

# RawTFNet: A Lightweight CNN Architecture for Speech Anti-spoofing

Yang Xiao\*, Ting Dang\*, Rohan Kumar Das†

\*The University of Melbourne, Australia

†Fortemedia Singapore, Singapore

Email: yxiao9550@student.unimelb.edu.au

**Abstract**—Automatic speaker verification (ASV) systems are often affected by spoofing attacks. Recent transformer-based models have improved anti-spoofing performance by learning strong feature representations. However, these models usually need high computing power. To address this, we introduce RawTFNet, a lightweight CNN model designed for audio signals. The RawTFNet separates feature processing along time and frequency dimensions, which helps to capture the fine-grained details of synthetic speech. We tested RawTFNet on the ASVspoof 2021 LA and DF evaluation datasets. The results show that RawTFNet reaches comparable performance to that of the state-of-the-art models, while also using fewer computing resources. The code and models will be made publicly available.

## I. INTRODUCTION

In recent years, deep learning technologies have greatly improved speech generation methods such as text-to-speech (TTS) [1] and voice conversion (VC) [2]. These technologies bring many advantages, especially for speech applications. However, they also create risks when used to generate fake speech that imitates a target speaker. Modern TTS and VC systems can now produce speech that sounds highly natural, which makes it easier to attack voice-based systems. As a result, automatic speaker verification (ASV) systems face growing threats from such synthetic speech [3]–[6]. To reduce these risks, researchers have focused on anti-spoofing methods that can detect fake voices. The ASVspoof challenge [7] has become an important platform in this field. It provides standard datasets, rules, and evaluation methods to help the research community develop and compare anti-spoofing systems.

Deep learning-based speech anti-spoofing methods can be divided into two main types: feature-based and end-to-end approaches. Feature-based methods use pre-designed representations such as linear-frequency cepstral coefficients (LFCC) [8], cochlear filter cepstral coefficient and instantaneous frequency [9] and constant-Q transform (CQT) based representations [10], [11]. These features help to capture the unique patterns of fake speech for detection.

On the other hand, end-to-end approaches work directly with raw audio signals [12]–[14]. They do not depend on handcrafted features and have shown strong generalization and performance. As a result, many recent studies focus on end-to-end models. For example, RawNet2 [15] uses time-domain convolution to learn features from raw waveforms. Attention-based models, like AASIST [16], apply graph attention networks to better capture both local and global speech

features across time and frequency. Similarly, Rawformer [17] combines convolutional layers with transformers, using self-attention to detect important clues for fake speech. Although these models perform well, they often need large computing resources. This makes them hard to use in real-world systems, especially where low power and fast processing are required. Therefore, it is important to find a balance between accuracy and efficiency for practical deployment.

To improve efficiency in speech processing tasks, researchers have proposed several lightweight convolutional neural network (CNN) architectures. For example, BC-ResNet [18] applies depthwise separable convolutions to reduce both model size and computational cost, while maintaining strong performance. Similarly, the TDNN [19] captures temporal structures in speech by applying convolution over time steps. These models are widely used in tasks such as ASV due to their good balance between speed and accuracy. However, although these architectures are efficient, they are not specifically tailored for speech anti-spoofing. Most of them concentrate on temporal features or use global representations, without focusing on the unique characteristics of synthetic speech. In practice, fake speech signals often contain subtle but important artifacts that are distributed across both time and frequency dimensions [20]. Standard CNNs, which process the input uniformly or emphasize only one dimension, may overlook these subtle patterns.

To handle the challenges of speech anti-spoofing, a task-specific model that can separate time and frequency information is needed. Inspired by time-frequency attention [21], we propose Time-Frequency Convolutions (TF-Convs), a 1D convolution-based module. It processes feature maps in two branches: one for time and another for frequency. This allows the model to capture detailed patterns in both dimensions while keeping computation low. We further propose a new anti-spoofing architecture (RawTFNet) that combines both efficiency and task-specific design. First, we use a Sinusoidal Convolution [22] layer and some SE-Res2Net blocks as the frontend, which effectively captures local patterns and enhances feature diversity. Then, the resulting feature map is processed by the TF-Convs module. Finally, the combined features are passed into the classifier. This architecture is designed to detect the fine-grained artifacts while keeping the model efficient. We tested RawTFNet on the ASVspoof 2021 LA and DF datasets. The results show that RawTFNet works directly

on raw features and reaches comparable accuracy to state-of-the-art (SOTA) models, while using only 70K parameters.

## II. PROPOSED RAWTFNET

To address the limitations of existing lightweight speech anti-spoofing models that do not explicitly separate time and frequency information, we propose a task-specific architecture for anti-spoofing called RawTFNet. Inspired by the time-frequency nature of speech signals, RawTFNet is designed to better capture fine-grained artifacts using dimension-aware processing. Fig. 1 (a) illustrates the overall structure of the proposed RawTFNet. In particular, we introduce a novel TF-Convs module that separates the analysis of time and frequency using 1D depthwise convolutions. Additionally, to improve the representation of local and contextual information, we integrate DWS SE-Res2Blocks that enhance feature diversity and channel attention with low computation.

### A. Frontend

The frontend in RawTFNet is designed to extract low-level spectro-temporal features directly from raw audio. We adopt a modified version of RawNet2 [15] to perform this task. First, the raw waveform is passed through a set of trainable Sinc filters [22] (with 70 filters in total), which act as learnable band-pass filters. This converts the signal into a two-dimensional feature map of frequency and time bins.

Next, the extracted features are refined through a series of convolutional layers. The resultant output passes through a ResNet block with 32 filters to capture early patterns. Then, it is processed by three SE-Res2Net blocks with 64 filters each, which enhance feature diversity using multiple receptive fields. To further improve efficiency, we replace standard 2D convolutions with DWS convolutions. This leads to the use of DWS SE-Res2Net blocks, as shown in Fig. 1 (b), which reduce computation while preserving performance. Each block also includes a squeeze-and-excitation (SE) module [23] to strengthen important channel-wise information. As a result, the encoder generates high-level feature maps  $x_{\text{HFM}} \in \mathbb{R}^{C \times F \times T}$ , where  $C$ ,  $F$ , and  $T$  represent the number of channels, frequency bins, and time steps, respectively.

### B. Time-Frequency Convolutions

To better detect the subtle patterns of synthetic speech from the high-level feature maps, we introduce Time-Frequency Convolutions (TF-Convs), a novel module designed specifically for anti-spoofing tasks. Unlike conventional CNN layers that treat all dimensions equally, TF-Convs separates the processing of time and frequency to capture finer details that are often overlooked. This design enables the model to focus on spoofing-related cues that may appear in different forms along the two axes.

As shown in Fig. 1, TF-Convs start with a  $1 \times 1$  transition convolution, which adjusts the number of channels from  $C$  to  $C'$ , transforming the input  $x_{\text{HFM}} \in \mathbb{R}^{C \times F \times T}$  into  $\mathbb{R}^{C' \times F \times T}$ . Then, a channel shuffle operation is applied to enhance feature mixing. The resulting tensor is evenly split along the channel

dimension into two parts: one for the frequency path  $x_f$  and one for the time path  $x_t$ , each of shape  $\mathbb{R}^{C'/2 \times F \times T}$ . Each branch is processed separately. The frequency path uses a  $3 \times 1$  depthwise convolution, followed by frequency-wise average pooling and a  $1 \times 1$  pointwise convolution. Similarly, the time path applies a  $1 \times 3$  depthwise convolution, followed by time-wise pooling and a pointwise convolution. These operations produce compact 1D features  $v_f$  and  $v_t$ .

Next comes the broadcasting operation [18] that copies a 1D feature (such as  $v_f$ ) across one dimension to match the size of the original 2D feature map (e.g.,  $F \times T$ ). This operation allows the model to add the compact 1D summary back to the full feature map in a way that maintains structure but reduces computation. The updated maps  $\hat{x}_f$  and  $\hat{x}_t$  are then created using element-wise addition between the original split maps and their corresponding broadcasted vectors, as shown in equations (1) and (2).

$$\hat{x}_f = \sum_{j=1}^T (x_{i,j,f} + v_{j,f}) \quad (1)$$

$$\hat{x}_t = \sum_{i=1}^F (x_{i,j,t} + v_{i,t}) \quad (2)$$

Finally, the two enhanced outputs are concatenated along the channel axis to form the final output  $y \in C' \times F \times T$ . All layers include batch normalization and ReLU activation to improve training stability. By separating and refining the time and frequency information, TF-Convs help the model learn more discriminative features for detecting fake speech.

### C. RawTFNet Architecture

In summary, the architecture of RawTFNet is designed to achieve strong spoofing detection accuracy with low computational cost. It is composed of three main components: a SinConv and DWS SE-Res2Net frontend, a stack of TF-Convs modules, and a Conv1D classifier. The frontend begins by extracting low-level spectro-temporal features from raw audio using a learnable Sinc filter, followed by DWS SE-Res2Blocks that enhance both local patterns and channel-wise information. These early layers allow the model to capture fine-grained acoustic cues that are critical for detecting synthetic speech.

After encoding, the features are refined through nine TF-Convs modules, each of which processes time and frequency dimensions separately. This separation allows the model to detect subtle artifacts in both directions. To improve abstraction, two  $2 \times 2$  max-pooling layers with a stride of 2 are inserted for intermediate downsampling. The final output is fed into a lightweight classifier composed of global average pooling, a 1D convolution layer, which outputs a binary score for bonafide or spoofed speech. To support different computing environments, a channel width parameter  $\tau$  is used to control the model's complexity. This makes RawTFNet- $\tau$  flexible and suitable for both edge devices and high-performance systems.

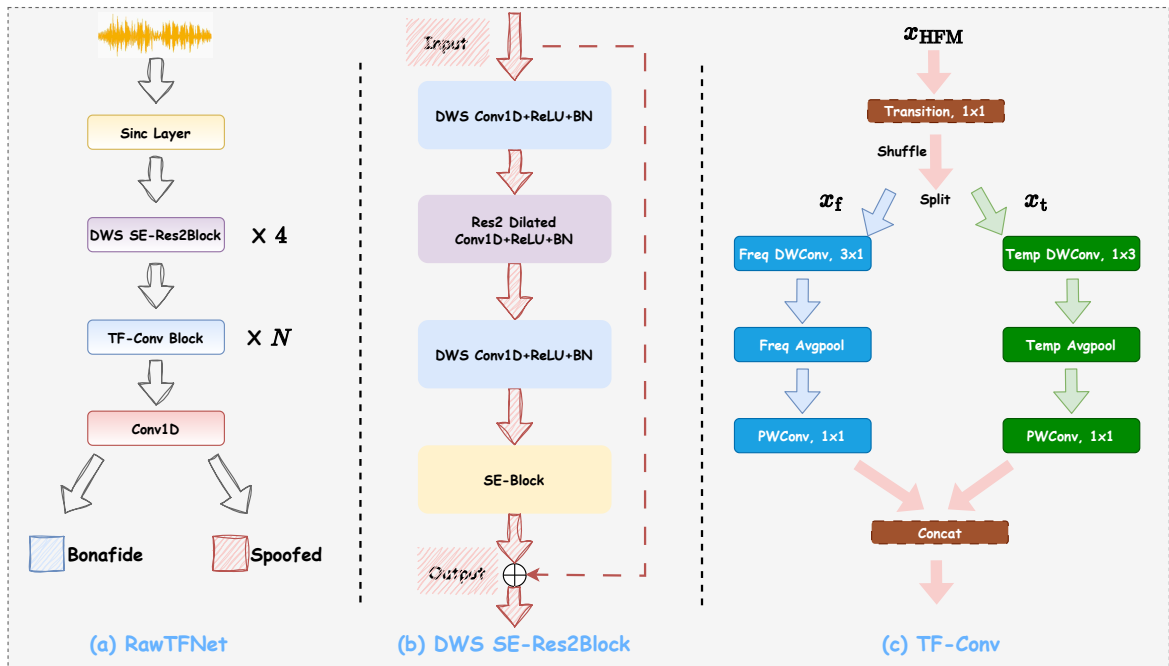


Fig. 1. Overview of the proposed RawTFNet architecture for speech anti-spoofing. (a) The model begins with a Sinc layer followed by four DWS SE-Res2Blocks to extract low-level features. These are then processed by stacked TF-Conv blocks and a final Conv1D layer for classification. (b) The DWS SE-Res2Block integrates depthwise separable convolutions, a Res2Net module with dilation, and a squeeze-and-excitation block to enhance channel-wise attention and efficiency. (c) The TF-Conv module separates feature maps along time and frequency dimensions using 1D depthwise convolutions, followed by pooling and pointwise convolution before merging. We fix the TF-Conv layer number ‘N’ to 9 in this work.

### III. EXPERIMENTAL SETUP

#### A. Datasets

To evaluate the effectiveness and generalization ability of the proposed RawTFNet, we conducted experiments on three benchmark datasets: ASVspoo 2021 LA, ASVspoo 2021 DF, and In-the-Wild. All models were trained using the ASVspoo 2019 LA training set [24], which provides labeled bonafide and spoofed speech samples for supervised learning.

The ASVspoo 2021 LA evaluation set contains about 180,000 utterances, including both real and spoofed speech generated through advanced TTS and VC systems [25]. The ASVspoo 2021 DF set consists of approximately 600,000 utterances, where audio samples are compressed using lossless codecs, reflecting real-world media storage conditions.

#### B. Metrics

We used two standard metrics to measure anti-spoofing performance: the Equal Error Rate (EER) and the minimum tandem detection cost function (min t-DCF) [26], both widely adopted in the ASVspoo Challenges. To assess model efficiency, we also report two system-level metrics: Parameters (PARAMS), which denotes the total number of model parameters, and MACS, referring to the number of million multiply-accumulate operations to measure computational cost.

#### C. Implementation Details

For training, we follow the baseline setup introduced in [15], using the ASVspoo 2019 LA dataset [24]. Each audio sample

TABLE I  
COMPARISON OF SOTA SINGLE ANTI-SPOOFING SYSTEMS ON ASVSPOO 2021 LA AND DF EVALUATION SETS.

Systems	PARAMS	MACS	2021LA		2021DF
			EER%	min t-DCF	EER%
RawGT-ST [27]	0.44M	17.4G	10.25	-	23.26
RawNet2 [15]	17.60M	<b>0.8G</b>	9.50	-	22.38
AASIST [16]	0.30M	8.9G	10.51	0.488	21.07
SE-Rawformer [17]	0.37M	6.1G	4.98	0.318	20.26
Rawformer-L [17]	0.29M	9.2G	6.83	0.361	-
RawTFNet-16 (Proposed)	<b>0.07M</b>	2.9G	<b>4.50</b>	<b>0.295</b>	18.53
RawTFNet-32 (Proposed)	0.17M	5.4G	5.05	0.321	<b>16.82</b>

is either cropped or concatenated to create a fixed-length input segment of approximately 4 seconds (64,000 samples). We adopt the Adam optimizer, with a weight decay of  $1 \times 10^{-4}$  to prevent overfitting. The model is trained using a batch size of 32 and a learning rate of  $1 \times 10^{-4}$ .

As the loss function, we use weighted cross-entropy, which accounts for class imbalance between bonafide and spoofed samples. Following [13], we apply RawBoost augmentation, specifically the ‘series: (1+2+3)’ combination (Algo4), to improve model robustness. All models are trained for 100 epochs, and for evaluation, we average the top five checkpoints based on validation performance to reduce variability and improve stability during testing.

### IV. RESULTS AND DISCUSSIONS

#### A. Comparison Study

Table I presents a comparison between the proposed RawTFNet-16 and RawTFNet-32 models and several SOTA

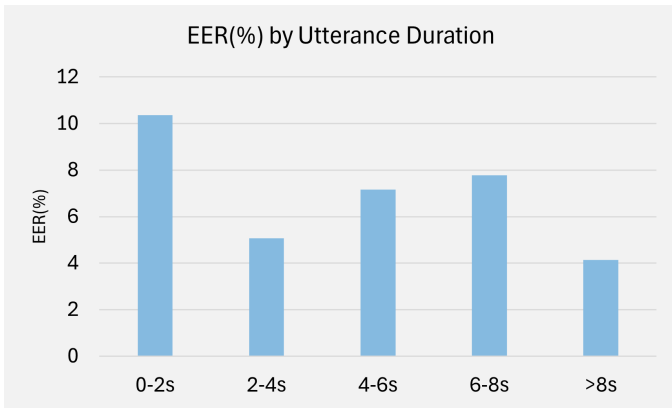


Fig. 2. EER (%) across different utterance duration ranges on the ASVspoof evaluation set. Each bar represents a specific duration group, and ‘n’ indicates the number of utterances in that group.

anti-spoofing systems on the ASVspoof 2021 LA and DF evaluation sets. Both variants of RawTFNet demonstrate a strong balance between detection performance and model efficiency. In terms of model size, RawTFNet-16 has only 0.07M parameters and 2.9G MACs, making it the most lightweight model in the comparison. RawTFNet-32 is slightly larger at 0.17M parameters and 5.4G MACs, but still significantly smaller and more efficient than models like RawNet2 (17.60M parameters with fewer MACs due to its heavy FC layers), AASIST (0.30M, 8.9G), and Rawformer-L (0.29M, 9.2G). These results indicate that both RawTFNet variants are highly suitable for resource-constrained and real-time environments, with RawTFNet-16 offering an extremely compact footprint.

In terms of detection performance, RawTFNet-32 achieves an EER of 5.05% and min t-DCF of 0.321 on the ASVspoof 2021 LA set. This is on par with SE-Rawformer (4.98%) but with much lower complexity. Meanwhile, RawTFNet-16 achieves an even lower EER of 4.50% and min t-DCF of 0.295, suggesting that smaller models can still maintain or even improve accuracy if designed with effective feature processing mechanisms. This performance surpasses larger models such as AASIST (EER 10.51%, min t-DCF 0.488), demonstrating the benefit of RawTFNet’s task-specific architecture. On the more challenging ASVspoof 2021 DF dataset, both RawTFNet variants again show strong generalization. RawTFNet-32 achieves the lowest EER of 16.82% among all systems, outperforming RawNet2 (22.38%), AASIST (21.07%), and SE-Rawformer (20.26%). RawTFNet-16 also performs well with an EER of 18.53%, showing that even the smallest version of our model generalizes effectively to unseen, codec-distorted speech.

Overall, these results confirm the strength of RawTFNet’s time-frequency separation strategy, depthwise convolutional design, and flexible scaling. Both RawTFNet-16 and RawTFNet-32 offer a compelling trade-off between performance and efficiency, making them suitable for diverse real-world deployments.

### B. Impact of Utterance Duration on Detection Performance

Fig. 2 shows how the EER of RawTFNet-32 varies across different utterance durations. The model performs worst on

TABLE II  
ABLATION STUDY ON RAWTFNET-32. ‘w/o’ INDICATES REMOVAL OF THE SPECIFIED COMPONENT FROM THE FULL MODEL. EER (%) IS REPORTED ON THE ASVspoof 2021 LA EVAL SET. PARAMS (K) AND MACs (M) REPRESENT THE NUMBER OF PARAMETERS AND MULTIPLY-ACCUMULATE OPERATIONS, RESPECTIVELY.

Model Variant	EER (%)	MACs (M)	Param (K)
RawTFNet-32 (Full)	<b>5.05</b>	5400	170
w/o freq branch	6.91	4550	150
w/o temp branch	7.08	4480	150
w/o shuffle	5.74	5400	170

very short clips (0–2s), with EER above 10%, likely due to insufficient acoustic context. Performance improves significantly for 2–4s inputs, reaching the lowest EER (~5%), matching the model’s training segment length. However, EER slightly increases for 4–6s and 6–8s ranges, possibly due to added variability or training bias. Interestingly, EER drops again for utterances longer than 8s, though this is based on fewer samples. These results suggest that utterance length influences spoof detection accuracy and highlight the potential of duration-aware training strategies.

### C. Ablation Study

To analyze the impact of each component in RawTFNet-32, we perform an ablation study on the ASVspoof 2021 LA set, as shown in Table II. Removing the frequency or temporal branch significantly increases the EER to 6.91% and 7.08%, respectively, confirming the importance of modeling both time and frequency information. Excluding the shuffle operation leads to a moderate drop in performance (EER: 5.74%) without affecting model size or cost, showing its positive role in feature interaction. These results demonstrate that all components contribute to RawTFNet’s effectiveness and efficiency.

## V. CONCLUSIONS

In this work, we introduced RawTFNet, a lightweight and effective model for speech anti-spoofing. By separating the processing of time and frequency features through the proposed TF-Convs, RawTFNet captures detailed cues in both dimensions while maintaining low computational cost. We further enhanced efficiency using DWS SE-Res2Blocks and a flexible channel width design. Experimental results on ASVspoof 2021 LA and DF evaluation datasets show that RawTFNet achieves competitive accuracy with fewer parameters and lower MACs compared to SOTA systems. Ablation studies confirm the importance of each component, especially the separate time-frequency paths. In summary, RawTFNet offers a robust and efficient framework, well-suited for deployment in real-world anti-spoofing scenarios.

## REFERENCES

- [1] O. Perrotin, B. Stephenson, S. Gerber, G. Bailly, and S. King, “Refining the evaluation of speech synthesis: A summary of the Blizzard Challenge 2023,” *Computer Speech & Language*, p. 101747, 2024.

- [2] Z. Yi, W.-C. Huang, X. Tian, *et al.*, “Voice conversion challenge 2020 — intra-lingual semi-parallel and cross-lingual voice conversion —,” in *Joint Workshop for the Blizzard Challenge and Voice Conversion Challenge 2020*, 2020, pp. 80–98.
- [3] Z. Wu, N. Evans, T. Kinnunen, J. Yamagishi, F. Alegre, and H. Li, “Spoofing and countermeasures for speaker verification: A survey,” *Speech Communication*, vol. 66, pp. 130–153, 2015.
- [4] Z. Wu, P. L. De Leon, C. Demiroglu, *et al.*, “Anti-spoofing for text-independent speaker verification: An initial database, comparison of countermeasures, and human performance,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 4, pp. 768–783, 2016.
- [5] R. K. Das, X. Tian, T. Kinnunen, and H. Li, “The attacker’s perspective on automatic speaker verification: An overview,” in *Proc. Interspeech*, 2020, pp. 4213–4217.
- [6] H. Yin, Y. Xiao, R. K. Das, *et al.*, “EnvSDD: Benchmarking environmental sound deepfake detection,” in *Proc. Interspeech*, 2025, 201–205.
- [7] Z. Wu, J. Yamagishi, T. Kinnunen, *et al.*, “ASVspoof: The automatic speaker verification spoofing and countermeasures challenge,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 4, pp. 588–604, 2017.
- [8] M. Sahidullah, T. Kinnunen, and C. Hanilçi, “A comparison of features for synthetic speech detection,” in *Proc. Interspeech*, 2015, pp. 2087–2091.
- [9] T. B. Patel and H. A. Patil, “Combining evidences from mel cepstral, cochlear filter cepstral and instantaneous frequency features for detection of natural vs. spoofed speech,” in *Proc. Interspeech 2015*, 2015, pp. 2062–2066.
- [10] M. Todisco, H. Delgado, and N. Evans, “Constant Q cepstral coefficients: A spoofing countermeasure for automatic speaker verification,” *Computer Speech & Language*, vol. 45, pp. 516–535, 2017.
- [11] R. K. Das, J. Yang, and H. Li, “Long range acoustic features for spoofed speech detection,” in *Interspeech 2019*, 2019, pp. 1058–1062.
- [12] Y. Chen, J. Yi, J. Xue, *et al.*, “RawBMamba: End-to-end bidirectional state space model for audio deepfake detection,” in *Proc. Interspeech*, 2024, pp. 2720–2724.
- [13] H. Tak, M. Kamble, J. Patino, M. Todisco, and N. Evans, “RawBoost: A raw data boosting and augmentation method applied to automatic speaker verification anti-spoofing,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 6382–6386.
- [14] Y. Xiao and R. K. Das, “XLSR-Mamba: a dual-column bidirectional state space model for spoofing attack detection,” *IEEE Signal Processing Letters*, vol. 32, pp. 1276–1280, 2025.
- [15] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. Evans, and A. Larcher, “End-to-end anti-spoofing with RawNet2,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6369–6373.
- [16] J.-w. Jung, H.-S. Heo, H. Tak, *et al.*, “AASIST: Audio anti-spoofing using integrated spectro-temporal graph attention networks,” in *Proc. IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2022, pp. 6367–6371.
- [17] X. Liu, M. Liu, L. Wang, K. A. Lee, H. Zhang, and J. Dang, “Leveraging positional-related local-global dependency for synthetic speech detection,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [18] B. Kim, S. Chang, J. Lee, and D. Sung, “Broadcasted residual learning for efficient keyword spotting,” in *Proc. Interspeech*, 2021, pp. 4538–4542.
- [19] V. Peddinti, D. Povey, and S. Khudanpur, “A time delay neural network architecture for efficient modeling of long temporal contexts,” in *Proc. Interspeech*, 2015, pp. 3214–3218.
- [20] Y. Xiao and R. K. Das, “TF-Mamba: A time-frequency network for sound source localization,” in *Proc. Interspeech*, 2025, 948–952.
- [21] Y. Cai, P. Zhang, and S. Li, “TF-SepNet: An efficient 1d kernel design in cnns for low-complexity acoustic scene classification,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 821–825.
- [22] M. Ravanelli and Y. Bengio, “Speaker recognition from raw waveform with sincnet,” in *IEEE Spoken Language Technology Workshop (SLT)*, 2018, pp. 1021–1028.
- [23] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proc. IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [24] A. Nautsch, X. Wang, N. Evans, *et al.*, “ASVspoof 2019: Spoofing countermeasures for the detection of synthesized, converted and replayed speech,” *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, no. 2, pp. 252–265, 2021.
- [25] J. Yamagishi, X. Wang, M. Todisco, *et al.*, “ASVspoof 2021: Accelerating progress in spoofed and deepfake speech detection,” in *Proc. The Automatic Speaker Verification Spoofing Countermeasures Workshop (ASVspoof)*, 2021, pp. 47–54.
- [26] T. Kinnunen, H. Delgado, N. Evans, *et al.*, “Tandem assessment of spoofing countermeasures and automatic speaker verification: Fundamentals,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2195–2210, 2020.
- [27] H. Tak, J.-w. Jung, J. Patino, M. Kamble, M. Todisco, and N. Evans, “End-to-end spectro-temporal graph attention networks for speaker verification anti-spoofing and speech deepfake detection,” in *Proc. ASVspoof 2021*, 2021, pp. 1–8.