PHASEFOOL: PHASE-ORIENTED AUDIO ADVERSARIAL EXAMPLES VIA ENERGY DISSIPATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Audio adversarial attacks design perturbations onto inputs that lead an automatic speech recognition (ASR) model to predict incorrect outputs. Current audio adversarial attacks optimize perturbations with different constraints (e.g. l_p-norm for waveform or the principle of auditory masking for magnitude spectrogram) to achieve their imperceptibility. Since phase is not relevant for speech recognition, the existing audio adversarial attacks neglect the influence of phase spectrogram. In this work, we propose a novel phase-oriented algorithm named PhaseFool that can efficiently construct imperceptible audio adversarial examples with energy dissipation. Specifically, we leverage the spectrogram consistency of short-time Fourier transform (STFT) to adversarially transfer phase perturbations to the adjacent frames of magnitude spectrogram and dissipate the energy that is crucial for ASR systems. Moreover, we propose a weighted loss function to improve the imperceptibility of PhaseFool. Experimental results demonstrate that PhaseFool can inherently generate full-sentence imperceptible audio adversarial examples with the 100% targeted success rate within 500 steps on average (9.24x speed-up over current state-of-the-art imperceptible counterparts), which is verified through a human study[1]. Most importantly, our PhaseFool is the first to exploit the phase-oriented energy dissipation in the audio adversarial examples rather than add perturbations on the audio waveform like most previous works.

1 INTRODUCTION

Adversarial examples are specifically designed instances that cause a machine-learning algorithm to produce a misclassification (Szegedy et al., 2013; Biggio et al., 2013). Since researchers have successfully created adversarial examples to make automatic speech recognition (ASR) system transcribe the audio as any different sentence (Carlini & Wagner, 2018; Vaidya et al., 2015; Carlini et al., 2016), audio adversarial attacks have become more and more popular. Those attack algorithms usually optimize perturbations added on the benign audio waveform with different constraints based on l_p-norm (Carlini & Wagner, 2018; Zhang et al., 2020). Although the perturbations introduced is small in magnitude, noises in the audio adversarial examples are usually perceptible. Later, algorithms based on psychoacoustic principle (Qin et al., 2019; Schönherr et al., 2019) successfully generate effectively imperceptible adversarial examples by hiding the adversarial perturbations to regions of magnitude spectrogram where it will not be heard by a human.

In current imperceptible targeted audio adversarial attacks, only waveform and magnitude spectrogram are taken into account leaving phase and its nature unexplored. Since the structure of natural speech is highly complicated, these two types of algorithms face several challenges:

- Waveform-oriented attacks are usually not imperceptible enough. Although constraining perturbations with l_2-norm or l_∞-norm (Carlini & Wagner, 2018; Liu et al., 2020), these attacks introduce audible noises to the adversarial examples. Later works focused on the temporal dependency (Zhang et al., 2020) successfully reduce human-perceptible noises but their perturbations are still audible, due to the uncorrelated relationship between psychoacoustic perception and l_p-norm metrics.

[1]Generated audio adversarial examples can be found at our demo website: `https://phasefool-demo.github.io/phasefool-demo/`.

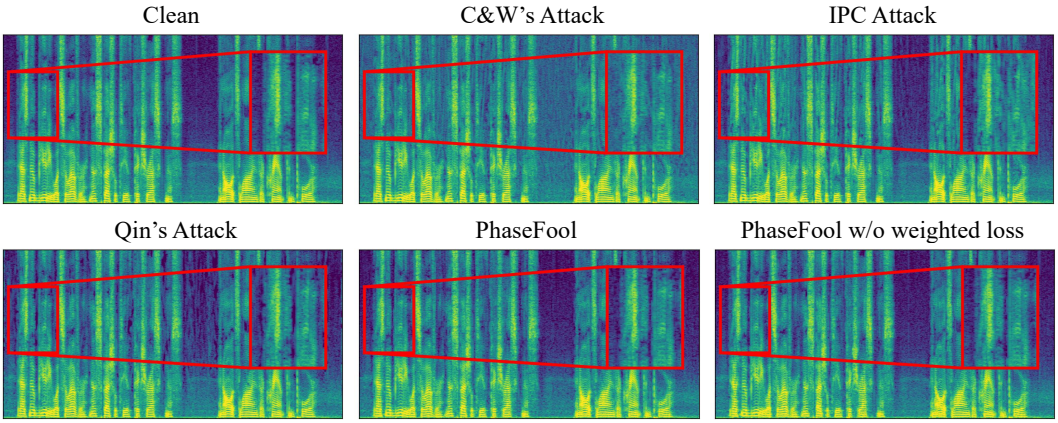

Figure 1: Visualizations of the ground-truth mel-spectrograms and the mel-spectrograms generated by several state-of-the-art attacks and our PhaseFool. Attacks based on the l_p-norm (e.g. C&W's attack (Carlini & Wagner, 2018) and IPC attack (Zhang et al., 2020)) introduce obvious perturbations to the mel-spectrogram and cause perceptible noises. Qin's attack leverage the psychoacoustic principle of auditory masking and hide noises in the inaudible regions while PhaseFool attack dissipates the energies that are important for classification while less important for human perception.

- The optimization of attacks based on the psychoacoustic principles is highly time-consuming and hard to converge. Due to the determining influence of magnitude spectrogram on human perception, the attack that leverages the psychoacoustic principle of auditory masking (Qin et al., 2019) successfully generates imperceptible audio adversarial examples. However, it needs two-stage optimization where the first stage of optimization focuses on finding a relatively small perturbation to fool the network (as is done in prior work (Carlini & Wagner, 2018)) and the second stage makes the adversarial examples imperceptible. The search process of the optimal balance between hiding the adversarial perturbations in the inaudible regions and fooling the ASR systems is empirically difficult. Generating a single full-sentence adversarial example requires approximately 5000 iterations until it finally converges to an imperceptible solution.

This research is motivated by a psychoacoustic principle that phase manipulation in one-channel audios imposes tiny effects on human perception (Vary & Eurasip, 1985). Besides, although phase is not relevant for speech recognition, the inherent spectrogram consistency of short-time Fourier transform (STFT) reveals that the adversarial perturbations in a frame of phase spectrogram will transfer to adjacent frames of magnitude spectrogram (Griffin & Lim, 1984). And these observation motivates us to exploit phase spectrogram for generating audio adversarial examples.

Considering the relationship between phase spectrogram and the spectrogram consistency of STFT, in this work, we propose a novel phase-oriented algorithm called PhaseFool that efficiently generates imperceptible audio adversarial examples via energy dissipation. We prove that the phase-oriented perturbations always reduce the magnitude of the spectrogram in the corresponding areas via spectrogram consistency, which is called energy dissipation. Leveraging such a phenomenon, we bring new directions for audio adversarial attacks and are able to perform a bi-directional attack that dissipates the crucial energy for ASR models and meanwhile, construct the perturbations for the adversarial attack. To the best of our knowledge, we are the first to consider phase spectrogram in the audio adversarial attacks. Our contributions can be summarized as follows:

- We leverage the spectrogram consistency of STFT to construct phase-oriented audio adversarial examples, which make further steps towards imperceptibility and efficiency.
- We further investigate the relationship between phase perturbations and magnitude spectrogram, which is called energy dissipation. Instead of the common l_p distance metrics or magnitude-oriented metrics based on the psychoacoustic principle, we propose a weighted loss function based on the phenomenon of energy dissipation to improve the imperceptibility of phase-oriented audio adversarial examples.

- Further investigating properties of PhaseFool, we observe additional insights about the vulnerability of ASR systems. It brings a better interpretation of audio adversarial attacks, i.e. the perturbation positions reveal that which part of one magnitude spectrogram is important but also fragile for the corresponding ASR system.

We conduct experiments on the LibriSpeech dataset (Panayotov et al., 2015) and implement a successful attack against a Transformer-based ASR system (Winata et al., 2020b). The results show that in terms of imperceptibility, PhaseFool greatly surpasses current state-of-the-art algorithms based on l_p-norm metrics and shows comparable imperceptibility with the attacks based on the psychoacoustic principle of auditory masking (verified by a human study). Furthermore, it can generate full-sentence audio adversarial examples with the 100% targeted success rate in around 30 seconds and achieves 9.24x speedup over the current state-of-the-art imperceptible counterpart. We attach the adversarial examples in the experiments and the implementation for reproducibility in the supplementary material.

2 BACKGROUNDS

In this section, we briefly overview the background of this work, including audio adversarial examples and phase-aware audio signal processing.

Audio Adversarial Examples. After the early works (Gong & Poellabauer, 2017; Cissé et al., 2017) towards neural networks based ASR systems successfully generate untargeted adversarial examples that produce incorrect but arbitrary transcriptions. Recently, more and more efforts have been made in the audio domain (Alzantot et al., 2018; Carlini & Wagner, 2018). Among them, one of the most outstanding studies is the iterative optimization-based attack (Carlini & Wagner, 2018), which introduces a two-stage attack and directly generates targeted adversarial examples using a gradient-descent minimization method. Nevertheless, the generated adversarial examples tend to include noticeable noise. Thus, two recent approaches (Qin et al., 2019; Schönherr et al., 2019) introduce hearing thresholds for designing inaudible perturbations by curbing the signal variation below the threshold of human perception. However, these imperceptible attacks require two stages of optimization and suffer from slow convergence rates.

Phase-aware Audio Signal Processing. The interest of utilizing phase information in the field of audio signal processing is constantly increasing (Mowlaee et al., 2016). The STFT phase has been widely used in signal enhancement (Miyahara & Sugiyama, 2014; Choi et al., 2019; Hu et al., 2020) and signal reconstruction (Mowlaee & Kulmer, 2015). The phase in the harmonic domain is used in the speech synthesis (Degottex & Erro, 2014; Ai & Ling, 2020) and the phase in perceptual bands is used in speech watermarking (Wang et al., 2021). The important information in the phase brings positive contributions to the above fields. Nevertheless, the existing audio adversarial attacks directly neglect the phase spectrogram, due to the fact that phase is not relevant for speech recognition. In this work, instead of introducing perturbations to the audio waveform like most of previous works, we propose to leverage the spectrogram consistency and generate phase-oriented perturbations for constructing imperceptible audio adversarial examples.

3 PHASE-ORIENTED AUDIO ADVERSARIAL EXAMPLES VIA ENERGY DISSIPATION

In this section, we first define the preliminaries of the phase-oriented audio adversarial attack. Then we investigate the relation between phase spectrogram and magnitude spectrogram under spectrogram consistency and prove that phase perturbations always reduce the magnitude of the spectrogram in the corresponding areas under spectrogram consistency, which is called energy dissipation (as is illustrated in Figure 1). Finally, we design a novel phase-oriented loss function based on energy dissipation for constructing imperceptible audio adversarial examples.

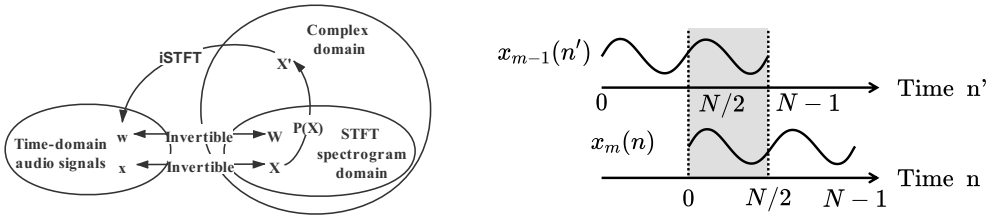

(a) (b)

Figure 2: The Illustration of spectrogram consistency. (a) demonstrates that if the spectrogram X is perturbed by $P(X)$, it usually locates outside of the STFT spectrogram domain making it non-invertible. (b) is the overlap-add process with 50% overlapped signals in the inverse STFT.

3.1 Preliminaries of Phase-oriented Audio Adversarial Attack

Denote the ASR classification model as $f : \{X\} \to Y$, with $X \in [0, 1]^{t\times c}$ being the audio space. $Y \in \{0, 1, ..., k\}^{n\times 1}$ is the sequence of the k-class output space corresponding to the transcription. Specifically, denote $x \in X$ as the input audio, $y \in Y$ as the ground-truth label, and $y_t \in Y$ as the target transcription, and $\mathcal{L}_{net}(f(x), y_t)$ be a loss function for the threat model that is minimized when $f(x) = y_t$. Then the general formulation of adversarial attack for ASR system is:

$$\underset{\delta}{\operatorname{argmin}} \quad \mathcal{L}_{net}(f(x+\delta), y_t) + \lambda \|\delta\|_p^p \ , \ \text{s.t.} \|\delta\|_\infty < \epsilon \ , \tag{1}$$

where $\mathcal{L}_{net}$ is the loss function used in the ASR model, $\lambda > 0$ is the trade-off parameter between the imperceptibility and adversarial loss (in some formulations λ = 0), ϵ control the maximum perturbation value, and p could be specified as different values, according to the attacker's requirement (e.g. $p = 2$ in C&W's attack (Carlini & Wagner, 2018)).

Denote $\phi(\cdot)$ as the function of the STFT and $\phi^{-1}(\cdot)$ as the function of the inverse STFT and $X = \phi(x) = |X| \cdot e^{i\theta_x}$, where $|X|$ and $\theta_x = \angle X$ representing the spectral magnitude and the instantaneous phase spectrogram, respectively. Denote $\mathcal{T}(X) = \phi(\phi^{-1}(X))$. We assume that window function w, window length N and window shift R is fixed and a signal can be perfectly reconstructed from its STFT spectrogram through the inverse STFT in this setting. We propose to generate perturbations on the instantaneous phase spectrogram and reconstruct the audio adversarial examples with the inverse STFT as the following formulation:

$$\underset{\delta}{\operatorname{argmin}} \quad \mathcal{L}_{net}(f(\mathcal{E}(\phi^{-1}(|X| \cdot e^{i\theta_x} \cdot e^{j\delta}))), y_t), \ \text{s.t.} -\pi \le \delta \le \pi \ , \tag{2}$$

where $\mathcal{E}(\cdot)$ is the feature extraction function for the ASR model and the $\phi^{-1}(|X| \cdot e^{i\theta_x} \cdot e^{j\delta})$ is the generated audio adversarial example.

As shown in Figure 2, the spectrogram $X' = |X| \cdot e^{i\theta_x} \cdot e^{j\delta}$ perturbed by $P(X)$ is usually located outside the STFT spectrogram domain. Considering the phenomenon of spectrogram consistency (Le Roux et al., 2010), it is easy to prove the following two conclusions: 1) $|X'| \neq |\mathcal{T}(X')|$ mathematically; 2) $\phi^{-1}(X')$ and $\phi^{-1}(\mathcal{T}(X'))$ represent the same time-domain audio signals. Since the STFT and iSTFT functions are differentiable, these conclusions enable us to adversarially manipulate the reconstructed magnitude spectrogram by introducing perturbations to the original phase spectrogram, which is called phase-oriented audio adversarial attack. Then we detail how the perturbations on a frame of phase spectrogram influence the adjacent frames of corresponding magnitude spectrogram and reveal the energy dissipation nature of the phase-oriented audio adversarial attack.

3.2 Relation Analysis between phase and magnitude

This subsection mainly discusses the essence of the relationship between phase perturbations and magnitude spectrogram frames. Reconstructing the time-domain signals after introducing perturbations to the phase spectrogram includes an overlap-add process. For simplicity, we assume that the window shift length R is equal to half of the window length N, which means the adjacent frames have 50% overlapped signals. As shown in Figure 2, the signals $x_{m-1}(n')$ and $x_m(n)$ in adjacent

frames $m-1$ and m overlap, where n' and n represent the time indices and $n' = n + \frac{N}{2}$. $x_{m-1}(n')$ and $x_m(n)$ are obtained by an inverse Fourier transform as the following:

$$x_{m-1}(n') = \phi^{-1}(X_{m-1}) = \frac{1}{N}\sum_{k=0}^{N-1} X_{m-1}[k]e^{j\frac{2\pi k n'}{N}} \ , \tag{3}$$

$$x_m(n) = \phi^{-1}(X_m) = \frac{1}{N}\sum_{k=0}^{N-1} X_m[k]e^{j\frac{2\pi k n}{N}} \ , \tag{4}$$

where j is the imaginary unit, $X_{m-1}[k]$ and $X_m[k]$ are spectrogram samples in frames $m-1$ and m at k-th frequency bin. Then we introduce perturbations δ_{m-1}, δ_m to the frame $m-1$, m and the corresponding signals become:

$$\hat{x}_{m-1}(n') = \frac{1}{N}\sum_{k=0}^{N-1} X_{m-1}[k]e^{j\frac{2\pi k n'}{N}} e^{j\delta_{m-1}[k]} \ , \tag{5}$$

$$\hat{x}_m(n) = \frac{1}{N}\sum_{k=0}^{N-1} X_m[k]e^{j\frac{2\pi k n}{N}} e^{j\delta_m[k]} \ . \tag{6}$$

Since the human auditory system is relatively less sensitive to the tiny time delay caused by phase rotation, the phase perturbations are inherently more imperceptible than the perturbations on the audio waveform. Then we further investigate how phase perturbations affect the magnitude spectrogram. The reconstructed signal $\tilde{x}_m(n)$ in the frame m is obtained as:

$$\begin{aligned}\tilde{x}_m(n) &= \hat{x}_{m-1}(n') + \hat{x}_m(n) \\ &= \frac{1}{N}\sum_{k=0}^{N-1} X_{m-1}[k]e^{j\frac{2\pi k n'}{N}} e^{j\delta_{m-1}[k]} + \frac{1}{N}\sum_{k=0}^{N-1} X_m[k]e^{j\frac{2\pi k n}{N}} e^{j\delta_m[k]} \ .\end{aligned} \tag{7}$$

We substitute $X_{m-1}[k]$ and $X_m[k]$ with the Fourier transform of $x_{m-1}(n')$ and $x_m(n)$, then the overlapping area of $\tilde{x}_m(n)$ in equation 7 will become:

$$\begin{aligned}\overline{x}_m(n) &= \frac{1}{N}\sum_{k=0}^{N-1}\left(\sum_{n'=0}^{N-1} x_{m-1}(n')e^{-j2\pi\frac{n'}{N}k}\right)e^{j2\pi\frac{n'}{N}k}e^{j\delta_{m-1}[k]} \\ &+ \frac{1}{N}\sum_{k=0}^{N-1}\left(\sum_{n=0}^{N-1} x_m(n)e^{-j2\pi\frac{n}{N}k}\right)e^{j2\pi\frac{n}{N}k}e^{j\delta_m[k]} \\ &= \frac{1}{N}\sum_{k=0}^{N-1}\left(\sum_{n'=0}^{N-1} x_{m-1}(n')e^{j\delta_{m-1}[k]} + \sum_{n=0}^{N-1} x_m(n)e^{j\phi_m[k]}\right) \\ &= \frac{1}{N}\sum_{k=0}^{N-1}\left(\sum_{n'=N/2}^{N-1} x_{m-1}(n')e^{j\delta_{m-1}[k]} + \sum_{n=0}^{N/2-1} x_m(n)e^{j\phi_m[k]}\right) \ .\end{aligned} \tag{8}$$

The overlapping area of the $x_{m-1}(n')$ and $x_m(n)$ in Figure 2 satisfy the following relationship based on spectrogram consistency:

$$x_{m-1}(t + \frac{N}{2}) = x_m(t), \ s.t. \ 0 \le t \le \frac{N}{2} \ . \tag{9}$$

With equation 9, we can deduce the final formulation for the reconstructed signal $\overline{x}_m(n)$ as following:

$$\begin{aligned}\overline{x}_m(n) &= \frac{1}{N}\sum_{k=0}^{N-1}\left(\sum_{n=0}^{N/2-1} x_m(n)e^{j\delta_{m-1}[k]} + \sum_{n=0}^{N/2-1} x_m(n)e^{j\delta_m[k]}\right) \\ &= \frac{1}{N}\sum_{n=0}^{N/2-1} x_m(n)\sum_{k=0}^{N-1}\left(e^{j\delta_{m-1}[k]} + e^{j\delta_m[k]}\right) \ ,\end{aligned} \tag{10}$$

where the clean audio can be represented by setting $\delta_{m-1}[k] = 0$ and $\delta_m[k] = 0$. Therefore, the ratio $\gamma_m[k]$ of the reconstructed signal's spectrogram energy to the clean audio's satisfies:

$$\gamma_m[k] = \frac{\left|e^{j\delta_{m-1}[k]} + e^{j\delta_m[k]}\right|}{|e^{j0} + e^{j0}|} = \frac{\sqrt{2 + 2\cos(\delta_{m-1}[k] - \delta_m[k])}}{2} \ , \tag{11}$$

where $\gamma_m[k]$ is ranged within $[0, 1]$. It means that the phase perturbations $\delta_{m-1}[k], \delta_m[k]$ always cause loss of the reconstructed magnitude spectrogram, which is called 'energy dissipation'. Till now, we reveal the bi-directional essence of phase-oriented audio adversarial attack: it dissipates the critical energy for the classification model and constructs the adversarial perturbations in the meantime.

3.3 Optimization with Energy Dissipation

Simply minimizing the l_p distortion of $(1 - \gamma)$ is a possible solution. However, the l_p based metrics yield more audible perturbations in the practical attacks due to the differences between l_p distance and human perceptibility. In this paper, we propose a weighted loss function to constrain phase perturbations in an imperceptible range. Since the perturbations of our attack always cause energy loss to the spectrogram, this function should dissipate the energy that is crucial for the classification model and meanwhile preserve the energy that is important for human perception. According to the characteristics of human auditory systems, we divide the spectrogram into three parts: 1) the signals that are important for human perception; 2) the signals that are relatively unimportant for human auditory systems; 3) the inaudible signals that are below the threshold in quiet; For 1) and 2), we manage to constrain their γ in different levels β_1 and β_2, respectively. For 3), we do not constrain the corresponding γ since the energy dissipation in these areas can not be perceived. We divide quiet signals 3) following Lin & Abdulla (2015). The signals 1) are the local maxima that have the largest magnitude within 3 bins in the spectrum frames (in order to preserve the signals in high frequency we do not adopt Bark scale) and the signals 2) are the rest of the signals in the spectrogram.

Weighted Loss function We formulate the construction of imperceptible phase-oriented adversarial examples as minimizing the following weighted loss function:

$$\mathcal{L} = \mathcal{L}_{net}(f(\mathcal{E}(\phi^{-1}(|X| \cdot e^{i\theta_x} \cdot e^{j\delta}))), y_t) + \lambda \cdot \mathcal{L}_{imp}(\gamma) \ , \tag{12}$$

$$\mathcal{L}_{imp}(\gamma) = \alpha_1 \cdot \sum_{m=0}^{N-1} \sum_{k=0}^{N-1} \mathcal{I}_1(max\{\beta_1 - \gamma_m[k], 0\}) + \alpha_2 \cdot \sum_{m=0}^{N-1} \sum_{k=0}^{N-1} \mathcal{I}_2(max\{\beta_2 - \gamma_m[k], 0\}) \ , \tag{13}$$

where $\mathcal{I}_1$ and $\mathcal{I}_2$ are the indicator functions that are equal to 1 when the $\gamma_m[k]$ is in the corresponding part in the spectrogram and $\alpha_1 = 1$, $\alpha_2 = 0.05$ are trade-off parameters. We initialize the parameter λ as 1×10^{-4} and update λ every 10 iterations according to the performance of the attack. Since the hinge loss is utilized to constrains the dissipation ratio γ of the part 1) and 2) above the levels $\beta_1 = 0.98$ and $\beta_2 = 0.90$, the tiny energy losses caused by the phase perturbations will be perceived as almost inaudible volume differences instead of the echoes or noises;

4 Experimental Setup

4.1 Datasets

We conduct experiments on the LibriSpeech dataset (Panayotov et al., 2015), which contains 1000 hours of 16KHz English speech derived from audiobooks. The train-clean-100 and train-clean-360 datasets are used to train the threat ASR model. Then we randomly select 100 audio examples from the test-clean dataset as source examples, and 100 separate transcriptions from the same dataset to be the targeted transcriptions. We ensure that each target transcription is around the same length as the original transcription because it is unrealistic to perturb a short audio sample (e.g., 10 words) to have a much longer transcription (e.g., 30 words). Examples of the original and targeted transcriptions are available in the supplementary material.

Method	WER↓	Q1↓	Q2↓
C&W's Attack	0%	77% / 33%	56%
IPC Attack	0%	70% / 40%	58%
Qin's Attack	0%	47% / 37%	8%
PhaseFool	0%	**46% / 35%**	9%

Table 1: Results of the user study for imperceptibility. Q1, Q2 are the two types of questions in the experiments. The results in the Q1 column before '/' are the scores of the adversarial examples and after '/' are the scores of the clean audio. '↓' means lower is more imperceptible.

4.2 Settings and Metrics

Implementation Details. We implement an end-to-end speech recognition model based on the low-rank transformer architecture following Winata et al. (2020a). The transformer model consists of four encoder layers and four decoder layers with the dim_{inner} of 2048, dim_{model} of 512, and dim_{emb} of 512. The number of attention heads is set to 8. The encoder module adopts a VGG net (Simonyan & Zisserman, 2015) with a 6-layer CNN architecture as the frontend feature extractor. The VGG net consists of convolutional layers that learn a universal representation of the spectrogram and generate input embedding. Since the feature extraction is differentiable, our phase attack can be easily extended to the ASR models that take other acoustic features as input. All experiments are carried out on 1 Nvidia GTX 3080 GPU with 10 GB memory. We use Adam (Kingma & Ba, 2015) as the optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\epsilon = 10^{-9}$ and the learning rate is set as 5×10^{-3}.

Threat Model. As is done in most previous works, we assume a white-box threat model, i.e. the adversary is allowed to compute gradients through the model in order to generate adversarial examples. Considering that there are many ways to convert a white-box model to a black-box model Chen et al. (2020) and to improve the robustness against distortions Yakura & Sakuma (2019); Liu et al. (2020) and defencesZhang et al. (2020), our work focus on the white-box model, as is done in most previous works.

Evaluation Metrics. We calculate the Word error rate (WER) metric which is defined as $WER = \frac{S+D+I}{N} \times 100\%$, where S, D, and I are the numbers of substitutions, deletions, and insertions, respectively, and N is the total number of words. Particularly, we choose the target transcription as the label for calculating the WER in the experiments.

5 Results

In this section, we evaluate the performance of PhaseFool in terms of imperceptibility, generation speedup, interpretation of attacks and ablation studies.

5.1 Imperceptibility Analysis

Experimental Design. We construct 100 full-sentence targeted adversarial examples sampled from the LibriSpeech test-clean dataset for our PhaseFool attack. For comparison, we then generate adversarial examples using the prior state-of-the-art works as baselines, including: 1) *C&W's attack (Carlini & Wagner, 2018)*, a two-stage adversarial attack for audio based on CTC loss and l_p-norm; 2) *IPC attack (Zhang et al., 2020)*, a new Iterative Proportional Clipping (IPC) algorithm that generates robust audio adversarial examples with temporal dependency; 3) *Qin's attack (Qin et al., 2019)*, the imperceptible attack that leverages the psychoacoustic principle of auditory masking. We ensure that all attacks have a 100% success rate and the target transcription is consistent among the experiments to exclude other interference factors.

The user study is conducted to verify the imperceptibility of these attacks. We recruit 24 users online and give them one of the two types of questions following (Qin et al., 2019), including: 1) clean or noisy, we give users a set of audio samples and ask them to tell us if they have any background noise (e.g., static, echoing, people talking in the background), which can represent the imperceptibility of attacks in practice; 2) identical or not, we present users with two audio samples and ask them if the

audio samples are identical, or if there are any differences, which is the most challenging experiment. The users are asked to listen to each sample with headphones on, and answer questions about a set of audio samples (one user can only answer one type of the questions). In all experiments, users can replay the audio samples multiple times when they are unsure of the answer.

Result Analysis. As shown in Figure 1, 33%∼40% of users believe that the clean audio samples contain noises. Only 47% and 46% of the users believe that the adversarial examples generated by the Qin's attack and PhaseFool attack contain some noises, which greatly surpass the l_p-norm based C&W's attack and IPC attack. In the more challenging experiment 2), there are only 8% and 9% users who can distinguish the adversarial examples of the Qin's attack and PhaseFool attack from the clean audio, respectively. It can be seen that our phase-oriented attack shows comparable imperceptibility with the attacks that leverage the psychoacoustic principle of auditory masking. However, our attack only needs 500 steps and one optimization stage to generate a full-sentence audio adversarial example rather than the 5000 step and two optimizations stage in Qin et al. (2019). We will further explore the generation speed of these attacks in the next subsection.

Method	**Steps**	**Efficiency(s)**
C&W's Attack	4000	211.87 ± 12.31
IPC Attack	3000	158.75 ± 10.99
Qin's Attack	5000	305.91 ± 8.02
PhaseFool	**500**	**33.10 ± 3.27**

Table 2: Comparison of the generation speed with 95% confidence intervals. Efficiency(s) means that the algorithm can generate a successful full-sentence audio adversarial examples within the listed seconds on average.

Method	**Steps**	**Efficiency(s)**	**Q1**	**Q2**
PhaseFool	500	33.10±3.27	46% / 35%	9%
PhaseFool with l_2 loss	500	32.54±3.50	73% / 35%	72%
PhaseFool w/o constraint	212	17.80±19.56	66% / 33%	46%

Table 3: The results of efficiency and imperceptibility evaluation in the ablation studies.

5.2 GENERATION SPEEDUP

We evaluate the generation speech of PhaseFool compared with several state-of-the-art audio adversarial attacks in the previous experiment. We generate 100 adversarial examples for each algorithm with the 100% success rate and ensure that the original audio samples and the target transcriptions are consistent among the experiments to exclude other interference factors. As shown in Table 2, it can be seen that PhaseFool can generate a successful full-sentence audio adversarial example in around 30 seconds on average, which is a 9.24x speedup compared with the imperceptible counterpart Qin's attack (Qin et al., 2019). PhaseFool is a one-stage attack without complex computations and only requires 500 steps to obtain a successful and imperceptible adversarial example, which demonstrates the superiority of the phase-oriented audio adversarial attack.

5.3 ABLATION STUDY

We evaluate the efficiency and imperceptibility in the ablation studies to explore the influences of the weighted loss function proposed in equation 13. As shown in Table 3, when we do not constrain the phase perturbations and optimize the example until it is successful, PhaseFool only needs 212 steps on average to construct an audio adversarial example. It can be seen that such examples are more imperceptible than the l_p-norm based attacks that need thousands of steps but introduce more noise than the psychoacoustic methods and PhaseFool with the weighted hinge loss in equation 13, which demonstrates the effectiveness of our proposed loss function. However, when we constrain the phase perturbations with l_2 loss, the noises in adversarial examples are even louder than the

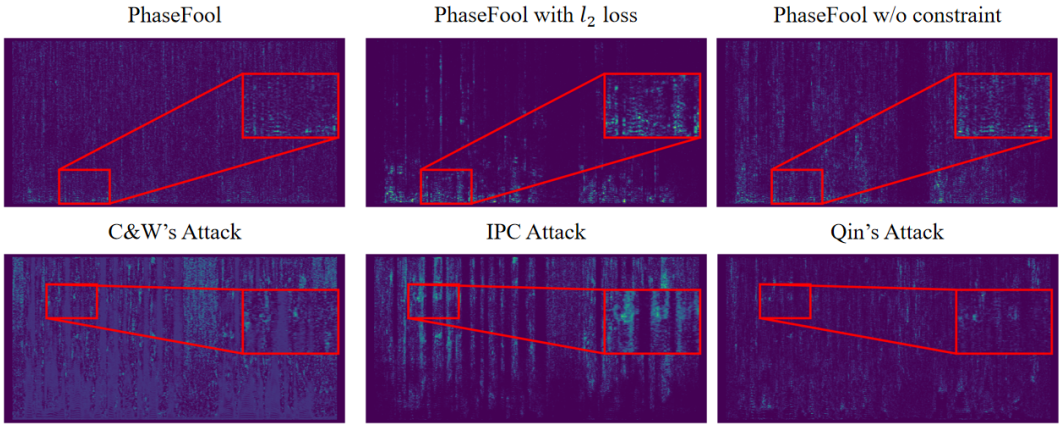

Figure 3: Pixel-wise difference of the generated audio adversarial examples in the mel-spectrogram domain. The perturbations of C&W's attack are disorderly distributed in the mel-spectrogram. Qin's attack and IPC attack tend to hide perturbations in the utterance regions while our PhaseFool attack generates perturbations around the harmonics in the spectrogram.

examples generated without constraint. It demonstrates the importance of curbing the energy dissipation ratio γ above the weighted level for different types of regions with the loss function proposed in subsection 3.3.

5.4 Interpretation of Attacks

For further investigating the interpretation of different attacks, we present the pixel-wise difference of the adversarial examples in the mel-spectrogram domain to visualize the learned adversarial perturbations in Figure 3. The perturbations of C&W's attack are disorderly distributed in the mel-spectrogram while Qin s attack and IPC attack tend to hide perturbations in the utterance regions. Our PhaseFool attack provides additional insights about the vulnerability of ASR models, i.e. the influences of different regions in the spectrogram for the audio adversarial attack. For example, as shown in Figure 3, the perturbations generated by our PhaseFool attack mainly occur on the positions around the harmonics in the mel-spectrogram. The perturbation positions reveal that which part of the spectrogram is important but also fragile for its prediction by the DNN model. With the l_2 loss function, PhaseFool significantly perturbs the mel-spectrogram regions that are important for human perception and cause audible echoes in the audio adversarial examples. When we removing the constraints of PhaseFool, the magnitude of perturbations in the mel-spectrogram grows obviously, which demonstrates the effectiveness of the proposed weighted loss.

6 Conclusions

In this work, we propose PhaseFool, a phase-oriented audio adversarial attack, which leverages energy dissipation to efficiently construct imperceptible audio adversarial examples. Leveraging such a phenomenon, we bring new directions for audio adversarial attacks and are able to perform a bi-directional attack that dissipates the crucial energy for ASR models and meanwhile, construct the perturbations for the adversarial attack. Moreover, we propose a weighted loss function to improve the imperceptibility of phase-oriented audio adversarial examples. Experiments on the LibriSpeech dataset (Panayotov et al., 2015) demonstrate that our PhaseFool significantly surpasses current state-of-the-art algorithms based on l_p-norm metrics and shows comparable imperceptibility with the attacks based on the psychoacoustic principle of auditory masking. Moreover, it can generate full-sentence adversarial examples with the 100% targeted success rate within 30 seconds approximately and achieves 9.24x speedup over the current state-of-the-art imperceptible counterparts.

For future work, we will continue to improve the imperceptibility of the phase-oriented audio adversarial examples, and apply PhaseFool to over-the-air and black-box settings for practical attacks. We will also investigate PhaseFool's capacity in online learning scenarios.

REFERENCES

Yang Ai and Zhen-Hua Ling. A neural vocoder with hierarchical generation of amplitude and phase spectra for statistical parametric speech synthesis. *IEEE ACM Trans. Audio Speech Lang. Process.*, 28:839–851, 2020. doi: 10.1109/TASLP.2020.2970241. URL https://doi.org/10.1109/TASLP.2020.2970241.

Moustafa Alzantot, Bharathan Balaji, and Mani Srivastava. Did you hear that? adversarial examples against automatic speech recognition. *arXiv preprint arXiv:1801.00554*, 2018.

Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Šrndić, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. Evasion attacks against machine learning at test time. In *Joint European conference on machine learning and knowledge discovery in databases*, pp. 387–402. Springer, 2013.

Nicholas Carlini and David A. Wagner. Audio adversarial examples: Targeted attacks on speech-to-text. In *2018 IEEE Security and Privacy Workshops, SP Workshops 2018, San Francisco, CA, USA, May 24, 2018*, pp. 1–7. IEEE Computer Society, 2018. doi: 10.1109/SPW.2018.00009. URL https://doi.org/10.1109/SPW.2018.00009.

Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David Wagner, and Wenchao Zhou. Hidden voice commands. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*, pp. 513–530, 2016.

Yuxuan Chen, Xuejing Yuan, Jiangshan Zhang, Yue Zhao, Shengzhi Zhang, Kai Chen, and XiaoFeng Wang. Devil's whisper: A general approach for physical adversarial attacks against commercial black-box speech recognition devices. In Srdjan Capkun and Franziska Roesner (eds.), *29th USENIX Security Symposium, USENIX Security 2020, August 12-14, 2020*, pp. 2667–2684. USENIX Association, 2020. URL https://www.usenix.org/conference/usenixsecurity20/presentation/chen-yuxuan.

Hyeong-Seok Choi, Jang-Hyun Kim, Jaesung Huh, Adrian Kim, Jung-Woo Ha, and Kyogu Lee. Phase-aware speech enhancement with deep complex u-net. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL https://openreview.net/forum?id=SkeRTsAcYm.

Moustapha Cissé, Yossi Adi, Natalia Neverova, and Joseph Keshet. Houdini: Fooling deep structured prediction models. *CoRR*, abs/1707.05373, 2017. URL http://arxiv.org/abs/1707.05373.

Gilles Degottex and Daniel Erro. A measure of phase randomness for the harmonic model in speech synthesis. In *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.

Yuan Gong and Christian Poellabauer. Crafting adversarial examples for speech paralinguistics applications. *CoRR*, abs/1711.03280, 2017. URL http://arxiv.org/abs/1711.03280.

Daniel Griffin and Jae Lim. Signal estimation from modified short-time fourier transform. *IEEE Transactions on acoustics, speech, and signal processing*, 32(2):236–243, 1984.

Yanxin Hu, Yun Liu, Shubo Lv, Mengtao Xing, Shimin Zhang, Yihui Fu, Jian Wu, Bihong Zhang, and Lei Xie. DCCRN: deep complex convolution recurrent network for phase-aware speech enhancement. In Helen Meng, Bo Xu, and Thomas Fang Zheng (eds.), *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*, pp. 2472–2476. ISCA, 2020. doi: 10.21437/Interspeech.2020-2537. URL https://doi.org/10.21437/Interspeech.2020-2537.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1412.6980.

Jonathan Le Roux, Hirokazu Kameoka, Nobutaka Ono, and Shigeki Sagayama. Fast signal reconstruction from magnitude stft spectrogram based on spectrogram consistency. In *Proc. DAFx*, volume 10, pp. 397–403, 2010.

Yiqing Lin and Waleed H Abdulla. Principles of psychoacoustics. In *Audio Watermark*, pp. 15–49. Springer, 2015.

Xiaolei Liu, Kun Wan, Yufei Ding, Xiaosong Zhang, and Qingxin Zhu. Weighted-sampling audio adversarial example attack. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 4908–4915, 2020.

Ryoji Miyahara and Akihiko Sugiyama. An auto-focusing noise suppressor for cellphone movies based on phase randomization and power compensation. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2199–2203. IEEE, 2014.

Pejman Mowlaee and Josef Kulmer. Phase estimation in single-channel speech enhancement: Limits-potential. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(8): 1283–1294, 2015.

Pejman Mowlaee, Rahim Saeidi, and Yannis Stylianou. Advances in phase-aware signal processing in speech communication. *Speech communication*, 81:1–29, 2016.

Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pp. 5206–5210. IEEE, 2015.

Yao Qin, Nicholas Carlini, Garrison W. Cottrell, Ian J. Goodfellow, and Colin Raffel. Imperceptible, robust, and targeted adversarial examples for automatic speech recognition. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5231–5240. PMLR, 2019. URL http://proceedings.mlr.press/v97/qin19a.html.

Lea Schönherr, Katharina Kohls, Steffen Zeiler, Thorsten Holz, and Dorothea Kolossa. Adversarial attacks against automatic speech recognition systems via psychoacoustic hiding. In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019. URL https://www.ndss-symposium.org/ndss-paper/adversarial-attacks-against-automatic-speech-recognition-systems-via-psychoacoustic-hiding/.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1409.1556.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

Tavish Vaidya, Yuankai Zhang, Micah Sherr, and Clay Shields. Cocaine noodles: exploiting the gap between human and machine speech recognition. In *9th {USENIX} Workshop on Offensive Technologies ({WOOT} 15)*, 2015.

Peter Vary and Member Eurasip. Noise suppression by spectral magnitude estimation—mechanism and theoretical limits—. *Signal processing*, 8(4):387–400, 1985.

Shengbei Wang, Weitao Yuan, Zhen Zhang, Jianming Wang, and Masashi Unoki. Synchronous multi-bit audio watermarking based on phase shifting. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*, pp. 2700–2704. IEEE, 2021. doi: 10.1109/ICASSP39728.2021.9414307. URL https://doi.org/10.1109/ICASSP39728.2021.9414307.

Genta Indra Winata, Samuel Cahyawijaya, Zhaojiang Lin, Zihan Liu, and Pascale Fung. Lightweight and efficient end-to-end speech recognition using low-rank transformer. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pp. 6144–6148. IEEE, 2020a. doi: 10.1109/ICASSP40776.2020.9053878. URL https://doi.org/10.1109/ICASSP40776.2020.9053878.

Genta Indra Winata, Samuel Cahyawijaya, Zhaojiang Lin, Zihan Liu, and Pascale Fung. Lightweight and efficient end-to-end speech recognition using low-rank transformer. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6144–6148. IEEE, 2020b.

Hiromu Yakura and Jun Sakuma. Robust audio adversarial example for a physical attack. In Sarit Kraus (ed.), *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pp. 5334–5341. ijcai.org, 2019. doi: 10.24963/ijcai.2019/741. URL https://doi.org/10.24963/ijcai.2019/741.

Hongting Zhang, P Zhou, Q Yan, and Xiao-Yang Liu. Generating robust audio adversarial examples with temporal dependency. In *Proc. Int. Joint Conf. Artif. Intell.*, pp. 1–5, 2020.