

# Improving Generalization for AI-Synthesized Voice Detection

Hainan Ren\*, Li Lin<sup>1</sup>, Chun-Hao Liu<sup>2</sup>, Xin Wang<sup>3</sup>, Shu Hu<sup>1†</sup>

<sup>1</sup> Purdue University

<sup>2</sup> Amazon

<sup>3</sup> University at Albany, SUNY

hnren666@gmail.com, lin1785@purdue.edu, chunhaol@amazon.com, xwang56@albany.edu, hu968@purdue.edu

## Abstract

AI-synthesized voice technology has the potential to create realistic human voices for beneficial applications, but it can also be misused for malicious purposes. While existing AI-synthesized voice detection models excel in intra-domain evaluation, they face challenges in generalizing across different domains, potentially becoming obsolete as new voice generators emerge. Current solutions use diverse data and advanced machine learning techniques (*e.g.*, domain-invariant representation, self-supervised learning), but are limited by predefined vocoders and sensitivity to factors like background noise and speaker identity. In this work, we introduce an innovative disentanglement framework aimed at extracting domain-agnostic artifact features related to vocoders. Utilizing these features, we enhance model learning in a flat loss landscape, enabling escape from suboptimal solutions and improving generalization. Extensive experiments on benchmarks show our approach outperforms state-of-the-art methods, achieving up to 5.12% improvement in the equal error rate metric in intra-domain and 7.59% in cross-domain evaluations.

**Code** — <https://github.com/Purdue-M2/AI-Synthesized-Voice-Generalization>

## Introduction

AI voice technology uses advanced models to generate natural human speech, mimicking tone and pronunciation to the extent that it is indistinguishable from real recordings. This technology has a wide range of applications, including voice assistants, audiobooks, voiceovers for videos and advertisements, and in fields such as healthcare and accessibility, where it can assist individuals with speech impairments. AI voice synthesis is rapidly evolving, using advanced variational auto-encoder (VAE) (Peng et al. 2020), generative adversarial network (GAN) models (Kumar et al. 2019; Yamamoto, Song, and Kim 2020; Lee et al. 2022), and diffusion models (Chen et al. 2020; Kong et al. 2020; Liu et al. 2023a) to generate audio waveforms from mel-spectrograms that are nearly indistinguishable from human recordings.

\*Work done as a remote student under Prof. Hu’s supervision

†Corresponding author (hu968@purdue.edu)

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

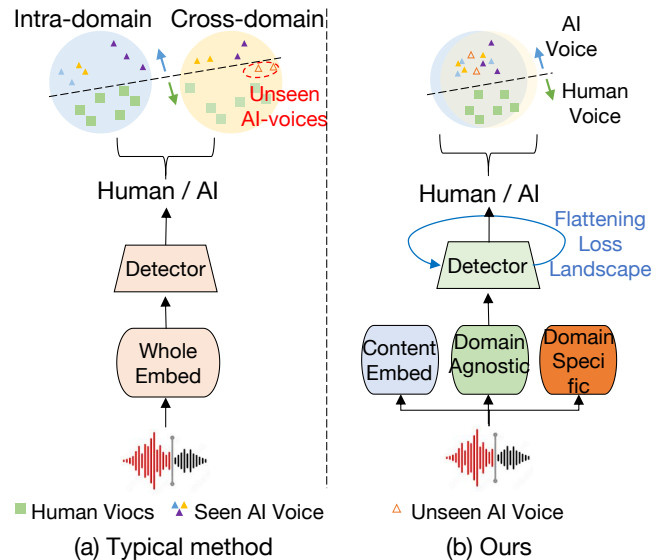


Figure 1: Comparison of audio deepfake detection methods. (a) The conventional method uses entire voice features and excels in distinguishing between human and AI-synthesized voices within familiar domains but struggles with voices generated from unseen vocoders, leading to inaccurate separation. (b) Our method achieves intra-&cross-domain detection by exposing domain-agnostic features for learning on a flattened loss landscape.

However, the widespread use of AI-synthesized voice has raised concerns about potential misuse, such as creating fake voice recordings, or ‘Audio Deepfakes’, for impersonation or fraud. For instance, in 2019, fraudsters used AI to mimic a CEO’s voice and stole over \$243,000 via a phone call (Forbes 2019). To address these concerns, research into detecting and mitigating AI-synthesized voice has become essential, with much of the progress (Barrington et al. 2023; Sun et al. 2023) driven by ASVspoof challenges and datasets (ASVspoof 2023). While these methods show promising results in intra-domain evaluations (*i.e.*, training and testing data come from the same vocoder), they face significant performance drops in cross-domain testing (*i.e.*, testing data is generated by an unseen vocoder).

Nüller et al. (Müller et al. 2022) provide a comprehensive evaluation of audio deepfake detectors’ generalization capabilities. They found that current detection methods, trained or designed based on the ASVspoof benchmark, perform poorly when detecting novel real-world audio deepfakes. To date, limited research has addressed the issue of generalization in detecting audio deepfakes. Previous efforts have focused on domain-invariant representation learning (Xie et al. 2023) and self-supervised learning (Wang and Yamagishi 2023). However, these approaches rely on predefined vocoders and are often influenced by extraneous factors like background noise and speaker identity.

To tackle these challenges, we propose a novel method that improves the generalization of AI-synthesized voice detection by addressing both feature-level and optimization-level factors (see Figure 1). We start by experimentally analyzing the complex interactions of features and how the sharpness of loss landscapes impacts generalization in AI-synthesized voice detection. We then introduce a novel disentanglement framework that combines multi-task and contrastive learning to extract domain-agnostic artifact features common across various vocoders. This process involves separating domain-agnostic from domain-specific artifacts (*i.e.*, artifacts linked to specific vocoders) and applying reconstruction regularization to maintain consistency between the original and reconstructed voices. To make the domain-agnostic features universally applicable and enhance generalization, we use content feature distribution as a benchmark and apply mutual information loss. This aligns the domain-agnostic features with the reference distribution. Finally, we optimize the model by flattening the loss landscape to avoid suboptimal solutions and further improve generalization. Our key contributions include:

1. We empirically analyze the factors from the feature and optimization levels affecting the generalization of AI-synthesized voice detection models.
2. We propose the first disentanglement framework to improve the generalization of AI-synthesized voice detection, focusing on both feature and optimization aspects. Specifically, we use disentanglement learning to extract domain-agnostic artifact features, which are utilized to enhance learning within a flattened loss landscape.
3. Our extensive experiments conducted on various prominent audio deepfake datasets demonstrate the effectiveness of our framework, which surpasses the performance of state-of-the-art methods in improving the generalization for cross-domain detection.

## Related Work

**Voice Synthesis.** Recent years witnessed remarkable improvements in synthesized speech quality, driven by deep learning (Tan et al. 2021). Two main approaches, Text-to-Speech (TTS) and Voice Conversion (VC), both utilize vocoders to synthesize voices. Early neural vocoders, like WaveNet (Oord et al. 2016) and WaveRNN (Kalchbrenner et al. 2018), used features directly for waveform generation. Recent ones, such as SC-WaveRNN (Paul, Pantazis, and Stylianou 2020) and MB WaveRNN (Yu et al. 2020) adopted

mel-spectrograms as input, but incurred slower inference times. To address this, alternative generative methods have emerged in waveform generation. These include *flow-based* methods like WaveGlow (Prenger, Valle, and Catanzaro 2019), FloWaveNet (Kim et al. 2019), and Squeeze-Wave (Zhai et al. 2020). *GAN-based* approaches namely MelGAN (Kumar et al. 2019), Parallel WaveGAN (Yamamoto, Song, and Kim 2020), and Fre-GAN (Lee et al. 2022). *VAE-based* techniques such as Wave-VAE (Peng et al. 2020). Additionally, *diffusion-based* methods like WaveGrad (Chen et al. 2020), DiffWave (Kong et al. 2020), and AudioLDM (Liu et al. 2023a).

**AI-Synthesized Voice Detection.** AI-synthesized voice detection is a fundamental task in machine learning, distinguishing authentic human speech from synthesized audio. The end-to-end approaches like RawNet2 (Tak et al. 2021) and RawFormer (Liu et al. 2023b) gained traction due to their competitive performance and efficiency. While these methods excel on intra-domain datasets, they often falter when dealing with cross-domain datasets. Some techniques (Salvi, Bestagini, and Tubaro 2023; Wang et al. 2023) Enhance generalization by identifying the most relevant signal segments or combining multi-view features to detect fake audio. However, they rely on predefined vocoders and consider the entire feature space, making them vulnerable to background noise and speaker identity. These challenges are the focal points of our work.

**Disentangled Representation Learning.** Disentanglement representation learning (Wang et al. 2022) is used in voice conversion, and voice style transfer during speech synthesis (Zhang et al. 2019; Aloufi, Haddadi, and Boyle 2020; Luong and Tran 2021; Champion, Jouviet, and Larcher 2022). Only Yadav et al. (2023) applies disentanglement representation learning in AI-synthesized voice detection. However, their approach—which extracts features directly from speech spectrograms—does not account for the artifact features characteristic of different vocoders. Additionally, it fails to address the domain-specific features that are intrinsic to different synthesis techniques. This leads to a limited generalization capability since it only eliminates the influence of the content features but may still overfit domain-specific patterns. In contrast, our proposed method effectively disentangles both domain-specific and domain-agnostic artifact features associated with various vocoders.

## Motivating Experiments

**Complex Entangled Information.** The generalization issue in AI-synthesized voice detection arises from two main factors. Firstly, many detectors overly emphasize irrelevant content, like identity and background noise. Secondly, different forgery techniques produce unique artifacts, as shown in Figure 2 (Left). The red circle shows unique large-scale artifacts in various AI-synthesized voices, easily detected by a vocoder-specific detector. However, detectors may become overly specialized in specific forgery, hindering generalization to unseen forgeries. To support this hypothesis, we extract features from the LibriSeVoc (Sun et al. 2023) dataset using RawNet2 (Tak et al. 2021) and Sun *et al.* (Sun et al. 2023). The UMAP visualizations (McInnes, Healy,

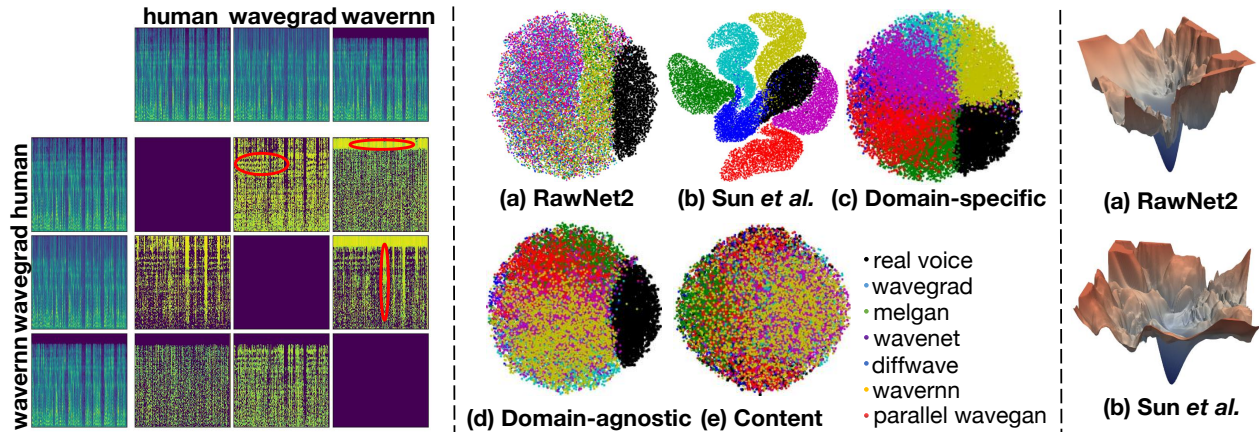


Figure 2: Experimental results for Motivation. **(Left)** The differences of mel-spectrogram between human voices and AI-synthesized ones (e.g., based on WaveGrad (Chen et al. 2020) and WaveRNN (Kalchbrenner et al. 2018) vocoders). The red circles highlight AI vocoder artifacts. More details about these differences can be found in Appendix. **(Middle)** The UMAP (McInnes, Healy, and Melville 2018) visualization of features from related methods and our framework on LibriSeVoc (Sun et al. 2023). The genuine and forged voices from six vocoders are separated in the latent space. The baselines (RawNet2 (Tak et al. 2021), Sun *et al.* (Sun et al. 2023)) and our domain-specific module can learn the domain-specific features, whereas the domain-agnostic module of our method captures the shared forgery features across different vocoders, and the content module exclusively captures forgery-irrelevant features. **(Right)** Visualization of loss landscape for RawNet2 and Sun *et al.* The sharp local and global minima could lead to models with poor generalization.

and Melville 2018) show forgery vocoders’ data clustering closely within baseline feature distribution (Figure 2 Middle (a) and (b)), while different vocoders’ data exhibit more distinct separations. This phenomenon is also observed in the domain-specific feature distribution from our method (Figure 2 (Middle (c))). Moreover, generalizable detectors should treat domain features from AI-synthesized voices equally but distinguish them from features of human voices (as illustrated in Figure 2 (Middle (d))). They should also treat content features equally, whether they are from human voices or AI-synthesized ones (as shown in Figure 2 (Middle (e))).

**Sharpness of Loss Landscape.** Existing DNN-based AI-synthesized voice detection models, such as RawNet2 (Tak et al. 2021) and Sun *et al.* (Sun et al. 2023), are highly over-parameterized and tend to memorize data patterns during training. This results in sharp loss landscapes with multiple minima (Figure 2 Right). Such sharpness presents challenges for models to locate the correct global minima for better generalization. Flattening the loss landscape is crucial to smooth the optimization path and enable robust generalization.

## Method

**Problem Setup.** Given a training dataset  $\mathcal{S} = \{(X_i, D_i, Y_i)\}_{i=1}^n$  of size  $n$ , where  $X_i \in \mathbb{R}^d$  represents the waveform of voice with feature dimension  $d$  and  $Y_i$  corresponds to the target label (e.g., synthetic voice or real voice). Moreover,  $D_i$  serves as the domain label, indicating the voice generation source of  $X_i$ . For instance, in the LibriSeVoc dataset (Sun et al. 2023),  $D_i \in \{\text{real}, \text{WaveNet} \text{ (Oord et al. 2016), WaveRNN} \text{ (Kalchbrenner et al. 2018), MelGAN} \text{ (Kumar et al. 2019), Parallel WaveGAN} \text{ (Yamamoto, Song, and Kim 2020), WaveGrad} \text{ (Chen et al. 2020), DiffWave} \text{ (Kong et al. 2020)}\}$ . Our goal is to train a

synthesized voice detector on  $\mathcal{S}$  for promising performance and high generalization to unseen synthesized voice data.

**Framework.** Our method, shown in Figure 3, consists of three main components: an encoder, a decoder, and two classification heads. The encoder contains two parts: a content encoder and an artifact encoder, each responsible for extracting content and artifact features. These two encoders share a common structure but operate with distinct sets of parameters. The decoder facilitates audio reconstruction by utilizing artifact features as a conditioning factor alongside content features. The classification module uses two heads: one for capturing domain-specific features and the other for extracting domain-agnostic features across different vocoders. The training process is guided by an optimization module designed to flatten the loss landscape. More details about the architecture can be found in Appendix.

## Disentanglement Process

We introduce a disentanglement learning module to extract vocoder-agnostic artifact features from the input voice for detection. To elaborate, consider a pair of voices ( $X_i, X_j$ ), where  $i, j \in \{1, \dots, n\}$  and  $i \neq j$ . Here,  $X_i$  represents the synthetic human voice (or real human voice), and  $X_j$  corresponds to a real human voice (or synthetic human voice). Our encoder, denoted as  $\mathbf{E}(\cdot)$ , comprises both a content encoder and an artifact encoder for the extraction of content features  $c$  and artifact features  $a$ . Notably, artifact features include domain-specific artifacts  $a^s$  (vocoder-specific features) and domain-agnostic artifacts  $a^g$  (common features across different vocoder models). The operation of the encoder can be represented as follows:  $c_i, a_i^s, a_i^g = \mathbf{E}(X_i)$ .

**Classification Loss.** To separate domain-specific artifacts from domain-agnostic ones, we use a simple approach with

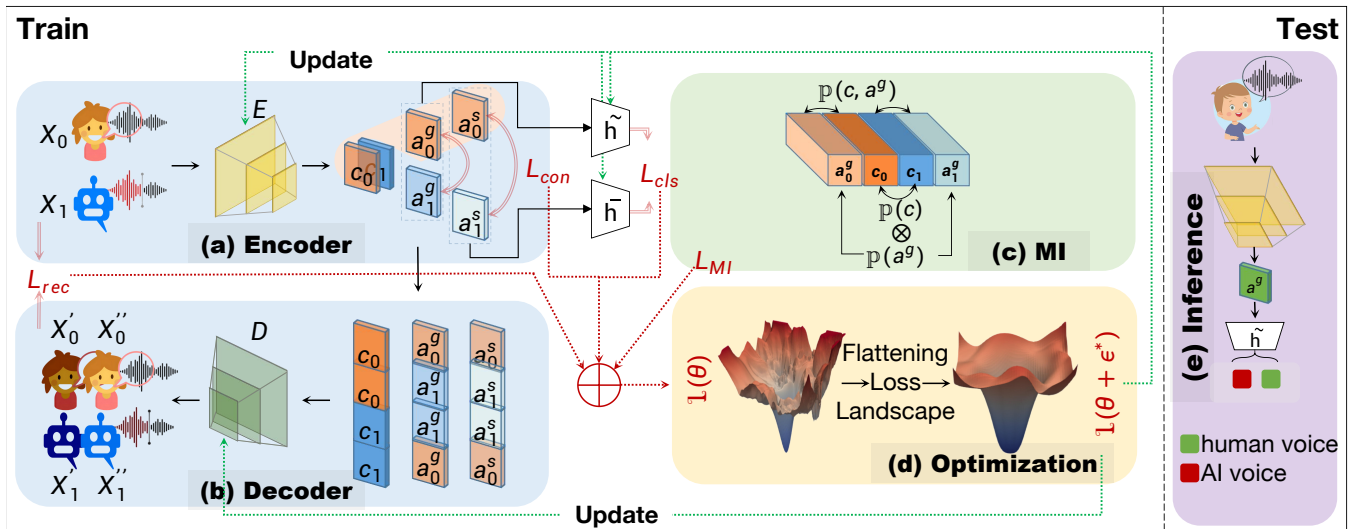


Figure 3: The overall architecture of our proposed approach. (a) In the encoder module, two RawNet2 (Tak et al. 2021)-structured backbones are added to the audio data signal to extract content and artifact features. Additionally, two headers further categorize artifact features into domain-agnostic and domain-specific categories. (b) In the decoder module, we use the content features as the base and fuse them with their own forgery features, as well as those from other samples for audio reconstruction. (c) Domain-agnostic features are made universally applicable by maximizing mutual information between domain-agnostic features and content features through joint and marginal distributions. (d) The sharpness-aware minimization (SAM) serves as an optimization technique to guide the model toward a flatter loss landscape to enhance its generalization. (e) For the inference, we take the predicted results of the domain-agnostic classification header.

multi-task learning, applying cross-entropy loss to each. The loss function can be formulated as follows:  $L_{cls} = C(\bar{h}(a_i^s), D_i) + \lambda_1 C(\bar{h}(a_i^g), Y_i)$ , where  $C(\cdot, \cdot)$  represents the CE loss.  $\bar{h}$  and  $\tilde{h}$  denote the classification heads for  $a_i^s$  and  $a_i^g$ , respectively.  $\lambda_1$  represents a hyperparameter. Training with this classification loss allows the encoder to learn both specific and shared artifact information, improving the model’s generalization capabilities.

**Contrastive Loss.** The classification loss mentioned above only considers individual voice information, neglecting the important global correlations between voices that can improve the encoder’s representation. Inspired by contrastive learning (Oord, Li, and Vinyals 2018; Yan et al. 2023; Lin et al. 2024), we address this by introducing a contrastive loss:  $L_{con} = [b + \|a_{\text{anchor}} - a_+\|_2 - \|a_{\text{anchor}} - a_-\|_2]_+$ , where  $a_{\text{anchor}}$  represents anchor artifact features of a voice, and  $a_+$  and  $a_-$  correspond to its positive counterpart from the same source and its negative counterpart from a different source, respectively. The hyperparameter  $b$  is introduced, and  $[\cdot]_+ = \max\{0, \cdot\}$  denotes a hinge function. We apply  $L_{con}$  to both domain-specific and domain-agnostic artifact features. For domain-specific features, the source is the neural vocoder type, and the contrastive loss drives the encoder to capture specific vocoder representations. For domain-agnostic features, the source is either synthetic or real human voice, encouraging the encoder to learn a generalizable representation independent of any particular vocoder.

**Reconstruction Loss.** To ensure the integrity of the extracted features and maintain consistency between the orig-

inal and reconstructed voices, we apply a reconstruction loss, defined as follows:  $L_{rec} = \|X_i - \mathbf{D}(c_i, a_i^s, a_i^g)\|_1 + \|X_i - \mathbf{D}(c_i, a_j^s, a_j^g)\|_1$ , where  $\mathbf{D}(\cdot, \cdot, \cdot)$  represents a decoder responsible for voice reconstruction based on the disentangled feature representations. In the  $L_{rec}$  loss, the first term is the self-reconstruction loss, which uses the input voice’s latent features to minimize reconstruction errors. The second term is the cross-reconstruction loss, which penalizes errors using the partner’s forgery feature. Together, these terms promote feature disentanglement.

**Mutual Information Loss.** To enhance generalization, it’s crucial to maintain consistent distributions of domain-agnostic features across different vocoders. This can be directly achieved by aligning their mutual relationships with the distributions of content features. While linear dependence/independence techniques (He et al. 2017, 2018) could be considered for this purpose, they often fail to capture the mutual relationships between content and domain-agnostic features in high-dimensional, nonlinear spaces. In contrast, mutual information (Kinney and Atwal 2014) is more effective for capturing arbitrary dependencies between variables.

Inspired by (Belghazi et al. 2018), we employ the Kullback-Leibler (KL) divergence (Joyce 2011), which is an equivalent form of mutual information, to quantify the dependencies between  $c$  and  $a^g$ . This is expressed as follows:  $\text{MI}(c; a^g) = \mathbb{D}_{\text{KL}}(\mathbb{P}(c, a^g) \| \mathbb{P}(c) \otimes \mathbb{P}(a^g))$ . In this equation,  $\mathbb{P}(\cdot, \cdot)$  represents the joint probability distribution,  $\mathbb{P}(\cdot)$  denotes the marginal probability distribution,  $\otimes$  signifies the product of the marginals, and  $\mathbb{D}_{\text{KL}}$  is the KL divergence.

Given that the probability densities  $\mathbb{P}(c, a^g)$  and

Methods	LibriSeVoc								ASVspoof2019							
	Seen vocoder				Unseen vocoder				Seen vocoder				Unseen vocoder			
	Avg	LSV	ASP	WF	Avg	ASP	FAVC	WF	Avg	ASP	LSV	Avg	ASP	FAVC	LSV	WF
LCNN (Lavrentyeva et al. 2019)	34.04	7.80	46.21	48.10	45.08	41.90	49.28	44.06	33.02	16.15	49.88	41.21	9.81	50.98	51.93	<b>52.13</b>
RawNet2 (Tak et al. 2021)	20.21	1.59	29.86	29.18	30.16	24.09	33.92	32.47	20.79	1.89	39.68	39.55	6.46	51.37	47.71	52.65
WavLM (Chen et al. 2022)	27.26	14.12	32.88	34.79	29.54	27.18	25.64	35.80	50.75	14.22	87.28	59.27	7.91	83.84	87.28	58.07
XLS-R (Babu et al. 2021)	33.47	11.21	45.37	43.83	45.11	51.23	42.91	41.18	53.65	9.04	98.26	74.85	6.40	94.91	98.26	99.82
Sun <i>et al.</i> (Sun et al. 2023)	18.67	3.79	22.77	29.45	27.86	24.42	25.52	33.65	22.18	3.92	40.44	41.59	8.38	54.78	50.59	52.62
<b>Ours</b>	<b>13.55</b>	<b>0.30</b>	<b>15.66</b>	<b>24.69</b>	<b>20.27</b>	<b>16.29</b>	<b>18.02</b>	<b>26.50</b>	<b>20.39</b>	<b>1.55</b>	<b>39.23</b>	<b>38.20</b>	<b>5.72</b>	<b>48.43</b>	<b>46.42</b>	52.24

Table 1: Comparisons of intra (seen)/cross (unseen)-domain generalization ability with comparable methods under EER (%). ‘LSV’, ‘ASP’, ‘WF’, and ‘FAVC’ stand for LibriSeVoc, ASVspoof2019, WaveFake, and FakeAVCeleb, respectively. ‘Avg’ represents the average EER (%) across the seen or unseen vocoders within each dataset.

$\mathbb{P}(c) \otimes \mathbb{P}(a^g)$  are not directly known, maximizing  $\mathbb{D}_{\text{KL}}(\mathbb{P}(c, a^g) || \mathbb{P}(c) \otimes \mathbb{P}(a^g))$  is a challenging task. However, we can maximize its lower bound  $L_{MI}$  using the Donsker-Varadhan representation (Donsker and Varadhan 1983), which can be expressed as:  $L_{MI} := \mathbb{E}_{x \sim \mathbb{P}(c, a^g)}[T(x)] - \log \mathbb{E}_{x \sim \mathbb{P}(c) \otimes \mathbb{P}(a^g)}[e^{T(x)}]$ , where  $T: \mathbb{R}^{d_c} \times \mathbb{R}^{d_{a^g}} \rightarrow \mathbb{R}$  represents a mutual information estimator. Inspired by (Hjelm et al. 2018),  $T$  can be absorbed into the encoder, combining the content feature  $c_i$  and domain-agnostic feature  $a_i^g$  in practice.

### Optimization with Flattening Loss Landscape

To sum up, the final loss function can be expressed as

$$\min_{\theta} \mathcal{L}(\theta) := \frac{1}{n} \sum_i [L_{cls} + \lambda_2 L_{con} + \lambda_3 L_{rec}] - \lambda_4 L_{MI}. \quad (1)$$

Here, we assume that the model weights of the entire framework are denoted as  $\theta$ . The hyperparameters  $\lambda_2$ ,  $\lambda_3$ , and  $\lambda_4$  are introduced to strike a balance between each term in the loss function. In practice, Eq. (1) can be solved by using a gradient descent approach to update  $\theta$ . To help the model avoid suboptimal solutions common in overparameterized DNNs and further improve generalization, we apply the SAM technique (Foret et al. 2020) to flatten the loss landscape. Specifically, this involves finding an optimal  $\epsilon^*$  to perturb  $\theta$  in a way that maximizes the loss, expressed as:

$$\epsilon^* = \arg \max_{\|\epsilon\|_2 \leq \gamma} \mathcal{L}(\theta + \epsilon) \approx \arg \max_{\|\epsilon\|_2 \leq \gamma} \epsilon^\top \nabla_{\theta} \mathcal{L} = \gamma \text{sign}(\nabla_{\theta} \mathcal{L}), \quad (2)$$

Here,  $\gamma$  is a hyperparameter that controls the perturbation magnitude, and  $\nabla_{\theta} \mathcal{L}$  is the gradient of  $\mathcal{L}$  with respect to  $\theta$ . The approximation term is derived using a first-order Taylor expansion, assuming  $\epsilon$  is small. The final equation results from solving a dual norm problem, using the sign function. Thus, the model weights are updated by solving:

$$\min_{\theta} \mathcal{L}(\theta + \epsilon^*). \quad (3)$$

The underlying idea is that perturbing the model in the direction of the gradient norm increases the loss value, thereby improving generalization. We optimize Eq. (3) using stochastic gradient descent, and the related algorithm is provided in the Appendix. Note that this is the first time to adapt SAM for AI-synthesized voice detection. Our ablation study demonstrates the effectiveness of this novel application of it in enhancing generalization.

## Experiment

### Experimental Settings

**Datasets.** To assess the generalization of our method, we tested it on various mainstream audio benchmarks, including LibriSeVoc (Sun et al. 2023), WaveFake (Frank and Schönherr 2021), ASVspoof 2019 (Lavrentyeva et al. 2019), and the audio segment of FakeAVCeleb (Khalid et al. 2021). More details of them and evaluations on other datasets (*e.g.*, ASVspoof2021 (Yamagishi et al. 2021)) are in Appendix.

**Baseline Methods.** To assess our method’s generalization capacity, we compared it with the following baselines: **1) LCNN** (Lavrentyeva et al. 2019) achieved the second-best performance in ASVspoof 2021 Speech Deepfake track. **2) RawNet2** (Tak et al. 2021) achieved the top performance in ASVspoof 2021 Speech Deepfake track. **3) WavLM** (Chen et al. 2022), developed by Microsoft, is a multilingual pre-trained model for general audio tasks. **4) XLS-R** (Babu et al. 2021), a robust cross-lingual speech model, excels in various domains, such as speech translation, speech recognition, and language identification. **5) Sun et al.** (Sun et al. 2023) is the latest method in detecting AI-synthesized voices by identifying the artifacts of vocoders in audio signals.

**Evaluation Metrics.** We evaluate our method using the Equal Error Rate (EER) metric, as commonly employed in previous studies (Frank and Schönherr 2021; Yamagishi et al. 2021) and baselines. A lower value means better performance. Furthermore, we add additional evaluations (Yamagishi et al. 2021) that include metrics AUC, False Acceptance Rate of Synthesized Voices, Rejection Rate of Real Voices, and Detection Cost Function in the Appendix, which also show the strong generalization of our method.

**Implementation Details.** Our encoders are based on RawNet2 (Tak et al. 2021) but exclude the last fully connected layer and utilize only features before it as encoder outputs. We employ the Adam (Kingma and Ba 2014) optimizer with a learning rate set to 0.0002 and a batch size of 16. Hyperparameters  $\lambda_1$ ,  $\lambda_2$ ,  $\lambda_3$ , and  $\lambda_4$  are set to 0.1, 0.3, 0.05, and 0.03, respectively. The margin  $b$  in  $L_{con}$  is set to 3. The  $\gamma$  in Eq. (2) is set to 0.07. We also use the original voice signal as input and apply the same data preprocessing as RawNet2 (Tak et al. 2021), padding all signals to the same size. More details can be found in the Appendix.

### Results

**Performance on Various Datasets.** To demonstrate the universality of our method, we expand our training set

Vocoders		LCNN	RawNet2	WavLM	XLS-R	Sun <i>et al.</i>	<b>Ours</b>
Seen	MelGAN	26.77	9.60	37.10	47.08	17.16	<b>2.93</b>
	PWGAN	59.21	43.81	31.05	40.08	36.71	<b>31.12</b>
	Avg	42.99	26.70	34.08	43.58	26.94	<b>17.03</b>
Unseen	WvGlow	28.21	4.14	36.47	31.34	12.87	<b>0.50</b>
	MBMGAN	46.06	37.21	31.77	44.13	37.91	<b>22.92</b>
	FBMGAN	48.69	44.65	37.08	45.42	46.85	<b>39.79</b>
	HiFiGAN	38.98	38.58	37.69	40.69	38.60	<b>29.76</b>
	Avg	40.49	31.14	35.75	40.40	34.06	<b>23.38</b>

Table 2: Detection EER (%) of cross-dataset evaluated on WaveFake Benchmark and trained on LibriSeVoc.

to include data from LibriSeVoc (Sun et al. 2023) and ASVspoof2019 (Lavrentyeva et al. 2019), respectively. Our model is subsequently evaluated on four distinct datasets: LibriSeVoc (Sun et al. 2023), ASVspoof2019 (Lavrentyeva et al. 2019), WaveFake (Frank and Schönherr 2021), and FakeAVCeleb (Khalid et al. 2021). We divide the test sets into two categories: seen vocoders from the same domain and unseen vocoders for cross-domain evaluation, based on the vocoder categories present in the training set. More details of dataset-vocoder partitions can be found in Appendix.

As shown in Table 1, our method excels baselines on the same domain and cross-domain scenarios. Specifically, our method outperforms Sun *et al.* (Sun et al. 2023) by nearly 5.12% (13.55% vs. 18.67% trained on LibriSeVoc) and 1.79% (20.39% vs. 22.18% trained on ASVspoof2019) in the intra-domain setting, and by approximately 7.59% (20.27% vs. 27.86% trained on LibriSeVoc) and 1.34% (38.20% vs. 39.55% trained on ASVspoof2019) in the cross-domain scene. Furthermore, despite the exceptional performance of the baselines in various audio downstream tasks, WavLM (Chen et al. 2022) and XLS-R (Babu et al. 2021) exhibit inferior performance in detecting AI-synthesized voices. While the LCNN method excels in performance for vocoder ‘WF’ on ASVspoof2019, our approach surpasses others in overall performance.

**Performance on Vocoders.** To assess our method on intra-domain and cross-domain vocoders, we train the model on LibriSeVoc (Sun et al. 2023) and evaluate it on WaveFake (Frank and Schönherr 2021). We categorize WaveFake into seen/unseen subsets and conduct testing on all vocoders within each subset. The results are reported in Table 2. We observe that our method not only achieves the best performance on seen vocoders but also exhibits a significant performance boost on unseen vocoders (23.38% vs. 34.06%).

**Performance Affected by Specific Vocoder.** To further investigate the model’s generalization capabilities and assess whether the models heavily rely on a specific vocoder, we remove one vocoder from LibriSeVoc (Sun et al. 2023) one by one, resulting in six sub-training sets. Then, we conduct tests on LibriSeVoc (Sun et al. 2023), ASVspoof2019 (Lavrentyeva et al. 2019), WaveFake (Frank and Schönherr 2021), and FakeAVCeleb (Khalid et al. 2021). Similarly, we categorize these benchmarks into seen and unseen domains and report the Equal Error Rate (EER) metric within these do-

Train Set	Methods	Seen vocoder				Unseen vocoder			
		Avg	LSV	ASP	WF	Avg	ASP	FAVC	WF
LSV w/o WaveNet	LCNN	33.9	8.3	45.1	48.4	46.2	40.6	54.1	44.0
	RawNet2	21.2	7.5	36.5	19.6	27.2	28.2	<b>23.0</b>	30.4
	Sun <i>et al.</i>	23.9	8.6	36.9	26.3	32.1	30.8	32.0	33.5
	<b>Ours</b>	<b>13.9</b>	<b>1.9</b>	<b>20.5</b>	<b>19.2</b>	<b>22.7</b>	<b>19.7</b>	23.3	<b>24.9</b>
LSV w/o WaveRNN	LCNN	36.5	9.8	49.9	50.0	47.7	50.0	43.5	49.7
	RawNet2	23.8	11.2	34.7	25.3	29.5	30.2	30.7	27.5
	Sun <i>et al.</i>	25.3	13.1	33.3	29.4	28.7	23.1	30.8	32.1
	<b>Ours</b>	<b>18.6</b>	<b>8.7</b>	<b>24.8</b>	<b>22.3</b>	<b>23.5</b>	<b>22.1</b>	<b>22.7</b>	<b>25.6</b>
LSV w/o WaveGrad	LCNN	33.1	7.8	43.9	47.5	44.6	42.1	45.8	45.9
	RawNet2	18.4	2.8	24.9	27.4	29.0	20.7	35.6	30.6
	Sun <i>et al.</i>	24.8	3.0	34.7	36.6	38.1	27.8	45.0	41.5
	<b>Ours</b>	<b>15.0</b>	<b>1.1</b>	<b>20.1</b>	<b>23.9</b>	<b>23.3</b>	<b>18.8</b>	<b>22.3</b>	<b>28.6</b>
LSV w/o DiffWave	LCNN	37.3	10.0	52.3	49.7	50.0	49.0	56.2	45.0
	RawNet2	24.5	2.6	39.5	31.4	34.6	29.7	40.3	33.8
	Sun <i>et al.</i>	29.3	6.7	38.7	42.4	33.4	29.4	33.4	37.4
	<b>Ours</b>	<b>16.4</b>	<b>1.2</b>	<b>17.3</b>	<b>30.6</b>	<b>27.5</b>	<b>19.8</b>	<b>31.5</b>	<b>31.3</b>
LSV w/o MelGAN	LCNN	36.4	12.5	49.8	46.8	45.6	48.2	47.6	41.1
	RawNet2	21.2	1.0	36.1	26.4	27.0	26.7	25.7	28.5
	Sun <i>et al.</i>	21.4	0.8	35.6	27.8	29.4	25.4	29.6	33.2
	<b>Ours</b>	<b>13.6</b>	<b>0.8</b>	<b>17.7</b>	<b>22.3</b>	<b>24.1</b>	<b>20.3</b>	<b>25.4</b>	<b>26.6</b>
LSV w/o PWaveGan	LCNN	36.0	7.9	50.0	50.3	47.0	46.5	50.0	44.6
	RawNet2	21.1	2.1	32.7	28.3	28.8	26.4	27.3	32.8
	Sun <i>et al.</i>	27.9	3.9	43.8	36.0	31.2	32.3	27.7	33.6
	<b>Ours</b>	<b>16.8</b>	<b>0.8</b>	<b>26.0</b>	<b>23.5</b>	<b>26.2</b>	<b>21.3</b>	<b>25.9</b>	<b>31.3</b>

Table 3: Comparisons of generalization ability with compared methods in EER (%). We use LSV (LibriSeVoc) as a train set, and test on four datasets: LSV, WF, ASP, and FAVC. The abbreviations ‘w/o WaveNet’ represent the voices generated by the WaveNet vocoder and are removed from LSV.

mains. As shown in Table 3, we observe that our method consistently achieves top performance on all six sub-training sets. Notably, Our feature decoupling approach generally performs better in unseen domains, demonstrating stable generalization in cross-domain scenarios.

## Ablation Study

**Analyzing the Framework Components in a Deconstructive Trajectory.** To evaluate the impact of each component in our proposed method, we conduct an analysis on several datasets. The results are reported in Table 4. Our modules consistently improve performance on seen or unseen vocoders. More detailed observations are as follows: **1)**  $V_A$  outperforms the baseline RawNet2 (Tak et al. 2021) on most datasets, especially on seen WF and unseen FAVC, revealing the efficacy of our reconstruction module. **2)**  $V_B$  and  $V_C$  achieve relatively similar results and further enhance the performance compared with  $V_A$  (e.g., 6.69% and 9.27% improvement of EER on seen ASP), indicating the necessity of our multi-task header and contrastive learning module. **3)** Implementing the mutual information module ( $V_D$ ) greatly improves the performance in cross-domain evaluation (e.g., 6.85% enhancement of EER on unseen FAVC compared

Method	Module						Seen vocoder				Unseen vocoder			
	RawNet2	Rec	Cls	Con	MI	SAM	Avg	LSV	ASP	WF	Avg	ASP	FAVC	WF
RawNet2	✓						20.21	1.59	29.86	29.18	30.16	24.09	33.92	32.47
$V_A$	✓	✓					19.16	0.53	31.84	25.11	28.57	27.49	26.35	31.87
$V_B$	✓	✓	✓				16.58	0.30	25.15	24.30	26.07	22.38	27.13	28.71
$V_C$	✓	✓	✓	✓			15.90	0.42	22.57	24.72	26.83	24.15	27.64	28.71
$V_D$	✓	✓	✓	✓	✓		15.71	0.38	25.62	<b>21.12</b>	23.23	23.52	20.79	<b>25.37</b>
<b>Ours</b>	✓	✓	✓	✓	✓	✓	<b>13.55</b>	<b>0.30</b>	<b>15.66</b>	24.69	<b>20.27</b>	<b>16.29</b>	<b>18.02</b>	26.50

Table 4: An analysis study of our key modules: RawNet2 (Tak et al. 2021), ‘Rec’(Reconstruction module), ‘Cls’(multi-task header), ‘Con’(contrastive learning module), ‘MI’(mutual information module), and ‘SAM’. These modules are constructed and trained under LSV and then tested on seen/unseen parts of LSV, ASP, WF, and FAVC.




Figure 4: (Left-Top) The virtualizations of model with MI ( $V_c$ ) and model without MI ( $V_d$ ). (Left-Bottom) The loss landscape visualization of our method with and without flattening the loss landscape. (Middle) (a) The effect of  $\lambda_4$  for balancing mutual information term. (b) The effect of  $\gamma$  in sharpness-aware minimization. (Right) The effect of the number of vocoders used in the train set. The reported average Equal Error Rate for seen (aEERs) and average Equal Error Rate for unseen (aEERu) domains.

with  $V_C$ ). To further illustrate the impact of mutual information, we present demographic-agnostic features in Figure 4 (Left-Top). It clearly shows that incorporating the mutual information module in the model leads to a more irregular distribution of fake features in the feature space, independent of the forgery techniques employed. **4)** Combining all components (Ours) achieves optimal results in both intra-domain and cross-domain evaluation, indicating the efficacy of the SAM approach. Figure 4 (Left-Bottom) illustrates SAM’s impact on the loss landscape, showing its ability to transform the rugged surface into a smoother one and enhance the model’s generalization capability.

**Effect of Hyperparameters.** With an increasing impact of mutual information ( $\lambda_4$ ), the model’s performance improves on both seen and unseen data. However, as  $\lambda_4$  continues to increase, the voice detection information within domain-agnostic features gradually diminishes, leading to a decline in the model’s ability to distinguish between Human and AI-synthesized voices. With increasing  $\lambda_4$ , voice detection information in domain-agnostic features diminishes, reducing the model’s ability to distinguish between human and AI-synthesized voices. Notably, we observe optimal performance when  $\lambda_4 = 0.03$  (Figure 4 (Middle (a))). As SAM intensity increases, model performance initially slightly decreases, followed by continuous improvement, and reaching optimal performance at  $\gamma = 0.07$  (Figure 4 (Middle (b))). This suggests that the model generalization benefits from a smoother loss landscape.

**Ablation on the number of vocoders in the train set.** To illustrate the impact of vocoder diversity in the training dataset on the model generalization. We create subsets of the

training data with different combinations of vocoder types, ranging from 1 to 6, sourced from LibriSeVoc. Trained models are evaluated on the similar seen/unseen manner, and the aEERs and aEERu are reported. Figure 4 (Right) indicates a clear trend: increasing vocoder diversity in training data enhances model generalizability.

## Conclusion

Existing AI-synthesized voice detection models are limited by predefined vocoders and sensitivity to factors like background noise and speaker identity. While excelling in intra-domain evaluation, they struggle to generalize across different domains with emerging voice generation advancements. To address this challenge, we propose a new disentanglement framework with a module extracting domain-agnostic features related to vocoders. Furthermore, we use these features to aid model learning in a flattened loss landscape, helping the model escape suboptimal solutions and enhance generalization. Extensive experiments on various datasets, compared to state-of-the-art methods, demonstrate the superior detection generalization of our framework. **Limitation:** One limitation of our method is its dependency on datasets including AI-synthesized voices generated by different types of vocoders. A few datasets contain synthesized voice from a single vocoder, limiting the extraction of domain-agnostic artifact features. **Future Work:** In future work, we aim to explore methods that can directly detect synthetic audio and improve detection generalization not limited to datasets containing multiple vocoders. We also plan to extend our approach to multi-modal tasks for audio and video detection.

## Acknowledgments

This work is supported by the U.S. National Science Foundation (NSF) under grant IIS-2434967 and the National Artificial Intelligence Research Resource (NAIRR) Pilot and TACC Lonestar6. The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of NSF and NAIRR Pilot.

## References

- Aloufi, R.; Haddadi, H.; and Boyle, D. 2020. Privacy-preserving voice analysis via disentangled representations. In *Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop*, 1–14.
- ASVspoof. 2023. Automatic Speaker Verification and Spoofing Countermeasures Challenge. In <https://www.asvspoof.org/>.
- Babu, A.; Wang, C.; Tjandra, A.; Lakhotia, K.; Xu, Q.; Goyal, N.; Singh, K.; von Platen, P.; Saraf, Y.; Pino, J.; et al. 2021. XLS-R: Self-supervised Cross-lingual Speech Representation Learning at Scale. *arXiv e-prints*, arXiv–2111.
- Barrington, S.; Barua, R.; Koorma, G.; and Farid, H. 2023. Single and Multi-Speaker Cloned Voice Detection: From Perceptual to Learned Features. *IEEE International Workshop on Information Forensics and Security*.
- Belghazi, M. I.; Baratin, A.; Rajeshwar, S.; Ozair, S.; Bengio, Y.; Courville, A.; and Hjelm, D. 2018. Mutual information neural estimation. In *International conference on machine learning*, 531–540. PMLR.
- Champion, P.; Jovet, D.; and Larcher, A. 2022. Are disentangled representations all you need to build speaker anonymization systems?
- Chen, N.; Zhang, Y.; Zen, H.; Weiss, R. J.; Norouzi, M.; and Chan, W. 2020. WaveGrad: Estimating Gradients for Waveform Generation. In *International Conference on Learning Representations*.
- Chen, S.; Wang, C.; Chen, Z.; Wu, Y.; Liu, S.; Chen, Z.; Li, J.; Kanda, N.; Yoshioka, T.; Xiao, X.; et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6): 1505–1518.
- Donsker, M. D.; and Varadhan, S. S. 1983. Asymptotic evaluation of certain Markov process expectations for large time. IV. *Communications on pure and applied mathematics*, 36(2): 183–212.
- Forbes. 2019. A Voice Deepfake Was Used To Scam A CEO Out Of \$243,000. In <https://www.forbes.com/sites/jessedamiani/2019/09/03/a-voice-deepfake-was-used-to-scam-a-ceo-out-of-243000/?sh=432e93d12241>.
- Foret, P.; Kleiner, A.; Mobahi, H.; and Neyshabur, B. 2020. Sharpness-aware Minimization for Efficiently Improving Generalization. In *International Conference on Learning Representations*.
- Frank, J.; and Schönherr, L. 2021. WaveFake: A Data Set to Facilitate Audio Deepfake Detection. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- He, R.; Wu, X.; Sun, Z.; and Tan, T. 2017. Learning invariant deep representation for nir-vis face recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- He, R.; Wu, X.; Sun, Z.; and Tan, T. 2018. Wasserstein CNN: Learning invariant features for NIR-VIS face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7): 1761–1773.
- Hjelm, R. D.; Fedorov, A.; Lavoie-Marchildon, S.; Grewal, K.; Bachman, P.; Trischler, A.; and Bengio, Y. 2018. Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*.
- Joyce, J. M. 2011. Kullback-leibler divergence. In *International Encyclopedia of Statistical Science*, 720–722.
- Kalchbrenner, N.; Elsen, E.; Simonyan, K.; Noury, S.; Casagrande, N.; Lockhart, E.; Stimberg, F.; Oord, A.; Dieleman, S.; and Kavukcuoglu, K. 2018. Efficient neural audio synthesis. In *International Conference on Machine Learning*, 2410–2419. PMLR.
- Khalid, H.; Tariq, S.; Kim, M.; and Woo, S. S. 2021. FakeAVCeleb: A Novel Audio-Video Multimodal Deepfake Dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Kim, S.; Lee, S.-G.; Song, J.; Kim, J.; and Yoon, S. 2019. FloWaveNet: A Generative Flow for Raw Audio. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 3370–3378. PMLR.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kinney, J. B.; and Atwal, G. S. 2014. Equitability, mutual information, and the maximal information coefficient. *Proceedings of the National Academy of Sciences*, 111(9): 3354–3359.
- Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2020. DiffWave: A Versatile Diffusion Model for Audio Synthesis. In *International Conference on Learning Representations*.
- Kumar, K.; Kumar, R.; De Boissiere, T.; Gestin, L.; Teoh, W. Z.; Sotelo, J.; De Brebisson, A.; Bengio, Y.; and Courville, A. C. 2019. Melgan: Generative adversarial networks for conditional waveform synthesis. *Advances in neural information processing systems*, 32.
- Lavrentyeva, G.; Tseren, A.; Volkova, M.; Gorlanov, A.; Kozlov, A.; and Novoselov, S. 2019. STC antispoofing systems for the AsVspoof2019 challenge. In *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, 1033–1037.
- Lee, S.-H.; Kim, J.-H.; Lee, K.-E.; and Lee, S.-W. 2022. FRE-GAN 2: Fast and Efficient Frequency-Consistent Audio Synthesis. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6192–6196.

- Lin, L.; He, X.; Ju, Y.; Wang, X.; Ding, F.; and Hu, S. 2024. Preserving fairness generalization in deepfake detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16815–16825.
- Liu, H.; Chen, Z.; Yuan, Y.; Mei, X.; Liu, X.; Mandic, D.; Wang, W.; and Plumbley, M. D. 2023a. AudioLDM: Text-to-Audio Generation with Latent Diffusion Models. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 21450–21474. PMLR.
- Liu, X.; Liu, M.; Wang, L.; Lee, K. A.; Zhang, H.; and Dang, J. 2023b. Leveraging Positional-Related Local-Global Dependency for Synthetic Speech Detection. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.
- Luong, M.; and Tran, V. A. 2021. Many-to-many voice conversion based feature disentanglement using variational autoencoder. In *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association*.
- McInnes, L.; Healy, J.; and Melville, J. 2018. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *ArXiv e-prints*.
- Müller, N.; Czempin, P.; Diekmann, F.; Froghyar, A.; and Böttinger, K. 2022. Does Audio Deepfake Detection Generalize? In *Proc. Interspeech 2022*, 2783–2787.
- Oord, A. v. d.; Dieleman, S.; Zen, H.; Simonyan, K.; Vinyals, O.; Graves, A.; Kalchbrenner, N.; Senior, A.; and Kavukcuoglu, K. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Paul, D.; Pantazis, Y.; and Stylianou, Y. 2020. Speaker Conditional WaveRNN: Towards Universal Neural Vocoder for Unseen Speaker and Recording Conditions. In *Proc. Interspeech 2020*, 235–239.
- Peng, K.; Ping, W.; Song, Z.; and Zhao, K. 2020. Non-Autoregressive Neural Text-to-Speech. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 7586–7598. PMLR.
- Prenger, R.; Valle, R.; and Catanzaro, B. 2019. Waveglow: A Flow-based Generative Network for Speech Synthesis. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3617–3621.
- Salvi, D.; Bestagini, P.; and Tubaro, S. 2023. Reliability Estimation for Synthetic Speech Detection. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.
- Sun, C.; Jia, S.; Hou, S.; and Lyu, S. 2023. AI-Synthesized Voice Detection Using Neural Vocoder Artifacts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop*, 904–912.
- Tak, H.; Patino, J.; Todisco, M.; Nautsch, A.; Evans, N.; and Larcher, A. 2021. End-to-end anti-spoofing with RawNet2. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6369–6373. IEEE.
- Tan, X.; Qin, T.; Soong, F. K.; and Liu, T. 2021. A Survey on Neural Speech Synthesis. *CoRR*, abs/2106.15561.
- Wang, C.; Yi, J.; Tao, J.; Zhang, C. Y.; Zhang, S.; and Chen, X. 2023. Detection of Cross-Dataset Fake Audio Based on Prosodic and Pronunciation Features. In *Proc. INTERSPEECH 2023*, 3844–3848.
- Wang, X.; Chen, H.; Tang, S.; Wu, Z.; and Zhu, W. 2022. Disentangled representation learning. *arXiv preprint arXiv:2211.11695*.
- Wang, X.; and Yamagishi, J. 2023. Can large-scale vocoded spoofed data improve speech spoofing countermeasure with a self-supervised front end? *arXiv preprint arXiv:2309.06014*.
- Xie, Y.; Cheng, H.; Wang, Y.; and Ye, L. 2023. Domain Generalization Via Aggregation and Separation for Audio Deepfake Detection. *IEEE Transactions on Information Forensics and Security*.
- Yadav, A. K. S.; Bhagtani, K.; Xiang, Z.; Bestagini, P.; Tubaro, S.; and Delp, E. J. 2023. Dsvae: Interpretable disentangled representation for synthetic speech detection. *arXiv preprint arXiv:2304.03323*.
- Yamagishi, J.; Wang, X.; Todisco, M.; Sahidullah, M.; Patino, J.; Nautsch, A.; Liu, X.; Lee, K. A.; Kinnunen, T.; Evans, N.; et al. 2021. ASVspoof 2021: accelerating progress in spoofed and deepfake speech detection. *arXiv preprint arXiv:2109.00537*.
- Yamamoto, R.; Song, E.; and Kim, J.-M. 2020. Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6199–6203. IEEE.
- Yan, Z.; Zhang, Y.; Fan, Y.; and Wu, B. 2023. Ucf: Uncovering common features for generalizable deepfake detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22412–22423.
- Yu, C.; Lu, H.; Hu, N.; Yu, M.; Weng, C.; Xu, K.; Liu, P.; Tuo, D.; Kang, S.; Lei, G.; Su, D.; and Yu, D. 2020. DurIAN: Duration Informed Attention Network for Speech Synthesis. In *Proc. Interspeech 2020*, 2027–2031.
- Zhai, B.; Gao, T.; Xue, F.; Rothchild, D.; Wu, B.; Gonzalez, J.; and Keutzer, K. 2020. SqueezeWave: Extremely Lightweight Vocoders for On-device Speech Synthesis. *ArXiv*, abs/2001.05685.
- Zhang, Y.-J.; Pan, S.; He, L.; and Ling, Z.-H. 2019. Learning latent representations for style control and transfer in end-to-end speech synthesis. In *IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE.