

VoiceCloak: A Multi-Dimensional Defense Framework Against Unauthorized Diffusion-Based Voice Cloning

Qianyue Hu¹, Junyan Wu¹, Wei Lu^{1*}, Xiangyang Luo²

¹ MoE Key Laboratory of Information Technology, School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou 510006, China

² State Key Laboratory of Mathematical Engineering and Advanced Computing, Zhengzhou 450002, China
huqy56@mail2.sysu.edu.cn, wujy298@mail2.sysu.edu.cn, luwei3@mail.sysu.edu.cn, luox_y_ieu@sina.com

Abstract

Diffusion Models (DMs) have achieved remarkable success in realistic voice cloning (VC), while they also increase the risk of malicious misuse. Existing proactive defenses designed for traditional VC models aim to disrupt the forgery process, but they have been proven incompatible with DMs due to the intricate generative mechanisms of diffusion. To bridge this gap, we introduce VoiceCloak, a multi-dimensional proactive defense framework with the goal of obfuscating speaker identity and degrading perceptual quality in potential unauthorized VC. To achieve these goals, we conduct a focused analysis to identify specific vulnerabilities within DMs, allowing VoiceCloak to disrupt the cloning process by introducing adversarial perturbations into the reference audio. Specifically, to obfuscate speaker identity, VoiceCloak first targets speaker identity by distorting representation learning embeddings to maximize identity variation, which is guided by auditory perception principles. Additionally, VoiceCloak disrupts crucial conditional guidance processes, particularly attention context, thereby preventing the alignment of vocal characteristics that are essential for achieving convincing cloning. Then, to address the second objective, VoiceCloak introduces score magnitude amplification to actively steer the reverse trajectory away from the generation of high-quality speech. Noise-guided semantic corruption is further employed to disrupt structural semantics captured by DMs, degrading output quality. Extensive experiments highlight VoiceCloak’s outstanding defense success rate against unauthorized diffusion-based voice cloning. Additional audio samples of VoiceCloak are available in demo pages.

Demo pages — <https://voice-cloak.github.io/VoiceCloak>

Introduction

Diffusion Models (DMs) (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020; Rombach et al. 2022) have recently emerged as powerful generative tools, achieving unprecedented success within realistic voice cloning (VC). Their iterative denoising process enables generating speech with remarkable naturalness, detail, and fidelity to human voice (Popov et al. 2022; Kong et al. 2020; Jeong et al. 2021;

*Corresponding author
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

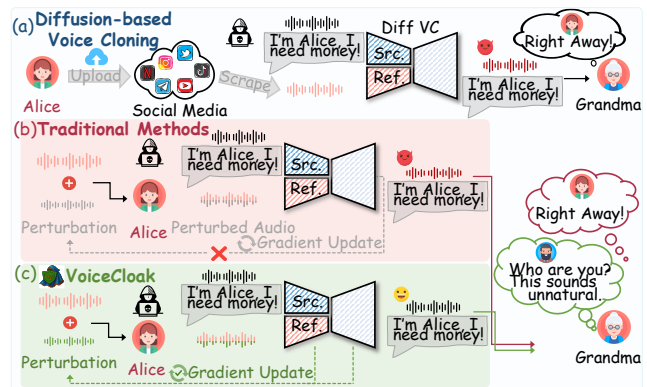


Figure 1: Illustration of diffusion-based voice cloning malicious misuse. (a) Voice forgery enables threats of fraud. (b) Traditional methods struggle due to ineffective disruptive gradients. (c) Audio protected by VoiceCloak resists high-fidelity cloning.

Shen et al. 2024). However, the open-source availability and ease of use of these models intensify concerns about potential misuse. Attackers can synthesize highly realistic voice replicas from short public audio clips, as depicted in Figure 1(a), enabling sophisticated fraud and circumvention of voiceprint authentication.

To counter such unauthorized use, two main defense paradigms are introduced, including forgery detection and proactive disruption. Reactive detection methods (Jung et al. 2022; Zhou and Lim 2021; Wu et al. 2024, 2025) identify forgeries after they are crafted, often too late to prevent harm. This highlights the need for proactive defenses that disrupt the synthesis process itself. Prior proactive work (Huang et al. 2021; Yu, Zhai, and Zhang 2023; Guo et al. 2025; Li et al. 2023; Qu et al. 2025) has focused on adding imperceptible adversarial perturbations to reference audio by compromising the functionality of either the voice decoder or the speaker identity encoder. However, existing defenses designed for prior architectures are largely ineffective against Diffusion Models (DMs). This incompatibility arises from two fundamental challenges: gradient vanishing and dynamic conditioning (Figure 1). Specifically, (1) the single forward pass gradient computation relied upon

by many defenses become unreliable or vanish within the multi-step denoising process of DMs and the corresponding deep computational graph, rendering such single-pass gradient information ineffective for disrupting the full generation trajectory (Kang, Song, and Li 2023). (2) Strategies targeting specific subnetworks (e.g. speaker or content encoders) fail because DMs often employ dynamic conditioning mechanisms, which means no single modules solely responsible for condition processing. Consequently, methods targeting individual components struggle to cause global disruption. These fundamental incompatibilities underscore the need for novel strategies tailored to this generative paradigm.

In light of the incompatibility of prior defenses with the diffusion paradigm, we introduce *VoiceCloak*, a novel proactive defense framework designed for two primary objectives against unauthorized voice cloning: **Speaker Identity Obfuscation** and **Perceptual Fidelity Degradation**. Driven by these two objectives, we conduct an analysis to identify and exploit corresponding intrinsic vulnerabilities within Diffusion Models (DMs). Based on this analysis, we design specific optimization objectives for the protective perturbation to effectively disrupt the synthesis process.

To achieve identity obfuscation, *VoiceCloak* first directly manipulates speaker representations within a universal embedding space, guided by psychoacoustic principles to maximize perceived identity distance and hinder the DM’s identity signature extraction. Second, recognizing that convincing mimicry depends on attention mechanisms in conditional guidance to align speaker style with content, we exploit this by introducing attention context divergence. This design prevents the attention mechanism from correctly utilizing contextual information, thereby disrupting the alignment required for accurate cloning.

Simultaneously, to achieve fidelity degradation, we focus on vulnerabilities within the core generative process itself. First, we employ Score Magnitude Amplification (SMA) to exploit the sensitivity of the iterative denoising trajectory which is crucial for realistic output, steering the generation path away from high-fidelity regions. Furthermore, acknowledging that the U-Net’s understanding of high-level semantics governs output naturalness, we utilize noise-guided semantic corruption to disrupt the capture of structural features to promote incoherence within the noise semantic space and degrade generation quality. These goal-driven strategies, which originate from adversarial analysis form a comprehensive defense. Extensive experiments confirm the superior defense efficacy against diffusion-based VC attacks under equivalent perturbation budgets.

Our contributions are summarized as follows:

- We propose *VoiceCloak*, a novel defense framework against Diffusion-Based VC that prevents unauthorized voice “theft” by exploiting intrinsic diffusion vulnerabilities through multi-dimensional adversarial interventions.
- We introduce auditory-perception-guided adversarial perturbations into speaker identity representations and disrupt the diffusion conditional guidance process to effectively distort identity information in synthesized audio.

- We present SMA, which controls the score function to divert the denoising trajectory, complemented by a semantic function designed to adversarially corrupt structural semantic features within the U-Net, degrading the perceptual fidelity of the forged audio.

Related Work

Audio Diffusion Models

Diffusion Models (Ruan et al. 2023; Kawar et al. 2023; Ruiz et al. 2023) a dominant force in generative modeling, demonstrating extraordinary performance across multimodal tasks, significantly advancing audio synthesis tasks like text-to-speech (Popov et al. 2021; Jeong et al. 2021) and unconditional audio generation (Kong et al. 2020; Liu et al. 2023). Particularly within voice cloning (VC), diffusion-based methods (Popov et al. 2022; Choi, Lee, and Lee 2024), mostly leveraging score-based formulations via stochastic differential equations (SDEs) (Song et al. 2020), now yield outputs with remarkable naturalness and speaker fidelity. While impressive, this state-of-the-art performance significantly heightens concerns regarding potential misuse, directly motivating proactive defense strategies such as the framework proposed herein.

Proactive Defense via Adversarial Examples

Beyond passive DeepFake detection, proactive defenses aim to preemptively disrupt malicious syntheses, by introducing adversarial perturbations to the original audio. Early work (Huang et al. 2021) demonstrated the feasibility of this approach but struggled to balance effectiveness with imperceptibility. Subsequent research focused on improving this trade-off. Strategies included using psychoacoustic models to enhance imperceptibility (Li et al. 2023), incorporating human-in-the-loop refinement for better balance (Yu, Zhai, and Zhang 2023), and employing GAN-based generators to improve efficiency (Dong et al. 2024).

Despite these advancements, prior proactive strategies were largely designed for earlier generative architectures like GANs. They often overlook the unique mechanisms and internal structures of Diffusion Models (DMs), limiting their applicability. Recognizing this critical gap, our work proposes a defense specifically tailored to the challenges of diffusion-based voice cloning.

Preliminaries

Score-based Diffusion

Score-based generative models define a continuous-time diffusion process using stochastic differential equations (SDEs) (Song et al. 2020). The forward process gradually perturbs clean data $\mathbf{x}_0 \sim p_0(\mathbf{x})$ into noise via the SDE:

$$d\mathbf{x} = f(\mathbf{x}, t)dt + g(t)d\mathbf{w}, \quad (1)$$

where $f(\mathbf{x}, t)$ and $g(t)$ are the drift and diffusion coefficients, and \mathbf{w} is a Wiener process. The corresponding **reverse-time SDE** that transforms x_T back into $p_0(x)$ can be expressed as:

$$d\mathbf{x} = [f(\mathbf{x}, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt + g(t)d\bar{\mathbf{w}}, \quad (2)$$

where $\bar{\mathbf{w}}$ is the reverse-time Wiener process and $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ represents the score-function. Then, the score-based diffusion model is trained to estimate the score function s_θ , and θ denotes the parameters of the denoiser:

$$\arg \min_{\theta} \lambda_t \mathbb{E}_{p_t(x)} \|s_\theta(\mathbf{x}_t, t) - \nabla \log p_t(\mathbf{x}_t | \mathbf{x}_0)\|_2^2, \quad (3)$$

where the expectation is taken over the data distribution $p_0(\mathbf{x}_0)$ and the transition kernel $p_t(\mathbf{x}_t | \mathbf{x}_0)$.

Adversarial Vulnerability Analysis

As mentioned before, the optimization of perturbation is guided by two core objectives: (1) Speaker Identity Obfuscation and (2) Perceptual Fidelity Degradation. The design of these perturbations stems from a targeted analysis to identify and exploit specific vulnerabilities within the Diffusion Model (DM) generative process itself.

A primary objective of VoiceCloak is speaker identity obfuscation. Convincing identity mimicry in Diffusion Models (DMs) depends on precise conditional control which is guided by acoustic details modeled from the reference audio x_{ref} (providing speaker ID). This guidance critically relies on mechanisms that align the target speaker’s acoustic characteristics from x_{ref} with the phonetic content from x_{src} (providing speech content). Specifically, the attention block is responsible for executing this alignment. We identify this crucial acoustic style-to-content mapping as a key vulnerability, as inaccurate alignment directly compromises the successful rendering of the target speaker’s identity. Consequently, disrupting this attention-driven pathway offers a strategy for identity obfuscation.

Complementary to identity obfuscation, the second objective is perceptual fidelity degradation, aimed at diminishing the usability of any synthesized speech. This involves targeting the core denoising process of DMs. The high fidelity of Diffusion-based VC relies on the model learning a precise reverse denoising trajectory to progressively refine noisy sample x_t towards the natural audio distribution. This reliance on trajectory precision presents an exploitable vulnerability. Therefore, adversarially diverting the trajectory can disrupt convergence towards the high quality audio region on the manifold.

Additionally, we target the U-Net’s internal feature representations as a vulnerability for degrading perceptual quality. Prior attribute editing work (Ronneberger, Fischer, and Brox 2015; Oh, Lee, and Lee 2024; Tumanyan et al. 2023; Wang et al. 2025) confirms that these hierarchical features within the U-Net are controllable and also encode crucial semantic and acoustic details that govern the coherence and naturalness of the synthesized speech. Therefore, adversarially corrupting these representations directly impair the model’s ability to synthesize natural-sounding speech, offering a complementary strategy for fidelity degradation. This analysis reveals critical vulnerabilities in DMs, providing targeted avenues for adversarial intervention.

Methodology

To provide the necessary background, the problem formulation is firstly established in Section 4.1. As shown in Figure

2, the VoiceCloak consists of four sub-modules, which will be introduced separately next.

Problem Formulation

Assume that malicious users can obtain reference audio x_{ref} of a target speaker. Leveraging open-source voice cloning models, they can synthesize speech which mimics the vocal characteristics of x_{ref} , denoted $\mathcal{VC}(x_{src}, x_{ref}, t)$. Our goal is to proactively safeguard x_{ref} against unauthorized voice cloning. We achieve this by introducing an imperceptible adversarial perturbation δ to create a protected version $x_{adv} = x_{ref} + \delta$. This perturbation is optimized to disrupt the diffusion synthesis process when conditioned on x_{adv} . Formally, we aim to find an optimal δ that maximizes the dissimilarity between the outputs generated using x_{ref} and x_{adv} :

$$\begin{aligned} & \arg \max_{\delta} \mathcal{D}(\mathcal{VC}(x_{src}, x_{ref}, t), \mathcal{VC}(x_{src}, x_{adv}, t)), \\ & \text{subject to } \|\delta\|_{\infty} \leq \epsilon, \end{aligned} \quad (4)$$

where $\mathcal{D}(\cdot)$ measures the output discrepancy, t is the diffusion timestep, and ϵ is the l_{∞} -norm budget for the perturbation δ . Our approach focuses on designing specific adversarial objectives that implicitly define \mathcal{D} by exploiting intrinsic vulnerabilities within the diffusion mechanism itself.

Adversarial Identity Obfuscation

Opposite-Gender Embedding Centroid Guidance Inspired by prior work showing dedicated speaker embeddings capture identity (Chen et al. 2024), directly manipulating these embeddings offers a direct approach, but its practical utility is hindered by poor transferability when attacking unknown or different encoder models.

Therefore, we explore leveraging speech Representation Learning (Chen et al. 2022b) to extract general acoustic representations that inherently capture speaker identity cues. Specifically, we select WavLM (Chen et al. 2022a) as our representation extractor, denoted as $\mathcal{R}(\cdot)$. By applying perturbations within this more general feature space, we aim for broader effectiveness against various models. As a baseline untargeted objective, we consider maximizing the embedding distance:

$$\mathcal{L}_{ID} = 1 - \text{Sim}(\mathcal{R}(x_{adv}), \mathcal{R}(x_{ref})), \quad (5)$$

where $\text{Sim}(x_1, x_2)$ represents the cosine similarity metric. While this untargeted objective effectively pushes x_{adv} away from the reference identity representation domain, it lacks specific directionality in high-dimensional space. So we further incorporate a targeted component guided by psychoacoustic principles.

Psychoacoustic studies (Kreiman and Sidtis 2011; Lavner, Gath, and Rosenhouse 2000) indicate that significant perceptual differences in speaker identity often exist between genders, linked to specific acoustic cues like F0 and formant structures. Leveraging this, we assume that guiding the adversarial embedding towards that of the opposite gender will likely create the strongest perceptual contrast, enhancing identity obfuscation.

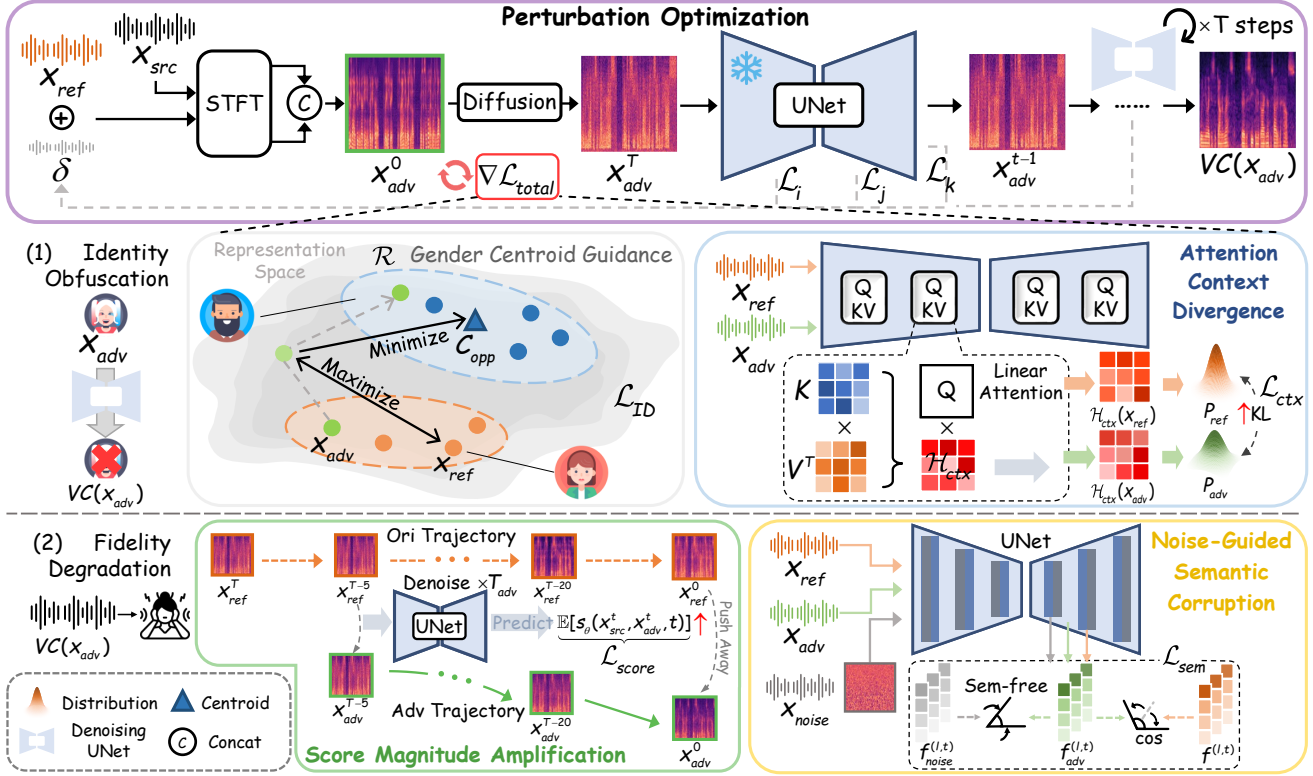


Figure 2: Overview of the proposed framework. Perturbation optimization is guided by gradients from \mathcal{L}_{total} , aggregating four targeting two objectives: (1) Identity Obfuscation (via Opposite-Gender Centroid Guidance and Attention Context Divergence) and (2) Perceptual Fidelity Degradation (Score Magnitude Amplification and Noise-Guided Semantic Corruption).

Based on this insight, we propose an auditory-perception-guided adversarial perturbation. Specifically, we first randomly select a speaker from the dataset whose gender is opposite to that of x_{ref} . To establish a representative identity embedding, we compute the centroid \mathcal{C}_{opp} by averaging the embeddings of all utterances in \mathcal{X} within the WavLM feature space \mathcal{R} :

$$\mathcal{C}_{opp} = \frac{1}{N} \sum_{x_i \in \mathcal{X}} \mathcal{R}(x_i), \quad (6)$$

$$\mathcal{L}_{ID} = -\text{Sim}(\mathcal{R}_{adv}, \mathcal{R}_{ref}) + \underbrace{\text{Sim}(\mathcal{R}_{adv}, \mathcal{C}_{opp})}_{\text{Gender}}, \quad (7)$$

where \mathcal{X} is the set of opposite gender utterances, N represents length of \mathcal{X} , \mathcal{R}_i is the embedding of x_i in space \mathcal{R} .

Finally, minimizing Eq. 7 directs the optimization process to simultaneously dissociate from the original speaker identity and converge towards a region selected based on psychoacoustic principles to maximize identity ambiguity.

Attention Context Divergence Building upon the motivation to disrupt conditional guidance for identity obfuscation (Section 3.2), we introduce Attention Context Divergence. This strategy targets the attention mechanism to interfere with its use of contextual information from x_{ref} .

Commonly in diffusion-based VC, conditional information from x_{ref} is integrated via Linear-attention layers (Katharopoulos et al. 2020). Within the U-Net, latent code representing the content x_{src}^t are linearly projected to Q matrix, whereas condition latents x_{ref}^t are projected to K and V . Linear-attention first computes a context matrix by aggregating values (V) weighted by keys (K):

$$\mathcal{H}_{ctx}(x_{ref}) = \text{Softmax}(\phi^{(l,t)}(x_{ref}^t)W_K^l)(\phi^{(l,t)}(x_{ref}^t)W_V^l)^T,$$

where \mathcal{H}_{ctx} represents the context hidden state, $\phi^{(l,t)}(\cdot)$ is the deep features of the l^{th} block in U-Net at timestep t , and W_K^l, W_V^l are the projection matrices. Then, the context interacts with queries (Q) to obtain the final attention output: $\mathcal{A}^l = (\mathcal{H}_{ctx})^T(W_Q^l \phi^{(l,t)}(x_{src}^t))$. This context representation provides a dynamic summary of the reference speaker’s stylistic features, weighted by their relevance to the current content queries during synthesis. By applying a softmax function, we obtain an explicit probability distribution over sequence positions which signifies the model’s information focus. Our strategy, therefore, is to maximize the Kullback-Leibler (KL) divergence (Kullback and Leibler 1951) between the context distribution derived from the original reference and the adversarial audio:

$$\mathcal{L}_{ctx} = D_{KL}(P_{ref} \parallel P_{adv}), \quad (8)$$

where $P_{adv} = \text{Softmax}(\mathcal{H}_{ctx}(x_{adv}))$.

Maximizing this divergence forces the attention pattern to deviate from the original, thereby hindering accurate style alignment.

To enhance the impact on identity, we direct the adversarial pressure on the U-Net’s downsampling path. The rationale is that these earlier layers primarily process coarser, lower-frequency features (Wang et al. 2024) strongly associated with speaker timbre and identity. Therefore, focusing our proposed loss \mathcal{L}_{ctx} on the attention layers within the U-Net’s downsampling path, the calculation of P_{adv} can be restated as:

$$P_{adv} = \text{Softmax}\left(\sum_l \mathcal{H}_{ctx}^l(x_{adv})\right), \quad (9)$$

where the layer index l iterates over the set Down ($l \in \text{Down}$), representing the U-Net’s downsampling blocks.

Perceptual Fidelity Degradation

Score Magnitude Amplification To degrade perceptual fidelity by exploiting the sensitivity of the denoising trajectory, we introduce Score Magnitude Amplification. This strategy directly targets the score function s_θ , which acts as the governing drift term for the reverse SDE (Eq. 3). Based on the insight that the magnitude of s_θ dictates the strength of the drift guiding the noise toward the data manifold. Consequently, our objective maximizes the magnitude of the predicted score s_θ :

$$\mathcal{L}_{score} = \mathbb{E}_{p_t(x), t \sim \mathcal{U}(1, T_{adv})} [\|s_\theta(x_{src}^t, x_{adv}^t, t)\|_2], \quad (10)$$

where $p_t(x)$ is the distribution of noisy samples $x^t \sim q(x^t|x^0)$, $\mathbb{E}[\cdot]$ calculates the average value, T_{adv} stands for adversarial timesteps which will be discussed in Section 4.4. Optimizing the above formula introduces an erroneous drift strength. Consequently, the denoising trajectory is forcefully diverted, resulting in a collapse in perceptual quality.

Furthermore, the iterative nature of the diffusion process may amplify these induced trajectory deviations. Perturbations introduced at earlier timesteps can propagate through subsequent steps. This error cumulative effect thus enhances the efficacy of our adversarial strategy.

Noise-Guided Semantic Corruption Following the motivation outlined in Section 3.2, we propose a semantic-level interference strategy. This strategy simultaneously forces the features of x_{adv} to diverge from its original reference x_{ref} , while regularizing them toward a “semantic-free” state.

Specifically, consider a network layer l within the frozen U-Net \mathcal{U}_θ and timestep t , we extract the original features $f^{(l,t)} = \mathcal{U}_\theta^l(x_{src}, x_{ref}, t)$ conditioned on x_{ref} , and $f_{adv}^{(l,t)} = \mathcal{U}_\theta^l(x_{src}, x_{adv}, t)$ corresponding to the adversarial version. Furthermore, to define a “semantic-free” target, we leverage the U-Net’s activation modes to unstructured information. We extract features $f_{noise}^{(l,t)} = \mathcal{U}_\theta^l(x_{noise}, x_{noise}, t)$ by feeding Gaussian white noise x_{noise} as both the source content and the reference condition. $f_{noise}^{(l,t)}$ can be considered to represent unstructured features and lack semantic information. The bidirectional objective aims to maximize the distance between $f_{adv}^{(l,t)}$ and $f^{(l,t)}$ while minimizing it with $f_{noise}^{(l,t)}$.

This objective encourages the adversarial features to abandon the original semantic structure and move towards a state of incoherence which can be formalized as:

$$\mathcal{L}_{sem} = 1 - \cos(f_{adv}^{(l,t)}, f^{(l,t)}) + \underbrace{\cos(f_{adv}^{(l,t)}, f_{noise}^{(l,t)})}_{\text{Sem-free}}, \quad (11)$$

where we employ the cosine distance metric $\cos(\cdot)$ as it emphasizes the structural feature similarity. The loss term \mathcal{L}_{sem} is strategically computed within the U-Net’s upsampling path, whose layers synthesize the high-frequency acoustic details essential for perceptual fidelity.

Joint Optimization of Defense Objectives

The final adversarial perturbation δ for VoiceCloak is optimized to simultaneously achieve our dual objectives (Section 4.2 and 4.3). The joint objectives of this comprehensive defense are formalized as follows

$$\begin{aligned} \mathcal{L}_{total} &= (\mathcal{L}_{ID}, \mathcal{L}_{ctx}, \mathcal{L}_{score}, \mathcal{L}_{sem}) \Lambda^\top, \\ \delta &:= \arg \max_{\delta} \mathcal{L}_{total}, \end{aligned} \quad (12)$$

where $\Lambda = (\lambda_{ID}, \lambda_{ctx}, \lambda_{score}, \lambda_{sem})$ controls the weight factors that balance the relative importance of these defenses.

Informed by findings that early diffusion timesteps primarily reconstruct low-frequency structural components (Yu et al. 2024), we target these initial steps (T_{adv}) to efficiently corrupt the fundamental integrity while minimizing computational cost.

Experiments

Experimental Setup

Datasets Experiments are conducted on the LibriTTS (Zen et al. 2019) and VCTK (Yamagishi, Veaux, and MacDonald 2019) datasets. We selected a gender-balanced audio subsets (479 utterances from LibriTTS, 500 from VCTK) to generate adversarial reference speech set $\mathcal{D}_{x_{adv}}$.

Baseline Methods We compare VoiceCloak against existing voice protection methods. We adopt the following methods for fair comparison: Attack-VC (Huang et al. 2021), VoicePrivacy (Chen et al. 2024), and VoiceGuard (Li et al. 2023). We also include a naive baseline: adding random Gaussian noise to x_{ref} .

Evaluation Metrics For identity protection, we report the Automatic Speaker Verification (ASV) acceptance rate, where a lower rate for protected outputs signifies more effective obfuscation. We define a comprehensive Defense Success Rate (DSR) to measure the achievement of both our objectives. A defense is considered successful if the protected output both fails speaker verification and exhibits low perceptual quality ($DS = s_{ASV} < \tau_{ASV} \wedge \text{NISQA}(y_{adv}) < \tau_q$), with thresholds $\tau_{ASV} = 0.25$ and $\tau_q = 3.0$. Additional metrics include Dynamic Time Warping (DTW) (Sakoe and Chiba 1978) and SSIM between y and y_{adv} spectrograms. Perturbation imperceptibility on x_{adv} versus x_{ref} is measured using PESQ (Rix et al. 2001), Mel-Cepstral Distortion (MCD) (Kubichek 1993) and SNR.

Datasets	Methods	Defense Effectiveness					Imperceptibility		
		DTW \uparrow	ASV \downarrow	SSIM \downarrow	NISQA \downarrow	DSR \uparrow	PESQ \uparrow	MCD \downarrow	SNR \uparrow
LibriTTS	Undefended	-	76.49%	-	3.96	-	-	-	-
	Random Noise l_∞	2.01	55.20%	0.31	3.72	16.00%	3.37	<u>1.35</u>	34.80
	Attack-VC	2.29	36.20%	0.31	<u>3.57</u>	30.40%	2.31	<u>3.71</u>	5.29
	VoicePrivacy	<u>2.26</u>	20.80%	0.30	3.60	26.80%	2.99	1.37	33.25
	VoiceGuard	2.08	16.49%	0.29	3.63	43.45%	2.15	4.39	10.58
	Ours	2.12	11.40%	0.27	2.36	71.40%	<u>3.22</u>	1.29	<u>33.53</u>
VCTK	Undefended	-	63.68%	-	3.41	-	-	-	-
	Random Noise l_∞	1.68	58.00%	0.35	3.16	11.38%	3.25	<u>1.38</u>	34.10
	Attack-VC	2.05	38.50%	0.33	2.82	26.20%	2.25	3.80	4.50
	VoicePrivacy	1.88	<u>30.28%</u>	0.35	<u>2.77</u>	<u>39.06%</u>	2.87	1.53	31.87
	VoiceGuard	1.87	31.42%	0.32	3.02	22.11%	2.05	4.57	10.00
	Ours	<u>1.93</u>	19.74%	0.29	2.51	63.41%	<u>3.09</u>	1.33	<u>32.41</u>

Table 1: Comparison of defense effectiveness and adversarial imperceptibility with SOTA methods. Higher values are better for metrics marked with \uparrow , and vice versa for those marked with \downarrow . The best result is marked in **BOLD**.

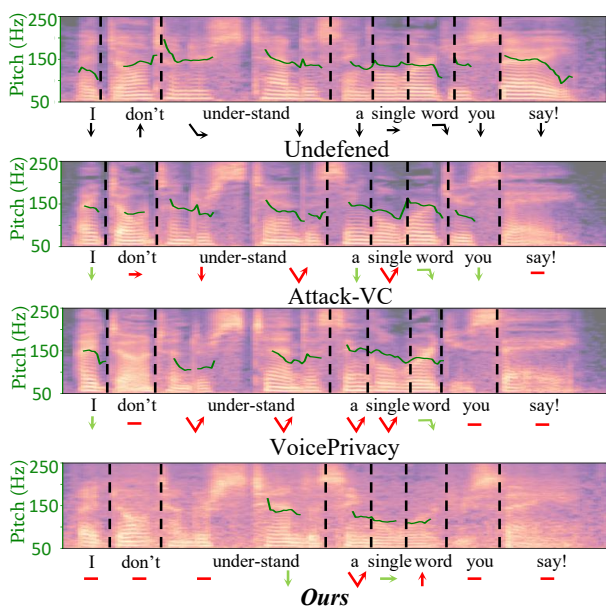


Figure 3: Mel spectrograms with F_0 pitch contours (green lines), and inferred intonation of the corresponding words. Arrows indicate perceived intonation shifts. (Intonation aligns with the ground truth, which is marked by green arrows, and diverges, which is marked by red arrows.)

Implementation Details We conduct our experiments mainly using DiffVC (Popov et al. 2022) as the target system. We set the number of optimization iterations to 50 and a step size $\alpha = 4 \times 10^{-5}$. And δ is constrained within an l_∞ -norm ball of $\epsilon = 0.002$. We set the adversarial and inference timesteps respectively $T_{adv} = 6$ and $T = 100$. The loss function combined identity and quality objectives with weights $\Lambda = (1.0, 4.5, 10, 0.85)$. All experiments are conducted on NVIDIA RTX 3090 GPU with a fixed random seed.

Comparison and Analysis

As shown in Table 1, VoiceCloak demonstrates exceptional defense efficacy, significantly outperforming all baselines




Figure 4: Protecting commercial speaker verification APIs (Iflytek, Azure) from spoofing attacks (lower are better).

with a DSR of over 71% on LibriTTS and 63% on VCTK. This success reflects our defense objectives: identity obfuscation is achieved by reducing speaker verification acceptance rates to 11%, while perceptual quality degradation is confirmed by low NISQA scores. This dual capability distinguishes VoiceCloak from defenses that may struggle with diffusion models. In contrast, the naive Random Noise baseline offers minimal protection. Regarding imperceptibility, VoiceCloak performs comparably to baselines, highlighting the superiority of targeting model vulnerabilities over merely increasing perturbation magnitude.

Figure 3 further visualizes the disruption results. Applying VoiceCloak protection corrupts the F_0 contour. The pitch curve is notably degraded, appearing blurred and unpredictable, accompanied by inconsistent intonation changes.

Protecting Commercial Systems We also evaluated the effectiveness in protecting commercial speaker verification (SV) APIs (Iflytek, Azure) to simulate real-world anti-spoofing scenarios. Successful protection aims to minimize the similarity score returned by the API. Figure 4 demonstrated VoiceCloak’s superior ability to decouple the protected audio from the original speaker’s identity.

User Study A user study involving 50 participants confirmed VoiceCloak’s perceptual efficacy. Listeners performed comparisons on two criteria: Timbre Dissimilarity and Naturalness Disruption. The aggregated results in Fig-




Figure 5: User perceptual study results. (a) Timbre Dissimilarity Preference. (b) Corresponding results for perceived Naturalness Disruption.

Settings		Defense Effectiveness				
\mathcal{L}_{ID}	\mathcal{L}_{ctx}	DTW \uparrow	ASV \downarrow	SSIM \downarrow	NISQA \downarrow	DSR \uparrow
-	-	1.96	46.82%	0.30	3.66	22.58%
✓	-	2.16	8.57%	0.30	3.57	27.74%
w/o Gender	-	2.25	19.92%	0.30	3.60	14.40%
-	✓	2.31	16.20%	0.27	2.96	62.57%
✓	✓	2.13	11.00%	0.27	2.85	69.20%

Table 2: Ablation study on the contribution of different settings for Identity Obfuscation. ✓ indicates the setting is used, "w/o" denotes the exclusion of the specified term.

Settings		Defense Effectiveness				
\mathcal{L}_{score}	\mathcal{L}_{sem}	DTW \uparrow	ASV \downarrow	SSIM \downarrow	NISQA \downarrow	DSR \uparrow
-	-	1.99	45.00%	0.31	3.09	20.20%
✓	-	2.42	31.80%	0.29	2.68	41.20%
-	✓	2.23	23.00%	0.27	2.44	60.60%
-	w/o Sem-free	2.28	26.36%	0.29	3.30	26.80%
✓	✓	2.22	23.60%	0.27	2.10	57.80%

Table 3: Ablation study for Perceptual Fidelity Degradation. Checkmark (✓) indicates the setting is used, "w/o" denotes the exclusion of the specified loss term.

ure 5, where "Neutral" indicates no perceived difference, revealed a consensus: our outputs are rated as more dissimilar in timbre and more unnatural, confirming its human-perceived effectiveness.

Ablation Study

Effectiveness of Adversarial Identity Obfuscation Table 2 presents an ablation study of our Adversarial Identity Obfuscation losses. First, \mathcal{L}_{ID} alone substantially reduces the ASV acceptance rate, demonstrating that perturbing the embedding space is an effective strategy for disrupting identity features. Furthermore, removing the opposite-gender guidance from \mathcal{L}_{ID} degrades performance, validating the efficacy of this directional strategy. The context divergence loss \mathcal{L}_{ctx} , also proves crucial, further lowering both ASV and NISQA by disrupting the attention mechanism's conditioning.

Effectiveness of Perceptual Fidelity Degradation Table 3 ablates the perceptual quality degradation losses. The results show that \mathcal{L}_{score} alone is effective, degrading quality (lower NISQA, higher DTW) by forcing the denoising trajectory away from high-fidelity regions. The semantic corruption loss, \mathcal{L}_{sem} , demonstrates an even stronger individual impact by directly corrupting internal U-Net features, thereby impairing the model's reconstruction of coherent,

Target Models	Defense Effectiveness				
	DTW \uparrow	ASV \downarrow	SSIM \downarrow	NISQA \downarrow	DSR \uparrow
DiffVC	2.12	11.40%	0.27	2.36	71.40%
DDDM-VC	1.67	16.80%	0.36	2.79	54.89%
DuTa-VC	2.41	13.77%	0.27	2.14	73.92%

Table 4: Transferability of the proposed defense method against unseen target models.

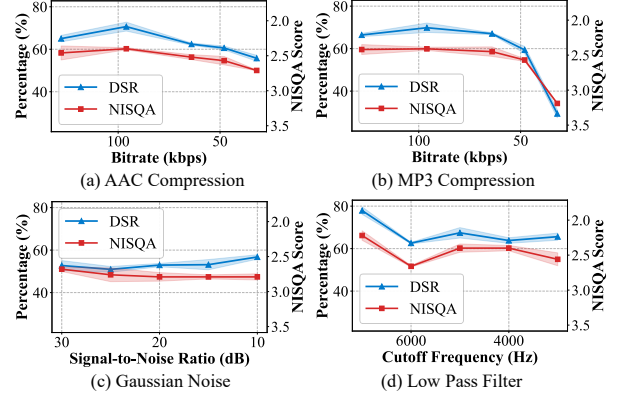


Figure 6: Resilience of VoiceCloak under more advanced robust scenario.

natural-sounding speech details. Critically, the necessity of guiding semantic features towards an incoherent state is evident from the "w/o Sem-free" variant.

Transferability

To evaluate the generalizability of VoiceCloak, we tested its performance on additional diffusion-based VC models: DuTa-VC (Wang et al. 2023) and DDDM-VC (Choi, Lee, and Lee 2024). As reported in Table 4, our method achieving an average DSR of 66.7%. We attribute this cross-model effectiveness to our strategy of targeting fundamental mechanisms, such as attention and score prediction that are shared across diffusion VCs.

Robustness

We evaluated the robustness of our method against four common distortions, with results presented in Figure 6. The results show consistent defense effectiveness, indicating resilience to real-world transformations and significantly outperforming the undefended baseline. Further details and analysis are available in extended version (Hu et al. 2025).

Conclusion

This paper introduced VoiceCloak, a defense against unauthorized diffusion VC. Our framework achieves superior defense by exploiting intrinsic vulnerabilities of DMs. Through strategies that disrupt attention driven conditional guidance, steer the denoising process and corrupt semantic space, VoiceCloak effectively undermines the synthesis process. Extensive experiments validate its efficacy, demonstrating success in simultaneously obfuscating speaker identity and degrading audio quality to mitigate mimicry threats.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No.62441237, No.62261160653), the Guangdong Provincial Key Laboratory of Information Security Technology (No. 2023B1212060026) and Henan Province Innovation Scientists and Technicians Troop Construction Projects of China (No.254000510007).

References

- Chen, S.; Chen, L.; Zhang, J.; Lee, K.; Ling, Z.; and Dai, L. 2024. Adversarial speech for voice privacy protection from personalized speech generation. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 11411–11415. IEEE.
- Chen, S.; Wang, C.; Chen, Z.; Wu, Y.; Liu, S.; Chen, Z.; Li, J.; Kanda, N.; Yoshioka, T.; Xiao, X.; et al. 2022a. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6): 1505–1518.
- Chen, S.; Wu, Y.; Wang, C.; Chen, Z.; Chen, Z.; Liu, S.; Wu, J.; Qian, Y.; Wei, F.; Li, J.; et al. 2022b. Unispeechsat: Universal speech representation learning with speaker aware pre-training. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6152–6156. IEEE.
- Choi, H.-Y.; Lee, S.-H.; and Lee, S.-W. 2024. Dddm-vc: Decoupled denoising diffusion models with disentangled representation and prior mixup for verified robust voice conversion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17862–17870.
- Dong, S.; Chen, B.; Ma, K.; and Zhao, G. 2024. Active Defense Against Voice Conversion through Generative Adversarial Network. *IEEE Signal Processing Letters*, 31: 706–710.
- Guo, Y.; Qu, Z.; Lu, W.; and Luo, X. 2025. Anti-Inpainting: A Proactive Defense Approach against Malicious Diffusion-based Inpainters under Unknown Conditions. *arXiv:2505.13023*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Hu, Q.; Wu, J.; Lu, W.; and Luo, X. 2025. VoiceCloak: A Multi-Dimensional Defense Framework against Unauthorized Diffusion-based Voice Cloning. *arXiv:2505.12332*.
- Huang, C.-y.; Lin, Y. Y.; Lee, H.-y.; and Lee, L.-s. 2021. Defending your voice: Adversarial attack on voice conversion. In *2021 IEEE Spoken Language Technology Workshop (SLT)*, 552–559. IEEE.
- Jeong, M.; Kim, H.; Cheon, S. J.; Choi, B. J.; and Kim, N. S. 2021. Diff-TTS: A Denoising Diffusion Model for Text-to-Speech. In *Proc. Interspeech 2021*, 3605–3609.
- Jung, J.-w.; Heo, H.-S.; Tak, H.; Shim, H.-j.; Chung, J. S.; Lee, B.-J.; Yu, H.-J.; and Evans, N. 2022. Aassist: Audio anti-spoofing using integrated spectro-temporal graph attention networks. In *ICASSP 2022-2022 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 6367–6371. IEEE.
- Kang, M.; Song, D.; and Li, B. 2023. Diffattack: Evasion attacks against diffusion-based adversarial purification. *Advances in Neural Information Processing Systems*, 36: 73919–73942.
- Katharopoulos, A.; Vyas, A.; Pappas, N.; and Fleuret, F. 2020. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, 5156–5165. PMLR.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6007–6017.
- Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2020. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*.
- Kreiman, J.; and Sidtis, D. 2011. *Foundations of voice studies: An interdisciplinary approach to voice production and perception*. John Wiley & Sons.
- Kubichek, R. 1993. Mel-cepstral distance measure for objective speech quality assessment. In *Proceedings of IEEE pacific rim conference on communications computers and signal processing*, volume 1, 125–128. IEEE.
- Kullback, S.; and Leibler, R. A. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1): 79–86.
- Lavner, Y.; Gath, I.; and Rosenhouse, J. 2000. The effects of acoustic modifications on the identification of familiar voices speaking isolated vowels. *Speech Communication*, 30(1): 9–26.
- Li, J.; Ye, D.; Tang, L.; Chen, C.; and Hu, S. 2023. Voice Guard: Protecting Voice Privacy with Strong and Imperceptible Adversarial Perturbation in the Time Domain. In *IJCAI*, 4812–4820.
- Liu, H.; Chen, Z.; Yuan, Y.; Mei, X.; Liu, X.; Mandic, D.; Wang, W.; and Plumbley, M. D. 2023. AudioLDM: Text-to-Audio Generation with Latent Diffusion Models. In *International Conference on Machine Learning*, 21450–21474. PMLR.
- Oh, H.-S.; Lee, S.-H.; and Lee, S.-W. 2024. Diffprosody: Diffusion-based latent prosody generation for expressive speech synthesis with prosody conditional adversarial training. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Popov, V.; Vovk, I.; Gogoryan, V.; Sadekova, T.; and Kudinov, M. 2021. Grad-tts: A diffusion probabilistic model for text-to-speech. In *International Conference on Machine Learning*, 8599–8608. PMLR.
- Popov, V.; Vovk, I.; Gogoryan, V.; Sadekova, T.; Kudinov, M. S.; and Wei, J. 2022. Diffusion-Based Voice Conversion with Fast Maximum Likelihood Sampling Scheme. In *International Conference on Learning Representations*.
- Qu, Z.; Lu, W.; Luo, X.; Wang, Q.; and Cao, X. 2025. ID-Guard: A Universal Framework for Combating Facial Manipulation via Breaking Identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–16.

- Rix, A. W.; Beerends, J. G.; Hollier, M. P.; and Hekstra, A. P. 2001. Perceptual evaluation of speech quality (PESQ)—a new method for speech quality assessment of telephone networks and codecs. In *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*, volume 2, 749–752. IEEE.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*, 234–241. Springer.
- Ruan, L.; Ma, Y.; Yang, H.; He, H.; Liu, B.; Fu, J.; Yuan, N. J.; Jin, Q.; and Guo, B. 2023. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10219–10228.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22500–22510.
- Sakoe, H.; and Chiba, S. 1978. Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing*, 26(1): 43–49.
- Shen, K.; Ju, Z.; Tan, X.; Liu, E.; Leng, Y.; He, L.; Qin, T.; Bian, J.; et al. 2024. NaturalSpeech 2: Latent Diffusion Models are Natural and Zero-Shot Speech and Singing Synthesizers. In *The Twelfth International Conference on Learning Representations*.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Tumanyan, N.; Geyer, M.; Bagon, S.; and Dekel, T. 2023. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1921–1930.
- Wang, F.; Tan, Z.; Wei, T.; Wu, Y.; and Huang, Q. 2024. Simac: A simple anti-customization method for protecting face privacy against text-to-image synthesis of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12047–12056.
- Wang, H.; Thebaud, T.; Villalba, J.; Sydnor, M.; Lammers, B.; Dehak, N.; and Moro-Velazquez, L. 2023. DuTa-VC: A Duration-aware Typical-to-atypical Voice Conversion Approach with Diffusion Probabilistic Model. In *Proc. Interspeech 2023*, 1548–1552.
- Wang, L.; Hu, Q.; Lu, W.; and Luo, X. 2025. Diffusion-based adversarial identity manipulation for facial privacy protection. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 11562–11571.
- Wu, J.; Lu, W.; Luo, X.; Yang, R.; Wang, Q.; and Cao, X. 2024. Coarse-to-fine proposal refinement framework for audio temporal forgery detection and localization. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 7395–7403.
- Wu, J.; Xu, W.; Lu, W.; Luo, X.; Yang, R.; and Guo, S. 2025. Weakly-supervised Audio Temporal Forgery Localization via Progressive Audio-language Co-learning Network. *arXiv preprint arXiv:2505.01880*.
- Yamagishi, J.; Veaux, C.; and MacDonald, K. 2019. CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit (version 0.92).
- Yu, H.; Chen, J.; Ding, X.; Zhang, Y.; Tang, T.; and Ma, H. 2024. Step vulnerability guided mean fluctuation adversarial attack against conditional diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6791–6799.
- Yu, Z.; Zhai, S.; and Zhang, N. 2023. Antifake: Using adversarial audio to prevent unauthorized speech synthesis. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 460–474.
- Zen, H.; Dang, V.; Clark, R.; Zhang, Y.; Weiss, R. J.; Jia, Y.; Chen, Z.; and Wu, Y. 2019. Libritts: A corpus derived from librispeech for text-to-speech. *arXiv preprint arXiv:1904.02882*.
- Zhou, Y.; and Lim, S.-N. 2021. Joint audio-visual deepfake detection. In *Proceedings of the IEEE/CVF International Conference on Computer Cision*, 14800–14809.