

# Can Emotion Fool Anti-spoofing?

Aurosweta Mahapatra<sup>1</sup>, Ismail Rasim Ulgen<sup>1</sup>, Abinay Reddy Naini<sup>2</sup>, Carlos Busso<sup>2</sup>,  
Berrak Sisman<sup>1,3</sup>

<sup>1</sup>Center for Language and Speech Processing, Johns Hopkins University, USA

<sup>2</sup>Language Technologies Institute, Carnegie Mellon University, USA

<sup>3</sup>Data Science and AI Institute (DSAI), Johns Hopkins University, USA

amahapa2@jhu.edu, iulgen1@jhu.edu, anaini@andrew.cmu.edu, busso@cmu.edu, sisman@jhu.edu

## Abstract

Traditional anti-spoofing focuses on models and datasets built on synthetic speech with mostly neutral state, neglecting diverse emotional variations. As a result, their robustness against high-quality, emotionally expressive synthetic speech is uncertain. We address this by introducing EmoSpoof-TTS, a corpus of emotional text-to-speech samples. Our analysis shows existing anti-spoofing models struggle with emotional synthetic speech, exposing risks of emotion-targeted attacks. Even trained on emotional data, the models underperform due to limited focus on emotional aspect and show performance disparities across emotions. This highlights the need for emotion-focused anti-spoofing paradigm in both dataset and methodology. We propose GEM, a gated ensemble of emotion-specialized models with a speech emotion recognition gating network. GEM performs effectively across all emotions and neutral state, improving defenses against spoofing attacks. We release the EmoSpoof-TTS Dataset<sup>1</sup>.

**Index Terms:** Anti-Spoofing, Speech Deepfake, Text-to-Speech, Ensemble Modeling, Emotion

## 1. Introduction

Anti-spoofing in the speech domain focuses on detecting spoofed speech produced through replay techniques, speech synthesis, and voice conversion [1]. Among these, the misuse of synthetic speech [1, 2] has become a growing concern, especially with advancements in emotionally expressive text-to-speech (TTS) models.

Recent TTS models can effectively capture a range of emotions, making synthetic speech increasingly realistic [3–8]. Furthermore, zero-shot TTS models [9–11] have demonstrated the ability to produce realistic synthetic speech from just a few seconds of reference audio while effectively capturing human emotions. While state-of-the-art (SOTA) anti-spoofing models [12, 13] perform relatively well in detecting neutral synthetic speech, their effectiveness against these emotional TTS models remains uncertain. This necessitates robust countermeasures against emotion-targeted attacks, where an attacker exploits the model’s sensitivity to a particular emotion.

Various challenges, such as the Automatic Speaker Verification and Spoofing Countermeasures Challenge (ASVspoof) [14] and the Audio Deep Synthesis Detection (ADD) Challenge [15, 16], have been introduced to raise awareness about anti-spoofing. These initiatives have led to the development of large-scale datasets, including ASVspoof 2019 [17], ASVspoof 2021 [18], and ASVspoof 2024 [19], which serve as benchmarks in this field. However, these datasets primarily contain

speech in a neutral state and do not specifically address the emotional aspects of speech, leaving a gap in the study of emotion-driven spoofing attacks.

The ASVspoof and ADD challenges, along with large-scale datasets, have significantly advanced the development of anti-spoofing models. Traditional models [20–22] rely on hand-crafted features for spoof speech detection. However, more recent end-to-end models like RawNet2 [12] and AASIST [13] process raw audio waveforms directly, achieving highly promising results by eliminating the need for manually engineered features. However, even these SOTA models have primarily been trained and evaluated on traditional datasets [17–19], which focus on neutral speech and lack diverse emotional variations. As a result, current anti-spoofing models are not explicitly designed to account for the impact of emotion on spoof detection, making them vulnerable to emotionally expressive synthetic speech.

In this work, we focus on the neglected emotional aspect of anti-spoofing by exploring whether the emotions can fool anti-spoofing systems. We reveal the limitations of current anti-spoofing paradigm in terms of lacking both suitable datasets and emotion tailored methods. To address this limitations, we build an emotionally rich anti-spoofing dataset and propose an anti-spoofing method specifically designed to tackle emotional speech to mitigate the risks associated with the emotion-targeted attacks.

We introduce EmoSpoof-TTS, a diverse dataset containing over 29 hours of emotionally expressive synthetic speech generated using recent TTS models. Our findings reveal that current anti-spoofing models are unreliable in scenarios involving emotional speech and exhibit performance disparities across different emotions. In addition, to improve the performance on emotionally expressive synthetic speech, we propose the Gated Ensemble Method (GEM). GEM leverages multiple emotion-specialized anti-spoofing models with a Speech Emotion Recognition (SER) model, which acts as a gating mechanism to enhance its performance in detecting high-quality, emotionally expressive speech while also addressing performance disparities across emotions. Our key contributions can be summarized as:

- We identify and highlight the significant limitations of existing anti-spoofing models on emotionally expressive synthetic speech and expose a new risk of emotion-targeted attacks.
- We introduce and release EmoSpoof-TTS, a corpus of emotionally expressive synthetic speech generated using recent TTS models.
- We propose the Gated Ensemble Method (GEM), which significantly improves emotional synthetic speech detection, highlighting the need to prioritize emotion in anti-spoofing design.

<sup>1</sup>EmoSpoof-TTS Dataset: <https://emospoof-tts.github.io/Dataset/>

## 2. Related Work

State-of-the-art anti-spoofing models, such as RawNet2 [12] and AASIST [13], are well known for their strong benchmark performance. These models employ an end-to-end architecture, eliminating the need for hand-crafted feature extraction and simplifying implementation. In this study, we select RawNet2 as our baseline anti-spoofing model due to its robust performance. As illustrated in Figure 1, RawNet2 comprises a Sinc-Net layer, residual blocks, a GRU layer, fully connected layers, and an output layer. Most prior studies utilize the publicly available pre-trained RawNet2 model, trained on the ASVspoof 2019 dataset [17]. We note that as TTS continues to advance, the ASVspoof datasets—despite their scale and popularity—lack emotional speech samples. This limitation raises concerns about whether high-performing anti-spoofing models remain effective against increasingly emotional synthetic speech.

Recently, EmoFake [23] was introduced as an emotional fake dataset generated via Emotional Voice Conversion (EVC) techniques. While valuable, EmoFake focuses on voice conversion and does not encompass emotional text-to-speech (TTS) synthesis. To bridge this gap, we introduce a new dataset, EmoSpoof-TTS, and utilize it to assess and enhance anti-spoofing robustness against advanced synthesis techniques and emotionally expressive speech. Further details are provided in Section 3.

## 3. Anti-spoofing vs Emotion

The state-of-the-art anti-spoofing models are primarily evaluated on traditional datasets, which lack emotionally expressive speech and do not account for recent advancements in speech synthesis. As a result, they are unreliable with emotional speech, making them vulnerable to emotion-targeted attacks. This section introduces EmoSpoof-TTS and compares anti-spoofing model performance on traditional datasets versus EmoSpoof-TTS.

### 3.1. EmoSpoof-TTS

We propose EmoSpoof-TTS to explore the impact of emotions in anti-spoofing. This dataset focuses on spoofed speech generated by zero-shot TTS models—StyleTTS2 [9], F5-TTS [10], and CosyVoice [11]—known for producing high-quality, realistic speech even with a short reference audio. These TTS models require emotional reference speech in order to synthesize emotional speech. In this work, we have utilized The Emotional Speech Database (ESD) [24] for this purpose. ESD dataset consists of 10 speakers expressing five emotions (Happiness, Anger, Sadness, Neutral, etc) with 350 utterances per emotion. For speech synthesis, we separate 50 utterances per speaker and emotion as reference speech. Using these reference speech, we synthesize 300 emotional speech samples which have the same content, speaker identity and emotion with the 300 bona-fide speech samples in ESD per speaker and emotion category.

We consider four basic emotions: *Happiness*, *Anger*, *Sadness*, and *Neutral state*. EmoSpoof-TTS contains a total of 36,000 synthesized speech samples from 4 emotions, 10 (5 male, 5 female) speakers and 3 TTS models [9–11] which is parallel to 12,000 bona-fide samples from ESD. EmoSpoof-TTS featuring multiple emotions, speakers, and TTS models, provides a valuable resource for studying emotional speech in anti-spoofing. We will publicly release the dataset to foster research in this area.

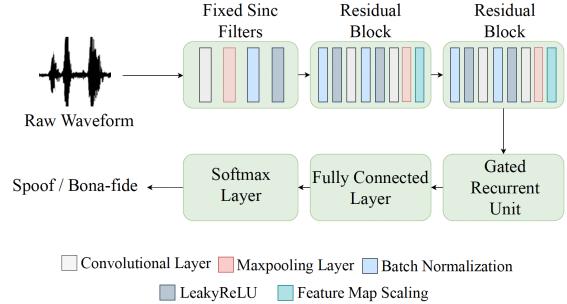


Figure 1: Architecture of RawNet2

Table 1: Performance of Pre-Trained RawNet2 model on test-set of Traditional Datasets (EER%)

Model	Test Dataset		
	ASVspoof 2019	ASVspoof 2021	ASVspoof 2024
RawNet2	4.60	14.73	40.67

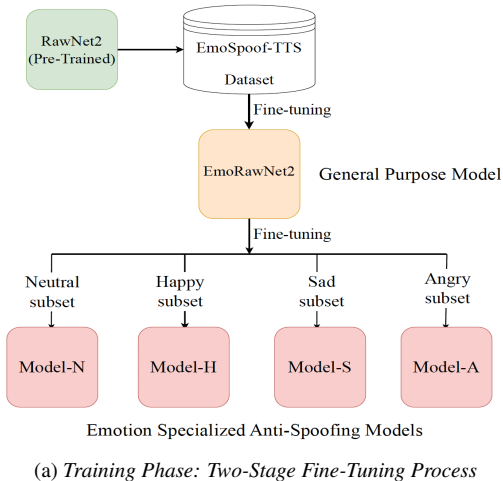
### 3.2. Analysis

We evaluate the pre-trained RawNet2 model, trained on the Logical Access (LA) partition of the ASVspoof 2019 dataset, as the traditional anti-spoofing method.

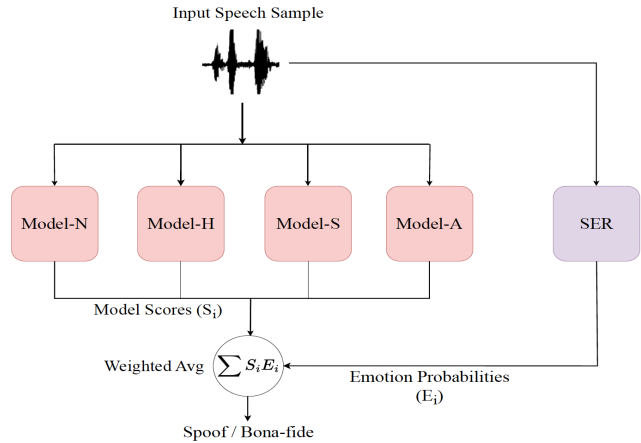
We first assess its performance using evaluation datasets from LA track for ASVspoof 2019, ASVspoof 2021, and the Track 1 for ASVspoof 2024. As shown in Table 1, RawNet2 performs well on ASVspoof 2019 and ASVspoof 2021, with low Equal Error Rate (EER) values. However, its performance drops significantly on ASVspoof 2024, a more recent dataset, highlighting its limitations in generalizing to evolving spoofing techniques.

We then evaluated the pre-trained RawNet2 model on EmoSpoof-TTS, using spoof samples from four speakers while maintaining speaker consistency across all three TTS models. The results, shown in Table 2, are reported for individual emotions and a combined category (Happiness, Anger, and Sadness, referred to as HAS). The HAS category allows us to assess the model’s performance on emotionally expressive speech compared to neutral speech, while individual emotions highlight detection performance across different emotional states.

The evaluation presented in Table 2 provides insights into the vulnerabilities of the pre-trained RawNet2 model when exposed to high-quality, emotionally expressive synthetic speech. The model struggles to detect fake samples from EmoSpoof-TTS, with EER values higher than those for the ASVspoof 2024 dataset across all TTS models. The EER for HAS across all three models is consistently higher than for Neutral. For instance, the EER for HAS in the case of StyleTTS2 is 44.03%, whereas for Neutral, it is 39%. This indicates lower confidence of the anti-spoofing model in distinguishing between spoofed and bona-fide speech when it is exposed to fake emotional speech. These results reinforce our claim that SOTA anti-spoofing models, despite performing well on traditional datasets like ASVspoof 2019 and ASVspoof 2021, struggle with speech samples from recent emotionally expressive TTS models, especially when handling emotional speech, making them unreliable in such scenarios. Additionally, our analysis reveals that EER values vary across emotion categories; for instance StyleTTS2, the EER for Happiness is 45.00%, for Sadness is 46.75%, and for Anger is 35.17%.



(a) Training Phase: Two-Stage Fine-Tuning Process



(b) Inference Phase: Computing the Weighted Average Prediction Score

Figure 2: Proposed Gated Ensemble Method (GEM)

Table 2: Performance of Pre-Trained RawNet2 model on EmoSpoof-TTS (EER%)

TTS Model	EmoSpoof-TTS					
	HAS	Neutral	Happy	Angry	Sad	Overall
StyleTTS2	44.03	39.00	45.00	35.17	46.75	43.06
F5-TTS	42.97	38.08	38.42	39.42	47.25	41.92
CosyVoice	42.44	42.17	39.58	39.75	47.50	42.65

### 3.3. Finetuning RawNet2 on EmoSpoof-TTS

Given the fact that, the pre-trained anti-spoofing method struggles to generalize to other datasets, we explored the performance when the model is trained with EmoSpoof-TTS. For this purpose, we fine-tuned the pretrained RawNet2 on the EmoSpoof-TTS using the train, validation and test subsets, which are described in Section 5.1. We structured the TTS models and speaker splits to ensure the training, validation, and test sets have no overlap, ensuring a fair evaluation. Results in Table 3 indicate a significant improvement in Emo-RawNet2’s performance over the pre-trained RawNet2 model. For instance, with StyleTTS2, the EER for HAS in the pre-trained RawNet2 is 44.03%, which drops to 10.36% for Emo-RawNet2. However, performance disparities across emotions persist, as reflected in the EER values: 8.00% for Happiness, 13.25% for Sadness, and 5.83% for the Neutral state.

This uneven performance across different emotions suggests that attackers could exploit these inconsistencies by targeting emotions to which the model is more sensitive, increasing the likelihood of successful attacks. Moreover, the results in Table 3 highlights that while fine-tuning an existing model improves overall performance, it fails to resolve the disparity across emotions. Therefore, defending against the vulnerability of anti-spoofing models to emotionally expressive speech does not only require a useful dataset but also anti-spoofing paradigms specifically tailored for detecting emotional speech.

## 4. Proposed Method

As discussed in Section 3, the performance of traditional anti-spoofing on emotional synthetic speech degrades significantly. Additionally, the uneven performance of anti-spoofing models across different emotions presents a security risk, as attackers

Table 3: Performance of Emo-RawNet2 on the EmoSpoof-TTS test set (EER%)

Model	EmoSpoof-TTS (test set)					
	HAS	Neutral	Happy	Angry	Sad	Overall
Emo-RawNet2	10.36	5.83	8.00	9.25	13.67	9.65

may exploit the model’s vulnerability to specific emotional categories. To mitigate these vulnerabilities and enhance robustness against emotion-targeted attacks, this section introduces a novel approach: the Gated Ensemble Method (GEM).

### 4.1. Emotion-specialized Anti-Spoofing models

To build emotion-specialized anti-spoofing models, we use Emo-RawNet2, our general-purpose model described in section 3.2. We fine-tune Emo-RawNet2 on emotion-specific subsets of the EmoSpoof-TTS dataset, resulting in four emotion-specialized models: Model-H (Happiness), Model-A (Anger), Model-S (Sadness), and Model-N (Neutral state).

### 4.2. Gated Ensemble Method (GEM)

During inference, we lack prior knowledge of the emotion in the input speech, so we must automatically select the corresponding specialized anti-spoofing model. We chose not to use a hard selection of models, as we believe that emotions are inherently interconnected, rather than orthogonal [25]. To account for the relationships between different emotions, we propose incorporating contributions from multiple emotion-specialized models. This approach ensures a more holistic analysis, improving the model’s ability to detect whether a sample is fake. Therefore, we introduce an ensemble of specialized models with a soft gating mechanism to achieve the optimal solution.

The core idea of GEM is to ensemble predictions from specialized models in a way that gives the highest weight to the model associated with the emotion of the input signal  $x$  during final decision-making. Soft gating enables this by allowing the most relevant model to make the final judgment while still incorporating the contributions from the other models.

For a given input  $x$ , each specialized model  $M_i$  produces a

score  $S_i$ :

$$S_i = M_i(x), \quad (1)$$

where  $S_i$  is the model’s spoof prediction score for  $x$ .

We employ a recent the SER system [26] to determine the emotion probabilities of  $x$ , as illustrated in Figure 1(b). The input  $x$  is fed to SER, which outputs the probabilities ( $E_i$ ) for each emotion category:

$$E_i = SER(x). \quad (2)$$

The final decision score  $y_s$  is computed as a weighted average of model scores, where weights are the emotion probabilities:

$$y_s = \sum_i S_i \cdot E_i. \quad (3)$$

We control the softness of this weights by applying softmax with temperature  $T$  to the output logits of the SER model. Based on  $y_s$ , the system determines whether the input signal is spoofed ( $y = 0$ ) or authentic ( $y = 1$ ).

## 5. Experiments and Results

### 5.1. Dataset

For the experiments in this paper, we partitioned the dataset into separate training, validation, and test sets based on speakers and TTS models to ensure a fair and unbiased evaluation:

- **EmoSpoof-TTS Train set:** Includes 4 emotion categories and 4 speakers, with data generated by *CosyVoice*, totaling 9,600 utterances (4,800 bona-fide and 4,800 spoofed).
- **EmoSpoof-TTS Validation set:** Includes 4 emotion categories and 2 speakers, with data generated by *F5-TTS*, totaling 4,800 utterances (2,800 bona-fide and 2,800 spoofed).
- **EmoSpoof-TTS Test set:** Includes 4 emotion categories and 4 speakers, with data generated by *StyleTTS2*, totaling 9,600 utterances (4,800 bona-fide and 4,800 spoofed).

To obtain Emo-RawNet2, we fine-tune the pre-trained RawNet2 on the training and validation sets of EmoSpoof-TTS. For the emotion-specialized models, we extracted emotion-specific subsets from these sets and fine-tuned Emo-RawNet2 separately for each emotion.

### 5.2. Implementation Details

We utilize the pre-trained RawNet2 with default parameters [12]. During the fine-tuning process to obtain the general-purpose model Emo-RawNet2, we used a batch size of 32 and trained the model for 50 epochs. For emotion-specific fine-tuning, the batch size was reduced to 8, and training was extended to 100 epochs. Learning rate was set to  $1e^{-4}$  in both the fine-tuning process. SER [26] is trained to classify 4 emotion categories (Neutral, Happiness, Anger, Sadness) on the MSP-Podcast dataset [27], using the default parameters and recipe provided. To ensure soft gating in the GEM, we set the temperature of the SER to 1.5.

### 5.3. Results

As a baseline, we consider the general-purpose model, Emo-RawNet2, which was fine-tuned with the full set of emotions from the EmoSpoof-TTS dataset, as described in Section 3.2. This allows us to analyze the effectiveness of our proposed

Table 4: Performance Comparison on EmoSpoof-TTS test set (EER%)

Model	EmoSpoof-TTS (test set)					
	HAS	Neutral	Happy	Angry	Sad	Overall
Model-N	7.89	<b>3.50</b>	7.92	7.67	7.08	6.79
Model-H	6.03	2.08	<b>6.67</b>	4.00	5.75	5.02
Model-A	11.36	9.58	13.25	<b>7.08</b>	12.83	11.02
Model-S	7.89	3.58	9.00	5.08	<b>7.67</b>	6.92
<b>GEM</b>	<b>6.22</b>	2.92	6.75	5.58	5.75	5.33

method, which ensembles specialized models using a soft gating mechanism. Results in Table 4 presents the performance of our emotion-specialized models Model-N, Model-H, Model-A, Model-S. Each specialized model not only enhances detection for its respective target emotions but also demonstrates improvements in detecting other emotions. These results indicate the non-orthogonality [25] of emotions, as improving the model’s performance on a particular emotion affects others.

The results presented in Table 4 demonstrate that GEM not only enhances the overall performance of the anti-spoofing model compared to Emo-RawNet2 but also mitigates the imbalance in spoof detection across different emotions. Its performance is notably more consistent across emotions compared to Emo-RawNet2 and individual emotion-specialized models. The results shows that proposed method (GEM) comes very close to the optimal performance of each specialized model for its respective emotion category leading to a much improved overall performance and more balanced performance across different emotions. Another key strength of GEM is that it enhances spoof detection in emotional speech without compromising performance on neutral speech. The above results demonstrate that handling spoof detection for all emotions together is a complex task for a single anti-spoofing model, resulting in suboptimal performance. In contrast, using multiple models specialized for each emotion not only improves emotion understanding but also captures inter-emotional relationships, leading to better overall anti-spoofing performance.

## 6. Conclusion

In this work we explored the effect of emotion on anti-spoofing, revealing emotion-related risks and the need of emotional aspect for this task in terms of both dataset and methodology. Our analysis showed current anti-spoofing paradigm with datasets and methods without emotional focus struggle with emotion-targeted attacks. To facilitate research in this domain, firstly we introduced EmoSpoof-TTS, a corpus created using recent text-to-speech (TTS) models to benchmark and enhance anti-spoofing performance on emotional synthetic speech. Furthermore, we proposed an anti-spoofing method specifically tailored for tackling emotional aspect of anti-spoofing. The proposed method, the Gated Ensemble Method (GEM), which leverages emotion-specialized models and a soft gating mechanism, improved anti-spoofing performance on emotional speech, validating the need for such methods. For future work, we aim to enhance anti-spoofing resilience to emotion-targeted attacks across diverse emotional states and TTS models while expanding EmoSpoof-TTS.

## 7. Acknowledgment

This work is supported by NSF CAREER award IIS-2338979.

## 8. References

- [1] A. Khan, K. M. Malik, J. Ryan, and M. Saravanan, "Battling voice spoofing: a review, comparative analysis, and generalizability evaluation of state-of-the-art voice spoofing countermeasures," *Artificial Intelligence Review*, vol. 56, no. Suppl 1, pp. 513–566, 2023.
- [2] M. Masood, M. Nawaz, K. M. Malik, A. Javed, A. Irtaza, and H. Malik, "Deepfakes generation and detection: State-of-the-art, open challenges, countermeasures, and way forward," *Applied intelligence*, vol. 53, no. 4, pp. 3974–4026, 2023.
- [3] Y. Guo, C. Du, X. Chen, and K. Yu, "Emodiff: Intensity controllable emotional text-to-speech with soft-label guidance," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [4] H. Tang, X. Zhang, N. Cheng, J. Xiao, and J. Wang, "Ed-tts: Multi-scale emotion modeling using cross-domain emotion diarization for emotional speech synthesis," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 12 146–12 150.
- [5] R. Shimizu, R. Yamamoto, M. Kawamura, Y. Shirahata, H. Doi, T. Komatsu, and K. Tachibana, "Prompttts++: Controlling speaker identity in prompt-based text-to-speech using natural language descriptions," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 12 672–12 676.
- [6] S.-H. Lee, S.-B. Kim, J.-H. Lee, E. Song, M.-J. Hwang, and S.-W. Lee, "Hierspeech: Bridging the gap between text and speech by hierarchical variational inference using self-supervised representations for speech synthesis," in *Advances in Neural Information Processing Systems*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022. [Online]. Available: <https://openreview.net/forum?id=awdyRVnfQKX>
- [7] M. Le, A. Vyas, B. Shi, B. Karrer, L. Sari, R. Moritz, M. Williamson, V. Manohar, Y. Adi, J. Mahadeokar, and W.-N. Hsu, "Voicebox: Text-guided multilingual universal speech generation at scale," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Online]. Available: <https://openreview.net/forum?id=gzCS252hCO>
- [8] S. Chen, C. Wang, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, "Neural codec language models are zero-shot text to speech synthesizers," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 33, pp. 705–718, 2025.
- [9] Y. A. Li, C. Han, V. Raghavan, G. Mischler, and N. Mesgarani, "Styletts 2: Towards human-level text-to-speech through style diffusion and adversarial training with large speech language models," in *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., vol. 36. Curran Associates, Inc., 2023, pp. 19 594–19 621. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/3eaad2a0b62b5ed7a2e66c2188bb1449-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/3eaad2a0b62b5ed7a2e66c2188bb1449-Paper-Conference.pdf)
- [10] Y. Chen, Z. Niu, Z. Ma, K. Deng, C. Wang, J. Zhao, K. Yu, and X. Chen, "F5-tts: A fairytaler that fakes fluent and faithful speech with flow matching," *arXiv preprint arXiv:2410.06885*, 2024.
- [11] Z. Du, Q. Chen, S. Zhang, K. Hu, H. Lu, Y. Yang, H. Hu, S. Zheng, Y. Gu, Z. Ma *et al.*, "Cosyvoice: A scalable multilingual zero-shot text-to-speech synthesizer based on supervised semantic tokens," *arXiv preprint arXiv:2407.05407*, 2024.
- [12] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. Evans, and A. Larcher, "End-to-end anti-spoofing with rawnet2," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6369–6373.
- [13] J.-w. Jung, H.-S. Heo, H. Tak, H.-j. Shim, J. S. Chung, B.-J. Lee, H.-J. Yu, and N. Evans, "Aasist: Audio anti-spoofing using integrated spectro-temporal graph attention networks," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6367–6371.
- [14] Z. Wu, T. Kinnunen, N. Evans, and J. Yamagishi, "Asvspoof 2015: Automatic speaker verification spoofing and countermeasures challenge evaluation plan," *Training*, vol. 10, no. 15, p. 3750, 2014.
- [15] J. Yi, R. Fu, J. Tao, S. Nie, H. Ma, C. Wang, T. Wang, Z. Tian, Y. Bai, C. Fan, S. Liang, S. Wang, S. Zhang, X. Yan, L. Xu, Z. Wen, and H. Li, "Add 2022: the first audio deep synthesis detection challenge," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 9216–9220.
- [16] J. Yi, J. Tao, R. Fu, X. Yan, C. Wang, T. Wang, C. Y. Zhang, X. Zhang, Y. Zhao, Y. Ren *et al.*, "Add 2023: the second audio deepfake detection challenge," *arXiv preprint arXiv:2305.13774*, 2023.
- [17] X. Wang, J. Yamagishi, M. Todisco, H. Delgado, A. Nautsch, N. Evans, M. Sahidullah, V. Vestman, T. Kinnunen, K. A. Lee, L. Juvela, P. Alku, Y.-H. Peng, H.-T. Hwang, Y. Tsao, H.-M. Wang, S. L. Maguer, M. Becker, F. Henderson, R. Clark, Y. Zhang, Q. Wang, Y. Jia, K. Onuma, K. Mushika, T. Kaneda, Y. Jiang, L.-J. Liu, Y.-C. Wu, W.-C. Huang, T. Toda, K. Tanaka, H. Kameoka, I. Steiner, D. Matrouf, J.-F. Bonastre, A. Govender, S. Ronanki, J.-X. Zhang, and Z.-H. Ling, "Asvspoof 2019: A large-scale public database of synthesized, converted and replayed speech," *Computer Speech Language*, vol. 64, p. 101114, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0885230820300474>
- [18] X. Liu, X. Wang, M. Sahidullah, J. Patino, H. Delgado, T. Kinnunen, M. Todisco, J. Yamagishi, N. Evans, A. Nautsch, and K. A. Lee, "Asvspoof 2021: Towards spoofed and deepfake speech detection in the wild," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 2507–2522, 2023.
- [19] X. Wang, H. Delgado, H. Tak, J.-w. Jung, H.-j. Shim, M. Todisco, I. Kukanov, X. Liu, M. Sahidullah, T. Kinnunen *et al.*, "Asvspoof 5: Crowdsourced speech data, deepfakes, and adversarial attacks at scale," *arXiv preprint arXiv:2408.08739*, 2024.
- [20] Z. Wu, R. K. Das, J. Yang, and H. Li, "Light convolutional neural network with feature genuinization for detection of synthetic speech attacks," *arXiv preprint arXiv:2009.09637*, 2020.
- [21] C.-I. Lai, N. Chen, J. Villalba, and N. Dehak, "Assert: Anti-spoofing with squeeze-excitation and residual networks," *arXiv preprint arXiv:1904.01120*, 2019.
- [22] K. Z. Mon, K. Galajit, C. O. Mawalim, J. Karnjana, T. Isshiki, and P. Aimmanee, "Spoof detection using voice contribution on lfcc features and resnet-34," in *2023 18th International Joint Symposium on Artificial Intelligence and Natural Language Processing (iSAI-NLP)*. IEEE, 2023, pp. 1–6.
- [23] Y. Zhao, J. Yi, J. Tao, C. Wang, and Y. Dong, "Emofake: An initial dataset for emotion fake audio detection," in *China National Conference on Chinese Computational Linguistics*. Springer, 2024, pp. 419–433.
- [24] K. Zhou, B. Sisman, R. Liu, and H. Li, "Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 920–924.
- [25] H.-C. Chou, C.-C. Lee, and C. Busso, "Exploiting co-occurrence frequency of emotions in perceptual evaluations to train a speech emotion classifier," *Interspeech 2022*, 2022.
- [26] L. Goncalves, A. N. Salman, A. R. Naini, L. M. Velazquez, T. Thebaud, L. P. Garcia, N. Dehak, B. Sisman, and C. Busso, "Odyssey 2024-speech emotion recognition challenge: Dataset, baseline framework, and results," *Development*, vol. 10, no. 9,290, pp. 4–54, 2024.
- [27] R. Lotfian and C. Busso, "Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings," *IEEE Transactions on Affective Computing*, vol. 10, no. 4, pp. 471–483, 2019.