

# AI-Synthesized Voice Detection Using Neural Vocoder Artifacts

Chengzhe Sun    Shan Jia    Shuwei Hou    Siwei Lyu  
University at Buffalo, State University of New York, NY, USA  
{csun22, shanjia, shuweiho, siweilyu}@buffalo.edu

## Abstract

*Advancements in AI-synthesized human voices have created a growing threat of impersonation and disinformation, making it crucial to develop methods to detect synthetic human voices. This study proposes a new approach to identifying synthetic human voices by detecting artifacts of vocoders in audio signals. Most DeepFake audio synthesis models use a neural vocoder, a neural network that generates waveforms from temporal-frequency representations like mel-spectrograms. By identifying neural vocoder processing in audio, we can determine if a sample is synthesized. To detect synthetic human voices, we introduce a multi-task learning framework for a binary-class RawNet2 model that shares the feature extractor with a vocoder identification module. By treating vocoder identification as a pretext task, we constrain the feature extractor to focus on vocoder artifacts and provide discriminative features for the final binary classifier. Our experiments show that the improved RawNet2 model based on vocoder identification achieves high classification performance on the binary task overall. Codes and data can be found at <https://github.com/csun22/Synthetic-Voice-Detection-Vocoder-Artifacts>.*

## 1. Introduction

In recent years, the rapid development of AI technologies, particularly deep learning, has resulted in a surge of synthetic media, commonly referred to as "DeepFakes." These media are highly realistic and can be challenging to distinguish from genuine content, making them a significant concern. While AI-synthesized still images and videos have received much attention, synthetic human voices have also undergone significant advances, achieving unprecedented quality and efficiency. These voices have the potential to revolutionize voice-based user interfaces for intelligent home assistants and wearable devices and can even help individuals who have lost their ability to speak due to conditions like strokes or Amyotrophic Lateral Sclerosis (ALS).

However, the increasing realism and availability of syn-

thetic human voices also pose significant risks. Scammers have used AI-synthesized voices to impersonate individuals and deceive others into transferring money or providing sensitive information. In one instance, a scammer used an AI-synthesized voice to impersonate a UK company's CEO and tricked an employee into transferring a large sum of money to the scammer's account [1]. Moreover, trolls on the internet have used free AI voice cloning tools to imitate the voices of celebrities and create content ranging from memes to virulent hate speech [2].

While methods to detect AI-synthesized images and videos have been extensively studied, methods to detect synthetic human voices have received less attention and are underdeveloped. This is because audio signals have different characteristics that make it difficult to apply image-based detection methods. Early detection methods often analyze statistical features unique to audio signals. For example, [3] compares higher-order statistics in the bi-spectral domain that capture local phase inconsistencies in synthetic voices. Recent studies [4–6] tend to use well-designed models for automatic and comprehensive feature learning to detect synthesized audio.

This work proposes a new approach to detecting synthetic human voices based on artifacts introduced by the neural vocoders used in the generation process. A neural vocoder is a specialized neural network that synthesizes audio waveforms from temporal-frequency representations like mel-spectrograms. Since neural vocoders are the final step in most AI-based audio synthesis models, it is unlikely that real audio signals will be processed with neural vocoders. Thus, the vocoder artifacts can provide cues to identify synthetic human voices.

We aim to highlight the distinct signal artifacts left by neural vocoders in synthetic audio signals. The foremost objective of this study is to explore the artifacts of vocoders. To do this, we constructed a dataset called *LibriSeVoc*, which controls for other factors and only probes for the vocoder signature. The dataset contains evenly distributed data for various vocoders. We used six different neural vocoders to create the LibriSeVoc dataset to reflect the diversity in architecture and mechanisms of neural vocoders.




Figure 1. The artifacts introduced by the neural vocoders to a voice signal. We show the mel-spectrogram of the original (top left) and the self-vocoded voice signal (top right five). Their differences corresponding to the artifacts introduced by the vocoder are shown at the bottom.

”Self-vocoding” samples were sourced from the same original audio signals to highlight the artifacts introduced by the vocoders. Figure 1 shows the differences in the mel-spectrogram of one original voice and its self-vocoded voice signals. Visible artifacts introduced by different neural vocoder models can be observed, which serve as the basis of our detection algorithm. While these artifacts may be subtle to visualize, this work demonstrates that they can be captured by a trained classifier.

To take advantage of the vocoder artifacts in detecting synthetic human voices, we developed a multi-task learning strategy. We used a binary classifier that shares the front-end feature extractor (e.g., RawNet2 [4]) with the vocoder identification module. This accommodates the insufficient number of existing real and synthetic human voice samples by including the self-vocoding samples in LibriSeVoc as additional training data. We treated vocoder identification as a pretext task to constrain the front-end feature extraction module to focus on vocoder-level artifacts and build highly discriminative features for the final binary classifier. Our experiments showed that our RawNet2 model achieved outstanding classification performance on our LibriSeVoc dataset and two public DeepFake audio datasets. We also evaluated our method under different post-processing scenarios and demonstrated good detection robustness to re-sampling and background noise.

The main contributions of our work are as follows:

- We propose to focus on neural vocoder artifacts as specific and interpretable features for detecting AI-synthesized audio;
- We designed a novel multi-task learning approach that combines a binary classification task with a vocoder identification module. This approach constrains the feature extractor to learn discriminative vocoder artifacts for detecting synthetic human voices;
- We provide LibriSeVoc as a dataset with self-vocoding samples created using six state-of-the-art vocoders to highlight and exploit the vocoder artifacts;
- Our proposed method was experimentally evaluated on three datasets and demonstrated its effectiveness.

Overall, our work provides a new and promising approach to detecting synthetic human voices by focusing on neural vocoder artifacts. Our multi-task learning strategy, together with the LibriSeVoc dataset, could serve as a valuable resource for future research in this area.

## 2. Related Works

In this section, we provide a literature review that is relevant to our research, including voice synthesis methods, state-of-the-art neural vocoder models, and existing AI-synthesized voice detection methods.

### 2.1. Human Voice Synthesis

The synthesis of human voice is a significant challenge in the field of artificial intelligence, with various practical applications such as voice-driven smart assistants and accessible user interfaces. Human voice synthesis can be classified into two general categories: text-to-speech (TTS) and voice conversion (VC). In this work, we focus on recent TTS and VC methods that use deep neural network models.

TTS systems transform input text into audio using the target voice and typically consist of three components: a text analysis module that converts the input text into linguistic features, an acoustic model that generates acoustic features in the form of a mel-spectrogram from the linguistic features, and a vocoder. Recent TTS models based on deep neural networks include WaveNet [7], Tacotron [8], Tacotron 2 [9], ClariNet [10], and FastSpeech 2s [11].

In contrast, VC models take a sample of one subject’s voice as input and create output audio of another subject’s voice of the same utterance. Recent VC models (e.g., [12–14]) usually work within the mel-spectrum domain and use deep neural network models to map between the mel-spectrograms of the input and output voice signals. These models use neural style transfer methods such as variational auto-encoder (VAE) or generative adversarial network (GAN) models to capture the utterance elements in the input voice and then combine them with the style of the output voice. The resulting mel-spectrogram is then reconstructed to an audio waveform using a neural vocoder. Both the TTS and VC models employ deep neural network models trained on large-scale human voice corpora.

## 2.2. Neural Vocoders

Vocoders are crucial components in both TTS and VC models as they synthesize output audio waveforms from mel-spectrograms. However, the transformation from audio waveforms to mel-spectrograms leads to the loss of information due to binning and filtering, making it difficult to recover the audio waveform from a mel-spectrogram. In recent years, deep neural network-based vocoders have been developed, significantly improving training efficiency and synthesis quality. There are three main categories of existing neural vocoders: autoregressive models, diffusion models, and GAN-based models.

Autoregressive models are probabilistic models that predict the distribution of each audio waveform sample based on all previous samples. However, since this process involves linear sample-by-sample generation, autoregressive models are slower than other methods. WaveNet [7], the first autoregressive neural vocoder, can also serve as a TTS or VC model depending on the input. WaveRNN [15] is another autoregressive vocoder that uses a single-layer recurrent neural network to efficiently predict 16-bit raw audio samples from mel-spectrogram slices.

Diffusion models are probabilistic generative models that run diffusion and reverse processes. The diffusion process is characterized by a Markov chain, which gradually adds Gaussian noise to an original signal until the noise is eliminated. The reverse process is a de-noising stage that removes the added Gaussian noise and converts a sample back to the original signal. WaveGrad [16] and DiffWave [17] are two notable examples of diffusion-based vocoder models. While diffusion models are the most time-efficient vocoders, their reconstruction qualities are inferior to autoregressive models, and the generated samples may contain higher levels of noise and artifacts.

GAN-based models follow the generative adversarial network (GAN) architecture [18], which employs a deep neural network generator to model the waveform signal in the time domain and a discriminator to estimate the quality of the generated speech. Mel-GAN [19] and Parallel WaveGAN [20] are the two most commonly used GAN-based neural vocoders. Recent works have shown that GAN-based vocoders outperform autoregressive and diffusion models in both generation speed and generation quality.

## 2.3. AI-synthetic Human Voice Detection

In recent years, detecting synthetic human voices has become crucial due to their potential misuse. While extensive research has focused on audio authentication for speech synthesis and replay attack detection [21, 22], detecting AI-generated audio with high realism and varying models is a developing field. One of the earliest methods for detecting AI-synthetic audio is bi-spectral analysis [3]. This method captures subtle inconsistencies in local phases of synthetic

human voices. Real human voice signals have random local phases due to audio waves transmitting and bouncing around in the physical environment, while synthetic human voices do not have these characteristics. Although these local phase inconsistencies cannot be detected by the human auditory system, they can be identified through bi-spectral analysis. Another method, known as DeepSonar [23], uses network responses of audio signals as the feature to detect synthetic audio. The ASVspoof Challenge 2021 evaluates additional state-of-the-art synthetic voice detection methods. The Gaussian mixture models CQCC-GMM [24], LFCC-GMM [24], a light convolutional neural network model LFCC-LCNN [24], and RawNet2 [4] have achieved the most reliable performance as primary baseline algorithms.

Recent studies have focused on improving the generalization capacity of fake audio detection. Various well-designed models have been proposed for Deepfake audio detection, such as the spectro-temporal graph attention network [4], unsupervised pretraining models [5], biometric characteristics verification model [25], and self-distillation framework [6]. However, Müller *et al.* [26] evaluated twelve architectures on their dataset with 37.9 hours of audio recording and found that related work performs poorly on real-world data, with some models even degenerating to random guessing. Thus, there is a high demand for developing efficient and effective models for AI-synthesized audio detection.

## 3. Method

We aim to detect synthetic human voices by identifying vocoder artifacts present in the audio signals. Since real human voice signals typically do not have vocoder artifacts, except for our self-vocoding signals that are specifically designed to have them, identifying the presence of vocoder artifacts is a key feature in detecting synthetic human voices.

To achieve this, let  $\mathbf{x}$  be the waveform of a human voice signal with a label  $y \in \{0, 1\}$ , where 0 corresponds to a real human voice and 1 corresponds to a synthetic human voice. Our goal is to build a classifier  $\hat{y} = F_{\theta}(\mathbf{x})$  that predicts the label of an input  $\mathbf{x}$ . We utilize the recent RawNet2 model [4] as the backbone for our classifier, as it was designed to operate directly on raw waveforms. This reduces the risk of losing information related to neural vocoder artifacts when compared to using pre-processed features such as mel-spectrograms or linear frequency cepstral coefficients (LFCCs).

The binary detection model can be constructed as a cascade of neural networks

$$F_{\theta}(\mathbf{x}) = B_{\theta_B}(R_{\theta_R}(\mathbf{x})) \quad (1)$$

where  $R_{\theta_R}(\mathbf{x})$  is the front-end RawNet2 model for feature extraction with its own set of parameters  $\theta_R$ ,  $B_{\theta_B}$  is a back-




Figure 2. Framework of the proposed synthesized voice detection method. Rather than simply learning discriminative features for binary classification, we incorporate a vocoder identification module with a multi-class classification loss to direct the feature extraction network toward prioritizing vocoder artifacts.

end binary classifier and  $\theta_B$  are its specific parameters, with  $\theta = (\theta_R, \theta_B)$ . We can train this classifier directly as in the previous work [27], by solving

$$\min_{\theta} \sum_{(\mathbf{x}, y) \in T} L_b(y, F_{\theta}(\mathbf{x})) \quad (2)$$

where  $L_b(y, \hat{y})$  could be any loss function for binary classification, for instance, the cross-entropy loss. The variable  $T$  refers to the training dataset that contains labeled real and synthetic human voice samples. However, this method assumes that there is a large number of synthetic human voice samples available, which is increasingly difficult to achieve due to the rapid advancement of synthesis technology. Additionally, this approach does not take into account the unique statistical properties of neural vocoders, which can be an essential indicator for synthetic audio signals.

To address the aforementioned problem, we propose a multi-task learning approach that combines binary classification with a vocoder identification task. This approach is designed to emphasize the importance of identifying vocoder-level artifacts in synthetic audio signals. Specifically, we augment our detection model with a vocoder identifier  $M_{\theta_M}$ , which categorizes a synthetic voice into one of the  $c$  possible neural vocoder models ( $c \in [0, C]$  where  $C \geq 2$ ). Our goal is to ensure that the feature extractor is trained to capture the distinct statistical characteristics of vocoders, making it more sensitive to these features. This approach is similar to self-supervised representation learning [28]. To this end, we form a new classification objective, as

$$\begin{aligned} & \min_{\theta_B, \theta_R} \lambda \sum_{(\mathbf{x}, y) \in T} L_b(y, B_{\theta_B}(R_{\theta_R}(\mathbf{x}))) \\ & + \min_{\theta_M, \theta_R} (1 - \lambda) \sum_{(\mathbf{x}, c) \in T'} L_m(c, M_{\theta_M}(R_{\theta_R}(\mathbf{x}))) \end{aligned} \quad (3)$$

In this equation,  $L_m$  is a multi-class loss function, and we use the softmax loss in our experiments.  $T'$  is a dataset containing synthetic human voices created with different neural

vocoders as corresponding labels. This dataset is much easier to create by performing “self-vocoding”, *i.e.*, creating synthetic human voices by running real samples through the mel-spectrogram transform and inverse, the latter performed with neural vocoders. We created such a dataset, LibriSeVoc, which will be described in detail in Section 4.1.  $\lambda$  is an adjustable hyper-parameter that controls the trade-off between the two loss terms.

The whole framework of our detection model is shown in Figure 2. Note that the two classification modules in the new learning objective function serve different roles. The first term pertains to binary classification and aims to distinguish between authentic and fake audio. Meanwhile, the second term focuses on vocoder identification, serving as a pretext task to direct the feature extractor’s attention toward vocoder-related artifacts. The two tasks share the feature extraction component so that the distinct features of the vocoders can be captured and transferred to the binary classification task.

## 4. Experiments

In this section, we present a series of experiments to evaluate the effectiveness of our proposed synthetic voice detection method. We start by introducing our LibriSevoc dataset, as well as two publicly available datasets that we use for evaluation. Then, we compare our method with state-of-the-art models on all three datasets, using both intra- and cross-dataset testing scenarios. Finally, we examine the robustness of our detection models to common post-processing operations.

### 4.1. Datasets

Three DeepFake audio datasets are considered in experiments, namely our LibriSeVoc, and two public datasets WaveFake [29] and ASVspoof 2019 [24].

**LibriSeVoc Dataset.** We have created a new open-source

Table 1. Data details of voice samples from real and each vocoder category in the LibriSeVoc dataset.

Model	Train-hour	Train-sample	Dev-hour	Dev-sample	Test-hour	Test-sample	Total-hour	Total-sample
Real	20.95	7,920	6.97	2,640	7.00	2,641	34.92	13,201
WaveNet (A01)	20.87	7,920	6.97	2,640	6.91	2,641	34.77	13,201
WaveRNN (A02)	20.95	7,920	7.00	2,640	6.94	2,641	34.90	13,201
WaveGrad (D01)	20.98	7,920	7.01	2,640	6.95	2,641	34.95	13,201
DiffWave (D02)	20.98	7,920	7.01	2,640	6.95	2,641	34.94	13,201
MelGAN (G01)	20.76	7,920	6.94	2,640	6.88	2,641	34.59	13,201
Parallel WaveGAN (G02)	20.76	7,920	6.94	2,640	6.88	2,641	34.59	13,201
<b>Total</b>	<b>146.25</b>	<b>55,440</b>	<b>48.84</b>	<b>18,480</b>	<b>48.51</b>	<b>18,487</b>	<b>243.66</b>	<b>92,407</b>

dataset, named LibriSeVoc, for the task of vocoder artifact detection. The statistical features of neural vocoders have not been extensively studied before, and the availability of large-scale datasets for the task of vocoder identification, particularly those that include multiple types of vocoders, is limited. We derived the LibriSeVoc dataset from the widely used LibriTTS speech corpus [30], which is often utilized in text-to-speech research [31–33]. The LibriTTS corpus is based on the Librispeech dataset [34], which contains samples extracted from audiobooks available on LibriVox [35].

We consider six state-of-the-art neural vocoders to generate speech samples in the LibriSeVoc dataset, namely, WaveNet and WaveRNN from the autoregressive vocoders, Mel-GAN and Parallel WaveGAN from the GAN-based vocoders, and WaveGrad and DiffWave from the diffusion-based vocoders. Specifically, we have 34.92 hours of real audio samples and have generated self-vocoded audio using six vocoders for each sample. A total of 208.74 hours of synthesized samples are created in the dataset.

Specifically, each vocoder synthesizes waveform samples from a given mel-spectrogram extracted from an original sample; we refer to this process as “self-vocoding.” By providing each vocoder with the same mel-spectrogram, we ensure that any unique artifacts present in the synthesized samples are attributable to the specific vocoder used to reconstruct the audio signal. The statistical information of the LibriSeVoc dataset is summarized below.

- The dataset contains 13,201 real audio samples and 79,206 fake audio samples from six vocoders (each with 13,201 audio samples).
- A total of 92,407 audio samples are included in the dataset, with audio length from 5 seconds to 34 seconds at 24kHz. Figure 3 shows the detailed audio length distribution in the dataset.
- We further split the whole dataset into three non-overlapped sets for training (55,440 samples), development (18,480 samples), and testing (18,487 samples) at the ratio of 6:2:2. More details can be found in Table 1.

**WaveFake Dataset.** This dataset [29] collects DeepFake audios from six vocoder architectures, including MelGAN,




Figure 3. Histogram of audio samples length in LibriSeVoc.

FullBand-MelGAN, MultiBand-MelGAN, HiFi-GAN, Parallel WaveGAN, and WaveGlow. It consists of approximately 196 hours of generated audio files derived from the LJSPEECH [37] dataset. Note that five of the six vocoder models in WaveFake are GAN-based vocoders. Differently, our LibriSeVoc dataset aims to consider a high diversity of vocoders for artifact extraction and covers three categories of widely-used vocoder structures, including autoregressive models, diffusion models, and GAN-based models.

**ASVspoof 2019 Dataset.** This dataset [36] is derived from the VCTK base corpus [38], which includes speech data captured from 107 speakers. It contains three major forms of spoofing attacks, namely synthetic, converted, and re-played speech. We labeled the samples from different vocoders as different classes in the training set for the multi-class loss calculation. More details about the dataset can be found in Table 2.

## 4.2. Implementation Details

We utilize the RawNet2 [4] model as the backbone for feature learning. The Adaptive Moment Estimation (Adam) [39] is used as our optimizer with a learning rate of 0.0001 and a batch size of 32. The loss weight  $\lambda$  is set as 0.5 in the experiment. To report the detection performance,

Table 2. Details of evaluation datasets.

Dataset	#Vocoder type	Frequency	Training size	Dev size	Testing size
LibriSeVoc	6	24kHz	55,440	18,480	18,487
WaveFake Dataset [29]	6	16kHz	64,000	16,000	24,800
ASVspoof 2019 [36]	6	16kHz	25,380	24,844	71,237

we calculate the Equal Error Rate (EER) following previous studies [24, 29, 40].

### 4.3. Baseline methods

**LFCC-LCNN [36].** The LFCC-LCNN method combines LFCC feature extraction with an LCNN classifier (DNN). LFCC-LCNN has been widely used since ASVspoof 2019. It achieved the second-best performance in the ASVspoof 2021 Speech Deepfake track during the ASVspoof challenge.

**RawNet2 [4].** The RawNet2 model is based on DNN speaker embedding extraction with the raw waveform as inputs. This powerful model uses a technique named feature map scaling which scales feature maps similar to squeeze-excitation. It performed the best in the ASVspoof 2021 Speech Deepfake track.

**WavLM [41].** The WavLM model from Microsoft is a self-supervised pre-trained multilingual model that can be used for a variety of downstream speech tasks. It is specifically designed to maintain speech content modeling based on masked speech prediction while also improving the potential for non-ASR tasks through speech denoising. Additionally, WavLM uses a gated relative position bias in its Transformer structure to better capture the sequence ordering of input speech.

**Wav2Vec2-XLS-R [42].** The Wav2Vec2-XLS-R model is a large-scale multilingual pre-trained model for speech-related tasks developed by Meta AI. It uses the wav2vec 2.0 objective for speech representation learning in 128 languages. This model can be fine-tuned for downstream tasks such as automatic speech recognition, translation, or speech classification.

### 4.4. Intra-dataset Evaluation

**Synthetic Human Voice Detection.** We first report the performance on the main task, *i.e.*, the classification of real and synthetic human voices on the LibriSeVoc and WaveFake datasets. The results in the first three columns of Table 3 show that our multi-task learning-based RawNet2 model with vocoder identification as a pretext task achieves the lowest EER of 0.13% on LibriSeVoc and 0.19% on WaveFake, obviously outperforming other baselines re-trained on each dataset.

**Evaluation on ASVspoof 2019.** We then evaluate our method on the ASVspoof 2019 dataset with spoofed and fake audios. The results in the last column of Table 3 show

Table 3. Detection EER (%) of intra-dataset testing on three datasets.

Methods	LibriSeVoc	WaveFake	ASVspoof
LFCC-LCNN [36]	0.14	<b>0.19</b>	11.60
RawNet2 [4]	0.17	0.32	6.10
WavLM [41]	0.45	2.92	6.94
Wav2Vec2-XLS-R [42]	1.54	2.33	13.48
<b>Ours</b>	<b>0.13</b>	<b>0.19</b>	<b>4.54</b>

that our model performed the best in detecting various audio spoofing attacks.

### 4.5. Cross-dataset Evaluation

In cross-domain testing, we evaluated the performance of pre-trained detection models trained on our LibriSeVoc dataset on the WaveFake dataset, which contained voice samples not included in the training set. Table 4 displays the comparison results with EER for each vocoder of the WaveFake dataset. It is evident that all the methods showed a noticeable degradation in performance during cross-domain testing, indicating poor generalization ability. Our proposed method achieved the lowest EERs in detecting unseen voice samples generated with the same vocoders as those in our LibriSeVoc dataset, specifically MelGAN and Parallel WaveGAN models. However, it is not surprising that the model exhibited poorer performance on other unseen vocoders. This highlights a limitation of our method, particularly when compared to the original Raw2net, as the vocoder artifacts are not generalized and are constrained by the training vocoder categories.

### 4.6. Robustness Evaluation

We evaluate the robustness of our detection method against common post-processing operations by constructing an augmented, degraded dataset from the LibriSeVoc test dataset. First, we resample the input speech to intermediate sampling rates (8kHz, 16kHz, 22.05kHz, 32kHz, and 44.1kHz) and then resample back to the original sampling rate (24 kHz). Additionally, we introduce background noise by adding a single pre-recorded crowd noise sample at three SNR values (8dB, 10dB, and 20dB). We randomly choose between the original, resampled, or noisy speech segments, with probabilities of 40%, 40%, and 20%, respectively. We use this augmented dataset to evaluate the robustness of our detection method against these post-processing operations.

The confusion matrix in Figure 4 further compares the

Table 4. Detection EER (%) of cross-dataset testing on WaveFake dataset.

Methods	Overall	Seen vocoder		Unseen vocoder			
		MelGAN	Parallel WaveGAN	WaveGlow	MultiBand MelGAN	FullBand MelGAN	HiFiGAN
LFCC-LCNN [36]	77.47	30.06	90.98	98.40	97.17	99.45	98.02
RawNet2 [4]	<b>25.98</b>	3.39	44.31	<b>0.58</b>	<b>23.34</b>	<b>43.82</b>	<b>33.38</b>
WavLM [41]	50	50	50	50	50	50	50
Wav2Vec2-XLS-R [42]	50	50	50	50	50	50	50
<b>Ours</b>	26.95	<b>2.16</b>	<b>35.53</b>	4.60	31.38	45.35	35.80




Figure 4. Confusion matrices evaluated on LibriSeVoc. Top: original testing set. Bottom: post-processed testing set. (OGA: ground truth, A01: WaveNet, A02: WaveRNN, D01: WaveGrad, D02: Diffwave, G01: MelGAN, G02: Parallel WaveGAN.)

detection and vocoder identification performance on the original LibriSeVoc set (with a detection EER of 0.13%) and on the post-processed dataset (with an EER of 2.73%). This shows that our method can extract discriminative vocoder-level features for vocoder identification and is also robust to common data post-processing operations.

## 5. Conclusion

In this work, we propose a novel approach for detecting synthetic human voices by identifying the traces of *neural*

*vocoders* in audio signals. To leverage the vocoder artifacts for synthetic human voice detection, we introduce a binary-class RawNet2 model that shares the front-end feature extractor with the one for vocoder identification. We employ a multi-task learning strategy where vocoder identification serves as a pretext task to constrain the front-end feature extraction module for building the final binary classifier. Our experiments demonstrate that our method achieves a high classification performance overall.

While our proposed approach has shown promising results, there is still room for improvement and several areas for future work. Firstly, we plan to expand the LibriSeVoc dataset to include a wider variety of real audio signals and environments. This will help to increase the diversity of the dataset and improve the generalization capability of our model. Secondly, although our method is effective in identifying synthetic audio through vocoder artifacts, we acknowledge that it is an indirect approach. In future work, we aim to explore methods that can directly differentiate real and synthetic audio by combining cues from vocoders and other signal features of audio DeepFakes. Finally, we will explore the use of our approach for detecting and preventing the misuse of synthetic audio in various applications, such as voice cloning or DeepFake videos.

## Social Impact Statement

The quality and efficiency of synthetic human voices generated by AI models have reached unparalleled heights. However, with the increasing realism and accessibility of these voices, there are significant risks involved. Fraudsters have exploited AI-generated voices to impersonate individuals and trick others into sharing sensitive information or transferring money. In this work, we propose a method that detects vocoder artifacts, which can expose AI-synthesized voices and mitigate the risks associated with them.

To ensure transparency, codes and data related with this work are made available as Open Source at <https://github.com/csun22/Synthetic-Voice-Detection-Vocoder-Artifacts>.

**Acknowledgment.** This work is supported by the Center for Identification Technology Research (CITeR) and the National Science Foundation under Grant No. 1822190.

## References

- [1] Forbes, “A Voice Deepfake Was Used To Scam A CEO Out Of \$243,000,” <https://www.cnn.com/2020/02/20/tech/fake-faces-deepfake/index.html>, 11 2019. 1
- [2] “4chan users embrace AI voice clone tool to generate celebrity hatespeech,” <https://www.theverge.com/2023/1/31/23579289/ai-voice-clone-deepfake-abuse-4chan-elevenlabs>, January 2023. 1
- [3] E. A. AlBadawy, S. Lyu, and H. Farid, “Detecting AI-synthesized speech using bispectral analysis.” in *CVPR Workshops*, 2019, pp. 104–109. 1, 3
- [4] H. Tak, J.-w. Jung, J. Patino, M. Kamble, M. Todisco, and N. Evans, “End-to-end spectro-temporal graph attention networks for speaker verification anti-spoofing and speech deepfake detection,” *arXiv preprint arXiv:2107.12710*, 2021. 1, 2, 3, 5, 6, 7
- [5] Z. Lv, S. Zhang, K. Tang, and P. Hu, “Fake audio detection based on unsupervised pretraining models,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 9231–9235. 1, 3
- [6] J. Xue, C. Fan, J. Yi, C. Wang, Z. Wen, D. Zhang, and Z. Lv, “Learning from yourself: A self-distillation method for fake speech detection,” *arXiv preprint arXiv:2303.01211*, 2023. 1, 3
- [7] A. van den Oord, S. Dieleman, H. Zen *et al.*, “WaveNet: A generative model for raw audio,” in *arXiv*, 2016. [Online]. Available: <https://arxiv.org/abs/1609.03499> 2, 3
- [8] Y. Wang, R. J. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. V. Le, Y. Agiomyrgiannakis, R. Clark, and R. A. Saurous, “Tacotron: A fully end-to-end text-to-speech synthesis model,” *CoRR*, vol. abs/1703.10135, 2017. [Online]. Available: <http://arxiv.org/abs/1703.10135> 2
- [9] Z. Wang, Y. Liu, and L. Shan, “CE-Tacotron2: End-to-end emotional speech synthesis,” in *2021 60th Annual Conference of the Society of Instrument and Control Engineers of Japan (SICE)*, 2021, pp. 48–52. 2
- [10] W. Ping, K. Peng, and J. Chen, “Clarinet: Parallel wave generation in end-to-end text-to-speech,” *arXiv preprint arXiv:1807.07281*, 2018. 2
- [11] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “FastSpeech 2: Fast and high-quality end-to-end text to speech,” *arXiv preprint arXiv:2006.04558*, 2020. 2
- [12] E. A. AlBadawy and S. Lyu, “Voice conversion using speech-to-speech neuro-style transfer,” *Proc. Interspeech 2017*, 2020. 2
- [13] L.-H. Chen, Z.-H. Ling, L.-J. Liu, and L.-R. Dai, “Voice conversion using deep neural networks with layer-wise generative training,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 12, pp. 1859–1872, 2014. 2
- [14] S. H. Mohammadi and A. Kain, “Voice conversion using deep neural networks with speaker-independent pre-training,” in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 19–23. 2
- [15] N. Kalchbrenner, E. Elsen, K. Simonyan *et al.*, “Efficient neural audio synthesis,” in *ICML*, 2018. 3
- [16] N. Chen, Y. Zhang, H. Zen *et al.*, “WaveGrad: Estimating gradients for waveform generation,” in *ICLR*, 2020. 3
- [17] Z. Kong, W. Ping, J. Huang *et al.*, “DiffWave: A versatile diffusion model for audio synthesis,” in *ICLR*, 2020. 3
- [18] I. Goodfellow, J. Pouget-Abadie, M. Mirza *et al.*, “Generative adversarial nets,” *NeurIPS*, vol. 27, 2014. 3
- [19] K. Kumar, R. Kumar, T. de Boissiere *et al.*, “MelGAN: Generative adversarial networks for conditional waveform synthesis,” *arXiv*, 2019. 3
- [20] R. Yamamoto, E. Song, and J.-M. Kim, “Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram,” in *ICASSP*, 2020. 3
- [21] Z. Wu, N. Evans, T. Kinnunen, J. Yamagishi, F. Alegre, and H. Li, “Spoofing and countermeasures for speaker verification: A survey,” *speech communication*, vol. 66, pp. 130–153, 2015. 3
- [22] H. A. Patil and M. R. Kamble, “A survey on replay attack detection for automatic speaker verification (asv) system,” in *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2018, pp. 1047–1053. 3
- [23] R. Wang, F. Juefei-Xu, Y. Huang, Q. Guo, X. Xie, L. Ma, and Y. Liu, “Deepsonar: Towards effective and robust detection of ai-synthesized fake voices,” in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 1207–1216. 3



- [24] M. Todisco, X. Wang, V. Vestman, M. Sahidullah, H. Delgado, A. Nautsch, J. Yamagishi, N. Evans, T. Kinnunen, and K. A. Lee, “ASVspoof 2019: Future horizons in spoofed and fake audio detection,” *arXiv preprint arXiv:1904.05441*, 2019. 3, 4, 6
- [25] A. Pianese, D. Cozzolino, G. Poggi, and L. Verdoliva, “Deepfake audio detection by speaker verification,” in *2022 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2022, pp. 1–6. 3
- [26] N. M. Müller, P. Czempin, F. Dieckmann, A. Froggyar, and K. Böttinger, “Does audio deepfake detection generalize?” *arXiv preprint arXiv:2203.16263*, 2022. 3
- [27] G. Yang, S. Yang, K. Liu *et al.*, “Multi-band melgan: Faster waveform generation for high-quality text-to-speech,” in *SLT workshop*, 2021, pp. 492–498. 4
- [28] L. Ericsson, H. Gouk, C. C. Loy, and T. M. Hospedales, “Self-supervised representation learning: Introduction, advances and challenges,” *IEEE Signal Processing Magazine*, vol. 32, 2022. 4
- [29] J. Frank and L. Schönherr, “Wavefake: a data set to facilitate audio deepfake detection,” *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS 2021)*, 2021. 4, 5, 6
- [30] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, “LibriTTS: A corpus derived from LibriSpeech for text-to-speech,” *arXiv preprint arXiv:1904.02882*, 2019. 5
- [31] J. Kim, S. Kim, J. Kong, and S. Yoon, “Glow-TTS: A generative flow for text-to-speech via monotonic alignment search,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 8067–8077, 2020. 5
- [32] R. Valle, K. Shih, R. Prenger, and B. Catanzaro, “Flowtron: an autoregressive flow-based generative network for text-to-speech synthesis,” *arXiv preprint arXiv:2005.05957*, 2020. 5
- [33] M. Chen, X. Tan, Y. Ren, J. Xu, H. Sun, S. Zhao, T. Qin, and T.-Y. Liu, “Multispeech: Multi-speaker text to speech with transformer,” *arXiv preprint arXiv:2006.04664*, 2020. 5
- [34] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “LibriSpeech: an ASR corpus based on public domain audio books,” in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210. 5
- [35] “Librivox – free public domain audiobooks,” <https://librivox.org/>. 5
- [36] G. Lavrentyeva, S. Novoselov, A. Tseren, M. Volkova, A. Gorlanov, and A. Kozlov, “Stc antispoofing systems for the asvspoof2019 challenge,” *arXiv preprint arXiv:1904.05576*, 2019. 5, 6, 7
- [37] K. Ito and L. Johnson, “The lj speech dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017. 5
- [38] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, “Superseded-cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit,” 2016. 5
- [39] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014. 5
- [40] J. Yamagishi, X. Wang, M. Todisco, M. Sahidullah, J. Patino, A. Nautsch, X. Liu, K. A. Lee, T. Kinnunen, N. Evans *et al.*, “ASVspoof 2021: accelerating progress in spoofed and deepfake speech detection,” *arXiv preprint arXiv:2109.00537*, 2021. 6
- [41] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, M. Zeng, X. Yu, and F. Wei, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, pp. 1–14, 10 2022. 6, 7
- [42] T. Arun Babu, C. Wang, A. Tjandra, K. Lakhota, Q. Xu, N. Goyal, K. Singh, P. Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, “Xls-r: Self-supervised cross-lingual speech representation learning at scale,” 11 2021. 6, 7