



Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation

Hemlata Tak¹, Massimiliano Todisco¹, Xin Wang², Jee-weon Jung³
Junichi Yamagishi² and Nicholas Evans¹

¹EURECOM, France, ²National Institute of Informatics, Japan

³Naver Corporation, South Korea

{tak,todisco,evans}@eurecom.fr, {wangxin,jyamagis}@nii.ac.jp, jeeveon.jung@navercorp.com

Abstract

The performance of spoofing countermeasure systems depends fundamentally upon the use of sufficiently representative training data. With this usually being limited, current solutions typically lack generalisation to attacks encountered in the wild. Strategies to improve reliability in the face of uncontrolled, unpredictable attacks are hence needed. We report in this paper our efforts to use self-supervised learning in the form of a wav2vec 2.0 front-end with fine tuning. Despite initial base representations being learned using only bona fide data and no spoofed data, we obtain the lowest equal error rates reported in the literature for both the ASVspoof 2021 Logical Access and Deepfake databases. When combined with data augmentation, these results correspond to an improvement of almost 90% relative to our baseline system.

1. Introduction

A persisting challenge in the design of spoofing countermeasures (CMs) for automatic speaker verification (ASV) is reliability in the face of diverse, unpredictable attacks [1]. ASV systems can be compromised by attacks belonging to a broad variety of different classes, e.g., converted voice, synthetic speech and replayed recordings. Even within each attack class, there is considerable potential variation, e.g., different algorithms or recording and replay device characteristics. An ideal spoofing detection solution should be robust to all such variation even if, in the wild, it is unpredictable. The acquisition of training data that is representative of spoofing attacks with near-boundless variability is obviously impracticable.

The ASVspoof initiative and challenge series have collected large databases of spoofed and bona fide utterances that are suitable for the training of spoofing countermeasures. To promote the development of generalisable countermeasures, namely detection solutions that cope well in the face of spoofing attacks not previously encountered, assessment is performed with experimental protocols and evaluation data that comprise spoofed utterances generated with a broad variety of different algorithms or techniques. The differences between training, development and evaluation data can lead to substantial differences in detection performance. For the most recent ASVspoof 2021 Logical Access (LA) evaluation [2], the equal error rate (EER) of the best performing baseline solution increased from 0.55% for the development set to 9.26% for the evaluation

set [2]. Submission results show better performance [3–9], but the fundamental gap between performance for development and evaluation data remains, indicating a persisting lack of generalisation.

Given that the training data used for ASVspoof challenges comprises spoofed utterances generated with a modest number of different attack algorithms (six in the case of the ASVspoof 2019 LA database), the lack of generalisation may be difficult to improve upon unless challenge rules are relaxed to allow training using external data. A relaxed training data policy would complicate comparisons between different systems and technology trained using different data – the playing field would no longer be level – though potential improvements to generalisation may make it worthwhile.

The question then is what external training data to use and how to use it? With the drive toward reproducible research, a number of speech synthesis and voice conversion algorithms are now openly available as open source. Additional training data, generated with different attack algorithms, can hence be produced readily. The number of algorithms remains limited, however, and can fundamentally never be fully representative of what can reasonably be expected in the wild. We have hence explored a different approach.

Motivated by (i) its proven application to the learning of general neural representations for a range of different tasks [11–19], (ii) evidence that fine-tuning with modest quantities of labelled data leads to state-of-the-art results, (iii) encouraging, previously reported results for anti-spoofing [20,21] and (iv) the appeal of one-class classification approaches [22], we have explored the use of self-supervised learning to improve generalisation. Our hypothesis is that better representations trained on diverse speech data, even those learned for other tasks and initially using only bona fide data (hence one-class), may help to reduce over-fitting and hence improve reliability and domain-robustness, particularly in the face of previously unseen spoofing attacks. Additionally, we hope that better trained representations are complementary to data augmentations techniques which are already known to improve generalisation [7, 23–27].

The principal contributions of this work are: (i) improved generalisation and domain robustness using a pre-trained, self-supervised speech model with fine-tuning; (ii) additional improvements using data augmentation showing complementary benefits to self-supervised learning; (iii) a new self-attention based aggregation layer which brings complementary improvements.

The first author is supported by the VoicePersonae project funded by the French Agence Nationale de la Recherche (ANR) and the Japan Science and Technology Agency (JST).




Figure 1: AASIST baseline framework reproduced from [10].




Figure 2: Front-end systems: (a) the baseline sinc-layer front-end; (b) the wav2vec 2.0 front-end.

2. Related work

Self-supervised learning (SSL) has attracted considerable attention in recent times. Existing works show that pre-trained models derived using SSL generalise well across a multitude of different tasks when a relatively simple classifier is learned from the resulting representations using only a modest quantity of labelled data [28]. A growing number of self-supervised speech models have been proposed. Examples include contrastive predictive coding (CPC) [16, 29], auto-regressive predictive coding [30], wav2vec [31], HuBERT [32, 33], wav2vec 2.0 [12, 34] and Wavlm [35], with all showing promising results for a variety of different speech processing tasks.

Two particularly popular approaches, HuBERT and wav2vec 2.0, have been applied to automatic speech recognition [12, 13], mispronunciation detection [36, 37], speaker recognition [38, 39] and emotion recognition [40]. The same techniques have been explored in the context of spoofing detection [20, 21]. Xie et al. [20] showed the benefit of using SSL with a Siamese network for spoofing detection. With learned representations coming from their combination, and without comparative assessments using representations derived via alternative means, the specific benefits of SSL are difficult to judge.

Wang et al. [21] compared different SSL based front-ends and back-end architectures and showed the importance of fine-tuning SSL models for spoofing detection. By replacing a linear frequency cepstral coefficient front-end with a wav2vec 2.0 front-end and by fine-tuning, they achieved relative reductions in the equal error rate (EER) of 68% and 79% for the ASVspoof 2021 LA and DF databases respectively. Nonetheless, the EERs remain at 5% and the additional or complementary benefit of data augmentation, which is known to be beneficial in both cases [3, 4, 7], was not explored. Results showed that the wav2vec 2.0 [12] front-end gives better generalised spoofing detection performance than a HuBERT model.

We have hence explored the wav2vec 2.0 XLS-R (0.3B) model¹ [13] in our work. wav2vec 2.0 XLS-R is a large-scale cross-lingually pre-trained model trained on diverse corpora including VoxPopuli data (VP-400K) [41], the multilingual LibriSpeech corpus (MLS) [42], CommonVoice (CV) [43], VoxLingua107 (VL) [44], and BABEL (BBL) [45] datasets. Together, they include speech data in 128 different languages from many different regions of the world. We explored the wav2vec 2.0 front-end with an integrated spectro-temporal graph attention network (AASIST) as a back-end [10]. The latter is described in the next section. Its coupling with a wav2vec 2.0 front-end is described in Section 4.

3. AASIST baseline system

The baseline system is an end-to-end, integrated spectro-temporal graph attention network named AASIST [10], illustrated in Fig. 1. It extracts representations directly from raw waveform inputs. As illustrated in Fig. 2-(a), AASIST uses a sinc convolutional layer based front-end [46]. It is initialised with 70 mel-scaled filters, each with a kernel size of 129 [47]. Through the addition of a channel dimension, the sinc layer output is fed to a post-processing layer and transformed to a spectro-temporal representation. These are fed to a RawNet2-based residual encoder, to learn a higher-level feature map $\mathbf{S} \in \mathbb{R}^{C \times F \times T}$ where C , F and T refer to the number of channels, spectral bins and time frames respectively.

Separate spectral and temporal representations are learned from \mathbf{S} using a max-pooling operation which is applied to the absolute values across either temporal or spectral dimensions in order to construct either a spectral input graph ($\mathcal{G}_s \in \mathbb{R}^{N_s \times d_s}$) or a temporal input graph ($\mathcal{G}_t \in \mathbb{R}^{N_t \times d_t}$). N_s and N_t are the set of graph nodes in spectral and temporal graphs respectively whereas d is the feature dimensionality of each node. Spec-

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

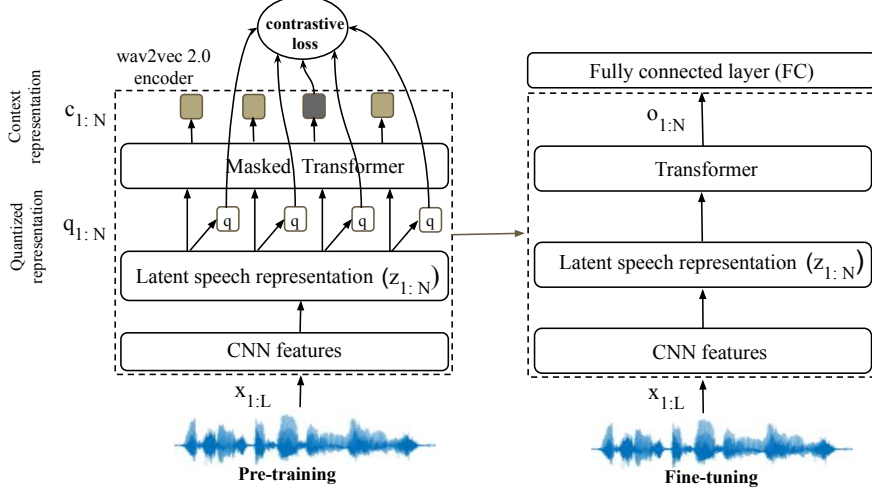


Figure 3: An overview of the pre-training and fine-tuning of the wav2vec 2.0 model, adapted from [13].

tral and temporal graphs \mathcal{G}_s and \mathcal{G}_t are modelled using a pair of parallel graph modules (grey boxes in Fig. 1), each comprising a graph attention network (GAT) [48] and a graph-pooling layer [49]:

$$\mathcal{G}_t = \text{graph_module}(\max_F(\text{abs}(\mathbf{S}))) \quad (1)$$

$$\mathcal{G}_s = \text{graph_module}(\max_T(\text{abs}(\mathbf{S}))) \quad (2)$$

A heterogeneous spectro-temporal graph (\mathcal{G}_{st}) is then formed by combining temporal (\mathcal{G}_t) and spectral (\mathcal{G}_s) graphs using a heterogeneous stacking graph attention layer (HS-GAL). Graph combination enables the concurrent modelling of heterogeneous graph representations with different node dimensions. An HS-GAL contains an attention mechanism modified in order to accommodate graph heterogeneity [50] and an additional stack node [28]. The latter acts to capture the relationships between artefacts spanning temporal and spectral domains. First, \mathcal{G}_t and \mathcal{G}_s are projected using an affine-transformation to another latent space with common dimensionality d_{st} before being fed into the HS-GAL which then constructs a combined heterogeneous graph \mathcal{G}_{st} .

HS-GALs are applied with a max graph operation (MGO) where two branches, each consisting of two HS-GALs, learn to detect different spoofing artefacts in parallel. Each HS-GAL is followed by a graph pooling layer and an element-wise maximum operation is applied to the branch outputs to produce another heterogeneous graph \mathcal{G}_{ST} . HS-GALs in each branch share a common stack node. The stack node of each preceding HS-GAL is fed to the following HS-GAL so that information in both temporal and spectral graphs is preserved.

The readout scheme (penultimate block in Fig. 1) uses node-wise maximum and average operations. The output of the readout layer is formed from the concatenation of five nodes. The first four nodes are derived by applying a maximum and average to spectral nodes (orange) and temporal nodes (blue) in \mathcal{G}_{ST} . The fifth is the copied stack node. A two-class prediction output (bona fide and spoofed) is finally generated using a hidden fully connected layer.

4. Self-supervised front-end

In this section we describe the replacement of the sinc-layer front-end with a wav2vec 2.0 model as illustrated in Fig. 2-(b). We describe both pre-training and fine-tuning to support downstream spoofing detection, both illustrated in Fig. 3.

4.1. Wav2vec 2.0 model

The wav2vec 2.0 pre-trained model is used to extract a sequence of feature representations $o_{1:N}$ from the raw input waveform $x_{1:L}$, where L is the number of samples. As shown in Fig. 3, the wav2vec 2.0 model consists of a convolutional neural network (CNN) and a transformer [28, 51] network. The former converts the input $x_{1:L}$ to a hidden feature sequence $z_{1:N}$ whereas the latter transforms $z_{1:N}$ to output sequence $o_{1:N}$. The ratio between L and N is dictated by the CNN stride of 20 ms (the default setting).

4.2. Pre-training

An illustration of the pre-training procedure following [12] is illustrated to the left in Fig. 3. Latent representations $z_{1:N}$ are quantised to representations $q_{1:N}$. Some portion of the latent representation $z_{1:N}$ is then masked and fed to the transformer which builds new context representations $c_{1:N}$. A contrastive loss for each masked time step n is then computed to measure how well the target q_n can be identified from among a set of distractors (i.e., $q_{n'}$ sampled from the other masked time steps where $n' \neq n$) given the corresponding context vector c_n . All work reported in this paper was performed with the wav2vec 2.0 XLS-R (0.3B) model [13]. We followed the example in the Fairseq project toolkit [52] to extract feature representations from self-supervised wav2vec 2.0 pre-trained model.²

4.3. Fine-tuning

Since pre-training is performed with only bona fide data (with no spoofed data), as per [21], spoofing detection performance is expected to improve with fine tuning using in-domain bona fide and spoofed training data. Our hypothesis is that fine tun-

²<https://github.com/pytorch/fairseq/tree/main/examples/wav2vec>

ing will protect against over-fitting and hence promote better generalisation. For *all* experiments reported in this paper, including those related to the ASVspoof 2021 LA dataset and the ASVspoof 2021 DF dataset, fine-tuning is performed using the ASVspoof 2019 LA training partition only. Whereas the 2021 LA data contains codec and transmission variation and the 2021 DF data contains compression variation, the 2019 LA data used for fine-tuning contains neither. During fine-tuning, the pre-trained wav2vec 2.0 XLS-R model is optimised jointly with the AASIST CM via back-propagation using the ASVspoof 2019 LA training set. This process is described in section 6.3 whereas the fine-tuning procedure is illustrated to the right in Fig. 3. It is performed using a weighted cross entropy objective function to minimize the training loss. In contrast to pre-training, input masking is not applied to hidden features $z_{1:N}$ during fine-tuning. Additionally, we add a fully connected layer on top of the wav2vec 2.0 transformer encoder output $o_{1:N}$ in order to reduce the representation dimension (top-right of Fig. 3).

4.4. Use with AASIST CM

The sinc layer shown in Fig. 2-(a) is replaced with the wav2vec 2.0 front-end shown in Fig. 2-(b). As before, the output $o_{1:N}$ is fed to a RawNet2-based residual encoder which is used to learn higher-level feature representations $\mathbf{S} \in \mathbb{R}^{C \times F \times T}$. Whereas the baseline system extracts temporal and spectral representations \mathbf{t} and \mathbf{f} from \mathbf{S} using a max-pooling operation, a self-attentive aggregation layer (described in Section 5) was found to improve performance of both front-ends. As shown in Fig. 1, temporal and spectral representations are then fed to the AASIST CM model to obtain a two-class prediction (bona fide and spoofed) in the same manner as described in Section 3. A summary of the wav2vec 2.0 front-end and downstream AASIST model configurations is presented in Table 1.

5. Self-attention based aggregation layer

Attention based pooling layers such as self-attentive pooling (SAP) and attentive statistical pooling (ASP) [53] has shown to be beneficial to the aggregation of frame-level features and the extraction of embeddings [54–58] for speaker recognition and verification tasks. We have also found that the introduction of a 2D self-attention based aggregation layer between the front-end and back-end helps to improve spoofing detection performance.

The new self-attentive aggregation layer is used to extract more attentive/relevant spectral and temporal representations. It helps to aggregate and assign higher attention weights through weighted summation to the most discriminative temporal and spectral features. We generate 2-D attention maps (an attention weight matrix) using a 2-D convolutional (conv2d) layer with *one* kernel-size rather than conventional conv1d based attention applied to a single domain. Weights are derived from representations \mathbf{S} processed by a conv2d layer followed by an activation & batch normalization (BN) layer, a 2-D convolutional layer and a softmax function to normalized the weights:

$$W = \text{Softmax}(\text{conv2d}(\text{BN}(\text{SeLU}(\text{conv2d}(\mathbf{S}))))), \quad (3)$$

where $\text{conv2d}(\cdot)$ denotes the 2-D convolution operation with an scaled exponential linear unit $\text{SeLU}(\cdot)$ as the activation function [59], and BN is batch normalisation [60]. Temporal and spectral representations are then extracted from the self-attentive aggregation layer according to:

$$\mathbf{t} = \sum_F \mathbf{S} \odot W, \quad (4)$$

Table 1: The wav2vec 2.0 and AASIST model architecture and configuration. Dimensions refer to (channels, frequency, time). Batch normalisation (BN) and scaled exponential linear unit (SeLU), beneath the dotted line, are applied to residual block outputs.

Layer	Input:64600 samples	Output shape
Data-aug	RawBoost	(64600)
SSL	wav2vec 2.0	(201,1024) (T,F)
front-end	FC (fine-tuning) transpose	(201,128) $o=(128,201)$ (F,T)
post-processing	add channel Maxpool-2D(3) BN & SeLU	(1,128,201) (1,42,67)
Res-block	$2 \times \left\{ \begin{array}{l} \text{Conv-2D}((2,3),1,32) \\ \text{BN \& SeLU} \\ \text{Conv-2D}((2,3),1,32) \end{array} \right\}$	(32,42,67)
Res-block	$4 \times \left\{ \begin{array}{l} \text{Conv-2D}((2,3),1,64) \\ \text{BN \& SeLU} \\ \text{Conv-2D}((2,3),1,64) \end{array} \right\}$	(64,42,67)

Spectral-attention		Temporal-attention
self att. agg. layer (\mathbf{f})=(64, 42)	self att. agg. layer (\mathbf{t})=(64, 67)	
$\mathcal{G}_s = (21(N_s), 64(d_s))$	$\mathcal{G}_t = (33(N_t), 64(d_t))$	
hetero. graph (\mathcal{G}_{st})	HS-GAL	(54(N_{st}),64(d_{st}))
HS-GAL→HS-GAL, stack node (32,26), (32,)	HS-GAL→HS-GAL, stack node (32,26),(32,)	
MGO (\mathcal{G}_{ST})	element-wise max.	(32,26), (32,)
readout	node-wise max. and avg. & concatenation	(160,)
Output	FC(2)	2

$$\mathbf{f} = \sum_T \mathbf{S} \odot W, \quad (5)$$

where \odot denotes element-wise multiplication. $W \in \mathbb{R}^{F \times T}$ is the 2-D attention normalised learnable weight matrix used in the self-attentive aggregation layer to calculate the weighted *sum* of the representation \mathbf{S} across time and frequency.

6. Experimental setup

Described in the following are the databases and metrics used in all reported experimental work, our use of data augmentation and specific, reproducible implementation details.

6.1. Databases and metrics

We used the training and development partitions of the ASVspoof 2019 LA database [61, 62] for training and validation. Evaluation was performed using the ASVspoof 2021 LA and domain mis-matched DF databases [2]. While both are generated from the same VCTK source database³, the LA database

³<http://dx.doi.org/10.7488/ds/1994>

contains codec and transmission variability whereas the DF database contains compression variability in addition to data stemming from sources other than the VCTK database [63]. The 2019 data used for training and validation contains neither. We use two evaluation metrics: the Equal Error Rate (EER) [64] and the Minimum Tandem Detection Cost Function (min t-DCF) [65]. We focus on the first for ease of interpretation and include the second since it is the default metric for ASVspoof challenges.

6.2. Data augmentation

Data augmentation (DA) is already known to reduce overfitting and hence to improve generalisation [3, 4, 6, 7, 27] and is particularly effective in the case of LA scenarios in which there is substantial variability stemming from, e.g., encoding, transmission and acquisition devices [2]. We are interested to determine whether self-supervised learning is complementary to DA. Unlike traditional DA techniques which enlarge the training dataset using *additional*, artificially generated utterances, and using the RawBoost⁴ DA tool [27], we add nuisance variability on-the-fly to the *existing* training data. RawBoost adds variation in the form of: i) linear and non-linear convolutive noise; ii) impulsive signal-dependent additive noise; iii) stationary signal-independent additive noise. Full details are available in [27].

DA is applied using exactly the same configuration and parameters reported in the original work [27]. It shows that a combination of linear and non-linear convolutive noise and impulsive signal-dependent additive noise strategies work best for the LA database. These augmentation strategies suit the convolutive and device related noise sources that characterise telephony applications. In contrast, for the DF database, DA works best using stationary signal-independent additive, randomly coloured noise, which match better the effects of compression [66] applied in generating the DF database. DA experiments and configuration are discussed further in Section 7.3.

6.3. Implementation details

Audio data are cropped or concatenated giving segments of approximately 4 seconds duration (64,600 samples). Graph pooling is applied with an empirically chosen pooling ratio of $k = 0.5$ for spectral and temporal graphs. We used the standard Adam optimiser [67] with a fixed learning rate of 0.0001 for experiments without the wav2vec 2.0 front-end. Since fine-tuning demands high GPU computation, experiments with wav2vec 2.0 were performed with a smaller batch size of 14 and a lower learning rate of 10^{-6} to avoid model over-fitting. As illustrated to right in Fig. 3, the fully connected layer after the wav2vec 2.0 SSL front end used 128 output dimensions. All other hyperparameters are the same for both front-ends which are both jointly optimised with the back-end classifier using back-propagation [68]. As is now common in the related literature [69, 70], we performed each experiments with three runs using different random seeds to initialize the network weights and report the results of the best performing seed and average results. All models were trained for 100 epochs on a single GeForce RTX 3090 GPU and all results are reproducible using open source code⁵ and with the same random seed and GPU environment.

⁴<https://github.com/TakHemlata/RawBoost-antispoofing>

⁵https://github.com/TakHemlata/SSL_Anti-spoofing

Table 2: Pooled EER and pooled min t-DCF results for the ASVspoof 2021 LA database evaluation set, for the sinc-layer and wav2vec 2.0 front-ends. SA refers to the self-attentive aggregation layer whereas DA refers to data augmentation. Results are the best (average) obtained from three runs of each experiment with different random seeds.

front-end	SA	DA	Pooled EER	Pooled min t-DCF
sinc-layer	×	×	11.47 (11.95)	0.5081 (0.5139)
wav2vec 2.0	×	×	6.15 (6.46)	0.3577 (0.3587)
sinc-layer	✓	×	8.73 (11.61)	0.4285 (0.5203)
wav2vec 2.0	✓	×	4.48 (6.15)	0.3094 (0.3482)
sinc-layer	✓	✓	7.65 (7.87)	0.3894 (0.3960)
wav2vec 2.0	✓	✓	0.82 (1.00)	0.2066 (0.2120)

7. Results

We present five sets of experiments. The first is a comparison of each front-end in terms of performance for the ASVspoof 2021 LA database. The second and third assess the complementary benefits coming from the new self-attention based aggregation layer and data augmentation. The fourth is an assessment performed on the ASVspoof 2021 DF database whereas the last is an assessment using a simplified CM solution.

7.1. Front-end comparison

Results for the AASIST baseline with the sinc-layer front-end (Section 3) and the same system with the wav2vec 2.0 front-end (Section 4) are presented in the first two rows of Table 2. These systems use neither the self-attentive aggregation layer nor data augmentation. The baseline EER of 11.47% is high and shows that the system is not robust to the codec and transmission variability which characterises the ASVspoof 2021 LA dataset. The same system using the wav2vec 2.0 front-end delivers an EER of 6.15%. While the relative reduction is almost 46%, the EER is still unacceptably high.

7.2. Self-attentive aggregation layer

Results for the same two front-end variants but using the self-attentive aggregation layer (SA) introduced in Section 5 are presented in rows 3 and 4 of Table 2. In both cases the EER drops substantially, to 8.73% for the sinc-layer front-end and to 4.48% for the wav2vec 2.0 frontend. In this case the wav2vec 2.0 front-end is responsible for a relative improvement of almost 50%.

7.3. Data augmentation

Results for the same two systems, both with the self-attentive aggregation layer (SA), and now also with data augmentation (DA), are shown in rows 5 and 6 of Table 2. DA reduces the EER only marginally from 8.73% to 7.65% in case of the sinc-layer front-end. To verify that this improvement is not due to random factors in neural network training (e.g., different, random initial network weights), we conducted a statistical analysis of the results following [69]. The results⁶ suggest that the improvement is statistically significant and is hence unlikely to be caused by factors other than DA. Its effect is more pronounced when using the wav2vec 2.0 front-end for which the

⁶Available in the form of an appendix available from arXiv at <https://arxiv.org/abs/2202.12233>

Table 3: As for Table 2 except for the ASVspooF DF database, evaluation set. Since there is no ASV in the DF scenario, there are no min t-DCF results. The last two lines show results for an LA-optimised DA configuration.

front-end	SA	DA	Pooled EER
sinc-layer	×	×	21.06 (22.11)
wav2vec 2.0	×	×	7.69 (9.48)
sinc-layer	✓	×	23.22 (25.08)
wav2vec 2.0	✓	×	4.57 (7.70)
sinc-layer	✓	✓	24.42 (25.38)
wav2vec 2.0	✓	✓	2.85 (3.69)
sinc-layer	✓	✓*	20.04 (20.50)
wav2vec 2.0	✓	✓*	6.64 (7.32)

EER decreases from 4.48% to 0.82%. This result is also statistically significant. This result corresponds to a relative improvement of almost 90% when compared to the baseline EER of 7.65%. To the best of our knowledge, this is the lowest EER reported for the ASVspooF 2021 LA database.

7.4. DeepFake results

Results for exactly the same experiments, but for the ASVspooF 2021 DeepFake (DF) database, are shown in Table 3. While neither SA, nor DA improve upon the baseline EER of 21.06%, consistency improvements are obtained for the wav2vec 2.0 front-end for which the EER drops from 7.69% to 2.85% using both SA and DA. This result is also statistically significant. To the best of our knowledge, this is the lowest EER reported for the ASVspooF 2021 DF database.

These results, while determined with the same wav2vec 2.0 front-end used for LA experiments, relate to a DA strategy optimised for the DF database (stationary signal-independent additive randomly coloured noise – see Section 6.2). Results for *exactly* the same setup, using the DA strategy optimised for LA (linear and non-linear convolutive noise and impulsive signal-dependent additive noise) are shown in the last two rows of Table 3. While the EER increases to 6.64%, this is still a competitive result and is 67% lower relative to the result of 20.04% for the sinc-layer front-end. Whereas a component of the DF database originates from the same VCTK database as the entire LA database, other components are sourced from multiple different corpora (Voice Conversion Challenge 2018 and 2020 databases) [2] including spoofed utterances generated with more than 100 different algorithms. With the ASVspooF 2019 LA training data containing neither codec or transmission variability (LA evaluation data), nor compression variability (DF evaluation), results show that the use of better pre-trained models leads to consistent improvements in generalisation, here being previously unseen spoofing attacks. Results for the DF database show that the benefit extends also to the case of domain mismatch.

7.5. Simplified CM solution

The last set of experiments were performed in order to gauge the relative importance of the AASIST and whether the improvements in generalisation are obtained for a simpler CM solu-

*Results using the DA strategy optimised for LA (linear and non-linear convolutive noise and impulsive signal-dependent additive noise).

Table 4: Pooled EER and pooled min t-DCF (LA only) results for the ASVspooF 2021 LA and DF databases (DBs), evaluation sets, using DB-optimised DA and the simplified back-end.

front-end	DA	DB	Pooled EER	Pooled min t-DCF
wav2vec 2.0	✓ [†]	LA	1.19	0.2175
wav2vec 2.0	✓ [†]	DF	4.38	-

tion. We removed the RawNet2-based encoder and replaced AASIST with a simple back-end comprising a max-pooling layer, a single graph module layer and a linear layer. Results for both ASVspooF 2021 LA and DF databases using optimised DA strategies for each are shown in Table 4. LA and DF results of 1.19% and 4.38% show that competitive EERs can be obtained using the fine-tuned wav2vec 2.0 front-end even with relatively less complex networks and that the benefits to generalisation are still complementary to those of DA.

8. Conclusions and discussion

We report in this paper our attempts to harness the power of self-supervised learning in the form of the popular wav2vec 2.0 front-end to learn more reliable representations to improve spoofing detection performance. We show that a well-trained, fine-tuned front-end, even when learned initially using massive quantities of only *bona fide* utterances, can improve generalisation substantially. Compared to a sinc-layer front-end, when coupled with a new self-attentive aggregation layer and data augmentation, the wav2vec 2.0 front-end delivers up to a 90% relative reduction in the equal error rate for the logical access spoofing detection task and up to an 88% relative reduction for a domain mis-matched Deepfake detection task in which spoofed utterances are generated with more than 100 different attack algorithms. Improvements stemming from the use of the self-supervised wav2vec 2.0 front-end are consistent for every experimental comparison and, to the best of the authors’ knowledge, results are the lowest reported EERs for both LA and DF databases to date.

We must nonetheless acknowledge that almost all results reported in the literature are derived using *fixed* training data, whereas those reported in this paper are derived from a model pre-trained using additional *external* data. Given that our results are obtained with different systems trained using different data, comparisons to the majority of results reported in the literature are obviously not fair. The scale of the improvements, however, up to a 90% relative reduction in EER, indicate the potential gain in performance that can be obtained with the use of additional, external training data and might suggest that the training data restrictions for ASVspooF evaluations might be relaxed. Nonetheless, the wav2vec 2.0 model is massively more complex when compared to the previous state-of-the-art solutions. Whether or not solutions with such footprints can be adapted to support practical applications remains to be seen.

Future work should investigate which particular characteristics of the self-supervised front-end are most beneficial. One objective of such work is to use the results of such analysis as a starting point to scale down the model size and complexity so that it might be more easily adopted for practical scenarios with modest computational resources.

[†]We used their respective optimised DA strategies for LA and DF as described in Section 6.2

9. References

- [1] A. Nautsch, X. Wang, et al., “ASVspoof 2019: spoofing countermeasures for the detection of synthesized, converted and replayed speech,” *IEEE Transactions on Biometrics, Behavior, and Identity Science (T-BIOM)*, vol. 3, 2021.
- [2] J. Yamagishi, X. Wang, et al., “ASVspoof 2021: accelerating progress in spoofed and deepfake speech detection,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [3] A. Tomilov, A. Svishchev, et al., “STC Antispoofing Systems for the ASVspoof 2021 Challenge,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [4] T. Chen, E. Houry, et al., “Pindrop Labs’ Submission to the ASVspoof 2021 Challenge,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [5] W. H. Kang, J. Alam, et al., “CRIM’s system description for the ASVspoof 2021 Challenge,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [6] J. Cáceres, R. Font, et al., “The Biometric Vox System for the ASVspoof 2021 Challenge,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [7] R. K. Das, “Known-unknown Data Augmentation Strategies for Detection of Logical Access, Physical Access and Speech Deepfake Attacks: ASVspoof 2021,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [8] X. Chen, Y. Zhang, et al., “UR Channel-Robust Synthetic Speech Detection System for ASVspoof 2021,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [9] X. Wang, X. Qin, et al., “The DKU-CMRI System for the ASVspoof 2021 Challenge: Vocoder based Replay Channel Response Estimation,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [10] J. Jung, H. Heo, et al., “AASIST: Audio Anti-Spoofing using Integrated Spectro-Temporal Graph Attention Networks,” in *Proc. ICASSP (to appear)*, 2022.
- [11] Aaron Van den Oord, Yazhe Li, and Oriol Vinyals, “Representation learning with contrastive predictive coding,” in *Proc. of NIPS*, 2018.
- [12] A. Baevski, Y. Zhou, et al., “Wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations,” in *Proc NIPS*, 2020, vol. 33.
- [13] A. Babu, C. Wang, et al., “XLS-R: Self-supervised cross-lingual speech representation learning at scale,” *arXiv preprint arXiv:2111.09296*, 2021.
- [14] S.-w. Yang, P.-H. Chi, et al., “SUPERB: Speech processing Universal PERFORMANCE Benchmark,” in *Proc. INTERSPEECH*, 2021.
- [15] A. Conneau, A. Baevski, et al., “Unsupervised cross-lingual representation learning for speech recognition,” in *Proc. INTERSPEECH*, 2021.
- [16] K. Kawakami, L. Wang, et al., “Learning Robust and Multilingual Speech Representations,” in *Proc. EMNLP*, 2020.
- [17] D. Cai, W. Wang, et al., “An iterative framework for self-supervised deep speaker representation learning,” in *Proc. ICASSP*, 2021.
- [18] A. Tjandra, D. Gon Choudhury, et al., “Improved language identification through cross-lingual self-supervised learning,” *arXiv preprint arXiv:2107.04082*, 2021.
- [19] S. Evain, H. Nguyen, et al., “LeBenchmark: A Reproducible Framework for Assessing Self-Supervised Representation Learning from Speech,” in *Proc. INTERSPEECH*, 2021.
- [20] Y. Xie, Z. Zhang, et al., “Siamese Network with wav2vec Feature for Spoofing Speech Detection,” in *Proc. INTERSPEECH*, 2021.
- [21] X. Wang and J. Yamagishi, “Investigating self-supervised front ends for speech spoofing countermeasures,” *arXiv preprint arXiv:2111.07725*, 2021.
- [22] F. Alegre, A. Amehraye, et al., “A one-class classification approach to generalised speaker verification spoofing countermeasures using local binary patterns,” in *Proc. IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, 2013.
- [23] T. Chen, A. Kumar, et al., “Generalization of Audio Deepfake Detection,” in *Proc. Speaker Odyssey Workshop*, 2020.
- [24] R. K. Das, J. Yang, et al., “Data Augmentation with Signal Companding for Detection of Logical Access Attacks,” in *Proc. ICASSP*, 2021.
- [25] Y. Zhang, G. Zhu, et al., “An Empirical Study on Channel Effects for Synthetic Voice Spoofing Countermeasure Systems,” in *Proc. INTERSPEECH*, 2021.
- [26] X. Chen, Y. Zhang, et al., “UR Channel-Robust Synthetic Speech Detection System for ASVspoof 2021,” in *Proc. ASVspoof 2021 Workshop*, 2021.
- [27] H. Tak, M. Kamble, et al., “RawBoost: A Raw Data Boosting and Augmentation Method applied to Automatic Speaker Verification Anti-Spoofing,” in *Proc. ICASSP (to appear)*, 2022.
- [28] J. Devlin, M.-W. Chang, et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *Proc. NAACL-HLT*, 2019.
- [29] A. v. d. Oord, Y. Li, et al., “Representation learning with contrastive predictive coding,” *CoRR*, vol. abs/1807.03748, 2018.
- [30] Y.-A. Chung, Y. Belinkov, et al., “Similarity Analysis of Self-Supervised Speech Representations,” in *Proc. ICASSP*, 2021.
- [31] S. Schneider, A. Baevski, et al., “Wav2vec: Unsupervised Pre-Training for Speech Recognition,” in *INTER-SPEECH*, 2019.
- [32] W.-N. Hsu, B. Bolte, et al., “HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units,” *arXiv preprint arXiv:2106.07447*, 2021.
- [33] W.-N. Hsu, Y.-H. H. Tsai, et al., “HuBERT: How much can a bad teacher benefit ASR pre-training?,” in *Proc. ICASSP*, 2021.
- [34] Q. Xu, A. Baevski, et al., “Self-training and pre-training are complementary for speech recognition,” in *Proc. ICASSP*, 2021.
- [35] S. Chen, C. Wang, et al., “WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing,” *arXiv preprint arXiv:2110.13900*, 2021.

- [36] X. Xu, Y. Kang, et al., “Explore wav2vec 2.0 for Mispronunciation Detection,” in *Proc. INTERSPEECH*, 2021.
- [37] L. Peng, K. Fu, et al., “A Study on Fine-Tuning wav2vec2.0 Model for the Task of Mispronunciation Detection and Diagnosis,” in *Proc. INTERSPEECH*, 2021.
- [38] N. Vaessen and D. A. van Leeuwen, “Fine-tuning wav2vec2 for speaker recognition,” *arXiv preprint arXiv:2109.15053*, 2021.
- [39] Z. Fan, M. Li, et al., “Exploring wav2vec 2.0 on speaker verification and language identification,” in *Proc. INTERSPEECH*, 2021.
- [40] L. Pepino, P. Riera, et al., “Emotion Recognition from Speech Using Wav2vec 2.0 Embeddings,” in *Proc. INTERSPEECH*, 2021.
- [41] C. Wang, M. Rivière, et al., “VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation,” in *Proc. of Annual Meeting of the Association for Computational Linguistics (ACL)*, 2021.
- [42] V. Pratap, Q. Xu, et al., “MLS: A Large-Scale Multilingual Dataset for Speech Research,” in *Proc. INTERSPEECH*, 2020.
- [43] R. Ardila, M. Branson, et al., “Common Voice: A Massively-Multilingual Speech Corpus,” in *Proc. LREC*, 2020.
- [44] J. Valk and T. Alumäe, “VoxLingua107: a dataset for spoken language recognition,” in *Proc. Spoken Language Technology Workshop (SLT)*, 2021.
- [45] M. JF Gales, K. M Knill, et al., “Speech recognition and keyword spotting for low-resource languages: BABEL project research at CUED,” in *Proc. Spoken language technologies for under-resourced languages workshop (SLTU)*, 2014.
- [46] M. Ravanelli and Y. Bengio, “Speaker recognition from raw waveform with SincNet,” in *IEEE Proc. SLT*, 2018.
- [47] H. Tak, J. Jung, et al., “End-to-end spectro-temporal graph attention networks for speaker verification anti-spoofing and speech deepfake detection,” in *Proc. ASVspoof 2021 workshop*, 2021.
- [48] P. Veličković, G. Cucurull, et al., “Graph attention networks,” in *Proc. ICLR*, 2018.
- [49] H. Gao and S. Ji, “Graph u-nets,” in *Proc. international conference on machine learning (PMLR)*, 2019.
- [50] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S Yu, “Heterogeneous graph attention network,” in *Proc. The World Wide Web (WWW) Conference*, 2019.
- [51] A. Vaswani, N. Shazeer, et al., “Attention is all you need,” in *Proc. NIPS*, 2017.
- [52] M. Ott, S. Edunov, et al., “fairseq: A Fast, Extensible Toolkit for Sequence Modeling,” in *Proc. NAACL-HLT 2019: Demonstrations*, 2019.
- [53] K. Okabe, T. Koshinaka, et al., “Attentive statistics pooling for deep speaker embedding,” in *Proc. INTERSPEECH*, 2018.
- [54] Y. Jung, S. M Kye, et al., “Improving Multi-Scale Aggregation Using Feature Pyramid Module for Robust Speaker Verification of Variable-Duration Utterances,” in *Proc. INTERSPEECH*, 2020.
- [55] J.-w. Jung, H.-S. Heo, et al., “Graph attention networks for speaker verification,” in *Proc. ICASSP*, 2021.
- [56] W. Xie, A. Nagrani, et al., “Utterance-level aggregation for speaker recognition in the wild,” in *Proc. ICASSP*, 2019.
- [57] J.-h. Kim, H.-j. Shim, et al., “RawNeXt: Speaker verification system for variable-duration utterances with deep layer aggregation and extended dynamic scaling policies,” in *Proc. ICASSP (to appear)*, 2022.
- [58] H.-j. Shim, J. Heo, et al., “Graph attentive feature aggregation for text-independent speaker verification,” in *Proc. ICASSP (to appear)*, 2022.
- [59] G. Klambauer, T. Unterthiner, et al., “Self-normalizing neural networks,” in *Proc. NIPS*, 2017.
- [60] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *Proc. International Conference on Machine Learning (ICML)*, 2015.
- [61] M. Todisco, X. Wang, et al., “ASVspoof 2019: Future Horizons in Spoofed and Fake Audio Detection,” in *Proc. INTERSPEECH*, 2019.
- [62] X. Wang, J. Yamagishi, et al., “ASVspoof 2019: a large-scale public database of synthesized, converted and re-played speech,” *Computer Speech & Language (CSL)*, vol. 64, 2020, 101114.
- [63] C. Veaux, J. Yamagishi, et al., “CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit,” 2017, <http://dx.doi.org/10.7488/ds/1994>.
- [64] N. Brümmer and E. de Villiers, “The BOSARIS toolkit: Theory, algorithms and code for surviving the new DCF,” *arXiv preprint arXiv:1304.2865*, 2013.
- [65] T. Kinnunen, H. Delgado, et al., “Tandem Assessment of Spoofing Countermeasures and Automatic Speaker Verification: Fundamentals,” *IEEE/ACM Transactions on Audio Speech and Language Processing (TASLP)*, vol. 28, 2020.
- [66] T. Painter and A. Spanias, “Perceptual coding of digital audio,” *Proceedings of the IEEE*, vol. 88, no. 4, 2000.
- [67] D. P Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [68] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016, <http://www.deeplearningbook.org>.
- [69] X. Wang and J. Yamagishi, “A Comparative Study on Recent Neural Spoofing Countermeasures for Synthetic Speech Detection,” in *Proc. INTERSPEECH*, 2021.
- [70] A. Kanervisto, V. Hautamäki, et al., “Optimizing Tandem Speaker Verification and Anti-Spoofing Systems,” *IEEE/ACM TASLP*, vol. 30, 2021.