

# Comprehensive Layer-wise Analysis of SSL Models for Audio Deepfake Detection

Yassine El Kheir<sup>1,2</sup>, Younes Samih<sup>3</sup>, Suraj Maharjan, Tim Polzehl<sup>1,2</sup>, Sebastian Möller<sup>2,1</sup>

<sup>1</sup>German Research Center for Artificial Intelligence (DFKI), Berlin, Germany

<sup>2</sup>Technical University of Berlin, Berlin, Germany

<sup>3</sup>IBM Research AI, Abu Dhabi, UAE

yassine\_el\_kheir@dfki.de, younes.samih@ibm.com, mhjsuraj@gmail.com

tim.polzehl@dfki.de, sebastian.moeller@dfki.de

## Abstract

This paper conducts a comprehensive layer-wise analysis of self-supervised learning (SSL) models for audio deepfake detection across diverse contexts, including multilingual datasets (English, Chinese, Spanish), partial, song, and scene-based deepfake scenarios. By systematically evaluating the contributions of different transformer layers, we uncover critical insights into model behavior and performance. Our findings reveal that lower layers consistently provide the most discriminative features, while higher layers capture less relevant information. Notably, all models achieve competitive equal error rate (EER) scores even when employing a reduced number of layers. This indicates that we can reduce computational costs and increase the inference speed of detecting deepfakes by utilizing only a few lower layers. This work enhances our understanding of SSL models in deepfake detection, offering valuable insights applicable across varied linguistic and contextual settings. Our models and code are publicly available at [https://github.com/Yasellej/SSL\\_Layerwise\\_Deepfake](https://github.com/Yasellej/SSL_Layerwise_Deepfake).

## 1 Introduction

Recent advancements in speech synthesis and voice conversion have resulted in high-fidelity synthetic audio, which can convincingly mimic real human voices (Kumar et al., 2023; Huang et al., 2023; Yi et al., 2023b). As a result, distinguishing between authentic speech and sophisticated deepfake audio has become increasingly challenging (Khan et al., 2022). This poses a significant threat, particularly for Automatic Speaker Verification (ASV) systems which are widely used for authentication in access control, telephone banking, and forensic investigations (Anjum and Swamy, 2017; Li et al., 2024). The potential for social and economic harm is evident, as highlighted by many incidents like the \$243,000 scam, which exploited voice mimicry

to deceive a CEO (Damiani, 2019). Such threats underscore the urgent need for effective audio deepfake detectors to ensure the security and trustworthiness of ASV systems.

In light of these challenges, SSL has emerged as a promising approach in speech processing. They are capable of extracting rich features from vast amounts of unlabeled audio data, which have enhanced the accuracy in speaker verification and speech recognition tasks (Mohamed et al., 2022; Borgholt et al., 2022; Yang et al., 2021; Tsai et al., 2022; Shon et al., 2022). Many studies show that features derived from SSL methods consistently outperform traditional acoustic features across various tasks (Babu et al., 2021; Xie et al., 2021; Martín-Doñas and Álvarez, 2022; Wang and Yamagishi, 2021). This suggests a new direction for enhancing the resilience of speech systems against deepfake threats.

While previous studies have highlighted the potential of SSL models in deepfake detection, a thorough layer-wise analysis of SSL models remains largely unexplored. Most current research focuses on deepfakes created exclusively in English within full-utterance audio scenarios, leaving several important questions unanswered: (1) How effectively do SSL models detect deepfakes across diverse scenarios, including multiple languages, partial, song, and scene deepfakes? (2) How do different layers of SSL models behave in these varying conditions, and is their performance consistent across different setups? (3) Which layers yield the most discriminative features for differentiating real audio from deepfakes?

In this paper, we address these gaps by undertaking a comprehensive analysis of SSL models across various settings. This includes (1) full speech utterance deepfake detection in English (En), Chinese (Zh), and Spanish (Es); (2) partial speech utterance detection in English and Chinese; and (3) detection song-based and scene-based (acoustic en-

vironment) deepfakes. By examining layer-wise contributions, we aim to provide insights into the effectiveness of SSL models in multiple languages and contexts.

Our main contributions in this paper are as follows:

- We show that the lower layers of the SSL models (Wav2Vec2, Hubert, and WavLM) play a more crucial role than the upper layers in detecting audio deepfakes. This is consistent across multiple languages and setups (full, partial deepfakes, song, scene) using different back-end classifiers.
- We demonstrate that utilizing only the lower layers of SSL models: 4-6 layers for Small models and 10-12 layers for Large models, achieves performance comparable to, or even surpassing, that of the full model while maintaining generalizability across diverse datasets and setups.

This research represents a pioneering effort to conduct a comprehensive layer-wise analysis of SSL models in the context of deepfake detection, encompassing a wide range of languages and settings.

## 2 Background

Traditional audio deepfake detection methods rely on hand-crafted acoustic features like MFCC, LFCC, or CQCC (Chakroborty et al., 2008; Alegre et al., 2013; Li et al., 2023; Sahidullah et al., 2015; Todisco et al., 2017). These methods require expert knowledge and may overlook essential discriminative information (Tak et al., 2020). As a result, while effective on some datasets, such methods struggle to generalize to new or unseen spoofing attacks due to the static nature of the features. This limitation leaves systems vulnerable to emerging spoofing techniques. Therefore, more robust and flexible anti-spoofing strategies are needed that can automatically learn and extract relevant features from raw speech data.

With advancements in deep neural network (DNN) approaches, DNN-based methods have generally been adopted in two ways. First, they can be used as back-end models with traditional hand-crafted features as front-ends (Zhang12 et al., 2021). In light of this, light convolutional neural network (LCNN) (Lavrentyeva et al., 2017; Tomilov et al., 2021), residual neural network

(ResNet) (Lai et al., 2019), and siamese convolutional neural network (Lei et al., 2020), have been successfully applied to anti-spoofing systems using MFCC, LFCC, or CQCC features. On the other hand, end-to-end DNN modeling techniques have been used to extract learnable embeddings directly from raw audio and automatically extract features relevant to anti-spoofing techniques. For instance, AASIST (Tak et al., 2021) used SincNet (Ravanelli and Bengio, 2018) to directly extract front-end features from raw audio, with fixed cut-off frequencies. Other structures as well, as explored in Bartusiak and Delp (2022); Shim et al. (2022); Khan et al. (2024), are applied directly to raw audio data or combined with hand-crafted features to construct deep embeddings through supervised training.

SSL models are large pre-trained frameworks that serve as the backbone for various speech tasks. These models excel in capturing high-quality representations that can be fine-tuned for specific downstream applications. In the field of speech deepfake detection, this approach has been widely adopted, leading to state-of-the-art results. Studies have shown that fine-tuning SSL models like Wav2Vec2 (Baeovski et al., 2020) with a classifier can significantly improve detection performance (Tak et al., 2022). Alternatively, fine-tuning the XLS-R (Babu et al., 2021) model and using features from the fifth layer’s hidden states instead of the last layer has also proven effective (Lee et al., 2022). Research highlights that utilizing hidden state features from various layers of pre-trained models can be highly beneficial for deepfake detection (Martín-Doñas and Álvarez, 2022).

Consequently, recent studies have focused on leveraging multi-layer features from SSL models to further enhance detection performance. Methods such as multi-fusion attentive classifiers (Guo et al., 2024), attentive merging techniques (Pan et al., 2024) with 10 & 12 layers, and their improved versions (Guragain et al., 2024) have been explored. Additionally, expert fusion techniques, such as the Mixture of Experts method, have been proposed to extract and integrate relevant features for fake audio detection from multiple layers of SSL models (Wang et al., 2024).

Recent advancements in multi-layer fusion techniques underscore the importance of a deeper examination of layer-wise contributions in SSL models. While previous research has largely focused on specific SSL models, particular classifiers, and predominantly English datasets (Lee et al., 2022;

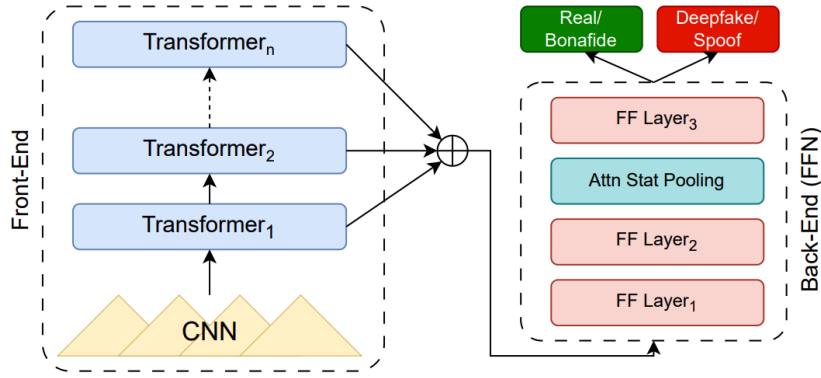


Figure 1: Layer-wise Contribution Framework. The framework consists of SSL models as front-ends to extract features and a back-end classifier. The front-end SSL models remain frozen during the experiments to evaluate the layer-wise feature contribution. FF stands for Linear feed-forward layer.

Pan et al., 2024), our study provides a broader perspective. Unlike prior work, our methodology extends beyond these limitations by comprehensively assessing deepfake detection tasks across various SSL models, languages, and scenarios. This approach enables a thorough analysis of layer-wise contributions. Our detailed analysis identifies the most critical layers for audio deepfake detection, enabling the development of computationally efficient models that maintain effectiveness across diverse datasets.

### 3 Methodology

Figure 1 shows our Layer-wise Contribution Framework. This framework provides insights into the layer-wise contribution of different transformer layers in detecting deepfakes. This framework employs SSL models as front-ends to extract features and a back-end classifier to perform classification (real vs deepfake). To systematically assess layer-wise contributions, we freeze the SSL model and evaluate the contribution of each transformer layer to the final detection performance.

Let  $\mathbf{x}$  represent the input audio signal, which is passed through an SSL model. This model is composed of  $L$  transformer layers, where the output of each layer  $l$  can be denoted as  $\mathbf{h}_l$ . The features  $\mathbf{h}_l$  extracted from each layer  $l \in \{1, 2, \dots, L\}$  are then weighted by a layer-wise importance factor  $w_l$ .

We follow a weighted aggregation of the extracted features, where the final representation  $\mathbf{h}_{\text{final}}$  used for classification is given by:

$$\mathbf{h}_{\text{final}} = \sum_{l=1}^L w_l \mathbf{h}_l \quad (1)$$

The weights  $w_l$  are learnable parameters, and their values are constrained to sum to one by applying a softmax function:

$$w_l = \frac{\exp(w_l)}{\sum_{k=1}^L \exp(w_k)} \quad (2)$$

This normalization ensures that the contributions from different transformer layers are appropriately scaled. The back-end classifier receives  $\mathbf{h}_{\text{final}}$  as input and generates the prediction  $\hat{y}$  for the detection task. This setup allows us to investigate which layers contribute most to the detection of audio deepfakes across various settings.

#### 3.1 Front-Ends: SSL Models

We experiment with 6 different front-end SSL models: small (base) and large versions of Wav2Vec2, Hubert (Hsu et al., 2021), and WavLM (Chen et al., 2022). These models are the most commonly used SSL models for audio deepfake detection. Wav2Vec2, Hubert, and WavLM also differ in their pre-training objectives.

**Wav2Vec2.0** Wav2Vec2 is a contrastive SSL model designed to learn speech representations by predicting latent features from masked audio regions. The architecture consists of two main components: a feature extractor and a stack of transformer layers. The feature extractor is a convolutional network (CNN) with 512 channels, designed to process audio segments of approximately 25ms from 16kHz sampled audio. The CNN compresses the input with strides and kernel sizes optimized to operate every 20ms. The model comes commonly in 2 configurations: the base version with

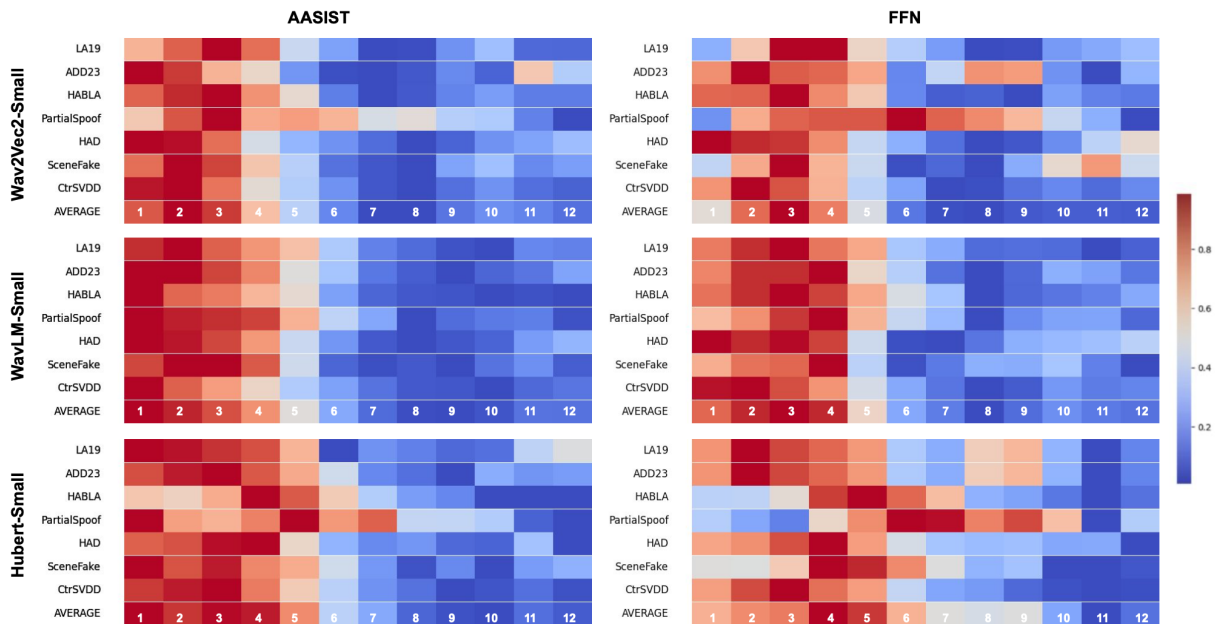


Figure 2: Heatmap of Normalized Layer-wise weights Across Various Datasets using Small SSL Models. AVERAGE row representing the average weights across datasets.

12 transformer layers and a hidden state dimension of 768, and the large version with 24 layers and a dimension of 1024. Wav2Vec2 is trained using Contrastive Predictive Coding (CPC) loss, which distinguishes between positive and negative samples by learning contextualized audio representations. This contrastive approach enables it to effectively model the underlying structure of speech data.

**Hubert** Hubert follows a similar architecture to Wav2Vec2.0, consisting of a convolutional feature extractor followed by a stack of transformer layers. However, it differs in its self-supervised training objective. Instead of contrastive loss, Hubert employs a masked prediction approach, where the model is trained to predict discrete target representations derived from clustered speech features. Initially, these targets are obtained through k-means clustering on MFCC features, and in later stages, the model learns from its embeddings. This iterative refinement enables Hubert to capture meaningful speech representations without requiring frame-level labels.

**WavLM** WavLM follows the same overall architecture as Wav2Vec2.0 and Hubert, consisting of a convolutional feature extractor and a stack of transformer layers. Like Hubert, it employs a masked prediction objective, but it extends this approach by incorporating speech enhancement techniques and

training on both clean and noisy speech. This additional training data improves its robustness in real-world scenarios, particularly for speaker-related tasks such as speaker verification and separation. Compared to Hubert, WavLM introduces an auxiliary denoising task, enabling better performance in diverse acoustic environments.

**Model Selection and Multilingual Setup** For our experiments, we use publicly available pre-trained models for both Wav2Vec2, Hubert, and WavLM. Specifically, we utilize multilingual versions of Hubert<sup>1</sup>, WavLM (both small and large)<sup>2</sup>, and Wav2Vec2 Large<sup>3</sup>. For the Wav2Vec2 small model, we opt for the English pre-trained variant due to the absence of a small multilingual pre-trained version.

### 3.2 Back-End Classifiers: FFN vs AASIST

We explore two different back-end models: (1) a simple lightweight feedforward neural network (FFN), and (2) the state-of-the-art AASIST (Jung et al., 2022) model. We experiment with these two back-end models to demonstrate the generalizability of layer-wise contributions across different classifier architectures. This ensures that our findings are not limited to a specific back-end structure.

<sup>1</sup>[mHuBERT-147](https://huggingface.co/facebook/mHuBERT-147)

<sup>2</sup>[unilm/blob/master/wavlm](https://github.com/facebook/unilm/blob/master/wavlm)

<sup>3</sup>[fairseq/blob/main/examples/wav2vec/xlsr](https://github.com/facebook/fairseq/blob/main/examples/wav2vec/xlsr)



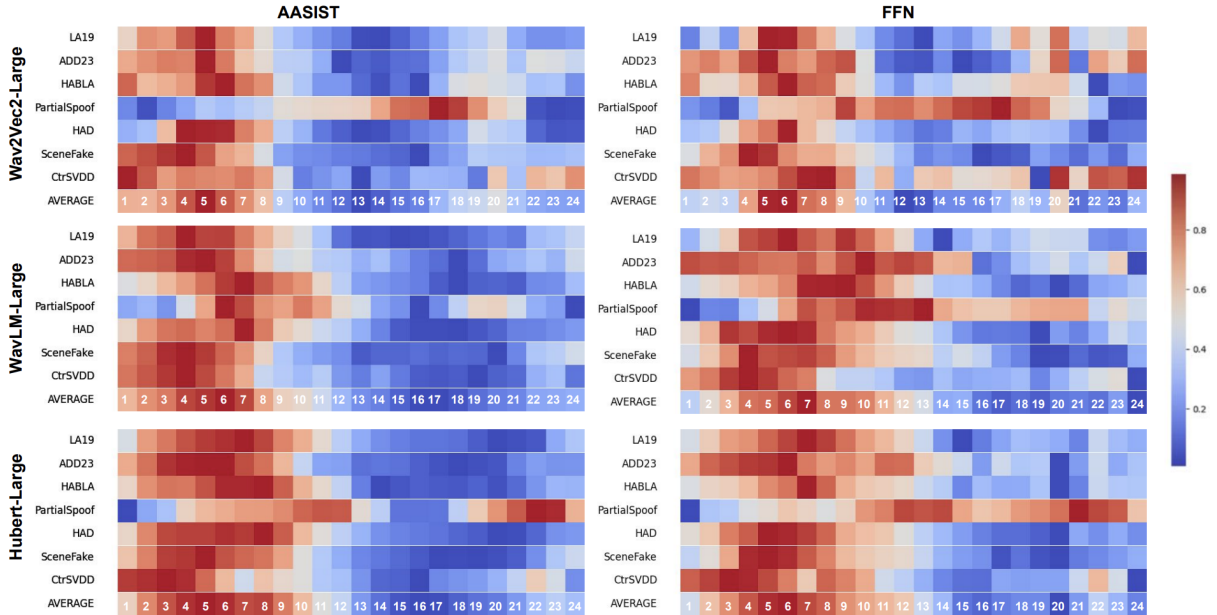


Figure 3: Heatmap of Normalized Layer-wise weights Across Various Datasets using Large SSL Models. AVERAGE row representing the average weights across datasets.

### 3.2.1 FFN

We use the lightweight back-end classifier approach from (Martín-Doñas and Álvarez, 2022) for our deepfake detection task as shown in Figure 1. Let  $\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_T]$  represent the weighted sequence of feature vectors extracted from the front-end SSL model, as defined in Equation 1. Each vector  $\mathbf{h}_t \in \mathbf{R}^d$ , where  $d = 768$  for SSL base models &  $d = 1024$  for SSL large models, is first passed through a batch normalization layer. The normalized vectors are then processed by two feedforward (FF) layers with a hidden dimension of 128, using the SELU activation function and dropout for regularization. Next, the processed vectors  $\mathbf{h}_t$  are aggregated using an attentive statistical pooling layer (Okabe et al., 2018), which computes a weighted mean and standard deviation over the time steps, producing a fixed-length representation for the entire utterance. This attentive pooling condenses the variable-length input sequence into a single vector of size 256. Finally, the pooled representation is projected into 2 classes.

### 3.2.2 AASIST

AASIST (Jung et al., 2022) is a state-of-the-art model that utilizes graph-based attention mechanisms to capture both spectral and temporal audio features. Its key components include: (1) Graph Attention Layer (GAT), which computes attention maps for spectral and temporal features using lin-

ear layers, batch normalization, and SELU activation; (2) Heterogeneous Graph Attention Layer (HtrgGAT), which processes and refines both spectral and temporal nodes; (3) a graph pooling layer, which selects the top- $k$  nodes based on attention scores; (4) residual blocks, which apply convolutional layers and SELU activation to process features; and (5) an attention mechanism that extracts spectral and temporal features from encoded data.

## 4 Experimental Design

### 4.1 Datasets

Table 2 summarizes datasets used in this study. These datasets are very diverse including both full utterance deep fakes (**Full Fake**) and partial utterance deepfakes (**Partial Fake**) in multiple languages: English (En), Chinese (Zh), and Spanish (Es). Additionally, **Scene**-based and **Song** deepfake setups are included. All datasets leverage cutting-edge TTS and VC generation techniques for creating audio deepfakes. Additional details on the different setups can be found in Appendix A, while Appendix B provides further information on the datasets.

### 4.2 Training setup

We employ the cross-entropy (CE) loss function and the Adam optimizer, setting the learning rate to  $1e^{-4}$  and a dropout rate of 0.2, with a batch size of 32 for training our models. During training,

Back-End	Front-end	#Layers	Datasets										
			LA19	LA21	DF21	ADD23.1.2.1	ADD23.1.2.2	HABLA	PartialSpoof	HAD <sup>†</sup>	SceneFake	CtrSVDD	
Wav2Vec2-Small	FFN	2	0.63	6.31	15.09	58.31	<b>67.14</b>	<b>9.80</b>	3.13	22.47	13.06	19.34	
		4	0.58	7.08	<b>12.1</b>	<b>60.07</b>	69.73	9.87	2.64	24.18	11.20	<b>17.10</b>	
		6	<b>0.48</b>	<b>5.53</b>	14.02	63.30	73.77	9.87	<b>2.58</b>	19.55	9.99	17.73	
		12	0.68	6.33	15.58	64.43	74.10	9.94	2.69	<b>17.63</b>	<b>9.26</b>	17.83	
	AASIST	2	0.63	5.09	15.01	62.59	66.85	<b>9.96</b>	2.97	37.43	12.48	<b>15.92</b>	
		4	<b>0.55</b>	5.69	<b>12.17</b>	<b>56.12</b>	<b>63.03</b>	10.02	<b>1.82</b>	36.41	9.87	20.42	
		6	0.92	6.48	12.50	57.78	65.62	10.09	2.11	34.65	10.87	16.80	
		12	1.02	<b>5.08</b>	15.80	58.59	64.91	10.11	2.25	<b>30.24</b>	<b>9.54</b>	16.67	
	WavLM-Small	FFN	2	3.41	27.13	17.98	57.27	63.22	9.91	3.36	31.32	12.27	19.18
			4	2.55	30.15	13.90	<b>57.19</b>	63.57	<b>9.88</b>	2.86	22.43	7.64	14.47
			6	<b>1.85</b>	<b>25.03</b>	<b>12.34</b>	57.46	<b>62.73</b>	9.99	<b>2.12</b>	21.74	4.07	<b>12.61</b>
			12	2.92	29.73	13.19	61.12	66.18	10.14	2.14	<b>19.00</b>	<b>2.44</b>	15.10
AASIST		2	0.55	<b>21.26</b>	17.70	64.33	70.01	10.08	3.00	28.5	8.45	17.92	
		4	<b>0.26</b>	21.51	12.42	<b>54.47</b>	<b>58.86</b>	10.14	2.18	31.03	5.79	15.89	
		6	<b>0.26</b>	25.26	<b>11.76</b>	55.85	60.02	<b>10.05</b>	<b>1.84</b>	30.87	3.80	14.83	
		12	0.57	39.87	15.48	58.55	66.05	10.10	1.97	<b>23.21</b>	<b>2.68</b>	<b>14.25</b>	
Hubert-Small		FFN	2	0.98	11.7	13.3	<b>56.46</b>	<b>64.14</b>	9.97	3.28	25.65	11.29	18.09
			4	0.60	6.78	9.40	57.92	65.09	9.94	2.33	18.67	4.83	17.33
			6	<b>0.44</b>	6.92	<b>9.39</b>	59.81	67.63	<b>9.86</b>	2.54	<b>15.99</b>	<b>3.25</b>	<b>16.24</b>
			12	0.63	<b>6.61</b>	11.01	64.53	72.18	9.89	<b>2.09</b>	18.49	4.33	16.51
	AASIST	2	0.68	5.5	15.11	57.69	<b>63.17</b>	10.03	3.37	32.62	8.70	17.73	
		4	<b>0.37</b>	5.42	11.53	<b>58.10</b>	66.38	10.04	2.05	<b>26.46</b>	<b>4.31</b>	14.72	
		6	0.39	4.55	<b>11.33</b>	61.07	70.54	<b>10.01</b>	1.96	29.72	4.52	<b>13.72</b>	
		12	0.96	<b>4.35</b>	12.37	60.25	66.11	10.08	<b>1.88</b>	29.65	6.31	18.73	

Table 1: Mean EER Results Across Datasets Using Full and Partial Transformer Layers of SSL Small Models. # Indicates the Number of Layers Used from the Front-End SSL Models. Best Average EER is Highlighted in yellow. <sup>†</sup>We trained on the HAD training set and reported results on PartialSpoof eval set.

Dataset	Language	Split
Full Fake		
ASVspoof 2019 (19LA) (Nautsch et al., 2021)	En	Train & Eval
ASVspoof 2021 (21LA)* (Yamagishi et al., 2021)	En	Eval
ASVspoof 2021 (21DF)* (Yamagishi et al., 2021)	En	Eval
ADD23 (Track 1.2)+ (Yi et al., 2023a)	Zh	Train & Eval
HABLA (Flórez et al.)	Es	Train & Eval
Partial Fake		
PartialSpoof (Zhang et al., 2021)	En	Train & Eval
Half-Truth (HAD) <sup>†</sup> (Yi et al., 2021)	Zh	Train
Song		
CtrSVDD (Zhang et al., 2024)	Multilingual	Train & Eval
Scene		
SceneFake (Yi et al., 2024)	En	Train & Eval

Table 2: Summary of datasets used in this study. \*Evaluated on models trained with LA19; + focused on Track 1.2 (Full Fake) with two eval sets, ADD23.1.2.1 and ADD23.1.2.2. <sup>†</sup>We trained on the HAD training set and reported results on PartialSpoof eval set.

we freeze the SSL model parameters and update only the back-end classifier’s parameters with the weights in Equation 1 initialized to ones. We train our models for 50 epochs, applying early stopping with patience of 10 epochs based on evaluation loss. We use a single NVIDIA H100 GPU to run our experiments. Each experiment is repeated 3 times with different seeds for reliable results and the average results are reported. We either crop or concatenate the audio data to create segments of approximately 4 seconds in duration (64,600 samples)<sup>4</sup>. We evaluate the model performance using the standard equal error rate (EER) metric.

<sup>4</sup>For **Partial Fake**, we used full audio

## 5 Results and Discussion

### 5.1 Layer-wise Contribution Analysis

After training SSL Model + Back-end classifiers with different datasets, we extract the layer-wise weights as outlined in Equation 2 to study the contribution of each layer to the task. For analysis, we first normalize these weights and then visualize them as heatmaps (we reported the average across 3 experiments for each dataset), as illustrated in Figures 2 and 3.

**Small SSL Models** Figure 2 presents the normalized layer-wise weights after training Wav2Vec2-Small, Hubert-Small, and WavLM-Small front-end models, using FFN and AASIST back-ends across various datasets. The results consistently show that the highest contributions come from lower layers (1-6) across all datasets and setups (speech, scene, and song), indicating that these layers capture the most discriminative features for the deepfake detection task. This pattern holds regardless of the back-end classifier used, suggesting that the classifier choice does not significantly impact the task reliance on lower layers. However, slight variations are observed in certain cases, such as PartialSpoof where intermediate layers contribute more significantly using FFN.

Looking at the average (AVERAGE) layer-wise weights across datasets, the trend becomes even clearer, with lower layers (1-4) consistently showing the highest contributions. This confirms the

Back-End	Front-end	#Layers	Datasets										
			LA19	LA21	DF21	ADD23.1.2.1	ADD23.1.2.2	HABLA	PartialSpoof	HAD <sup>†</sup>	SceneFake	CtrSVDD	
Wav2Vec2-Large	FFN	4	0.37	19.77	14.45	55.59	69.45	9.92	4.43	35.19	9.09	17.65	
		6	0.18	13.35	7.74	51.08	65.51	9.89	3.81	24.89	4.02	14.19	
		8	<b>0.15</b>	9.82	5.85	45.82	60.75	<b>9.91</b>	2.92	21.92	3.01	14.73	
		10	0.17	6.43	5.27	46.78	60.98	9.99	2.91	19.64	<b>2.73</b>	13.61	
		12	0.21	5.96	4.20	<b>45.17</b>	<b>57.76</b>	9.96	2.91	18.42	3.65	14.31	
		24	0.21	<b>4.63</b>	<b>4.14</b>	46.29	59.91	10.01	<b>2.48</b>	<b>16.83</b>	2.96	<b>13.18</b>	
	AASIST	4	0.69	13.41	18.77	61.76	70.28	10.05	8.11	43.73	6.40	19.3	
		6	0.54	9.61	12.40	52.83	66.10	<b>10.04</b>	5.78	33.27	4.15	15.04	
		8	0.90	11.05	9.80	46.56	61.11	10.05	3.56	30.58	4.38	14.37	
		10	0.80	8.65	10.51	46.57	<b>59.28</b>	10.05	3.66	27.09	<b>3.08</b>	14.12	
		12	<b>0.21</b>	<b>7.95</b>	9.38	<b>45.83</b>	60.05	<b>10.04</b>	2.99	<b>24.74</b>	3.58	14.19	
		24	0.73	10.24	<b>8.12</b>	49.69	61.31	10.11	<b>1.87</b>	26.51	4.25	<b>13.34</b>	
	WavLM-Large	FFN	4	1.32	13.17	15.97	57.85	69.87	<b>9.99</b>	5.99	27.05	5.14	21.27
			6	0.43	8.82	12.36	55.01	68.51	10.03	5.05	24.93	4.89	19.88
			8	0.33	8.00	9.32	52.11	68.02	10.03	3.57	19.02	4.19	16.62
			10	<b>0.23</b>	8.30	7.83	54.27	69.84	10.00	3.57	14.17	3.75	15.98
			12	0.30	<b>4.95</b>	8.29	<b>51.96</b>	<b>67.98</b>	<b>9.99</b>	3.32	13.59	3.27	<b>15.84</b>
			24	0.38	6.97	<b>6.24</b>	56.54	68.19	<b>9.99</b>	<b>2.63</b>	<b>9.90</b>	<b>3.17</b>	16.84
AASIST		4	0.89	8.33	19.63	54.17	68.11	<b>10.06</b>	12.85	39.79	3.61	17.25	
		6	0.44	10.14	13.88	55.05	67.69	10.11	8.45	41.48	<b>2.96</b>	15.84	
		8	<b>0.38</b>	8.04	12.72	53.57	68.03	10.11	7.96	37.36	3.31	16.14	
		10	0.57	<b>6.77</b>	14.71	54.56	68.72	10.15	5.43	<b>36.21</b>	3.49	<b>13.74</b>	
		12	0.52	7.79	12.98	<b>48.50</b>	<b>65.10</b>	10.14	9.09	43.26	4.85	14.00	
		24	0.57	8.17	<b>9.29</b>	51.21	67.32	10.12	<b>3.53</b>	43.43	3.96	15.51	
Hubert-Large		FFN	4	0.44	12.40	12.27	53.30	65.08	<b>9.96</b>	4.09	38.25	8.90	17.22
			6	<b>0.19</b>	10.34	8.69	52.77	65.82	9.89	3.63	29.58	4.71	14.67
			8	<b>0.19</b>	9.57	6.09	49.37	63.84	9.90	3.13	21.57	3.80	12.19
			10	0.22	7.25	5.24	<b>48.61</b>	62.74	9.89	2.81	16.20	4.02	12.49
			12	0.20	5.70	<b>4.47</b>	49.36	<b>62.43</b>	<b>9.96</b>	2.57	15.59	3.34	12.71
			24	0.28	<b>4.81</b>	5.12	52.11	63.85	<b>9.95</b>	<b>2.33</b>	<b>13.20</b>	<b>2.97</b>	<b>12.01</b>
	AASIST	4	0.70	12.26	16.98	53.29	63.53	10.06	6.54	34.97	4.79	17.18	
		6	0.64	10.58	15.07	48.54	61.47	<b>10.04</b>	6.38	29.40	5.38	13.58	
		8	0.32	6.70	12.59	<b>45.35</b>	58.91	10.08	4.19	29.45	4.14	11.17	
		10	<b>0.23</b>	7.84	10.64	45.4	59.06	10.05	2.50	27.29	4.47	<b>10.52</b>	
		12	0.27	<b>5.41</b>	10.29	46.27	<b>56.97</b>	10.06	2.47	<b>17.05</b>	<b>3.11</b>	11.04	
		24	0.47	6.65	<b>8.46</b>	45.44	59.84	10.09	<b>2.37</b>	22.67	3.64	11.16	

Table 3: Mean EER Results Across Datasets Using Full and Partial Transformer Layers of SSL Large Models. # Indicates the Number of Layers Used from the Front-End SSL Models. Best Average EER is Highlighted in yellow. <sup>†</sup>We trained on the HAD training set and reported results on PartialSpoof eval set.

dominant role of lower layers in the deepfake performance of small SSL models.

**Large SSL Models** Figure 3 shows the normalized layer-wise weights for Wav2Vec2-Large, Hubert-Large, and WavLM-Large front-end models. Unlike the concentrated lower-layer contributions seen with small SSL models, large SSL models display a lesser concentrated pattern of layer importance, especially with the FFN back-end. For Wav2Vec2-Large, Hubert-Large, and WavLM-Large, lower layers (2-13) still contribute significantly across most datasets when using both FFN and AASIST back-ends. However, a shift towards middle layers (12-21) is observed in PartialSpoof. When looking at the average (AVERAGE) layer-wise weights, this trend remains stable across SSL models and back-end classifiers, with the highest contributions emerging in lower layers (1-12), particularly around the 4-7<sup>th</sup> layers.

**Focus on Local Features Across Models** Across all models, there is a consistent reliance on the lower layers, particularly layers 1-6 in Small SSL models and layers 1-12 in Large SSL models. This trend suggests that the models focus on local features, which aligns with how audio deepfake de-

tectors identify artifacts left by VC and TTS algorithms. These artifacts often reside in specific frequency sub-bands or temporal segments, highlighting the significance of local feature extraction (Sriskandaraja et al., 2016; Tak et al., 2020).

Particularly, in our investigation of the PartialSpoof data, we observed significant disfluency in certain speech segments, indicating that the model tends to focus more on the transitions between fake and real segments as shown in (Liu et al., 2024). This may explain the model’s increased reliance on intermediate and upper SSL layers, which are more effective at capturing disfluent speech patterns, as demonstrated in (Shih et al., 2024).

## 5.2 Performance of Reduced-Layer SSL Models

This section investigates the capabilities of SSL models when limited to a reduced number of transformer layers. We evaluate Wav2Vec2-Small, Hubert-Small, and WavLM-Small using 2, 4, and 6 layers, as well as Wav2Vec2-Large, Hubert-Large, and WavLM-Large using 4, 6, 8, 10, and 12 layers. These configurations are guided by our earlier layer-wise analysis. The reduced-layer models are trained under the same conditions as the

Wav2Vec-Small				WavLM-Small				Hubert-Small				AVG <sub>out</sub>
#Layers	FFN	AASIST	AVG <sub>in</sub>	#Layers	FFN	AASIST	AVG <sub>in</sub>	#Layers	FFN	AASIST	AVG <sub>in</sub>	
<b>2</b>	21.53	22.89	22.21	<b>2</b>	24.51	24.18	24.35	<b>2</b>	21.49	21.46	21.48	22.68
<b>4</b>	<b>21.46</b>	21.61	<b>21.54</b>	<b>4</b>	22.46	<b>21.26</b>	21.86	<b>4</b>	19.29	<b>19.94</b>	<b>19.62</b>	21.01
<b>6</b>	21.68	21.78	21.73	<b>6</b>	<b>20.99</b>	21.45	<b>21.22</b>	<b>6</b>	<b>19.21</b>	20.78	20.00	<b>20.98</b>
<b>12</b>	21.85	<b>21.42</b>	21.64	<b>12</b>	22.20	23.27	22.74	<b>12</b>	20.63	21.07	20.85	21.74
Wav2Vec-Large				WavLM-Large				Hubert-Large				
#Layers	FFN	AASIST	AVG <sub>in</sub>	#Layers	FFN	AASIST	AVG <sub>in</sub>	#Layers	FFN	AASIST	AVG <sub>in</sub>	
<b>4</b>	23.59	25.25	24.42	<b>4</b>	22.76	23.47	23.12	<b>4</b>	22.19	22.03	22.11	23.22
<b>6</b>	19.47	20.98	20.23	<b>6</b>	20.99	22.6	21.80	<b>6</b>	20.03	20.11	20.07	20.70
<b>8</b>	17.49	19.24	18.37	<b>8</b>	19.12	21.76	20.44	<b>8</b>	17.97	18.29	18.13	18.98
<b>10</b>	16.85	18.38	17.62	<b>10</b>	18.79	21.44	20.12	<b>10</b>	16.95	17.80	17.38	18.37
<b>12</b>	16.26	<b>17.90</b>	<b>17.08</b>	<b>12</b>	<b>17.95</b>	21.62	19.79	<b>12</b>	<b>16.63</b>	<b>16.29</b>	<b>16.46</b>	<b>17.78</b>
<b>24</b>	<b>16.06</b>	18.62	17.34	<b>24</b>	18.09	<b>21.31</b>	<b>19.7</b>	<b>24</b>	16.66	17.08	16.87	17.97

Table 4: Reported average EER across datasets using different Wav2Vec2, Hubert, and WavLM SSL models with different back-end layers.  $AVG_{in}$  represents the average EER across datasets for each SSL model, calculated by first averaging the EER for different back-ends (FFN and AASIST) within the same reduced-layer model.  $AVG_{out}$  indicates the overall average of  $AVG_{in}$ , calculated separately for SSL Small and SSL Large models. #L indicates the number of transformer layers.

initial full models described in Section 4.2, and their performance is assessed across the same datasets and setups. We introduce a notation where "SSLModel-X" represents a model with X layers (e.g., Wav2Vec2-Small-4 denotes Wav2Vec2-Small with 4 layers).

**Performance Comparison Across Deepfake Scenarios** Our analysis demonstrates that reduced-layer SSL models perform on par with, or even better than, full-layer models in various deepfake detection scenarios. As shown in Tables 1 and 3, no single configuration consistently outperforms across all datasets and back-end classifiers. While full models excel in some cases, reduced models—especially 4- and 6-layer configurations for small models and 8-, 10-, and 12-layer configurations for large models—achieve comparable and superior results in others.

For example, Wav2Vec2-Small-6+FFN achieves a lower EER of **0.48%** on LA19 compared to **0.68%** for the full model, with better generalization on LA21 and DF21. WavLM-Small-6+FFN outperforms the full model on 6 out of 10 datasets. The WavLM-Small-4+AASIST configuration yields an EER of **54.47%** on ADD23.1.2.1, compared to **58.55%** for the full model, indicating the effectiveness of reduced layers. Hubert-Small shows similar trends, with reduced-layer models performing either comparably or better than the full model. For instance, the Hubert-Small-6+ASSIST model scores **13.72%** on the CtrSVDD dataset, outperforming the full model’s **18.73%**. However, on the SceneFake datasets, full models for both Wav2Vec-Small and WavLM-Small consistently outperform

the reduced-layer versions, although the difference is minimal, ranging from **0.5%** to **1.5%** EER.

For Large SSL models, the Wav2Vec2-Large-12+AASIST configuration significantly reduces the EER to **0.21%** on LA19, outperforming the full model while maintaining as well generalization on LA21 and DF21. Specifically, it shows a **2.29%** gain on LA21, with only a slight **1.26%** reduction on DF21. The WavLM-Large-12+FFN configuration achieves **15.84%** EER on CtrSVDD, outperforming the full model. Hubert-Large follows a similar trend, with the 8-layer model reaching **0.19%** EER on LA19, and Hubert-Large-12+FFN achieving the best result on DF21 with an EER of **4.47%**.

Across critical datasets such as ADD23, reduced-layer models consistently outperform full models: Wav2Vec2-Large-12+FFN achieves **45.17%** EER on ADD23.1.2.1, WavLM-Large-12+FFN reaches **51.96%**, and Hubert-Large-10 hits **48.61%**. In contrast, for PartialSpoof, full models consistently outperform the reduced-layer SSL models, aligning with our observation in Figure 3, where the upper-middle layers contribute most to this dataset. Nonetheless, despite the reduced-layer SSL models not significantly lowering the results, the differences between the best reduced-layer models and full models remain minimal, with EER differences ranging from **0.24%** to **1.9%**.

### 5.3 Optimal Layer Configuration Across Datasets

Our results reveal that using 4-6 layers for the Small models (Wav2Vec-Small, Hubert-Small, WavLM-Small) and 10-12 layers for the Large



models (Wav2Vec2-Large, Hubert-Large, WavLM-Large) consistently delivers strong average performance across various datasets. As illustrated in Table 4,  $\text{AVG}_{in}$  represents the average EER across datasets for each SSL model, computed by averaging the EER for the different back-end configurations (FFN and AASIST) within the same reduced-layer model. The  $\text{AVG}_{out}$  column provides the overall average of  $\text{AVG}_{in}$ , calculated separately for Small and Large SSL models. Notably, Wav2Vec2-Small-6, Hubert-Small-6, and WavLM-Small-6 achieve an  $\text{AVG}_{out}$  of **20.98%**, while Wav2Vec2-Small-4, Hubert-Small-4, and WavLM-Small-4 reach **21.01%**, outperforming both the full SSL model and other reduced-layer configurations. Likewise, the 12-layer configuration for the Large models also demonstrates optimal performance, with an  $\text{AVG}_{out}$  of **17.78%**, showing its robustness across datasets. However, in some cases, the full model outperforms other configurations, such as with Wav2Vec2-Small-AASIST, Wav2Vec2-Large-FFN, and WavLM-Large-AASIST. Despite this, the reduced-layer models come close in performance, with minimal degradation in EER, ranging from a **0.31%** to **0.72%** difference.

## 6 Small vs Large SSL Models

When comparing Small and Large SSL models across various datasets, no significant performance difference is observed for most datasets, including LA19, LA21, HABLA, PartialSpoof, and SceneFake. However, for more recent datasets, where deepfake models are more advanced, such as DF21, ADD23.1.2.1, and ADD23.1.2.2, larger architectures are essential for capturing subtle deepfake artifacts and improving generalization across languages, as demonstrated with HAD<sup>5</sup>. For instance, Tables 1 and 3 show a **10.95%** difference in EER is observed when comparing the best performance of Wav2Vec2-Small on ADD23.1.2.1 with Wav2Vec2-Large. A **6.2%** difference in EER is also noted between WavLM-Small and WavLM-Large on ADD23.1.2.2. On the other hand, Hubert-Large achieves its best performance on CtrSVDD with an EER of **10.53%**, lagging behind the best Hubert-Small configuration by **3.19%** EER difference.

<sup>5</sup>Models trained on HAD training set but evaluated on PartialSpoof eval set.

## 7 Conclusion

This study provides a comprehensive layer-wise analysis of SSL models for audio deepfake detection, covering multilingual datasets and a range of deepfake types, including Full Fake, Partial Fake, and Song- and Scene-based scenarios. Our results show that using only first 4-6 layers for Small models and 10-12 layers for Large models consistently delivers optimal performance across datasets, reducing parameters by at least half. This significantly lowers computational costs while maintaining competitive EER scores.

## 8 Limitations

Our study provides a comprehensive layer-wise analysis of SSL models for deepfake detection across languages and contexts, but it has some limitations. We did not investigate the latest Audio Language Models and Speech tokenizers used for audio deepfake generation (Xie et al., 2024), which might lead to different artifacts and model behavior. Additionally, we focused solely on speech-based SSL models, excluding acoustic SSL models. Future work could explore these models, especially for tasks involving non-speech artifacts, like those in the SceneFake dataset. We also plan to investigate alternative fusion strategies, such as attentional fusion, beyond a simple weighted-sum layer-wise combination.

## 9 Acknowledgments

We thank the anonymous reviewers of NAACL 2025 for their helpful and constructive feedback. This work has been supported by the German Federal Ministry of Education and Research as part of the projects news-polygraph (03RU2U151C) and VERANDA (16KIS2047).

## References

- Federico Alegre, Asmaa Amehraye, and Nicholas Evans. 2013. A one-class classification approach to generalised speaker verification spoofing countermeasures using local binary patterns. In *2013 IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–8. IEEE.
- Z Khamar Anjum and R Kumara Swamy. 2017. Spoofing and countermeasures for speaker verification: A review. In *2017 International Conference on Wireless Communications, Signal Processing and Networking (WiSPNET)*, pages 467–471. IEEE.

- Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick Von Platen, Yatharth Saraf, Juan Pino, et al. 2021. Xls-r: Self-supervised cross-lingual speech representation learning at scale. *arXiv preprint arXiv:2111.09296*.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460.
- Emily R Bartusiak and Edward J Delp. 2022. Transformer-based speech synthesizer attribution in an open set scenario. In *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 329–336. IEEE.
- Lasse Borgholt, Jakob Drachmann Havtorn, Joakim Edin, Lars Maaløe, and Christian Igel. 2022. A brief overview of unsupervised neural speech representation learning. *arXiv preprint arXiv:2203.01829*.
- Sandipan Chakraborty, Anindya Roy, and Goutam Saha. 2008. Improved closed set text-independent speaker identification by combining mfcc with evidence from flipped filter banks. *International Journal of Electronics and Communication Engineering*, 2(11):2554–2561.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- Jesse Damiani. 2019. A voice deepfake was used to scam a ceo out of \$243,000. *Forbes Magazine*.
- Pablo Andrés Tamayo Flórez, Rubén Manrique, and Bernardo Pereira Nunes. Habla: A dataset of latin american spanish accents for voice anti-spoofing.
- Yinlin Guo, Haofan Huang, Xi Chen, He Zhao, and Yuehai Wang. 2024. Audio deepfake detection with self-supervised wavlm and multi-fusion attentive classifier. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12702–12706. IEEE.
- Anmol Guragain, Tianchi Liu, Zihan Pan, Hardik B Sailor, and Qiongqiong Wang. 2024. Speech foundation model ensembles for the controlled singing voice deepfake detection (ctrsvdd) challenge 2024. *arXiv preprint arXiv:2409.02302*.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhota, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM transactions on audio, speech, and language processing*, 29:3451–3460.
- Wen-Chin Huang, Lester Phillip Violeta, Songxiang Liu, Jiatong Shi, and Tomoki Toda. 2023. The singing voice conversion challenge 2023. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 1–8. IEEE.
- Jee-weon Jung, Hee-Soo Heo, Hemlata Tak, Hye-jin Shim, Joon Son Chung, Bong-Jin Lee, Ha-Jin Yu, and Nicholas Evans. 2022. Aasist: Audio anti-spoofing using integrated spectro-temporal graph attention networks. In *ICASSP 2022-2022 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 6367–6371. IEEE.
- Awais Khan, Khalid Mahmood Malik, and Shah Nawaz. 2024. Frame-to-utterance convergence: A spectral-temporal approach for unified spoofing detection. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 10761–10765. IEEE.
- Awais Khan, Khalid Mahmood Malik, James Ryan, and Mikul Saravanan. 2022. Voice spoofing countermeasures: Taxonomy, state-of-the-art, experimental analysis of generalizability, open challenges, and the way forward. *arXiv preprint arXiv:2210.00417*.
- Yogesh Kumar, Apeksha Koul, and Chamkaur Singh. 2023. A deep learning approaches in text-to-speech system: a systematic review and recent research perspective. *Multimedia Tools and Applications*, 82(10):15171–15197.
- Cheng-I Lai, Nanxin Chen, Jesús Villalba, and Najim Dehak. 2019. Assert: Anti-spoofing with squeeze-excitation and residual networks. *arXiv preprint arXiv:1904.01120*.
- Galina Lavrentyeva, Sergey Novoselov, Egor Malykh, Alexander Kozlov, Oleg Kudashev, and Vadim Shchemelinin. 2017. Audio replay attack detection with deep learning frameworks. In *Interspeech*, pages 82–86.
- Jin Woo Lee, Eungbeom Kim, Junghyun Koo, and Kyogu Lee. 2022. Representation selective self-distillation and wav2vec 2.0 feature exploration for spoof-aware speaker verification. *arXiv preprint arXiv:2204.02639*.
- Zhenchun Lei, Yingen Yang, Changhong Liu, and Jihua Ye. 2020. Siamese convolutional neural network using gaussian probability feature for spoofing speech detection. In *Interspeech*, pages 1116–1120.
- Menglu Li, Yasaman Ahmadiadli, and Xiao-Ping Zhang. 2024. Audio anti-spoofing detection: A survey. *arXiv preprint arXiv:2404.13914*.
- Wei Li, Jichen Yang, and Pei Lin. 2023. Investigation of the influence of blocks on the linear spectrum for synthetic speech detection. *Electronics Letters*, 59(9):e12797.

- Tianchi Liu, Lin Zhang, Rohan Kumar Das, Yi Ma, Ruijie Tao, and Haizhou Li. 2024. How do neural spoofing countermeasures detect partially spoofed audio? *arXiv preprint arXiv:2406.02483*.
- Juan M Martín-Doñas and Aitor Álvarez. 2022. The vicomtech audio deepfake detection system based on wav2vec2 for the 2022 add challenge. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 9241–9245. IEEE.
- Abdelrahman Mohamed, Hung-yi Lee, Lasse Borgholt, Jakob D Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff, Shang-Wen Li, Karen Livescu, Lars Maaløe, et al. 2022. Self-supervised speech representation learning: A review. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1179–1210.
- Andreas Nautsch, Xin Wang, Nicholas Evans, Tomi H Kinnunen, Ville Vestman, Massimiliano Todisco, Héctor Delgado, Md Sahidullah, Junichi Yamagishi, and Kong Aik Lee. 2021. Asvspoof 2019: spoofing countermeasures for the detection of synthesized, converted and replayed speech. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(2):252–265.
- Koji Okabe, Takafumi Koshinaka, and Koichi Shinoda. 2018. Attentive statistics pooling for deep speaker embedding. *arXiv preprint arXiv:1803.10963*.
- Zihan Pan, Tianchi Liu, Hardik B Sailor, and Qionqiong Wang. 2024. Attentive merging of hidden embeddings from pre-trained speech model for anti-spoofing detection. *arXiv preprint arXiv:2406.10283*.
- Mirco Ravanelli and Yoshua Bengio. 2018. Speaker recognition from raw waveform with sincnet. In *2018 IEEE spoken language technology workshop (SLT)*, pages 1021–1028. IEEE.
- Md. Sahidullah, Tomi Kinnunen, and Cemal Hanilçi. 2015. A comparison of features for synthetic speech detection. In *Proc. Interspeech 2015*, pages 2087–2091.
- Yao Shi, Hui Bu, Xin Xu, Shaoji Zhang, and Ming Li. 2020. Aishell-3: A multi-speaker mandarin tts corpus and the baselines. *arXiv preprint arXiv:2010.11567*.
- Yi-Jen Shih, Zoi Gkalitsiou, Alexandros G Dimakis, and David Harwath. 2024. Self-supervised speech models for word-level stuttered speech detection. *arXiv preprint arXiv:2409.10704*.
- Hye-jin Shim, Jungwoo Heo, Jae-Han Park, Ga-Hui Lee, and Ha-Jin Yu. 2022. Graph attentive feature aggregation for text-independent speaker verification. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7972–7976. IEEE.
- Suwon Shon, Ankita Pasad, Felix Wu, Pablo Brusco, Yoav Artzi, Karen Livescu, and Kyu J Han. 2022. Slue: New benchmark tasks for spoken language understanding evaluation on natural speech. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7927–7931. IEEE.
- Kaavya Sriskandaraja, Vidhyasaharan Sethu, Phu Ngoc Le, and Eliathamby Ambikairajah. 2016. Investigation of sub-band discriminative information between spoofed and genuine speech. In *Interspeech*, pages 1710–1714.
- Hemlata Tak, Jose Patino, Andreas Nautsch, Nicholas Evans, and Massimiliano Todisco. 2020. An explainability study of the constant q cepstral coefficient spoofing countermeasure for automatic speaker verification. *arXiv preprint arXiv:2004.06422*.
- Hemlata Tak, Jose Patino, Massimiliano Todisco, Andreas Nautsch, Nicholas Evans, and Anthony Larcher. 2021. End-to-end anti-spoofing with rawnet2. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6369–6373. IEEE.
- Hemlata Tak, Massimiliano Todisco, Xin Wang, Jee-weon Jung, Junichi Yamagishi, and Nicholas Evans. 2022. Automatic speaker verification spoofing and deepfake detection using wav2vec 2.0 and data augmentation. *arXiv preprint arXiv:2202.12233*.
- Massimiliano Todisco, Héctor Delgado, and Nicholas Evans. 2017. Constant q cepstral coefficients: A spoofing countermeasure for automatic speaker verification. *Computer Speech & Language*, 45:516–535.
- Anton Tomilov, Aleksei Svishchev, Marina Volkova, Artem Chirkovskiy, Alexander Kondratev, and Galina Lavrentyeva. 2021. Stc antispoofing systems for the asvspoof2021 challenge. In *Proc. ASVspoof 2021 Workshop*, pages 61–67.
- Hsiang-Sheng Tsai, Heng-Jui Chang, Wen-Chin Huang, Zili Huang, Kushal Lakhota, Shu-wen Yang, Shuyan Dong, Andy T Liu, Cheng-I Jeff Lai, Jiatong Shi, et al. 2022. Superb-sg: Enhanced speech processing universal performance benchmark for semantic and generative capabilities. *arXiv preprint arXiv:2203.06849*.
- Xin Wang and Junichi Yamagishi. 2021. Investigating self-supervised front ends for speech spoofing countermeasures. *arXiv preprint arXiv:2111.07725*.
- Zhiyong Wang, Ruibo Fu, Zhengqi Wen, Jianhua Tao, Xiaopeng Wang, Yuankun Xie, Xin Qi, Shuchen Shi, Yi Lu, Yukun Liu, et al. 2024. Mixture of experts fusion for fake audio detection using frozen wav2vec 2.0. *arXiv preprint arXiv:2409.11909*.
- Yang Xie, Zhenchuan Zhang, and Yingchun Yang. 2021. Siamese network with wav2vec feature for spoofing speech detection. In *Interspeech*, pages 4269–4273.

Yuankun Xie, Yi Lu, Ruibo Fu, Zhengqi Wen, Zhiyong Wang, Jianhua Tao, Xin Qi, Xiaopeng Wang, Yukun Liu, Haonan Cheng, Long Ye, and Yi Sun. 2024. [The codefake dataset and countermeasures for the universal detection of deepfake audio](#). *Preprint*, arXiv:2405.04880.

Junichi Yamagishi, Xin Wang, Massimiliano Todisco, Md Sahidullah, Jose Patino, Andreas Nautsch, Xuechen Liu, Kong Aik Lee, Tomi Kinnunen, Nicholas Evans, et al. 2021. *Asvspoof 2021: accelerating progress in spoofed and deepfake speech detection*. In *ASVspoof 2021 Workshop-Automatic Speaker Verification and Spoofing Countermeasures Challenge*.

Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhota, Yist Y Lin, Andy T Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, et al. 2021. *Superb: Speech processing universal performance benchmark*. *arXiv preprint arXiv:2105.01051*.

Jiangyan Yi, Ye Bai, Jianhua Tao, Haoxin Ma, Zhengkun Tian, Chenglong Wang, Tao Wang, and Ruibo Fu. 2021. *Half-truth: A partially fake audio detection dataset*. *arXiv preprint arXiv:2104.03617*.

Jiangyan Yi, Ruibo Fu, Jianhua Tao, Shuai Nie, Haoxin Ma, Chenglong Wang, Tao Wang, Zhengkun Tian, Ye Bai, Cunhang Fan, et al. 2022. *Add 2022: the first audio deep synthesis detection challenge*. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 9216–9220. IEEE.

Jiangyan Yi, Jianhua Tao, Ruibo Fu, Xinrui Yan, Chenglong Wang, Tao Wang, Chu Yuan Zhang, Xiaohui Zhang, Yan Zhao, Yong Ren, et al. 2023a. *Add 2023: the second audio deepfake detection challenge*. *arXiv preprint arXiv:2305.13774*.

Jiangyan Yi, Chenglong Wang, Jianhua Tao, Chu Yuan Zhang, Cunhang Fan, Zhengkun Tian, Haoxin Ma, and Ruibo Fu. 2024. *Scenefake: An initial dataset and benchmarks for scene fake audio detection*. *Pattern Recognition*, 152:110468.

Jiangyan Yi, Chenglong Wang, Jianhua Tao, Xiaohui Zhang, Chu Yuan Zhang, and Yan Zhao. 2023b. *Audio deepfake detection: A survey*. *arXiv preprint arXiv:2308.14970*.

Lin Zhang, Xin Wang, Erica Cooper, Junichi Yamagishi, Jose Patino, and Nicholas Evans. 2021. *An initial investigation for detecting partially spoofed audio*. *arXiv preprint arXiv:2104.02518*.

You Zhang, Yongyi Zang, Jiatong Shi, Ryuichi Yamamoto, Jionghao Han, Yuxun Tang, Tomoki Toda, and Zhiyao Duan. 2024. *Svdd challenge 2024: A singing voice deepfake detection challenge evaluation plan*. *arXiv preprint arXiv:2405.05244*.

Yuxiang Zhang<sup>12</sup>, Wenchao Wang<sup>12</sup>, and Pengyuan Zhang<sup>12</sup>. 2021. *The effect of silence and dual-band fusion in anti-spoofing system*. In *Proc. Interspeech*.

## A Setup Description

In this study, various types of audio deepfakes are utilized to evaluate the performance of SSL. These deepfakes are categorized as follows:

**Full Fake:** The entire utterance is generated using voice conversion (VC) or text-to-speech (TTS) technologies. No part of the original speech remains intact.

**Partial Fake:** Only specific segments or words from the original utterance are replaced with VC or TTS-generated content. The remainder of the utterance is real.

**Song:** These setups involve utterances accompanied by background music, with corresponding speech content either real or deepfaked.

**Scene:** In these cases, the speech itself is real, but the background environment or context has been altered (e.g., a different ambiance or noise is introduced to simulate a new scene).

## B Datasets

In this section, we provide information about different datasets used in the study:

### B.1 ASVSpooF19 (LA19):

Our work focuses on the LA subset of the ASVSpooF19 database, which contains both bonafide (real) and spoofed speech (deepfake) data. The spoofed data is generated using 17 different TTS and Voice VC systems. While the TTS and VC systems were trained on data from the VCTK<sup>6</sup> database, there is no overlap with the data in the 2019 database. Of the 17 systems, 6 are categorized as known attacks, and the remaining 11 are classified as unknown. The training and development sets consist exclusively of known attacks, while the evaluation set includes 2 known attacks and 11 unknown spoofing attacks. Among the 6 known attacks, 2 are VC systems and 4 are TTS systems.

### B.2 ASVSpooF21 (LA21 & DF21):

The ASVSpooF21 evaluation dataset contains bonafide and spoofed speech transmitted over various telephony systems, including VoIP and PSTN. While no additional noise is introduced, transmission introduces variability and artifacts from both spoofing and encoding. Unlike previous LA tasks, LA21 includes acoustic propagation, with all data featuring reverberation and noise. The evaluation

<sup>6</sup><https://datashare.ed.ac.uk/handle/10283/2651>



set includes recordings from the same 48 speakers (21 male, 27 female) as ASVspoof 2019. The DF21 task simulates an attack where the attacker has access to a victim’s voice data, such as audio from social media. The dataset is sourced from the VCTK corpus and additional undisclosed datasets.

### **B.3 ADD23 (Track1.2):**

ADD23 Track 1.2 focuses on detecting fake utterances, particularly those generated in Track 1.1 (from the ADD23 challenge). The training and development sets are the same as those used in Track 3.2 of ADD 2022 (Yi et al., 2022), comprising both real and fake utterances based on the AISHELL-3 dataset (Shi et al., 2020). The evaluation phase includes two distinct datasets, referred to as ADD23.1.2.1 and ADD23.1.2.2.

### **B.4 HABLA:**

The HABLA dataset is a Spanish-language anti-spoofing corpus that includes accents from Argentina, Colombia, Peru, Venezuela, and Chile. It contains more than 22,000 genuine speech samples from both male and female speakers across these five countries, alongside 58,000 spoofed samples generated using six different speech synthesis methods. We use the train/dev/eval split provided by the dataset authors.

### **B.5 PartialSpoof:**

The PartialSpoof database is derived from the ASVspoof 2019 LA dataset, which includes 17 types of spoofed data generated by advanced speech synthesizers, voice converters, and hybrid systems. The same bonafide data from the ASVspoof19 LA corpus is used to create partially spoofed audio by following these steps: first, voice activity detection (VAD) algorithms are applied, and based on the detected boundaries, a randomly selected segment of a bonafide utterance are replaced with a spoofed segment.

### **B.6 Half-Truth (HAD):**

The HAD (Half-Truth Audio Detection) dataset is focused on detecting partially fake audio, where only a small portion of an utterance—such as individual words—is altered using advanced TTS systems. It is based on the AISHELL-3 speech corpus, which includes multi-speaker recordings designed for TTS model training. The dataset provides the challenging task of identifying subtle manipula-

tions in audio where most of the content remains genuine.

### **B.7 CtrSVDD:**

CtrSVDD is a diverse dataset focused on detecting deepfake singing vocals. It consists of 47.64 hours of bona fide singing and 260.34 hours of deepfake vocals generated using state-of-the-art synthesis techniques. Deepfake songs span 14 different methods, including 7 speech synthesis and 7 voice conversion (VC) techniques. The dataset includes 164 singer identities and provides a total of 188,486 deepfake song clips and 32,312 genuine song clips, with an average clip length of 5.02 seconds. CtrSVDD offers a rich resource for studying deepfake detection in singing, making it a valuable tool for advancing research in this area.

### **B.8 SceneFake:**

SceneFake is a dataset developed for scene fake audio detection, focusing on manipulated audio samples generated by altering the acoustic scene of real utterances using advanced speech enhancement technologies. Unlike traditional datasets that modify timbre, prosody, or linguistic content, SceneFake specifically addresses the challenge of detecting audio where the background environment has been tampered with. The dataset includes a variety of real utterances paired with their manipulated counterparts, and initial experiments indicate that existing models trained on datasets like ASVspoof19 struggle to reliably detect scene fake utterances.