

MOS-FAD: IMPROVING FAKE AUDIO DETECTION VIA AUTOMATIC MEAN OPINION SCORE PREDICTION

Wangjin Zhou¹, Zhengdong Yang¹, Chenhui Chu¹, Sheng Li², Raj Dabre², Yi Zhao³, Kawahara Tatsuya¹

¹Graduate School of Informatics, Kyoto University, Sakyo-ku, Kyoto, Japan

²National Institute of Information and Communications Technology (NICT), Kyoto, Japan

³Kuaishou Technology, Beijing, China

ABSTRACT

Automatic Mean Opinion Score (MOS) prediction is employed to evaluate the quality of synthetic speech. This study extends the application of predicted MOS to the task of Fake Audio Detection (FAD) as we expect that MOS can be used to assess how close synthesized speech is to the natural human voice. We propose MOS-FAD, where MOS can be leveraged at two key points in FAD: training data selection and model fusion. In training data selection, we demonstrate that MOS enables effective filtering of samples from unbalanced datasets. In the model fusion, our results demonstrate that incorporating MOS as a gating mechanism in FAD model fusion enhances overall performance.

Index Terms— MOS prediction, self-supervised learned (SSL) model, model fusion, fake audio detection (FAD)

1. INTRODUCTION

Recent developments in text-to-speech (TTS) [1, 2, 3, 4] and voice conversion (VC) [5] have made it feasible to create a human-like speech, which could be misused for malevolent purposes such as spoofing attacks. As a countermeasure, there is a growing emphasis on fake audio detection (FAD), which aims to distinguish fake audio from real audio [6, 7, 8, 9, 10, 11].

Efforts for FAD tasks have focused on enhancing the acoustic front end to improve the effectiveness of FAD systems. Research has shown that well-designed acoustic features can effectively distinguish fake audio from real audio [12, 13, 14, 15]. Concurrently, there are also studies concentrated on designing effective classification models to distinguish between real and fake audios [16, 17, 18, 19, 20, 21, 22, 14, 15].

Unlike the approach mentioned earlier, our core idea is to improve fake audio detection by leveraging external knowledge. Considering that MOS (Mean Opinion Score) is used to assess the quality of synthesized speech, MOS will be promising tool for detecting fake audio. Moreover, compared to the binary classification of fake audio detection, MOS offers a finer-grained evaluation criterion that can yield additional information.

In this paper, we adopt the self-supervised learning (SSL) models [23], widely studied and performed well in various speech classification tasks [24], as the foundational architecture for constructing the FAD classifier. To ensure the robustness of our final results and reduce dependence on the performance of a single SSL model, we conduct experiments with various SSL models. Building upon this SSL-based architecture, we explore the integration of MOS scores into FAD tasks through various methods. This investigation ultimately leads to the development of our innovative MOS-FAD framework, where MOS scores are used as gating mechanisms in model

construction. This framework has achieved state-of-the-art performance.

We also extend the use of MOS scores to guide the selection of training data. This approach addresses two key challenges: Firstly, it alleviates the long-tail issue associated with imbalanced training data. Secondly, by utilizing MOS scores for data selection, the model can be trained on samples that pose a greater challenge in distinguishing between real and fake audio, ultimately enhancing the model's classification capabilities.

2. RELATED WORK

2.1. Automatic MOS Predictor

The Mean Opinion Score (MOS) [25] is a widely used subjective quality evaluation criterion for synthesized speech. In MOS assessments, listeners rate speech sample quality on a scale ranging from 1 (poor) to 5 (excellent). The average of these ratings is used to gauge overall speech quality. However, traditional MOS evaluations can be expensive and time-consuming.

Given these challenges, there has been a growing interest in training automatic MOS predictors. Notably, the VoiceMOS Challenge [26] has played a pivotal role in advancing this field in recent years.

The state-of-the-art automatic MOS predictors employ neural network-based frameworks and use large-scale synthetic speech data for training a model to predict MOS score [27, 28, 29, 30, 31, 32]. Another solution is to introduce self-supervised learning (SSL) models [22, 33, 34, 33, 35], which have already been demonstrated in the MOS prediction task. Since the choice of data and various model configurations used for pre-training will affect the performance, fusing diverse models with different architectures or training strategies [32, 31] could help achieve better results on the MOS prediction task.

2.2. Fake Audio Detection

The current FAD research focuses on acoustic front-end and classification models.

Research has shown that well-designed acoustic features can effectively distinguish fake audio from real audio. The traditional Mel frequency spectrum coefficients (MFCC) were improved by using constant Q-transform instead of the short-time Fourier transform [12]. The linear frequency spectrum coefficients (LFCC) [13] was proposed to replace the Mel-scale filters with linear filters, making it more focused on high-frequency band features. Wav2vec pre-trained models [14, 15] were used as the feature extractors, which obtain more robust acoustic features.

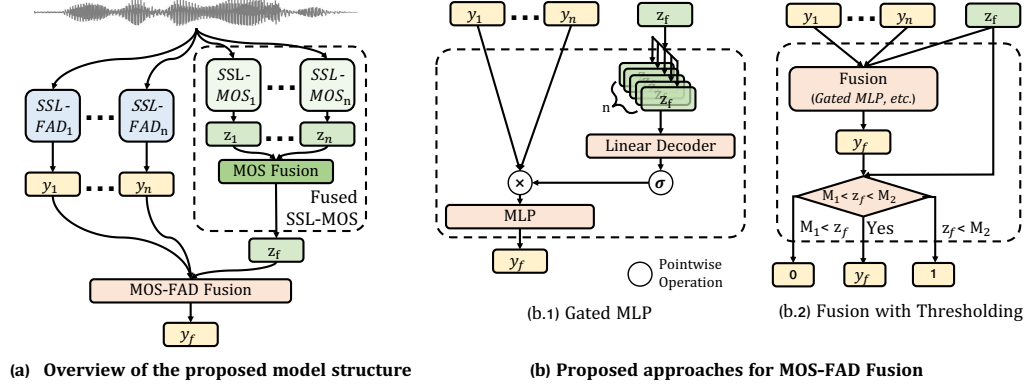


Fig. 1. Proposed model structure.

An effective classification model can also help distinguish between real and fake audio. With the development of deep neural networks, the classifier has been upgraded from convolutional neural networks (CNN) [16, 17], light convolution neural network (LCNN) [18], residual network (ResNet) [19], differentiable architecture search (DARTS) [20, 21] to state-of-the-art Transformer-based models, e.g., Wav2vec2.0 [22, 14, 15].

3. PROPOSED APPROACH

As illustrated in Fig. 1(a), our proposed method consists of three components: SSL-FADs, Fused SSL-MOS sub-system, and MOS-FAD Fusion. In this section, we introduce each of them, respectively.

3.1. SSL-FAD

SSL-FADs are SSL-based models that predict FAD scores. This work adopts the idea in [30], which adds a mean pooling layer and a fully connected layer after the feature extractor of an SSL model to construct an SSL-FAD model. In this paper, we adopt 7 different SSL models (Wav2Vec 2.0 Base, Wav2Vec 2.0 Large, Wav2Vec 2.0 (LV-60), HuBERT Base, WavLM Base, WavLM Base+ and WavLM Large) to construct 7 SSL-FADs, as shown in Table 2. By n SSL-FADs, FAD scores $\mathcal{Y} = \{y_i \mid y_i \in [0, 1], i = 1, 2, \dots, n\}$ are computed.

3.2. Fused SSL-MOS

The fused SSL-MOS is an automatic MOS predictor proposed in [31], which consists of several SSL-MOS models and a MOS fusion. Similar to SSL-FAD, SSL-MOS is constructed by adding a mean pooling layer and a fully connected layer after the feature extractor of the SSL model. The SSL models employed in constructing SSL-MOSs are identical to those used in constructing SSL-FAD. The MOS model fuser is a 2-layer neural network consisting of a fully connected layer without a bias for capturing the weighted information, and a linear function for obtaining the residual information between the ground truth and the predicted scores after a fully connected layer. The MOS fusion takes the MOS scores $\mathcal{Z} = \{z_i \mid z_i \in [0, 5], i = 1, 2, \dots, n\}$ predicted by n SSL-MOS models as inputs and outputs the final MOS score $z_f \in [0, 5]$.

3.3. MOS-FAD Fusion

The MOS-FAD fusion performs the fusion of FAD scores predicted by the SSL-FADs and the MOS scores predicted by the Fused SSL-MOS. It outputs the prediction $y_f \in [0, 1]$ for the FAD task.

We conducted experiments using both Multi-Layer Perceptron (MLP) and LightGBM [36] as potential candidates for MOS-FAD. Through a comparative analysis of their performance, we concluded that MOS can be effectively employed as a gating mechanism within the fusion model. Building upon this insight, we introduced Gated MLP, as depicted in Fig.1(b.1). Moreover, to further leverage the prior information from MOS scores, the Fusion with thresholding is proposed, as shown in Fig.1(b.2).

a. MLP: In our experiment, the MLP is a two-layer neural network with one sigmoid activation function and three hidden neurons. The linear transformation in the first layer is performed without a bias.

b. LightGBM: Light Gradient Boosting Machine (LightGBM) [36] is an ensemble model that trains a series of decision trees sequentially in a leaf-wise fashion and combines them. During sequential training, the decision trees use the error from the previous tree to adjust their learning and eventually minimize the loss function. Therefore, weak learners of the decision trees can be combined as a high-performance model.

c. Gated MLP: Instead of feeding all the predicted scores (FAD scores and MOS scores) to the MLP, we design a special gate layer for the predicted MOS score. The gates control how much information from the FAD scores should be fed to the later MLP. The first step is to decode the MOS scores into a tensor of the same length as the FAD scores using a Linear Decoder layer. Then, the values in this tensor are normalized to the range of 0 to 1 using a *Sigmoid* activation function. These outputs are treated as gate g . Then, predicted FAD scores $y_1, y_2 \dots$ are pointwisely multiplied by g , and the results are used as inputs to the MLP. The MLP component in Gated MLP remains consistent with the description provided in 3.3.a.

d. Fusion with Thresholding: Building upon the definition of MOS scores, we introduce a straightforward screening mechanism named as ‘fusion with thresholding’. The underlying concept is to categorize samples with excessively low MOS scores as fake audio and those with excessively high MOS scores as genuine audio. We assign a score of 0 to samples with MOS scores below the threshold M_1 , designating them as fake speech, and assign a score of 1 to samples with MOS scores above the threshold M_2 as real speech. For the remaining samples, FAD scores are determined based on the outcomes

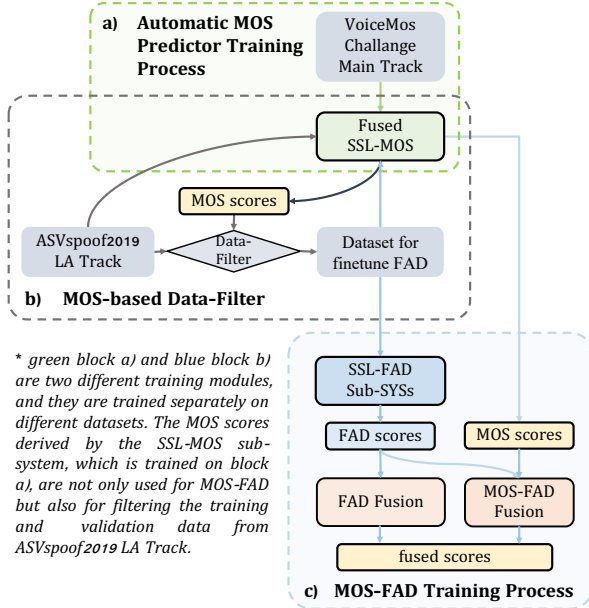


Fig. 2. Training process of our proposed model.

of specific fusion models. It is worth highlighting that this fusion with thresholding approach is applicable to any FAD model. In our experiments, all fake speech of training and validation datasets received MOS scores below 4.0, while 98.81% of real speech received MOS scores above 2.5. Consequently, we set the threshold values as $M_1 = 2.5$ and $M_2 = 4.0$.

4. EXPERIMENT

4.1. Datasets

a. VoiceMOS dataset: We used the VoiceMOS Challenge 2022 [26] main track dataset for training automatic MOS predictor. The main track data contains only English and consists of 4, 974 samples for training, 1, 066 samples for validation, and 1, 066 examples for testing. We quantized the main track’s MOS scores ranging from 1 to 5 into 33 intervals with a step of 0.125.

b. ASVspoof dataset: ASVspoof datasets were leveraged for the FAD task. First, we utilized the training and validation datasets from the ASVspoof2019 [9] LA dataset to train our proposed framework. Next, we employed the ASVspoof2021 challenge [8] evaluation set of the DF track to evaluate the effectiveness of our approach. This DF evaluation set comprises a total of 611, 829 samples. It shows audio coding and compression artifacts, with approximately 600k of audio processing with various commercial audio codecs.

4.2. Experimental Settings

All neural network models were trained using a fixed learning rate of 0.001, Stochastic Gradient Descent as the optimizer, and *CrossEntropy* as the loss function. Training would stop if there was no decrease in the validation loss for 20 consecutive epochs. During the training of LightGBM, the following parameter configuration was employed: {'objective': 'binary', 'metric': 'auc', 'num_leaves': 16, 'max_bin': 25, 'max_depth': 4, 'learning_rate': 0.1}.




Fig. 3. The MOS score distributions in ASVspoof2019 LA track.

Table 1. The number of samples of ASVspoof2019 LA Track with or without the data-filter. With the data-filter we select data by choosing data with MOS scores ranging between 3.0 and 4.0.

	Training Dataset			Validation Dataset		
	Total	Real	Fake	Total	Real	Fake
w/o filter	25,380	2,580	22,800	22,438	2,548	22,438
w/ filter	5,034	2,568	2,466	5,198	2,533	2665

4.3. Training Process

Fig. 2 shows the training details of our proposed approach. The training process can be divided into two independent modules, the Automatic MOS Predictor Training Process and the MOS-FAD Training Process, and an additional data processing module, the MOS-based data-filter, for the FAD task training dataset.

a. Automatic MOS Predictor Training Process: As depicted in the green block in Fig. 2, we employed the MOS fusion model, as described in Section 3.2, to generate MOS scores. This model undergoes training using the VoiceMos Challenge 2022 Main Track dataset. It is important to note that this training process is conducted separately from the MOS-FAD training process, and the training specifics for SSL-MOS closely adhere to those detailed in [31].

b. MOS-based Data-Filter: The MOS-based data-filter serves as a selector for training FAD task models, with two primary objectives. Firstly, it aims to alleviate the issue of sample imbalance, as evident in Table 1, where the original training set exhibits a significant imbalance, with nearly ten times more fake speeches than real ones. Secondly, it enhances model training effectiveness, as illustrated in Fig. 3. The distribution of MOS scores in the LA track dataset suggests that samples with MOS scores below 2.5 can be confidently categorized as fake, while those with scores exceeding 4.0 can be considered as real. Consequently, the FAD model only requires training on data that falls within the range of uncertainty (overlap part in the MOS distribution graph) regarding its authenticity. In summary, to achieve a balanced sample distribution and enhance training effectiveness, we opted for samples from the ASVspoof2019 LA dataset with MOS scores between 3.0 and 4.0 as the new training and validation datasets for the FAD task. As depicted in Table 1, this filtering approach results in an approximate balance between real and fake speeches.

c. MOS-FAD Training Process: We used the filtered dataset to finetune SSL-FADs and train their fusions (FAD Fusion and MOS-FAD Fusion) for FAD tasks (as shown in Fig. 2). The MOS-FAD Fusion in Fig. 2 represents a model fusion approach that combines both MOS scores and FAD scores. Concurrently, The FAD Fusion in Fig. 2 serves as a control group, enabling us to evaluate the importance of incorporating MOS scores into the fusion model.

Table 2. EER on 7 FAD scores of ASVspoof2021 DF Eval Track by each SSL-FAD which trained with or without MOS Data-Filter.

SSL Base	DF Eval	
	w/ filter	w/o filter
W2V Base	0.2514	0.5510
HuBERT Base	0.2353	0.4449
W2V large	0.1891	0.4794
W2V Large(LV-6.0)	0.2395	0.3776
WavLM Base	0.2167	0.5126
WavLM Base+	0.2088	0.5310
WavLM Large	0.2453	0.5881

4.4. Inference

To assess the authenticity of samples in the ASVspoof2021 DF evaluation set, we first employ the Fused SSL-MOS system to generate its MOS scores. Concurrently, we use individual SSL-FAD systems to derive the corresponding FAD scores. Subsequently, the MOS-FAD Fusion process combines these FAD and MOS scores to yield a unified score for each sample. It’s important to note that a lower fused score indicates a higher likelihood that the sample is fake.

5. RESULTS AND ANALYSIS

5.1. Effectiveness of MOS-based Data-Filter

We conducted experiments by training SSL-FADs directly using the original training dataset and compared their performance with models trained on the filtered dataset, as displayed in Table 2. The results demonstrate that models trained after data filtering outperform those trained on the unfiltered dataset, indicating the effectiveness of filtered data in training.

5.2. Effect of MOS in Fusion

Table 3 shows the performance of different fusion models. Among them, the ‘Score Fusion’ investigates how to combine 7 different FAD scores and 1 of MOS scores, while the ‘Embedding Fusion’ validates the effectiveness of ‘Score Fusion’.

a. Score Fusion: Only using FAD scores, both FAD Fusion methods fail to outperform the Wav2Vec Large-based SSL-FAD model. This could be attributed to the high correlations among FAD scores, indicating multicollinearity, which makes it challenging to improve performance by score fusion.

In contrast, within the MOS-FAD Fusions that incorporate MOS scores, LightGBM outperforms MLP and any SSL-FADs. Considering LightGBM’s tree-like structure, we reason that MOS plays a gating role in the fusion process. Building on this hypothesis, we introduce the Gated MLP, which outperforms LightGBM, providing an evidence of MOS as a gating mechanism.

For further improvement, we introduced a more rigorous gate filtering mechanism named ‘fusion with thresholding,’ drawing upon prior MOS knowledge. Among MOS-FAD Fusions, the adoption of fusion with thresholding yields performance improvements. Notably, both LightGBM with thresholding and Gated MLP with thresholding surpass the previous SOTA performance. Remarkably, Gated MLP with thresholding demonstrates an impressive 13.6% reduction in EER.

b. Embedding Fusion: We substituted the MOS and FAD scores with embeddings generated by the SSL model, except for the MOS

Table 3. EER of each fusion on ASVspoof2021 DF Eval Track. Scores stand for the fusion of FAD and MOS scores. Emb stands for the fusion of FAD and MOS embeddings. † is significantly better than the former SOTA when significance level p is set at 0.01.

Model Type	Model	DF Eval	
		Score	Emb
FAD Fusion	MLP	0.2110	0.2049
	LightGBM	0.2221	\
MOS-FAD Fusion	MLP	0.3160	0.1763
	LightGBM	0.1732	\
	Gated MLP	0.1661	0.1890
	MLP + threshold	0.2921	0.1509†
	LightGBM + threshold	0.1476†	\
	Gated MLP + threshold	0.1351†	0.1509†
Former SOTA in ASVspoof 2021 (T23 in [8])		0.1564	

z_f in ‘fusion with thresholding.’ In the MLP approach of FAD Fusion, all embeddings from SSL-FADs are concatenated into a 6,144-dimension tensor, which is then input into the MLP. In the MLP approach of MOS-FAD Fusion, embeddings from both SSL-FAD and SSL-MOS are concatenated to create a 12,288-dimension tensor as input for the MLP. In the case of Gated MLP, the concatenated 6,144-dimension embeddings, generated by SSL-MOS, undergo a linear layer followed by a Sigmoid activation layer to produce gate units of 6,144 dimensions. These gate units control how much information from the concatenated 6,144-dimensional embeddings, produced by SSL-FADs, is passed into the MLP. Notably, the number of hidden units in the MLP described in this section is half of the input dimension.

The results show that the embeddings-based MLP without MOS data failed to surpass the performance of that with MOS data and the best SSL-FAD model, highlighting the necessity of incorporating MOS information. Furthermore, embeddings-based fusions with thresholding outperform their respective baseline models, which further validates the appropriateness of MOS as the gating mechanism in fusion.

Finally, when the significance level p is set to 0.01, the best performance of ‘Score Fusion’ significantly outperforms that of ‘Embeddings Fusion.’ A possible reason is that the gate logic in score fusion serves as a more powerful mechanism for combinations than the combinations learned through the model in embedding fusion. This confirms the effectiveness of ‘Score Fusion.’

6. CONCLUSION

In this paper, we have proposed MOS-FAD, a robust approach for FAD via automatic MOS predictor. We used the MOS-based data-filter to resample the training and validation dataset, which led to effective training for SSL-FADs. Additionally, our results have shown that utilizing MOS as a gating mechanism in the fusion with FAD scores improves performance, based on which we proposed a novel fusion model, Gated MLP with thresholding. This MOS-FAD fusion method, significantly decreases the EER by 13.6%, compared with the best system reported so far. Given that the misuse of fake audio generation can potentially have negative ramifications on society, we hope that this study will encourage further research into this rapidly evolving and crucial field and help circumvent any problems.

7. REFERENCES

- [1] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, and et al., “Tacotron: Towards end-to-end speech synthesis,” in *Proc. INTERSPEECH*, 2017.
- [2] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, “Fast-speech: Fast, robust and controllable text to speech,” in *Proc. NeurIPS*, 2019.
- [3] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, “FastSpeech 2: Fast and high-quality end-to-end text to speech,” in *Proc. ICLR*, 2020.
- [4] C. Yu, H. Lu, N. Hu, M. Yu, C. Weng, K. Xu, P. Liu, D. Tuo, S. Kang, G. Lei, and et al., “DurIAN: Duration informed attention network for multimodal synthesis,” in *Proc. INTERSPEECH*, 2020.
- [5] B. Sisman, J. Yamagishi, S. King, and H. Li, “An overview of voice conversion and its challenges: From statistical modeling to deep learning,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2020.
- [6] J. Yi, Y. Bai, J. Tao, H. Ma, Z. Tian, C. Wang, T. Wang, and R. Fu, “Half-Truth: A Partially Fake Audio Detection Dataset,” in *Proc. INTERSPEECH*, 2021, pp. 1654–1658.
- [7] H. Ma, J. Yi, J. Tao, Y. Bai, Z. Tian, and C. Wang, “Continual Learning for Fake Audio Detection,” in *Proc. INTERSPEECH*, 2021, pp. 886–890.
- [8] J. Yamagishi, X. Wang, M. Todisco, M. Sahidullah, J. Patino, A. Nautsch, X. Liu, K. Lee, T. Kinnunen, N. Evans, and et al., “ASvspoof 2021: accelerating progress in spoofed and deepfake speech detection,” in *Proc. ASvspoof Workshop*, 2021.
- [9] M. Todisco, X. Wang, V. Vestman, M. Sahidullah, H. Delgado, A. Nautsch, J. Yamagishi, N. Evans, T. Kinnunen, and K. Lee, “ASvspoof 2019: Future Horizons in Spoofed and Fake Audio Detection,” in *Proc. INTERSPEECH*, 2019, pp. 1008–1012.
- [10] T. Kinnunen, M. Sahidullah, H. Delgado, M. Todisco, N. Evans, J. Yamagishi, and K. Lee, “The ASvspoof 2017 Challenge: Assessing the Limits of Replay Spoofing Attack Detection,” in *Proc. INTERSPEECH*, 2017, pp. 2–6.
- [11] J. Yi, R. Fu, J. Tao, S. Nie, H. Ma, C. Wang, T. Wang, Z. Tian, Y. Bai, C. Fan, and et al., “Add 2022: the first audio deep synthesis detection challenge,” in *Proc. IEEE-ICASSP*, 2022, pp. 9216–9220.
- [12] 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.
- [13] M. Sahidullah, T. Kinnunen, and C. Hanilçi, “A comparison of features for synthetic speech detection,” in *Proc. INTERSPEECH*, 2015, pp. 2087–2091.
- [14] Y. Xie, Z. Zhang, and Y. Yang, “Siamese network with wav2vec feature for spoofing speech detection,” in *Proc. INTERSPEECH*, 2021, pp. 4269–4273.
- [15] J. Martín-Doñas and A. Álvarez, “The vicomtech audio deepfake detection system based on wav2vec2 for the 2022 add challenge,” in *Proc. IEEE-ICASSP*, 2022, pp. 9241–9245.
- [16] C. Lai, N. Chen, J. Villalba, and N. Dehak, “ASSERT: Anti-Spoofing with Squeeze-Excitation and Residual Networks,” in *Proc. INTERSPEECH*, 2019, pp. 1013–1017.
- [17] M. Alzantot, Z. Wang, and M. Srivastava, “Deep Residual Neural Networks for Audio Spoofing Detection,” in *Proc. INTERSPEECH*, 2019, pp. 1078–1082.
- [18] G. Lavrentyeva, S. Novoselov, E. Malykh, A. Kozlov, O. Kudashev, and V. Shchemelinin, “Audio replay attack detection with deep learning frameworks,” in *Proc. INTERSPEECH*, 2017, pp. 82–86.
- [19] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016, pp. 770–778.
- [20] H. Liu, K. Simonyan, and Y. Yang, “DARTS: differentiable architecture search,” in *Proc. ICLR*, 2019.
- [21] W. Ge, M. Panariello, J. Patino, M. Todisco, and N. Evans, “Partially-connected differentiable architecture search for deepfake and spoofing detection,” in *Proc. INTERSPEECH*, 2021.
- [22] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, 2020, vol. 33, pp. 12449–12460.
- [23] A. Oord, Y. Li, and O. Vinyals, “Representation learning with contrastive predictive coding,” in *Proc. NeurIPS*, 2018.
- [24] S. Yang and et al., “SUPERB: Speech processing universal performance benchmark,” in *Proc. INTERSPEECH*, 2021, pp. 1194–1198.
- [25] R. Huang and Y. Li, “MOS: towards scaling out-of-distribution detection for large semantic space,” in *Proc. CVPR*, 2021, pp. 8710–8719.
- [26] W. Huang, E. Cooper, Y. Tsao, H. Wang, T. Toda, and J. Yamagishi, “The voicemos challenge 2022,” in *Proc. INTERSPEECH*, 2022.
- [27] C. Lo, S. Fu, W. Huang, X. Wang, J. Yamagishi, Y. Tsao, and H. Wang, “Mosnet: Deep learning based objective assessment for voice conversion,” in *Proc. INTERSPEECH*, 2019.
- [28] Y. Choi, Y. Jung, and H. Kim, “Neural mos prediction for synthesized speech using multi-task learning with spoofing detection and spoofing type classification,” in *Proc. IEEE-SLT*, 2021, pp. 462–469.
- [29] W. Huang, E. Cooper, J. Yamagishi, and T. Toda, “LDNet: Unified listener dependent modeling in mos prediction for synthetic speech,” in *Proc. IEEE-ICASSP*, 2022.
- [30] E. Cooper, W. Huang, T. Toda, and J. Yamagishi, “Generalization ability of mos prediction networks,” *arXiv preprint arXiv:2110.02635*, 2021.
- [31] Z. Yang, W. Zhou, C. Chu, S. Li, R. Dabre, R. Rubino, and Y. Zhao, “Fusion of self-supervised learned models for mos prediction,” in *Proc. INTERSPEECH*, 2022.
- [32] T. Saeki, D. Xin, W. Nakata, T. Koriyama, S. Takamichi, and H. Saruwatari, “Utmos: Utokyo-sarulab system for voicemos challenge 2022,” in *Proc. INTERSPEECH*, 2022.
- [33] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, and et al., “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, pp. 1505–1518, 2021.
- [34] W. Hsu, B. Bolte, Y. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [35] A. Baevski, W. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, “Data2vec: A general framework for self-supervised learning in speech, vision and language,” in *Proc. ICML*, 2022.
- [36] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T. Liu, “Lightgbm: A highly efficient gradient boosting decision tree,” in *Proc. NeurIPS*, 2017, vol. 30.