



Genuine-Focused Learning using Mask AutoEncoder for Generalized Fake Audio Detection

Xiaopeng Wang^{1,2}, Ruibo Fu^{1,*}, Zhengqi Wen¹, Zhiyong Wang^{1,2}, Yuankun Xie³, Yukun Liu², Jianhua Tao^{4,5}, Xuefei Liu¹, Yongwei Li¹, Xin Qi^{1,2}, Yi Lu^{1,2}, Shuchen Shi⁶

¹ Institute of Automation, Chinese Academy of Sciences ² School of Artificial Intelligence, University of Chinese Academy of Sciences ³ School of Information and Communication Engineering, Communication University of China ⁴ Department of Automation, Tsinghua University ⁵ Beijing National Research Center for Information Science and Technology, Tsinghua University ⁶ Shanghai Polytechnic University

wangxiaopeng22@mails.ucas.ac.cn, ruibo.fu@nlpr.ia.ac.cn.

Abstract

The generalization of Fake Audio Detection (FAD) is critical due to the emergence of new spoofing techniques. Traditional FAD methods often focus solely on distinguishing between genuine and known spoofed audio. We propose a Genuine-Focused Learning (GFL) framework guided, aiming for highly generalized FAD, called GFL-FAD. This method incorporates a Counterfactual Reasoning Enhanced Representation (CRER) based on audio reconstruction using the Mask AutoEncoder (MAE) architecture to accurately model genuine audio features. To reduce the influence of spoofed audio during training, we introduce a genuine audio reconstruction loss, maintaining the focus on learning genuine data features. In addition, content-related bottleneck (BN) features are extracted from the MAE to supplement the knowledge of the original audio. These BN features are adaptively fused with CRER to further improve robustness. Our method achieves state-of-the-art performance with an EER of 0.25% on ASVspoof2019 LA.

Index Terms: fake audio detection, genuine-focused learning, counterfactual reasoning enhanced representation, mask autoencoder

1. Introduction

Fake Audio Detection (FAD) is vital for ensuring information security and preventing fraud. With the emergence of new spoofing techniques[1], developing robust FAD measures for Automatic Speaker Verification (ASV) [2] systems has become increasingly urgent. FAD systems must have the ability to generalize to adapt to the evolving and unpredictable nature of spoofing attacks in the real world. To promote the development of generalized research on FAD, the ASVspoof [3, 4] initiative and its series of challenges have amassed a substantial database of genuine and spoofed utterances. Of particular note, the test sets contain numerous spoofing methods not present in the training sets, providing valuable resources for evaluating the performance and generalization challenges of FAD systems.

Currently, the detection of spoofed audio is mainly based on a two-step approach involving feature extraction and classification tasks. The choice of front-end feature extraction is critical for the model's generalizability. Traditional approaches involve manual extraction of features such as first-order spectral features [5, 6, 7, 8], second-order spectral features [9, 10, 11], and methods that fuse spectral features [12, 13]. These approaches have somewhat improved the robustness of FAD; however, their capability to detect unknown spoofing methods is still limited, and their generalizability needs further enhancement.

* corresponding author

In recent years, pre-trained models based on first-order spectral features have demonstrated remarkable noise robustness and generalization capabilities in FAD tasks [14, 15, 16, 17]. They use an encoder-only architecture, trained with contrastive loss for the purpose of self-supervised learning [18, 19]. Their outstanding performance is due to pre-training on a large corpus of genuine data with background noise. In contrast, pre-trained models based on second-order spectral features use a different architecture [20, 21, 22], typically an encoder-decoder structure where the decoder supports the encoder during self-supervised training. However, similar to first-order spectral features, when applied to FAD tasks, these models also function merely as feature extractors, without further modeling the genuine audio patterns. These patterns are present in the FAD training data. In addition, research on better exploiting second-order spectral features pre-training for transfer learning to FAD tasks is still very limited.

In this study, we introduce a novel Genuine-Focused Learning (GFL) framework guided to enhance the generalization capability of FAD, called GFL-FAD. Unlike previous self-supervised models, we adopt an encoder-decoder architecture. To magnify the fine differences between genuine and spoofed audio, a Counterfactual Reasoning Enhanced Representation (CRER) method is proposed in the light of the audio reconstruction mechanism, where Mask AutoEncoder (MAE) [23] architecture is employed for accurately modeling genuine audio features. To reduce the influence of spoofed audio during training, we propose the Genuine Audio Reconstruction Loss (GAR Loss) to maintain the focus on learning genuine audio features, thereby deepening the understanding of the essence of genuine audio and improving the detection ability of the model against unknown spoofing techniques. In addition, content-related bottleneck (BN) features are extracted from MAE to supplement the knowledge of the original audio, and we use an attention mechanism to effectively fuse the BN features of MAE with CRER, which improves the robustness of the model. Evaluated in the ASVspoof2019 Logical Access (LA) scenario, our approach achieves state-of-the-art performance with an Equal Error Rate (EER) of 0.25%.

- We propose the GFL-FAD architecture, focusing on genuine audio to amplify the differences between genuine and spoofed audio, thereby enhancing detection capabilities.
- We propose the GAR Loss to further focus on genuine learning and better capture the distribution of genuine features, thereby improving the generalizability of the model.
- We use an attention mechanism to effectively fuse the BN features of MAE with CRER, thereby increasing the robustness of the model.

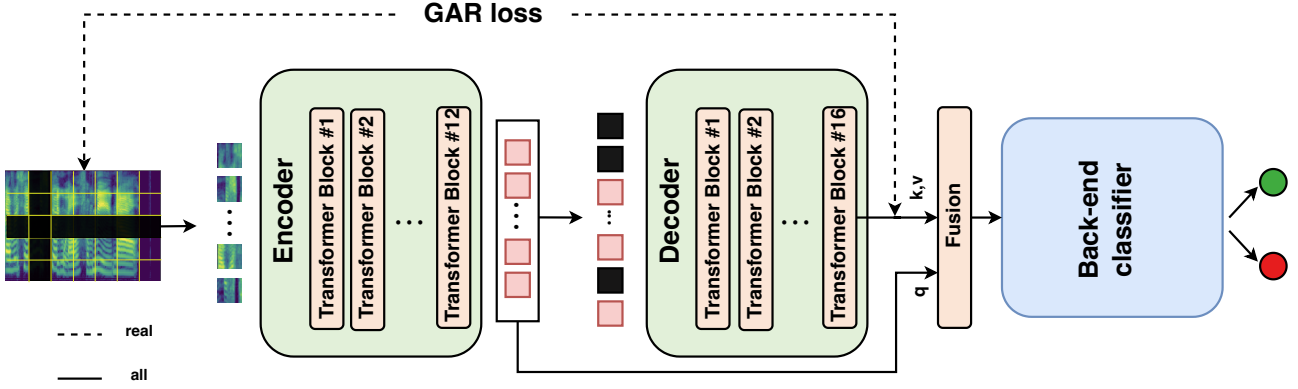


Figure 1: Illustration of the overall architecture for the proposed GFL-FAD. We start by segmenting patches from the spectral representation. A portion of these patches are encoded by the MAE encoder, resulting in BN features. The remaining patches are masked and concatenated with the BN features before being fed to the MAE decoder for spectrogram reconstruction. We refer to the reconstructed spectral features as CRER. If the input audio is genuine, an additional GAR Loss is computed to maintain focus on modeling genuine audio knowledge. Through fusion, the CRER is combined with the BN features and passed to the back-end classification network for final detection.

2. Proposed method

The overall framework of our approach is shown in Figure 1. Unlike previous self-supervised models, which are primarily used as feature extractors, we jointly train the decoder of the MAE to enhance the discrimination between genuine and spoofed audio (Section 2.1). To improve the method’s focus on genuine, we introduce the GAR Loss (Section 2.2). In addition, to preserve the original information, we use a fusion module to combine the BN features of MAE with CRER (Section 2.3) for the final classification task.

2.1. Counterfactual Reasoning Enhanced Representation

The CRER is extracted through the decoder of the MAE. Our MAE architecture is based on the Audio Masked AutoEncoder (audio-MAE)[21], with the encoder comprising 12 layers of ViT-Base (ViT-B), and the decoder being a 16 layers transformer that includes a local attention mechanism. The spectrogram input to the MAE must first be preprocessed, including segmented into small patches, each of which is mapped to a C -dimensional embedding vector. All embedding vectors are represented as $X = [x_1, x_2, \dots, x_N] \in \mathbb{R}^{C \times N}$, where $N = F \times T$ is the number of patches, F is the number of sub-bands, and T is the number of temporal segments. Position embedding is added to each patch to capture its location and the relative positional relationships between patches. The formula is as follows:

$$X^p = X + E_{pos} \quad (1)$$

The sequence $X^p = [x_1^p, x_2^p, \dots, x_N^p] \in \mathbb{R}^{C \times N}$ consists of patch embedding with integrated positional information, while $E_{pos} \in \mathbb{R}^{C \times N}$ denotes the positional embedding for all patches. After obtaining X^p , certain areas of time and frequency domains are randomly masked, and the unmasked patches are fed into the encoder to encode the audio representations directly as content-related BN features. The masked patches, along with the encoded audio representations, are then input into the decoder for spectrogram reconstruction. The decoded features serve as the CRER.

2.2. Genuine Audio Reconstruction Loss

A key component of our training strategy is the inclusion of the GAR Loss, specifically designed for GFL, to enhance the model’s ability to discriminate between genuine and spoofed audio. The formula is as follows:

$$L_{GAR} = \frac{\sum_{i=1}^N \text{mask}_i \cdot L_{\text{recon},i}}{\sum_{i=1}^N \text{mask}_i} \quad (2)$$

where mask_i denotes the authenticity of sample i (label 1 is the genuine sample), and $L_{\text{recon},i}$ denotes the reconstruction loss for the i -th sample. The formula for the reconstruction loss is as follows:

$$L_{\text{recon},i} = \text{MSE}(\text{prediction}_i, \text{input}_i) \quad (3)$$

We use the Mean Squared Error (MSE) as the reconstruction loss, aiming to minimize the MSE between the CRER and the original spectrogram. In the reconstruction loss, we calculate the loss only for the masked patch blocks, rather than the entire spectrogram. Additionally, to better optimize the model, we standardize the reconstructed CRER by its mean and variance.

2.3. Fusion Strategy for Feature Characteristics

In our fusion strategy, we employ the encoder segment of the Transformer model [24] to integrate the BN features from MAE with CRER. We utilize the multi-head attention mechanism of the Transformer, which conducts several self-attention operations concurrently, to highlight the unique characteristics and focal points of each feature set. The attention computation is formulated as follows:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (4)$$

Here, d_k denotes the dimension of the keys. CRER and BN features are first mapped to the same dimension through two separate linear layers. the embeddings of BN features to serve as queries (Q), with CRER acting as both keys (K) and values

(V). This configuration allows for a targeted fusion of features, emphasizing the significant aspects of genuine audio present in the BN features while enhancing them with the nuanced distinctions provided by CRER.

The multi-head attention mechanism generates h distinct representations of (Q, K, V) , each undergoing a scaled dot-product attention calculation. The outcomes from each head are then concatenated and passed through a feed-forward layer, as described by:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h)W^O \quad (5)$$

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V) \quad (6)$$

This fusion strategy emphasizes the strengths of each feature set, ensuring that the integrated representation benefits from the detailed resolution of the BN features and the discriminative power of the CRER.

2.4. Loss Function

In our framework, the total loss function L_{total} is a blend of Cross-Entropy Loss L_{CE} for classification and Genuine Audio Reconstruction Loss L_{GAR} for audio reconstruction, formulated as:

$$L_{\text{total}} = L_{\text{CE}} + \alpha \cdot L_{\text{GAR}} \quad (7)$$

Here, L_{CE} evaluates the model’s ability to classify audio as genuine or spoofed, directly impacting the accuracy of fake audio detection. L_{GAR} focuses on the reconstruction of genuine audio, ensuring that the model retains the ability to accurately reproduce the characteristics of genuine audio signals. The coefficient α , initially set to 0.01, is a non-learnable parameter that we investigate in ablation experiments to understand its impact on performance. This optimization enhances the model’s focus on classification for spoofed audio while maintaining the integrity of genuine audio reconstruction.

3. Experimental setup

3.1. Dataset details

Our experiments were conducted using the ASVspoof2019 LA dataset, which is recognized for its extensive collection of genuine and spoofed audio samples, produced via advanced text-to-speech and voice conversion techniques. The evaluation set is particularly noted for encompassing a broader range of spoofing attacks, featuring 13 distinct scenarios (A07-A19). This dataset and the distribution of attacks across the partitions are summarized in Table 1. Additionally, acknowledging the influence of initialization on the performance of spoofing detection systems, the results presented are averaged over multiple runs with varied random seeds, ensuring a reliable and consistent performance evaluation.

Table 1: Overview of the ASVspoof2019 LA dataset

Partition	genuine	Spoof	Attacks
Training	2,580	22,800	A01 - A06
Development	2,548	22,296	A01 - A06
Evaluation	7,355	63,882	A07 - A19

3.2. Implementation details

In our experimental setup, a raw waveform comprising 64,600 frames (approximately 4 seconds) was acquired for analysis. We use a window size of 25 ms with a hop length of 10 ms to transform the waveform into 128 mel-bank features. We initialized our MAE model using the pre-trained parameters of audio-MAE [21]. We used AASIST [5] as our back-end classifier. Subsequently, our proposed GFL-FAD method was implemented and trained using the PyTorch framework for 100 epochs. The training process was conducted on an NVIDIA RTX 4090 GPU, utilizing a batch size of 16 to efficiently process the data. To optimize the training procedure, we adopted the AdamW [25] optimizer with a learning rate of 5×10^{-6} . Additionally, a cosine annealing learning rate decay scheme was employed to dynamically adjust the learning rate throughout the training process, ensuring optimal model convergence and performance.

For the evaluation of our proposed system, we utilized two principal metrics: the EER and the Minimum Tandem Detection Cost Function (min t-DCF) [26], with a particular emphasis on the EER for its direct interpretability. The min t-DCF is also reported as it represents the standard for the ASVspoof challenges, providing a comprehensive assessment in conjunction with ASV systems.

4. Results

4.1. Performance Comparison

We visualized the high-dimensional representations of the model at its best performance using T-SNE, as shown in Figure 2. The visualization confirms the accurate fitting of genuine audio by the model, highlighting the effectiveness of our approach in learning the patterns of genuine data. However, the visualization also shows that some spoofed samples were incorrectly classified as genuine. This may be because the artifacts in the spectra of these spoofed samples are not pronounced, making it difficult for the model to accurately detect them even when the differences are amplified.

In our performance analysis, we listed the performance of SOTA methods based on first-order and second-order spectral features in terms of EER and min t-DCF metrics on the ASVspoof2019 dataset, as shown in Table 2. F0-SENet34 [31] achieved the best performance by fusing the fundamental




Figure 2: Visualization of High-Dimensional Representations of Genuine and Spoofed Audio Samples using T-SNE

Table 2: Performance Comparison of First-Order and Second-Order Spectral Methods With and Without Pre-Training in ASVspoof 2019 LA Evaluation

System	Front-end	min t-DCF	EER (%)
No pre-training			
LCNN-LSTM-sum [11]	LFCC	0.0524	1.92
Attention+Resnet [27]	FFT	0.0510	1.87
MCG-Res2Net50 [10]	CQT	0.0520	1.78
FFT-L-SENet [28]	FFT	0.0368	1.14
Raw PC-DARTS [29]	Raw waveform	0.0517	1.77
Res-TSSDNet [8]	Raw waveform	0.0481	1.64
RawGAT-ST [7]	Raw waveform	0.0335	1.06
S ² pecNet [12]	Raw & LFCC	0.0240	0.84
AASIST [5]	Raw waveform	0.0275	0.83
Graph-ST [9]	LFB	0.0166	0.58
DFSincNet [30]	Raw waveform	0.0176	0.52
FO-SENet34 [31]	STFT	0.0143	0.43
pre-training			
HuRawNet2 [32]	Raw waveform	0.1393	1.96
Wav2vec2+LGF [33]	Raw waveform	-	1.28
Wav2vec2+AASIST [34]	Raw waveform	-	0.90
Wav2vec2+LLGF [33]	Raw waveform	-	0.86
Wav2vec2+VIB [35]	Raw waveform	0.0107	0.40
Wav2vec2+ASP [15]	Raw waveform	-	0.31
Ours	LFB	0.0071	0.25

frequency with real and imaginary spectrograms in scenarios without pre-training. In scenarios with pre-trained models, the studies were based on first-order spectral features pre-training, with the Wav2vec2+ASP [15] method achieving the best performance. Our GFL-FAD, based on second-order spectral features pre-training, achieved the SOTA performance among all of the above methods. In addition, we compared the performance of our GFL-FAD method with that of using only the encoder model pre-trained with audio-MAE in scenarios using the same classifier, which also demonstrated the effectiveness of our proposed approach. The specific performance details are provided in the ablation experiments.

4.2. Ablation experiments

The ablation experiments performed on the ASVspoof2019 LA dataset demonstrate the significant contributions of each component within our GFL-FAD approach, as detailed in Table 3. Removing the GAR Loss component slightly increases the EER to 0.39% and the min t-DCF to 0.0115, indicating the importance of GAR Loss within the framework. Furthermore, the absence of the BN features (EN) leads to an increase in the EER to 0.41%, demonstrating the positive impact of these features on system performance. It is noteworthy that not using the CRER features (DE), which essentially treats the pre-trained model encoder as a mere feature extractor, results in the worst performance. We also explored the optimal setting of the hyperparameter alpha, which balances detection and GAR Loss. The results in Table 4 show that the best performance is achieved when alpha is set to 0.01. In addition, we investigated the effect of different mask ratios (time + frequency) on FAD, as shown in Figure 3. The study revealed that excessively high mask ratios prevent MAE from learning effective audio representations due to a significant loss of spectral information. Conversely, excessively low mask ratios oversimplify the reconstruction task

and prevent effective learning by the MAE. Ultimately, the best performance was achieved with a mask ratio of 0.3.




Figure 3: Performance of GFL-FAD at Different Mask Ratios

Table 3: Ablation Experiments of Our Contributions

Method	EER(%)	min t-DCF
GFL-FAD	0.25	0.0071
w/o GAR	0.39	0.0115
w/o EN	0.41	0.0119
w/o DE	0.56	0.0184

Table 4: Selection of hyperparameter α

α	EER(%)	min t-DCF
1	0.54	0.0153
0.1	0.41	0.0122
0.01	0.25	0.0071

5. Conclusion

In this study, we present GFL-FAD, a novel framework that uses GFL to enhance the generalization capabilities of FAD. This approach marks the first application of an encoder-decoder structure, specifically the MAE, to FAD tasks. Our method enhances the subtle differences between real and spoofed audio through the CRER, which is informed by an audio reconstruction mechanism. To minimize the impact of spoofed audio during training, we propose the GAR loss, which ensures focused learning of genuine audio features. This deepens our understanding of the nature of genuine audio and improves the model's ability to detect novel spoofing methods. In addition, we extract BN features from the MAE encoder and fuse them with CRER using an attention mechanism to improve the robustness of the model. Our GFL-FAD demonstrates state-of-the-art performance with an EER of 0.25% in the challenging ASVspoof2019 LA scenario. In the future, we aim to explore more effective ways to model genuine audio patterns and extend our methodology to other speech anti-spoofing contexts, such as the detection of partially spoofed speech.

6. Acknowledgements

This work is supported by the National Natural Science Foundation of China (NSFC) (No. 62101553, No. 62306316, No. U21B20210, No. 62201571).

7. References

- [1] R. Fu, J. Tao, Z. Wen, J. Yi, and T. Wang, "Focusing on attention: prosody transfer and adaptative optimization strategy for multi-speaker end-to-end speech synthesis," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6709–6713.
- [2] D. A. Reynolds, "An overview of automatic speaker recognition technology," in *2002 IEEE international conference on acoustics, speech, and signal processing*, vol. 4. IEEE, 2002, pp. IV–4072.
- [3] M. Todisco, X. Wang, V. Vestman, M. Sahidullah, H. Delgado, A. Nautsch, J. Yamagishi, N. Evans, T. Kinnunen, and K. A. Lee, "Asvspoof 2019: Future horizons in spoofed and fake audio detection," in *INTERSPEECH 2019-20th Annual Conference of the International Speech Communication Association*, 2019.
- [4] X. Wang, J. Yamagishi, M. Todisco, H. Delgado, A. Nautsch, N. Evans, M. Sahidullah, V. Vestman, T. Kinnunen, K. A. Lee *et al.*, "Asvspoof 2019: A large-scale public database of synthesized, converted and replayed speech," *Computer Speech & Language*, vol. 64, p. 101114, 2020.
- [5] 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)*. IEEE, 2022, pp. 6367–6371.
- [6] 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 *ASVSPOOF 2021, Automatic Speaker Verification and Spoofing Countermeasures Challenge*. ISCA, 2021, pp. 1–8.
- [7] H. Tak, J. weon 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. 2021 Edition of the Automatic Speaker Verification and Spoofing Countermeasures Challenge*, 2021, pp. 1–8.
- [8] G. Hua, A. B. J. Teoh, and H. Zhang, "Towards end-to-end synthetic speech detection," *IEEE Signal Processing Letters*, vol. 28, pp. 1265–1269, 2021.
- [9] F. Chen, S. Deng, T. Zheng, Y. He, and J. Han, "Graph-based spectro-temporal dependency modeling for anti-spoofing," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [10] X. Li, X. Wu, H. Lu, X. Liu, and H. Meng, "Channel-Wise Gated Res2Net: Towards Robust Detection of Synthetic Speech Attacks," in *Proc. Interspeech 2021*, 2021, pp. 4314–4318.
- [11] X. Wang and J. Yamagishi, "A comparative study on recent neural spoofing countermeasures for synthetic speech detection," *Interspeech 2021*, 2021.
- [12] P. Wen, K. Hu, W. Yue, S. Zhang, W. Zhou, and Z. Wang, "Robust Audio Anti-Spoofing with Fusion-Reconstruction Learning on Multi-Order Spectrograms," in *Proc. INTERSPEECH 2023*, 2023, pp. 271–275.
- [13] C. Fan, J. Xue, S. Dong, M. Ding, J. Yi, J. Li, and Z. Lv, "Subband fusion of complex spectrogram for fake speech detection," *Speech Communication*, vol. 155, p. 102988, 2023.
- [14] Y. Xie, H. Cheng, Y. Wang, and L. Ye, "Domain generalization via aggregation and separation for audio deepfake detection," *IEEE Transactions on Information Forensics and Security*, 2023.
- [15] J. W. Lee, E. Kim, J. Koo, and K. Lee, "Representation Selective Self-distillation and wav2vec 2.0 Feature Exploration for Spoof-aware Speaker Verification," in *Proc. Interspeech 2022*, 2022, pp. 2898–2902.
- [16] Y. Xie, H. Cheng, Y. Wang, and L. Ye, "Learning a self-supervised domain-invariant feature representation for generalized audio deepfake detection," in *Proc. INTERSPEECH*, vol. 2023, 2023, pp. 2808–2812.
- [17] Z. Cai and M. Li, "Integrating frame-level boundary detection and deepfake detection for locating manipulated regions in partially spoofed audio forgery attacks," *Computer Speech & Language*, vol. 85, p. 101597, 2024.
- [18] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [19] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. 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.
- [20] A. Baade, P. Peng, and D. Harwath, "MAE-AST: Masked Autoencoding Audio Spectrogram Transformer," in *Proc. Interspeech 2022*, 2022, pp. 2438–2442.
- [21] P.-Y. Huang, H. Xu, J. Li, A. Baevski, M. Auli, W. Galuba, F. Metze, and C. Feichtenhofer, "Masked autoencoders that listen," *Advances in Neural Information Processing Systems*, vol. 35, pp. 28 708–28 720, 2022.
- [22] Y. Gong, C.-I. Lai, Y.-A. Chung, and J. Glass, "Ssast: Self-supervised audio spectrogram transformer," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 10, 2022, pp. 10 699–10 709.
- [23] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 16 000–16 009.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [25] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2018.
- [26] T. Kinnunen, H. Delgado, N. Evans, K. A. Lee, V. Vestman, A. Nautsch, M. Todisco, X. Wang, M. Sahidullah, J. Yamagishi *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. Ling, L. Huang, J. Huang, B. Zhang, and P. Li, "Attention-based convolutional neural network for asv spoofing detection," in *Interspeech*, 2021, pp. 4289–4293.
- [28] J. Xue, C. Fan, Z. Lv, J. Tao, J. Yi, C. Zheng, Z. Wen, M. Yuan, and S. Shao, "Audio deepfake detection based on a combination of f0 information and real plus imaginary spectrogram features," in *Proceedings of the 1st International Workshop on Deepfake Detection for Audio Multimedia*, 2022, pp. 19–26.
- [29] W. Ge, J. Patino, M. Todisco, and N. Evans, "Raw differentiable architecture search for speech deepfake and spoofing detection," in *ASVSPOOF 2021, Automatic Speaker Verification and Spoofing Countermeasures Challenge*. ISCA, 2021, pp. 22–28.
- [30] B. Huang, S. Cui, J. Huang, and X. Kang, "Discriminative frequency information learning for end-to-end speech anti-spoofing," *IEEE Signal Processing Letters*, vol. 30, pp. 185–189, 2023.
- [31] J. Xue, C. Fan, Z. Lv, J. Tao, J. Yi, C. Zheng, Z. Wen, M. Yuan, and S. Shao, "Audio deepfake detection based on a combination of f0 information and real plus imaginary spectrogram features," in *Proceedings of the 1st International Workshop on Deepfake Detection for Audio Multimedia*, 2022, pp. 19–26.
- [32] L. Li, T. Lu, X. Ma, M. Yuan, and D. Wan, "Voice deepfake detection using the self-supervised pre-training model hubert," *Applied Sciences*, vol. 13, no. 14, p. 8488, 2023.
- [33] X. Wang and J. Yamagishi, "Investigating Self-Supervised Front Ends for Speech Spoofing Countermeasures," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2022)*, 2022, pp. 100–106.
- [34] H. Tak, M. Todisco, X. Wang, J. weon Jung, J. Yamagishi, and N. Evans, "Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2022)*, 2022, pp. 112–119.
- [35] Y. Eom, Y. Lee, J. S. Um, and H. R. Kim, "Anti-Spoofing Using Transfer Learning with Variational Information Bottleneck," in *Proc. Interspeech 2022*, 2022, pp. 3568–3572.