



# Representation Selective Self-distillation and wav2vec 2.0 Feature Exploration for Spoof-aware Speaker Verification

Jin Woo Lee<sup>1</sup>, Eungbeom Kim<sup>2</sup>, Junghyun Koo<sup>1</sup>, and Kyogu Lee<sup>1,2,3</sup>

<sup>1</sup> Department of Intelligence and Information, Seoul National University

<sup>2</sup> Interdisciplinary Program in Artificial Intelligence, Seoul National University

<sup>3</sup> AI Institute, Seoul National University

{jinwlee, eb.kim, dg22302, kglee}@snu.ac.kr

## Abstract

Text-to-speech and voice conversion studies are constantly improving to the extent where they can produce synthetic speech almost indistinguishable from bona fide human speech. In this regard, the importance of countermeasures (CM) against synthetic voice attacks of the automatic speaker verification (ASV) systems emerges. Nonetheless, most end-to-end spoofing detection networks are black-box systems, and the answer to what is an effective representation for finding artifacts remains veiled. In this paper, we examine which feature space can effectively represent synthetic artifacts using wav2vec 2.0, and study which architecture can effectively utilize the space. Our study allows us to analyze which attribute of speech signals is advantageous for the CM systems. The proposed CM system achieved 0.31% equal error rate (EER) on ASVspoof 2019 LA evaluation set for the spoof detection task. We further propose a simple yet effective spoofing aware speaker verification (SASV) method, which takes advantage of the disentangled representations from our countermeasure system. Evaluation performed with the SASV Challenge 2022 database show 1.08% of SASV EER. Quantitative analysis shows that using the explored feature space of wav2vec 2.0 advantages both spoofing CM and SASV.

**Index Terms:** speech anti-spoofing, spoofing aware speaker verification

## 1. Introduction

Automatic speaker verification (ASV) system aims to automatically verify the identity of individual speakers for given utterances. To verify that the input sample is directly uttered by the target speaker, it is essential that ASV systems intrinsically reject any kind of spurious attempts, e.g., non-target utterances, synthetic speeches, or converted voices [1]. Although recent ASV systems have shown prominent results, such as achieving an equal error rates (EERs) less than 1% in the *in-the-wild* scenarios [2, 3], there is room for improvement in being prepared for spoofing attacks [4].

Text-to-speech (TTS) [5, 6, 7] and voice conversion (VC) [8, 9, 10] studies have advanced to the extent that they could produce perceptually almost indistinguishable from bona fide human speech in terms of naturalness and similarity. In order to deceive ASV systems, attackers aim to precisely mimic the target speaker's voice identity through replayed recordings, synthetic speeches, or cloned voices. From this concern, various studies have been proposed to develop the ASV system to have robust reliability for spoofing [4].

Early anti-spoofing studies were developed based on hand-crafted features such as spectral centroid frequency coefficients [11], or constant-Q cepstral coefficients [12]. However, it has

empirically been shown that spoofing artifacts lie at the sub-band level [11, 13], and that front-ends with high spectral resolution in the same band detect artifacts in reliably well [14, 15]. Since cepstral analysis compresses information on the entire spectral domain rather than emphasizing sub-band level features, the importance of research on advanced front-ends for spoof countermeasure networks compared to the conventional cepstral processing has emerged [16].

Recent studies proposed end-to-end anti-spoofing systems that operate on raw waveforms [17, 18, 19]. These end-to-end models can also be viewed as a category of methods that replace the aforementioned feature extraction with neural nets with trainable parameters. Several studies utilized sinc convolution-based network [20] as a front-end of anti-spoofing model [17, 21].

More recent studies investigated the effect of self-supervised front-ends as speech spoofing countermeasures [22, 23]. Self-supervised learning has established itself as a powerful framework to learn general data representations from unlabeled data [24, 25, 26]. Tak *et al.* [22] used wav2vec 2.0 model as front-end of existing AASIST [17] countermeasure network. They used output features of XLS-R [27] to improve the performance of AASIST. Wang *et al.* [23] investigated several self-supervised models as feature extractors of the spoof detection task. However, not many end-to-end anti-spoofing studies analyzed to find clues about the attributes that are important for finding artifacts.

Inspired by [28], where they investigated the most relevant features of each wav2vec 2.0 Transformer layer to the characteristics of speech signals, we examine features extracted from each layer of the wav2vec pre-trained model for the spoofed speech detection task. Contrary to the approaches that studied the use of self-supervised models as front-end of spoofing countermeasures, we seek to find answers to the following questions:

- Which attribute is most related to the features advantageous for spoofing detection?
- How much simpler the back-end model can be behind the improved feature extractor?
- How can the countermeasure network actually help building a secure speaker verification against spoofing?

To answer these questions, we first study feature spaces that, when spoofed speeches are projected, reveals their artifacts more clearly. After that, we examine whether there is no performance degradation by replacing the back-end model with a simple one upon using the feature extractor. We further propose a simple system inspired by representation distillation [29] and self-distillation [30] for spoof-aware speaker verification utilizing the improved countermeasure system.

## 2. Methods

### 2.1. Self-supervised representation for spoofing detection

To explore an effective representation space for spoofing detection, we divide the CM system into two parts, front-end and back-end. We define the front-end by the part that projects the audio input to the feature space, and the back-end by the rest of the network that generates the embeddings for the final classification. We study the feature space for spoofing countermeasure, using a pre-trained wav2vec 2.0 model as a front-end. The architecture of wav2vec 2.0 consists of a convolutional encoder to map raw audio to latent speech representations and stacks of Transformer blocks that build context representations [24].

The wav2vec model is trained to predict speech representations by solving a contrastive task, which requires identifying the quantized latent speech representation of the convolutional encoder for a time step within a set of distractors. Through this self-supervised framework, the model builds context representations by capturing dependencies over the entire sequence of latent representations without any labels. Such context can be correlated with various features including audial features, fluency, pronunciation, or semantic level text features [28]. Also, it has empirically been observed that, in controlled settings of clean and native speech inputs, wav2vec 2.0 model outperforms Mockingjay [26] and BERT [25] on audio and fluency feature representations [28]. Assuming the attacks are attempted using clean native speeches, we fix wav2vec 2.0 pre-trained model and examine whether each Transformer block’s output serves as an effective feature representation for the countermeasure network.

In addition, we experimented with how lightweight the back-end model can be underneath the feature space of each front-end. Through this experiment, we analyze how advantageous each feature extractor is in distinguishing artifacts, even though the constraints for the back-end architecture are considerably relaxed. We adopt AASIST as a baseline of the back-end model, which is equipped with a front-end based on SincNet [20] and a back-end based on heterogeneous stacking graph attention mechanism [17]. Unless otherwise stated, we denote AASIST to mean only the back-end system of [17]. As simple back-end models, we examine a 3-layer multilayer perceptron (MLP), and an attentive statistics pooling (ASP) layer [2] equipped with a fully-connected (FC) layer. The attentive statistics pooling layer calculates the mean and standard deviations of the front-end features and allows the model to select the frames that are deemed to be relevant via the attention mechanism.

### 2.2. Spoof-aware speaker verification

Our SASV system consists of three neural networks: a speaker verification network  $E$  based on ECAPA-TDNN [2], a spoofing countermeasure network  $C$  based on section 2.1, and a newly proposed representation selective self-distillation (RSSD) module based on representation distillation. While we keep  $E$  and  $C$  to be independently trained, RSSD is trained to selectively transform the speaker representations by  $E$  using the countermeasure embeddings of  $C$ .

Let  $\mathbf{e}_e = E(\mathbf{x}_e)$  and  $\mathbf{e}_t = E(\mathbf{x}_t)$  denote speaker embeddings of enrollment and test speaker, respectively. Denote countermeasure embedding of the test utterance by  $\mathbf{c}_t = C(\mathbf{x}_t)$ . We measure spoof-aware speaker similarity by

$$\text{sim}(\mathbf{e}_e, g(f_\theta(\mathbf{c}_t), \mathbf{e}_t)), \quad (1)$$

where  $\text{sim}(\mathbf{a}, \mathbf{b}) = \langle \mathbf{a}, \mathbf{b} \rangle / \|\mathbf{a}\| \|\mathbf{b}\|$  denotes cosine similarity.

RSSD network consists of two modules: a feature transform layer  $f_\theta$  and a gating operator  $g$ . The layer  $f_\theta$  transform the embedding  $\mathbf{c}_t$  of test speech  $\mathbf{x}_t$ , to modulate the speaker embedding of  $\mathbf{x}_t$ . Then, gate operator  $g$  outputs self-distilled representation using the two embeddings  $f_\theta(\mathbf{c}_t)$  and  $\mathbf{e}_t$ . Finally, we measure spoof-aware speaker similarity as equation (1).

During training, our objective function  $\mathcal{L}$  is computed for given triad of labeled utterances  $(\mathbf{x}_e, \mathbf{x}_t, y)$  as

$$\mathcal{L}(\mathbf{x}_e, \mathbf{x}_t) = \mathbb{1}_{[y=1]} \mathcal{L}_{\text{r-distill}}(\mathbf{x}_t) + \mathbb{1}_{[y \neq 1]} \mathcal{L}_{\text{spooof}}(\mathbf{x}_e, \mathbf{x}_t), \quad (2)$$

where the label  $y = 1$  if  $\mathbf{x}_t$  is bona fide, and  $y = 0$  otherwise.  $\mathcal{L}_{\text{r-distill}}$  denotes representation self-distillation loss, and  $\mathcal{L}_{\text{spooof}}$  denotes spoofing countermeasure loss.

#### 2.2.1. Representation self-distillation loss

RSSD distills the information through representations extracted from the pre-trained speaker verification network. This is the difference from conventional knowledge distillation which distill the information through probabilistic outputs from teacher networks [31]. The representation self-distillation loss  $\mathcal{L}_{\text{r-distill}}$  of bona fide example  $\mathbf{x}$  is defined as

$$\mathcal{L}_{\text{r-distill}}(\mathbf{x}_t) = -\text{sim}(\mathbf{e}_t, g(f_\theta(\mathbf{c}_t), \mathbf{e}_t)), \quad (3)$$

where  $\mathbf{e}_t = E(\mathbf{x}_t)$  denotes representation from the pre-trained speaker verification network, and  $\mathbf{c}_t = C(\mathbf{x}_t)$  denotes representation from the pre-trained countermeasure network. Parameters of the feature transformation layer  $f_\theta(\cdot)$  are optimized to preserve the bona fide speaker representation from the pre-trained network. The transformed countermeasure representation  $f_\theta(C(\cdot))$  determines whether to preserve speaker representation  $E$  using the gating operator  $g$  or not.

#### 2.2.2. Spoofing countermeasure loss

The spoofing loss  $\mathcal{L}_{\text{spooof}}$  between spoof example  $\mathbf{x}_t$  and enrollment example  $\mathbf{x}_e$  is defined as

$$\mathcal{L}_{\text{spooof}}(\mathbf{x}_e, \mathbf{x}_t) = \text{sim}(\mathbf{e}_e, g(f_\theta(\mathbf{c}_t), \mathbf{e}_t)), \quad (4)$$

where  $\mathbf{e}_e = E(\mathbf{x}_e)$  denotes speaker representation of the enrollment speaker. In contrast to bona fide examples, the desired representation space should not distill the information of given spoof speaker representations so that transformed representation  $g(f_\theta(\mathbf{c}_t), \mathbf{e}_t)$  drifts apart from original enrollment speaker representation  $\mathbf{e}_e$ . The feature transformation layer  $f_\theta$  changes spoof countermeasure representation  $\mathbf{c}_t$  to displace spoof speaker representation far from enrollment speaker representation.

## 3. Experimental Setup

### 3.1. Implementation details

For all of the experiments, we use ASVspoof 2019 LA database [4]. It consists of bona fide speeches collected from VCTK corpus [32], and spoofed speeches generated using a variety of TTS and VC methods. The training and development sets were constructed using 6 different algorithms, while the evaluation set consists of non-overlapping 13 different methods [33]. The training, development, and evaluation datasets consist of 20, 20, and 67 unique speakers, respectively. For more specific details on the experimental dataset, we refer to [4, 33].

Figure 1: Breakdown of CM-EERs tested using XLSR-53 front-end with ASP back-end, for the spoofing countermeasure task.




Table 1: CM-EERs for spoofing countermeasure networks. All self-supervised (SS) front-ends below consist of 317M number of parameters. ‘Inter.’ indicates whether an intermediate feature is used, other than the final output of Transformer.

System	Front-end	SS	Inter.	EER
RawGAT-ST [21]	SincNet [20]	✗	-	1.06
AASIST-L [17]	SincNet [20]	✗	-	0.99
AASIST [17]	SincNet [20]	✗	-	0.83
LGF [23]	XLSR-53 [34]	✓	✗	1.28
LLGF [23]	W2V-Large [24]	✓	✗	0.86
Ours (MLP)	XLSR-53 [34]	✓	✓	0.80
Ours (AASIST)	XLSR-53 [34]	✓	✓	0.40
Ours (ASP)	XLSR-53 [34]	✓	✓	<b>0.31</b>

### 3.1.1. Feature exploration of spoofing countermeasure

We conduct experiments to replace the front-end of AASIST by XLSR-53 [34]: a wav2vec 2.0 model pre-trained on 56k hours of speech in 53 languages. We use the XLSR-53 Large model implemented in fairseq<sup>1</sup> [35], equipped with 24 Transformer blocks with dimension of 1024. For the back-end model comparison, we use an MLP that consists of 3 FC layers with dimension 1024 and leaky ReLUs between each layer. We also examine the use of ASP layer that pools the feature representation into statistics of dimension 2048. The statistics are linearly projected to represent an embedding of dimension 160.

### 3.1.2. Spoof-aware speaker verification

We use 3 fully connected layers with leaky ReLU activation for  $f_\theta$  and Hadamard product for the gating operator  $g$ . We train RSSD with batch size 32 using Adam optimizer [36] for 20 epochs with the learning rate  $10^{-4}$ . For evaluation, we use a model with the best scores on the development set. We adopt two methods presented in the SASV challenge [1] as our baselines. As with the baselines, we also leverage the pre-trained model of state-of-the-art ASV network ECAPA-TDNN<sup>2</sup>.

## 3.2. Evaluation metric

We use EER (%) as the evaluation metric for every experiment, where we denote the EER of the CM and ASV tasks as CM-EER and SV-EER, respectively. On the SASV task, we identify the EER for the spoofed input as SPF-EER. The performance of RSSD is evaluated using SASV-EER [1], which is a combination of SPF-EER and SV-EER. SASV-EER shows how well a system can distinguish bona fide and target signals from spoofed or (zero-effort) non-target ones.

<sup>1</sup><https://github.com/pytorch/fairseq/tree/main/examples/wav2vec>

<sup>2</sup><https://huggingface.co/speechbrain/spkrec-ecapa-voxceleb>

## 4. Results

### 4.1. Spoofing countermeasure

Table 1 compares EERs of CM systems with different combinations of front-end and back-end modules. For the same back-end system with AASIST, we achieve a 43% performance improvement just by replacing the front-end from SincNet to XLSR-53. Moreover, a simple 3-layer MLP back-end with XLSR-53 front-end performed even better than the original AASIST. Yet, this observation involves a more in-depth analysis than simply replacing the front-end with a self-supervised model. Compared to the best-reported performance of [23], where they used the output of the final Transformer layer of XLSR-53 or wav2vec 2.0 Large model, our approach of selecting an intermediate, 5<sup>th</sup>, layer unveils the effective feature space for CM by outperforming the previous studies. This result suggests the importance of exploring the feature space of each wav2vec 2.0 layer.

#### 4.1.1. Feature exploration

Figure 1 shows EERs of our CM system with XLSR-53 front-end and ASP back-end. Among 24 Transformer layers of XLSR-53, using features from 5<sup>th</sup> layer shows the lowest pooled EER. This observation can be interpreted by connecting with the empirical analysis of [28] that audial features (e.g., local jittering or shimmering) are highly related to the features from 2<sup>nd</sup> out of 12 Transformer layers. Hence, it can be inferred that synthetic artifacts are more highly exposed by the audial features than other characteristics such as fluency, pronunciation, and text semantics, where [28] reported they are best learned from the 6<sup>th</sup> to 9<sup>th</sup> layers.

Our system shows strength against TTS attacks (A07-12, A16) compared to VC attacks (A13-15, A17-19), except for A10. This can be attributed to vocoders that also synthesized spoofed training data or are difficult to produce utterances sufficiently similar to genuine speech in the current setting [33]. A10 was relatively more difficult to figure out regardless of back-end models, which is similar to the report of [33] that the difference between A10 and bona fide utterances is marginally significant. A10 is the only approach to use WaveRNN [37] vocoder.

For VC attacks, we observe that EER can vary depending on the quality of the input source rather than the vocoder. We could easily identify A13 which used a publicly available TTS system as the input source of the VC system, but relatively higher EER was scored for A14 and A15 which are of similar approach but used a commercial TTS synthesis engine. The vocoder of A15 is the same as that used for the training data A01 which was included in the training data for the back-end. A13 and A17 use the same vocoder which is based on waveform filtering, but the EER of A17, which receives genuine human voice as a conversion source, is measured higher.

Table 2: Pooled CM-EERs for CM task. The result of the model with sinc-AASIST is taken from the best EER reported in [17]. All EERs are rounded to the first decimal place.

Back-end	sinc	XLSR-53						
		1 <sup>st</sup>	3 <sup>rd</sup>	5 <sup>th</sup>	7 <sup>th</sup>	9 <sup>th</sup>	13 <sup>th</sup>	17 <sup>th</sup>
MLP	-	2.2	2.1	0.8	1.4	1.0	1.7	2.2
AASIST	0.8	2.9	0.8	0.4	0.5	0.5	0.7	0.9
ASP	1.0	1.0	0.7	0.3	0.5	0.4	0.5	1.0

Table 2 compares CM-EERs for various front-ends. For all choices of back-end models, using 5<sup>th</sup> layer of XLSR-53 achieved the best performance. We observe that a single ASP layer outperforms AASIST, in the case of using XLSR-53 as the front-end. It can be inferred that the attentive pooling of self-supervised representations is effective for spoofing detection.

#### 4.1.2. Back-end model selection for self-supervised front-ends

Scatter plots of t-SNE for CM embeddings are illustrated in Figure 2. For the original AASIST with sinc front-end, we observe that the embeddings corresponding to A17 (gray) and A18 (olive) are entangled with those of bona fide (red) utterances. Considering that the two attack types which recorded the highest EER of AASIST are A17 and A18 [17], this observation can be considered to be consistent with the breakdown EER of original AASIST. Embeddings extracted using XLSR-53 as front-end show better disentanglement of bona fide utterances against A17 and A18.

Above all, ASP with XLSR-53 front-end shows the best disentanglement. Although no information on attack types was included in the training, ASP even distinguishes the embeddings for attack types A09, A11, and A13. A13 is known to deceive CM systems relatively well, even though it is easily identified by humans [33], and we can easily distinguish it down to 0.00% EER. Interestingly, the embeddings corresponding to the VC systems (A17-19) are clustered together (bottom of the XLSR-ASP t-SNE). Similar to the observation of [33], VC systems that converted synthetic speeches (A13-15) tend to be grouped with TTS systems.

Table 2 shows that with help of an appropriate choice of feature extractor, the performance could outperform the existing system even with a simpler back-end system. For all back-end systems, using the output of the 5<sup>th</sup> layer as a feature show the best performance. Also, we find that changing the back-end of AASIST to a 3-layer MLP did not cause significant performance degradation under the appropriate front-end selection.




Figure 2: t-SNE samples of countermeasure embeddings for different types of spoofing attacks in the evaluation set. The caption below each image indicates the front and back-end pair used to extract the embeddings. The scattered embeddings are displayed in different colors depending on the attack type, and the embeddings for the bona fide utterances are marked in red.

Table 3: Pooled EERs with RSSD for SASV task. For the same ASV and CM subsystem, RSSD significantly outperforms the baselines. Improving the CM network further develops RSSD.

System	CM Front	CM Back	SV	SPF	SASV
Baseline 1	SincNet	AASIST	35.32	0.67	19.31
Baseline 2	SincNet	AASIST	11.48	0.78	6.37
RSSD	SincNet	AASIST	1.41	0.76	1.15
RSSD	XLSR-53	AASIST	1.34	0.60	1.11
RSSD	XLSR-53	ASP	1.32	0.59	1.08

However, even with a fewer number of parameters than the rest, ASP effectively improves the countermeasure, by even outperforming the others for most front-ends. It can be inferred that with a suitable choice of feature extractor, attentive pooling operation can outperform other back-ends.

## 4.2. Spoof-aware speaker verification

Table 3 compares our method with the SASV 2022 challenge baselines. By replacing the baseline SASV systems with RSSD using the original AASIST as CM network, we achieve 1.15% SASV-EER, surpassing the baselines by a large margin. We further increase the performance of RSSD by replacing the front-end of the CM network with XLSR-53, and achieve the best performance of 1.08% SASV-EER by adopting ASP as a back-end. It is also worth mentioning that ECAPA-TDNN achieves 1.63% SV-EER for ASV task and the original AASIST achieves 0.83% CM-EER for CM task. Since RSSD outperforms each system by solving SASV task, it can be inferred that RSSD not only distills the representations of the subsystems, but also completes the embedding space to improve the performance of the subtasks.

## 5. Conclusion

This paper explores the use of wav2vec 2.0 as the feature extractor of the spoofing countermeasures. We investigate from which layer of XLSR-53 the features obtained are advantageous for defending spoof attacks. Quantitative analysis shows that using the features of the 5<sup>th</sup> layer of XLSR-53 is effective for spoofing detection, outperforming the existing sinc convolution-based front-end model. We further find that for such feature representations, the back-end system with a simple ASP layer outperformed AASIST and MLP. With the findings of wav2vec 2.0 features for spoofing detection, we propose RSSD for SASV to fully utilize the state-of-the-art models from both ASV and spoofing detection. By selectively performing speaker verification based on self-distillation using spoof countermeasures and speaker embeddings, RSSD achieves SASV-EER of 1.08%. Empirical results show that RSSD takes advantage of the state-of-the-art speaker verification and spoofing detection models with efficient time and memory complexity.

## 6. Acknowledgement

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2022-0-00320, Artificial intelligence research about cross-modal dialogue modeling for one-on-one multi-modal interactions, 90%) and [NO.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University), 10%].

## 7. References

- [1] J.-w. Jung, H. Tak, H.-j. Shim, H.-S. Heo, B.-J. Lee, S.-W. Chung, H.-G. Kang, H.-J. Yu, N. Evans, and T. Kinnunen, "Sasv challenge 2022: A spoofing aware speaker verification challenge evaluation plan," *arXiv preprint arXiv:2201.10283*, 2022.
- [2] B. Desplanques, J. Thienpondt, and K. Demuynck, "Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," *arXiv preprint arXiv:2005.07143*, 2020.
- [3] J. Thienpondt, B. Desplanques, and K. Demuynck, "Integrating frequency translational invariance in tdnns and frequency positional information in 2d resnets to enhance speaker verification," *arXiv preprint arXiv:2104.02370*, 2021.
- [4] 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," *arXiv preprint arXiv:1904.05441*, 2019.
- [5] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [6] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio *et al.*, "Tacotron: Towards end-to-end speech synthesis," *arXiv preprint arXiv:1703.10135*, 2017.
- [7] R. Prenger, R. Valle, and B. Catanzaro, "Waveglow: A flow-based generative network for speech synthesis," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3617–3621.
- [8] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, "Autovc: Zero-shot voice style transfer with only autoencoder loss," in *International Conference on Machine Learning*. PMLR, 2019, pp. 5210–5219.
- [9] S.-H. Lee, J.-H. Kim, H. Chung, and S.-W. Lee, "Voicemixer: Adversarial voice style mixup," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [10] H.-S. Choi, J. Lee, W. Kim, J. Lee, H. Heo, and K. Lee, "Neural analysis and synthesis: Reconstructing speech from self-supervised representations," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [11] M. Sahidullah, T. Kinnunen, and C. Hanilçi, "A comparison of features for synthetic speech detection," 2015.
- [12] M. Todisco, H. Delgado, and N. W. Evans, "A new feature for automatic speaker verification anti-spoofing: Constant q cepstral coefficients," in *Odyssey*, vol. 2016, 2016, pp. 283–290.
- [13] J. Yang, R. K. Das, and H. Li, "Significance of subband features for synthetic speech detection," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 2160–2170, 2019.
- [14] J.-w. Jung, H.-j. Shim, H.-S. Heo, and H.-J. Yu, "Replay attack detection with complementary high-resolution information using end-to-end dnn for the asvspoof 2019 challenge," *arXiv preprint arXiv:1904.10134*, 2019.
- [15] H. Tak, J. Patino, A. Nautsch, N. Evans, and M. Todisco, "An explainability study of the constant q cepstral coefficient spoofing countermeasure for automatic speaker verification," *arXiv preprint arXiv:2004.06422*, 2020.
- [16] H. Tak, J. Patino, A. Nautsch, N. Evans, and M. Todisco, "Spoofing attack detection using the non-linear fusion of sub-band classifiers," *arXiv preprint arXiv:2005.10393*, 2020.
- [17] 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," *arXiv preprint arXiv:2110.01200*, 2021.
- [18] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. Evans, and A. Larcher, "End-to-end anti-spoofing with rawnet2," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6369–6373.
- [19] Y. Ma, Z. Ren, and S. Xu, "Rw-resnet: A novel speech anti-spoofing model using raw waveform," *arXiv preprint arXiv:2108.05684*, 2021.
- [20] M. Ravanelli and Y. Bengio, "Speaker recognition from raw waveform with sincnet," in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 1021–1028.
- [21] 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," *arXiv preprint arXiv:2107.12710*, 2021.
- [22] H. Tak, M. Todisco, X. Wang, J.-w. Jung, J. Yamagishi, and N. Evans, "Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation," *arXiv preprint arXiv:2202.12233*, 2022.
- [23] X. Wang and J. Yamagishi, "Investigating self-supervised front ends for speech spoofing countermeasures," *arXiv preprint arXiv:2111.07725*, 2021.
- [24] 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.
- [25] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [26] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, "Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6419–6423.
- [27] A. Babu, C. Wang, A. Tjandra, K. Lakhota, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino *et al.*, "Xls-r: Self-supervised cross-lingual speech representation learning at scale," *arXiv preprint arXiv:2111.09296*, 2021.
- [28] J. Shah, Y. K. Singla, C. Chen, and R. R. Shah, "What all do audio transformer models hear? probing acoustic representations for language delivery and its structure," *arXiv preprint arXiv:2101.00387*, 2021.
- [29] Y. Tian, D. Krishnan, and P. Isola, "Contrastive representation distillation," *arXiv preprint arXiv:1910.10699*, 2019.
- [30] L. Zhang, J. Song, A. Gao, J. Chen, C. Bao, and K. Ma, "Be your own teacher: Improve the performance of convolutional neural networks via self distillation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 3713–3722.
- [31] G. Hinton, O. Vinyals, J. Dean *et al.*, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, vol. 2, no. 7, 2015.
- [32] J. Yamagishi, C. Veaux, K. MacDonald *et al.*, "Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit (version 0.92)," 2019.
- [33] 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.
- [34] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, "Unsupervised cross-lingual representation learning for speech recognition," *arXiv preprint arXiv:2006.13979*, 2020.
- [35] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "fairseq: A fast, extensible toolkit for sequence modeling," *arXiv preprint arXiv:1904.01038*, 2019.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [37] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. Oord, S. Dieleman, and K. Kavukcuoglu, "Efficient neural audio synthesis," in *International Conference on Machine Learning*. PMLR, 2018, pp. 2410–2419.