

DIN-CTS: Low-Complexity Depthwise-Inception Neural Network with Contrastive Training Strategy for Deepfake Speech Detection

Lam Pham*, Dat Tran*, Phat Lam, Florian Skopik,
Alexander Schindler, Silvia Poletti, David Fischinger, Martin Boyer,

{Lam.Pham, Florian.Skopik, Alexander.Schindler, Silvia.Poletti, David.Fischinger, Martin.Boyer}@ait.ac.at,
datttse170375@fpt.edu.vn, phat.lamhcmutddk21@hcmut.edu.vn

Abstract—In this paper, we propose a deep neural network approach for deepfake speech detection (DSD) based on a low-complexity Depthwise-Inception Network (DIN) trained with a contrastive training strategy (CTS) (This work is a part of our DERAME FAKES and EUCINF projects). In this framework, input audio recordings are first transformed into spectrograms using Short-Time Fourier Transform (STFT) and Linear Filter (LF), which are then used to train the DIN. Once trained, the DIN processes bonafide utterances to extract audio embeddings, which are used to construct a Gaussian distribution representing genuine speech. Deepfake detection is then performed by computing the distance between a test utterance and this distribution to determine whether the utterance is fake or bonafide. To evaluate our proposed systems, we conducted extensive experiments on the benchmark dataset of ASVspoof 2019 LA. The experimental results demonstrate the effectiveness of combining the Depthwise-Inception Network with the contrastive learning strategy in distinguishing between fake and bonafide utterances. We achieved Equal Error Rate (EER), Accuracy (Acc.), F1, AUC scores of 4.6%, 95.4%, 97.3%, and 98.9% respectively using a single, low-complexity DIN with just 1.77 M parameters and 985 M FLOPS on short audio segments (4 seconds). Furthermore, our proposed system outperforms the single-system submissions in the ASVspoof 2019 LA challenge, showcasing its potential for real-time applications.

Items— deepfake audio, spectrogram, feature extraction, classification model.

I. INTRODUCTION

Thanks to advancements in deep learning technologies, speech generation systems now beneficially support a wide range of real-world applications including text-to-speech for individuals with speech disorders, voice chatbots in call centers, cross-linguistic speech translation, etc. However, deep learning has also facilitated the creation of fake speech for malicious purposes, such as spoofing attacks, raising serious security concerns. As a result, deepfake speech detection (DSD) has recently gained significant attention from the research community. The state-of-the-art systems, which have been proposed for the task of deepfake speech detection, approach neural network architectures and deep learning techniques [1], [2], [3]. To achieve a high-performance

DSD system, ensembles of input features or models are leveraged. In particular, multiple input features of LFCC, PSCC, LLFB were used in [4]. Similarly, authors in [5] made uses of MEL, CQT, GAM, LFCC, and Wavelet based spectrograms. Regarding model ensembling, two approaches are commonly used. The first involves fusing the individual results of multiple models, as seen in [6] with the use of LCNN and ResNet. The second approach integrates multiple branches within a single network architecture, where feature maps extracted from different pre-trained models are combined. For an example, feature maps from XLS-R, WavLM, and Hubert are merged in [7]. Although ensemble models prove effective to achieve high performance (i.e., the best systems proposed in deepfake speech detection challenges of ASVspoof 2019 [8], 2021 [9], and 2024 [10] leveraged ensemble models), this method leads an issue of a high-complexity model with a large number of trainable parameters and FLOPS. This poses a challenge for applying ensemble models to real-world applications that require a real-time inference or are constrained by hardware limitations. Recently, single models with encoder-decoder based architectures have become popular [1]. In these systems, encoder architectures leverage pre-trained models such as Whisper [11], WavLM [12], or Wave2Vec2.0 [13] (i.e., these models were trained on large-scale datasets of human speech in advance) to extract general feature maps. Meanwhile, decoders present diverse architectures such as Multilayer Perceptron [5], GAN-based architecture [14], Multi-feature attention [15], Graph Attention Network [16] to explore the feature maps extracted from the encoders. Although this approach proposes single models for the DSD task, leveraging pre-trained models as encoders still results in a highly complex system. For examples, the smallest Whisper model presents 39 M parameters and the largest Whisper model has 1550 M parameters. Meanwhile, the pre-trained Wave2Vec2.0 BASE and LARGE models presents around 95 M and 300 M of parameters, respectively. Additionally, this approach mainly focuses on exploring encoder and decoder architectures rather than analyzing training strategies which enforce the model to separate the distributions of bonafide and fake utterances.

To tackle these mentioned limitations, we propose a deep-learning-based model for the deepfake speech detection

L. Pham, F. Skopik, A. Schindler, S. Poletti, D. Fischinger, and M. Boyer are with Austrian Institute of Technology, Vienna, Austria
D. Tran is with FPT University, Vietnam
P. Lam is with HCM University of Technology, Vietnam
(*) Main and equal contribution into the paper.




Fig. 1. The proposed Depthwise-Inception Network architecture and Contrastive Training Strategy

(DSD) and highlight the contributions: (1) We first propose a novel and low-complexity deep neural network for DSD task that is inspired by depthwise convolution and inception architectures, referred to as the Depthwise-Inception Network (DIN). (2) To train DIN model, we propose a contrastive training strategy (CTS) that proves effective to separate distribution of bonafide and fake utterances. (3) By combining DIN model and the CTS method, we achieved a low-complexity and high-performance DSD system.

II. DEPTHWISE-INCEPTION NETWORK ARCHITECTURE AND CONTRASTIVE TRAINING STRATEGY

The detailed architecture of the proposed deep-learning-based model is denoted in Fig. 1 and comprises two main parts: spectrogram feature extraction and deep learning model.

A. Spectrogram Feature Extraction

The input utterances are first split into 4-second audio segments. This segment length generally provides sufficient context to capture important features and allows faster training and inference for applications that require real-time detection. Then, Short Time Fourier Transform (STFT) and Linear Filter (LF) are applied on 4-second segments to generate STFT-LF spectrograms. By setting window size, hop size, and linear filter number of 1024, 512, and 64 respectively, we obtain STFT-LF spectrograms of 128×128 . Then, each spectrogram is concatenated with its first and second derivatives to create a 3-channel representation, resulting in a size of $3 \times 128 \times 128$. Finally, we apply the SpecAug data augmentation [17] to the spectrograms before feeding into the backend classification model.

B. Depthwise-Inception Network Architecture

By comparing pairs of bonafide and fake utterances, we observed that the power distribution across the frequency axis in fake spectrograms follows a regular pattern, whereas in bonafide utterances, it is more variable. This inspired

us to make use of inception network layers which are effective to learn the minor difference in local regions of the spectrograms. Additionally, we leverage depthwise convolution layers and pointwise convolution layers rather than the traditional convolution layers to reduce the model size, but still preserve the power of convolution computation to be able to generate distinct feature maps between fake and bonafide utterances. By combining depthwise convolution and inception layers, we construct the novel Depthwise-Inception Network (DIN). Fig. 1 illustrates the proposed DIN architecture, which can be divided into two main components: the backbone and the heads. The backbone as shown in the right side of the Fig. 1 presents a dense layer (Conv [4x4], BatchNorm (BN) - Gaussian Error Linear Units (GELU)), followed by four Dew-Inc blocks, and a final Global Max Pooling (GM Pooling) layer. For each Dew-Inc block, it presents inception and residual based architecture with four branches of inception-convolution layers with different convolutional kernels (i.e., $[1 \times 1]$, $[3 \times 3]$, $[3 \times 1]$, $[5 \times 1]$) and one residual shortcut. The output of the backbone is fed into two heads: Softmax head and Contrastive head. Softmax head presents two dense layers. The first dense layer presents Fully Connected layer (FC), followed by BN and GELU. The second dense layer comprises FC and Softmax layers. Meanwhile, Contrastive head performs two dense layers which share the same configuration of FC, BN, and GELU.

C. Contrastive Training Strategy

To train the proposed DIN model, we introduce a contrastive training strategy with three training stages. In the first stage, we train DIN model with multiple classes and multiple loss functions. While bonafide utterance is considered as one class, fake generators are the other classes. Training with multiple classes enforces DIN model to learn distinct features among bonafide and fake utterances.

To further enforce the training process, we propose three loss functions. Let N be the number of spectrograms

in each training batch. Then, the output feature vector of the backbone, the Softmax head and the Contrastive head are indicated as $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$, the output feature maps of Softmax head and Contrastive head as $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$ and $\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N\}$, respectively. The first loss function L_1 is A-Softmax loss as shown in Equation (1).

$$L_1 = \frac{1}{N} \sum_{n=1}^N -\log \left(\frac{\exp(s\phi(\theta_c^n))}{\exp(s\phi(\theta_c^n)) + \sum_{j \neq c} \exp(s \cos \theta_j^n)} \right) \quad (1)$$

where $\phi(\theta) = (-1)^k \cos(m\theta) - 2k$, m is the margin, s is the scalar factor; θ_c^n is the angle between the embedding $\mathbf{y}_n$ and the class weight $\mathbf{w}_{y_n}$, c is the label of the embedding $\mathbf{y}_n$. This loss is applied to the feature maps $\mathbf{Y}$ for detecting multiple classes of bonafide and fake generators. We use A-Softmax loss [18] rather than the traditional Entropy loss to maximize angles between feature maps from different classes. m and s are empirically set to 4 and 30, respectively.

The second loss L_2 is a type of contrastive loss which is applied for self-supervised learning [19]. To apply this loss, we consider only two fake speech generators, namely Text-To-Speech (TTS) generator and Voice Conversion (VC) generator. The loss, which is applied on the feature maps $\mathbf{Z}$, aims to minimize the distances among feature maps of the same classes $\mathbf{Z}_C$ and maximize the distances among feature maps of the different classes of $\mathbf{Z}_C$ and $\mathbf{Z}_J$, where $\mathbf{Z}_C$ and $\mathbf{Z}_J$ are subsets of $\mathbf{Z}$ with $C + J = N$. The loss is defined as:

$$L_2 = \frac{1}{N} \sum_{n=1}^N \frac{1}{C} \sum_{c=1}^C \log \left(\frac{\exp(\mathbf{z}_n \cdot \mathbf{z}_c / \tau)}{\exp(\mathbf{z}_n \cdot \mathbf{z}_c / \tau) + \sum_{j=1}^J \exp(\mathbf{z}_n \cdot \mathbf{z}_j / \tau)} \right) \quad (2)$$

where τ is empirically set to 0.01.

A key consideration is that, with the advancement of deep learning, DSD systems should be continuously updated to keep pace with emerging fake speech generators. However, given the rapid proliferation of new generators, this is often impractical. As a result, DSD models frequently encounter a large number of unseen fake speech they were not trained on, posing a challenge to their effectiveness. To tackle this issue, we develop a loss function which focuses on bonafide rather than fake speech. In particular, the loss function aims to minimize the variance of the distribution of feature maps of bonafide utterances. The loss is presented by Equation (3)

$$L_3 = \frac{1}{K} \sum_{k=1}^K \|\mathbf{x}_k - \mathbf{c}\|_2^2 \quad (3)$$

where $\mathbf{c}$ is the central feature map of bonafide utterances, $\mathbf{x}_k$ is a bonafide feature map obtained from the backbone, K is the number of bonafide utterances in the batch of N spectrograms ($K < N$). The central feature map $\mathbf{c}$ is the average of bonafide feature maps in each training batch, but it is recomputed from all bonafide feature maps in the entire training set every 5 epoches.

By combining three loss function, we obtain the final loss L to train the DIN architecture as the Equation (4)

$$L = \alpha L_1 + \beta L_2 + \gamma L_3 \quad (4)$$

TABLE I
PROPOSED CONTRASTIVE TRAINING STRATEGY

Algorithm 1: Contrastive Training Strategy
Input: A set of T spectrograms split into batches of $\mathbf{B}_t = \{\mathbf{I}_1, \mathbf{I}_2, \dots, \mathbf{I}_N\}$
Output: Trained model for fake/bonafide speech detection.
Components:
- The backbone E to extract a set of embeddings $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ from batch $\mathbf{B}_t$
- The Softmax Head SH to extract embeddings $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$ from $\mathbf{X}$
- The Contrastive Head CH to extract embeddings $\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N\}$ from $\mathbf{X}$
- A set of 4 losses represented by functions: $\mathcal{H} = \{H_1, H_2, H_3, H_4\}$
The first stage:
for $e = 1$ to Training Epochs do:
for $t = 1$ to T/N do:
- Extract embeddings $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ from batch $\mathbf{B}_t$:
$\mathbf{X} \leftarrow E(\mathbf{B}_t)$
$\mathbf{Y} \leftarrow SH(\mathbf{X})$
$\mathbf{Z} \leftarrow CH(\mathbf{X})$
- Compute losses:
$L_1 \leftarrow H_1(\mathbf{Y})$
$L_2 \leftarrow H_2(\mathbf{Z})$
$L_3 \leftarrow H_3(\mathbf{X})$
- Compute final loss:
$L \leftarrow \alpha L_1 + \beta L_2 + \gamma L_3$
- Backward pass and update weights of backbone and heads
end for
end for
The second stage:
for $e = 1$ to Finetune Epochs do:
for $t = 1$ to T/N do:
- Extract embeddings $\mathbf{X}$ from batch $\mathbf{B}_t$:
$\mathbf{X} \leftarrow E(\mathbf{B}_t)$
- Compute Entropy loss:
$L_{\text{Entropy}} \leftarrow H_4(\mathbf{X})$
- Backward pass and update weights
end for
end for
The third stage:
- Extract bonafide embeddings $\mathbf{X}_b$ from all batches $\mathbf{B}_t$:
$\mathbf{X}_b \leftarrow E(\mathbf{B}_t)$
- Compute distribution of bonafide embeddings:
$\mathbf{X}_b \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$

where α , β , and γ are empirically set to 0.2, 0.4, 0.4, respectively.

In the second stage, we replace the A-Softmax and Contrastive heads by a new head, referred to as the Entropy head with only a FC layer and a Softmax layer. While the Entropy head is fine-tuned with a high learning rate, the trainable parameters in the backbone are fine-tuned with a low learning rate at this stage.

In the third stage, we fed only bonafide utterances into the pre-trained DIN model obtained from the second stage, extracting the output feature maps of the backbone that are denoted by the $\mathbf{X}_b$ in Fig. 1. Then, the mean $\boldsymbol{\mu}$ and the co-variance matrix $\boldsymbol{\Sigma}$ are computed from the feature maps $\mathbf{X}_b$. In other words, we obtain the Gaussian distribution of bonafide utterances $\mathbf{X}_b \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. In summary, the proposed contrastive training strategy is described in Table I.

D. Inference Process

In the inference process, a testing utterance is fed into the pre-trained DIN model in the second stage to extract the output feature map of the backbone $\mathbf{x}^t$. The Mahalanobis distance d between the testing feature map $\mathbf{x}^t$ and the

TABLE II
PERFORMANCE COMPARISON BETWEEN BASELINE (RESNET18) AND OUR PROPOSED DEEP LEARNING MODELS ON ASV2019-LA

Network & Settings	Acc $\uparrow$	F1 $\uparrow$	AUC $\uparrow$	EER $\downarrow$	Parameters $\downarrow$	FLOPS $\downarrow$
(Baseline) ResNet18 w/ Entropy head & Entropy loss	92.0%	95.4%	96.6%	8.0%	11.18M	1192M
(M1) DIN backbone w/ Entropy head & Entropy loss	92.1%	95.5%	97.5%	7.9%	1.77M	985M
(M2) DIN backbone w/ multiple heads, multiple losses (stage 1&2)	94.6%	97.0%	98.4%	5.3%	1.77M	985M
(M3) DIN backbone w/ multiple heads, multiple losses (3 stages)	95.4%	97.3%	98.9%	4.6%	1.77M	985M




Fig. 2. Histogram of Mahalanobis distances between the training bonafide distribution (‘Train’ subset) and the test bonafide & fake utterances (‘Evaluating’ subset) with the model M1 on the left side and the model M3 on the right side

Gaussian distribution of bonafide utterances $\mathcal{N}(\mu, \Sigma)$ is computed by Equation (5) to decide whether the testing utterance is fake or bonafide.

$$d\{x^t, \mathcal{N}(\mu, \Sigma)\} = \sqrt{(x^t - \mu)^T \Sigma^{-1} (x^t - \mu)} \quad (5)$$

III. EXPERIMENTS AND RESULTS

A. Datasets and evaluation metrics

We evaluate the proposed models on the Logic Access dataset of ASVspoof 2019 challenge (ASV2019-LA). The models are trained on the ‘Train’ subset and evaluated on ‘Development’ subset from ASV2019-LA. These subsets comprise bonafide utterances and fake utterances from six speech generators (2 VC and 4 TTS systems), referred to as A01 to A06. Finally, the models are tested on the ‘Evaluation’ subsets from ASV2019-LA. This subset comprises bonafide utterances and fake utterances from 13 speech generators, referred to as A07 to A19, which are different from fake generators in ‘Train’ and ‘Development’ subsets. We follow the ASVspoof challenge guidelines and use the Equal Error Rate (EER) as the primary evaluation metric for the proposed models. Additionally, we report Accuracy (Acc.), F1 score, and AUC to facilitate performance comparison among the models.

B. Experimental settings

The proposed DSD system has been developed within the Pytorch framework. The architecture has been trained for 60 epochs in total, using Titan RTX 24GB GPU. The first 50 epochs have been used for the Stage 1 of the training. Then, the model is finetuned on two classes (fake/bonafide) for the

remaining 10 epoches in Stage 2. The Adam method [20] is used for the optimization.

C. Results and discussion

To evaluate our proposed model, we first construct a baseline leveraging ResNet18 backbone for feature map extraction and an Entropy head for classification (This Entropy head presents the architecture mentioned in stage 2 of the training process in section II-C). The baseline is trained on two classes (fake and bonafide) using Cross-Entropy loss. We also evaluated three other models, referred to as M1, M2, and M3. M1 uses the architecture in the stage 2 which presents the depthwise-inception backbone as shown in right side of Fig. 1 and the Entropy head. This model is train from scratch on two classes (fake and bonafide) using Cross-Entropy loss. We compare M1 with the baseline to evaluate the role of proposed DIN architecture. M2 presents the architecture in the stage 1. Then, this model is fine-tuned in stage 2, but it is not applied in the stage 3. We evaluate the role of contrastive training strategy with multiple classes and multiple loss functions in M2 model. Finally, M3 is the proposed full model of Fig. 1 including the entire three-stage training strategy which leads to the estimation of the bonafide distribution in stage 3.

Experimental results in Table II indicate that M1 with DIN architecture outperforms the ResNet18 baseline over all metrics. While M1 presents a complexity with 1.77 M parameters and 985 M FLOPS, ResNet baseline shows a much higher complexity with 11.18 M parameters and 1192 M FLOPS. When the proposed contrastive training strategy




Fig. 3. TSNE-based visualization of the output feature maps extracted from the backbone (embedding X extracted from model M1 on the left side and extracted from model M3 on the right side) among testing bonafide and fake utterances ('Evaluation' subset with A07 to A19 generators)

in stages 1 and 2 is applied to M2, it significantly improves the EER performance by 2.6%. Applying stage 3 where the bonafide distribution is used to measure the Mahalanobis distances among utterances helps to further improve all metric scores.

The visualization shown in Fig. 2 indicates that the proposed three-stage contrastive training strategy is effective to separate the distributions of fake and bonafide utterances. This leads a smaller overlapping region between bonafide and fake utterances in M3 model compared with M1 model. The visualization shown in Fig. 3, which uses TSNE from Scikit-Learn with $n_components$, $perplexity$, ann_iter parameters set to 2, 30, 300 respectively, again proves that the contrastive training strategy is effective to separate bonafide utterances and fake utterances from different generators.

As a result, we achieve the single model M3 with the best EER score of 4.6%, which outperforms submitted single systems in the challenge of ASV-2019 LA task (Top-4 single system in the challenges: T32 (4.92%), T04 (5.74%), T01 (6.01%), and T58 (6.14%) in [21]). With a low complexity of 1.77 M parameters and Acc., F1, AUC scores of 95.4%, 97.3%, and 98.9%, the model M3 shows strong potential for real-time deployment in DSD systems.

IV. CONCLUSION

This paper has presented a deep-learning-based model for the task of deepfake speech detection. By combining the depthwise-inception network architecture and the three-stage contrastive training strategy, we achieve a low-complexity and high-performance single model (M3) which shows great for real-time DSD applications. The future work, which focus on cross-dataset and cross-language evaluation, evaluate how general and robust is the proposed model.

ACKNOWLEDGMENTS

EUCINF project is co-funded by European Union under grant agreement N°101121418. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

Defame Fakes is funded by the Austrian security research programme KIRAS of the Federal Ministry of Finance (BMF).

REFERENCES

- [1] Lam Pham et al., "A comprehensive survey with critical analysis for deepfake speech detection," *arXiv preprint arXiv:2409.15180*, 2024.
- [2] Menglu Li, Yasaman Ahmadiadli, and Xiao-Ping Zhang, "Audio anti-spoofing detection: A survey," *arXiv preprint arXiv:2404.13914*, 2024.
- [3] Zahra Khanjani, Gabrielle Watson, and Vandana P Janeja, "Audio deepfakes: A survey," *Frontiers in Big Data*, vol. 5, pp. 1001063, 2023.
- [4] Woo Hyun Kang, Jahangir Alam, and Abderrahim Fathan, "Crim's system description for the asvspoof2021 challenge," in *Proc. INTERSPEECH*, 2021, pp. 100–106.
- [5] Lam Pham, Phat Lam, Truong Nguyen, Huyen Nguyen, and Alexander Schindler, "Deepfake audio detection using spectrogram-based feature and ensemble of deep learning models," in *Proc. IS2*, 2024, pp. 1–5.
- [6] Anton Tomilov, Aleksei Svishchev, Marina Volkova, Artem Chirkovskiy, Alexander Kondratyev, and Galina Lavrentyeva, "Stc antispoofing systems for the asvspoof2021 challenge," in *Proc. INTERSPEECH*, 2021, pp. 61–67.
- [7] Yujie Yang, Haochen Qin, Hang Zhou, Chengcheng Wang, Tianyu Guo, Kai Han, and Yunhe Wang, "A robust audio deepfake detection system via multi-view feature," in *Proc. ICASSP*, 2024, pp. 13131–13135.
- [8] Xin Wang et al., "Asvspoof 2019: A large-scale public database of synthesized, converted and replayed speech," *Computer Speech & Language*, vol. 64, pp. 101114, 2020.
- [9] Junichi Yamagishi et al., "Asvspoof 2021: accelerating progress in spoofed and deepfake speech detection," in *Proc. INTERSPEECH*, 2021, pp. 10744–10753.
- [10] Xin Wang et al., "Asvspoof 5: Crowdsourced speech data, deepfakes, and adversarial attacks at scale," in *Proc. INTERSPEECH*, 2024, pp. 1008–1012.
- [11] Alec Radford et al., "Robust speech recognition via large-scale weak supervision," in *Proc. ICML*, 2023, pp. 28492–28518.
- [12] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al., "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [13] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli, "wav2vec 2.0: a framework for self-supervised learning of speech representations," in *Proc. NIPS*, 2020.
- [14] Junlong Deng, Yanzhen Ren, Tong Zhang, Hongcheng Zhu, and Zongkun Sun, "VFD-net: Vocoder fingerprints detection for fake audio," in *Proc. ICASSP*, 2024, pp. 12151–12155.
- [15] Yinlin Guo, Haofan Huang, Xi Chen, He Zhao, and Yuehai Wang, "Audio deepfake detection with self-supervised wavlm and multi-fusion attentive classifier," in *Proc. ICASSP*, 2024, pp. 12702–12706.
- [16] Hemlata Tak, Massimiliano Todisco, Xin Wang, Jee-weon Jung, Junichi Yamagishi, and Nicholas Evans, "Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation," in *Proc. INTERSPEECH*, 2022, pp. 112–119.
- [17] D. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *preprint arXiv:1904.08779*, 2019.
- [18] Yutian Li, Feng Gao, Zhijian Ou, and Jiasong Sun, "Angular softmax loss for end-to-end speaker verification," in *Proc. ISCSLP*, 2018, pp. 190–194.
- [19] Patrick Feeney and Michael C Hughes, "Sincere: Supervised information noise-contrastive estimation revisited," *arXiv preprint arXiv:2309.14277*, 2023.
- [20] P. K. Diederik and B. Jimmy, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2015.
- [21] Massimiliano Todisco et al., "Asvspoof 2019: Future horizons in spoofed and fake audio detection," in *Proc. INTERSPEECH*, 2019, pp. 1008–1012.