increases (*i.e.*, the noise becomes smaller), more volunteers consider the AE as a normal one. Specifically, when the SNR is lower than 8dB, over 84% of volunteers can hear talking in the AE audios, and more than 31% of volunteers can recognize the commands in the AEs. The results validate that if the SNR is low, it is easy for the users to notice or even recognize the AEs. Since the audio AEs need to be stealthy enough to fool the human, AEs with high SNRs are considered more in practice.

Table 11: Evaluation results on human perception of AEs with different SNRs.

SNR (dB)	Normal (%)	Noise (%)	Talking (%)	Once-recognize (%)	Twice-recognize (%)
16	36.00	38.33	25.67	3.33	4.33
14	27.67	35.33	37.00	6.67	7.33
12	14.33	27.00	58.67	14.00	17.00
10	8.67	15.00	75.33	26.33	29.00
8	4.33	11.00	84.67	31.00	31.00
6	0.00	4.67	95.33	43.67	43.67

Figure 11: SROAs after different numbers of query attempts.

F SUPPLEMENTARY EVALUATION OF NI-OCCAM

As a supplement to Section 4.3, here we provide the detailed results of our NI-Occam on 10 target phrases in Table 12. Besides, we also conduct an experiment on the impact of the number of attempts. In previous experiments on NI-Occam, the default setting for the number of query attempts is 3 times, and Devil’s Whisper [32] adopted a number of 30 attempts. Theoretically, more attempts will lead to higher SROAs. To illustrate this statement, we conduct an experiment with a range of [1, 30] query attempts using NI-Occam, and the results are given in Figure 11. As shown, at the very start, the increasing number of query attempts can help increase the SROA. However, when the number of attempts is larger than 17, the SROA will remain at 70% and no longer increases, indicating a performance limit. In practice, however, more attempts also require longer attacking periods, making the attack less stealthy. Hence, we suggest the attacker conduct as few attempts as possible.

G EVALUATION OF OCCAM AND NI-OCCAM ON LARGE COMMAND SETS

As illustrated before, the target phrases in Devil’s Whisper [32] are restricted to several phrases for a substitute model. Different from Devil’s Whisper, Occam is not restricted to the number of targeted phrases, where the adversary can express freely. To illustrate this point, we evaluate Occam and NI-Occam on 60 targeted phrases. In the experiment of Occam, we generate 60 audio AEs with 60 different targeted voice commands by querying Alibaba SSR 30,000 times and feed them into it. The target commands and SNRs of the AEs are presented in Table 13. The 60 AEs have an average SNR

Table 12: Detailed results of NI-Occam against voice control devices.

Command	Apple Siri		iFlytek		Microsoft Cortana		Google Assistant		Amazon Echo	
	Attack	SNR (dB)	Attack	SNR (dB)	Attack	SNR (dB)	Attack	SNR (dB)	Attack	SNR (dB)
Call my wife	✓	11.57	✓	7.61	✓	7.61	✓	11.57	✓	8.76
Open the door	✗	-	✓	8.57	✗	-	✗	-	✗	-
Play music	✓	9.70	✓	9.70	✓	9.70	✓	9.70	✓	11.51
Send the text	✗	-	✗	-	✓	11.51	✗	-	✗	-
Take a picture	✗	-	✗	-	✗	-	✗	-	✗	-
Turn off the light	✓	8.61	✓	8.61	✓	8.61	✓	8.61	✓	8.61
Open the website	✓	8.66	✓	8.50	✓	8.50	✗	-	✗	-
Make it warmer	✗	-	✗	-	✗	-	✗	-	✗	-
Navigate to my home	✓	11.54	✓	11.54	✓	11.54	✓	11.54	✗	-
Turn on airplane mode	✓	8.80	✗	-	✗	-	✗	-	✓	8.80
Average	6/10	9.81	6/10	9.09	6/10	9.58	4/10	10.36	4/10	9.42

Table 13: Performance of Occam and NI-Occam on large target command sets.

Commands	Occam	NI-Occam		Commands	Occam	NI-Occam	
	Alibaba	Cortana	Siri		Alibaba	Cortana	Siri
watch TV	17.40	9.74	9.74	clear notification	9.98	8.76	8.76
open the box	11.11	8.74	-	when do you get up	10.96	-	-
flip a coin	11.92	8.81	8.81	ask me a question	24.97	9.18	-
set a timer	18.15	8.6	8.6	where is my home	16.21	10.23	10.23
sing a song	14.26	9.29	9.29	where is my hotel	14.42	9.73	-
crystal ball	9.24	-	-	mischief managed	20.6	8.51	8.51
find a hotel	10.31	9.58	8.73	turn off all alarms	14.71	8.86	8.86
set an alarm	25.01	10.09	10.09	turn the volume up	20.37	8.28	8.28
surprise me	28.88	10.36	10.36	show me the menu	28.69	8.61	8.61
take a selfie	17.83	8.56	-	where is this place	11.52	9.27	-
reading a book	9.5	-	-	what is the weather	14.09	-	9.03
close twitter	16.25	-	-	show me my flights	15	9.14	-
classical music	14.59	8.83	9.25	listen to voice mail	17.73	7.65	-
call my uncle	14.71	-	8.8	show me the money	16.83	9.11	9.11
tell me a joke	22.6	10.27	10.27	show me my alarms	12.78	10.12	10.12
tell me a story	24.67	8.87	8.87	when is my birthday	31.49	9.28	9.28
turn on the TV	25.07	9.15	9.15	turn on the window	10.58	-	8.53
set a reminder	11.87	8.55	9.62	turn off the computer	25.2	-	-
clean my room	24.04	9.76	9.76	make me a sandwich	19.14	-	7.37
what time is it	38.07	8.84	9.41	turn off the window	9.07	7.71	-
record a video	9.36	8.48	-	where is my package	20.74	8.45	-
spin the wheel	18.01	-	-	turn the volume down	21.6	9.68	9.68
open instagram	9.71	8.59	-	share the new version	21.45	-	-
turn on the camera	11.23	8.62	-	give me a compliment	14.23	-	-
cancel an alarm	13.11	-	-	show me my messages	25.8	8.93	8.93
where is my car	22.24	8.54	9.89	sing me happy birthday	17.09	9.34	9.34
how old are you	9.09	-	-	what is the temperature	13.17	9.08	9.08
mute the volume	13.99	-	9.32	turn on the coffee maker	20.08	-	-
turn on bluetooth	26.24	10.18	10.18	Oscar winner of this year	9.85	-	-
turn off bluetooth	15.04	9.51	9.51	what would you recommend	16.01	7.72	-

Note that, (i) this table shows the SNRs of the AEs generated from Occam and NI-Occam, and “-” denotes this AE fails. (ii) In total, the SRoAs of Occam against Alibaba SSR API, NI-Occam against Cortana, and NI-Occam against Apple Siri are 100%, 71.7%, and 58.3%, respectively.

of 17.27dB, and all of them can be recognized as the target commands successfully. For NI-Occam, we also generate 60 audio AEs with 60 different voice commands and play the AEs to a Samsung C9 running Cortana and an iPhone 11 running Siri. As shown in Table 13, in total, 43 AEs have successfully attacked Cortana, *i.e.*, the SRoA is 71.7%. For Siri, there are 35 effective AEs that can be recognized as the target commands, *i.e.*, the SRoA is 58.3%. The results indicate that both Occam and NI-Occam can perform well on large command sets.

H RESULTS OF POSSIBLE COUNTERMEASURES

Due to the space limitation, we provide the detailed results of four countermeasures against our attacks in this section. Table 14 shows the performance of our attacks against different countermeasures, including local smoothing, downsampling and temporal dependency based approach, and Table 15 presents the performance of our NI-Occam against adversarial training.

In our experiment of adversarial training, we adversarially train Kaldi (the aforementioned Mini LibriSpeech model) on the Mini LibriSpeech dataset¹⁴. We set p to ∞ , iterations to 10, and the step size α to $\epsilon/5$ for PGD, which are the same with [50, 51]. Since the prior work [51] demonstrates that adversarial training with $\epsilon > 0.01$ will be not able to converge, we set ϵ to 0.002, 0.004 and 0.006. For evaluation, we also generate 10 audio AEs from Mini LibriSpeech model.

Table 14: Performance of our attacks against different countermeasures.

Countermeasures	Setting	Attack Methods	
		Occam	NI-Occam
Local smoothing	$h = 1$	2/10	4/10
	$h = 3$	0/10	4/10
Downsampling	DSR = 12kHz	0/10	3/10
	DSR = 10kHz	0/10	2/10
Temporal dependency	$k = \text{Rand}(0.2, 0.8)$	AUC Score	
		100%	68%

Note that in these experiments, we test Occam against Alibaba SSR and NI-Occam against Cortana. DSR means the downsampling rate.

Table 15: Performance of NI-Occam against adversarial training.

	ϵ	WER [‡]	SRoA [†]
Standard training	-	10.69	10/10
Adversarial training ^{\$}	0.002	19.82	3/10
	0.004	31.54	1/10
	0.006	58.17	0/10

Note that, (i) [‡]: WER (Word Error Rate) is a common metric to measure the performance of speech recognition systems. WER calculates the ratio of the minimum number of word-level operations, including substitutions, deletions and insertions, required to convert the transcribed text into the target text, to the total number of words. A lower WER indicates a higher accuracy of ASR systems. (ii) ^{\$}: Adversarial training is conducted on the Mini LibriSpeech dataset that contains 5 hours of training data and 2 hours of test data. (iii) [†]: Audio AEs are generated using NI-Occam in Kaldi (Mini LibriSpeech model).

¹⁴<https://www.openslr.org/31/>