



# Improved DeepFake Detection Using Whisper Features

Piotr Kawa<sup>1</sup>, Marcin Plata<sup>1</sup>, Michał Czuba<sup>1</sup>, Piotr Szymański<sup>1</sup>, Piotr Syga<sup>1</sup>

<sup>1</sup>Wrocław University of Science and Technology, Wrocław, Poland

{piotr.kawa, marcin.plata, michal.czuba, piotr.szymanski, piotr.syga}@pwr.edu.pl

## Abstract

With a recent influx of voice generation methods, the threat introduced by audio DeepFake (DF) is ever-increasing. Several different detection methods have been presented as a countermeasure. Many methods are based on so-called front-ends, which, by transforming the raw audio, emphasize features crucial for assessing the genuineness of the audio sample. Our contribution contains investigating the influence of the state-of-the-art Whisper automatic speech recognition model as a DF detection front-end. We compare various combinations of Whisper and well-established front-ends by training 3 detection models (LCNN, SpecRNet, and MesoNet) on a widely used ASVspoof 2021 DF dataset and later evaluating them on the DF In-The-Wild dataset. We show that using Whisper-based features improves the detection for each model and outperforms recent results on the In-The-Wild dataset by reducing Equal Error Rate by 21%.

**Index Terms:** audio DeepFake, DeepFake detection, feature extraction, Whisper

## 1. Introduction

Audio DeepFakes (DF) is a collection of deep learning techniques that create artificial speech. These methods may involve creating entirely new sentences using Text-To-Speech or Voice-Cloning (aiming at mimicking speech patterns of a specific person or sounding *natural* to a human listener) or transferring the qualities of the victim's voice to the attacker's speech, referred to as Voice-Conversion [1]. With the increasing sophistication of deep learning techniques, it has become relatively easy to create audio DeepFakes that are difficult to distinguish from bona fide recordings. Such malicious activities can cause significant harm, including compromising the security of systems protected by speaker recognition and contributing to spreading fake news or defaming an individual's reputation. As a result, developing effective DF detection techniques has become increasingly critical for ensuring the integrity and trustworthiness of audio-based communication systems. This need has been reflected in the recent growth of methods based on deep neural networks that assess the validity of utterances [2, 3, 4, 5, 6, 7].

Detecting DeepFake audio is a problem analogous to speech spoofing [8, 9]. Despite the superficial similarity, they differ in the targets they deceive: spoofing aims to fool speaker verification systems, whereas DF targets humans. One can easily point out attacks typical for one of the areas (e.g., replay attack for spoofing) that are not present in the other — therefore, these areas are considered separate.

Audio feature extraction is crucial in many applications, like speech recognition or speaker identification. Many existing approaches to DeepFake detection focus on some extracted

features instead of a raw waveform. That makes the extraction method vital to DF detection and generates motivation for in-detail investigation. Feature extraction aims to identify an audio signal's key characteristics and emphasize them. Mel-frequency cepstral coefficients (MFCC) [10] and linear-frequency cepstral coefficients (LFCC) [11] are some of the most widely used methods. MFCCs are based on the human auditory system's non-linear frequency response. In contrast, LFCCs are designed to address some of the limitations of MFCCs, such as their insensitivity to low-frequency information and lack of robustness to noise.

Whisper [12] is a state-of-the-art automatic speech recognition (ASR) system. It was trained on 680,000 hours of content. Due to the data's diversity and magnitude, it has shown to be robust against broad spectra of background interferences, accents, and languages. Its name refers to the family of the models differing, i.a., by width and sizes of layers. Whisper is based on off-the-shelf encoder-decoder Transformer architecture [13]. Its encoder is based on two convolutional layers, each processed by a GeLU activation function [14]. The information is later modified by adding position embeddings [13]. The encoder ends with a series of the pre-activation residual attention blocks [15], followed by the normalization layer.

In this work, we harness the feature extraction capabilities of the pre-trained Whisper's encoder not to capture speech properties later used for ASR but to investigate its performance in DF detection. We use it along with three detection models to verify if neural network-provided features might help in DF detection. We selected Whisper for the evaluation due to its effectiveness in speech recognition, which comes from the large and diversified speech corpora that it was trained on. As such, we infer that Whisper's features would ignore most of the naturally occurring artefacts and help identify artificially modified speech samples. In particular, it would help with the problem of generalization, which refers to the poor efficacy of the models on the data outside of the training set's distribution — currently one of the most challenging problems in DF detection. Our experiments cover the smallest available Whisper version — *tiny.en*. We aim at minimal overhead to ensure that the provided solution can be widely used in production environments. In addition, the model was trained strictly on English data, which is the main language among DF detection datasets. Please note that bigger Whisper versions were shown to yield even more satisfying results in various tasks (e.g. *large* performs up to 3x better [12] in speech recognition or translation than the *tiny.en* model). This allows us to expect that the results can be enhanced even further. Note that Whisper was trained on human speech samples — a bona fide and thus highly-biased set in the sense of DF detection. However, a similar approach was proposed in [7]. The authors used a front-end based on wav2vec 2.0 [16] that

was originally designed as an unsupervised pre-trained model for representations and used in the task of speech recognition. wav2vec was also trained on data considered bona fide. Fine-tuning this front-end led to a substantial increase in the results and generalization [17, 18]. We choose Whisper for our approach as it significantly improves over wav2vec, not only in the results reported but also in the data scale used for training and self-supervision (over 16x more). While this may play a lesser role for speech recognition as the ratio of importance between the audio model and language model can differ between approaches, we treat Whisper as an audio encoder only and thus expect to see an impact of a much larger dataset used in training.

The codebase related to our research can be found on GitHub: [github.com/piotrkawa/deepfake-whisper-features](https://github.com/piotrkawa/deepfake-whisper-features).

## 2. Detection models and datasets

We consider four models — three processing spectrogram-like features: LCNN [2], MesoNet [4] (MesoInception-4 variant), and SpecRNet [3] as well as RawNet3 [19] that analyzes raw audio. The models consist of respectively 467,425 (LCNN), 28,486 (MesoNet), 277,963 (SpecRNet), and 15,496,197 (RawNet3) parameters. SpecRNet used in our comparison differs from its original implementation — to enable processing of the higher-dimensions front-ends we add an adaptive 2D average pooling after the last SeLU layer [20]. A similar scenario occurred in MesoNet, where we add adaptive 1D average pooling right before the penultimate fully connected layer. In the case of LCNN, we increase the size from 160 to 768 of input features and hidden features of two bi-LSTM layers and the input features of the last Linear layer.

For the spectrogram-based models, we consider 3 front-ends: LFCC, MFCC, and the output of the Whisper ASR encoder. We additionally evaluate the concatenated front-ends of cepstral-coefficients with Whisper features. The intuition behind it follows [21] — concatenation of different front-ends may yield better results. We use LFCC and MFCC based on the window and hop lengths of 400 and 160; they are composed of 128 coefficients. We concatenate front-ends in a second dimension with its delta and double delta features. This results in the shape data (128 \* 3, 3000). In our experiments, we use *tiny.en* variant of the Whisper model. Its encoder contains 7,632,384 parameters and outputs data of shape (376, 1500). To match it with the size of the other front-ends (required for using it in the concatenated front-end setting), we replicate one of the dimensions achieving a tensor of size (376, 3000).

The dataset used in the paper consists of 125,000 samples randomly selected from ASVspoof 2021 (DF) [22] and all 31,779 samples of DeepFakes In-The-Wild [23]. The decision is motivated by the general scarcity of DF datasets, of which ASVspoof is among the largest and most popular. In contrast, the latter dataset consists of samples reflecting real-world scenarios (being gathered from the Internet). To emulate the scenario in which architectures are developed using training on the most popular datasets, they should be effective in the actual environment while determining the authenticity of new samples, possibly distorted by noise. Even though several models achieve high efficacy on popular datasets like ASVspoof or WaveFake [24], the investigation in [23] showed that those methods do not generalize well to unknown, real-world samples. The EER of LCNN evaluated on the In-The-Wild dataset increased by up to 1000%, and even more for RawNet3. Naturally, one of the solutions might be similar to the one presented in [21], mixing the datasets used so that multiple creation meth-

ods may be recognized. Such an approach may be infeasible in practical scenarios, where new DF creation methods should also be detected. To mimic the practical scenario in our investigation, we decided to train the models on a well-established DF dataset, as an end-user would, and later test it on the dataset that reflects a possible real-world sample to verify.

## 3. Experimental setup

Each sample underwent a standardized preprocessing procedure. It covered resampling to 16 kHz mono-channel, removing silences that were longer than 0.2 s and padding (by repetition), or trimming samples to 30 s of content. For two reasons, we decided on the input length of the 30 s instead of the typical length of about 4 s ([5, 21, 24, 3]). Firstly, works like [23] show that analyzing longer utterances yields better results. Secondly, Whisper takes as an input 30 s of content: samples are trimmed or padded with zeros. Instead of padding it with zeros, we decided to fill the whole input tensor with speech.

We trained models on a random subset of 100,000 training and 25,000 validation samples of the ASVspoof 2021 DF dataset. We used a subset of this dataset for two reasons: we wanted to make our solution possible to train on a single GPU in about 24 hours; moreover — we did not anticipate a significant gain from the samples’ quantity for an architecture like Whisper *tiny*. We addressed the disproportion between bona fide and fake classes with oversampling. We used a learning rate of  $10^{-4}$  and a weight decay of  $10^{-4}$  for all spectrogram-based models. RawNet3 used a learning rate of  $10^{-3}$  and a weight decay of  $5 \cdot 10^{-5}$ . We trained models with a binary cross-entropy function for 10 epochs with a batch size of 8. Training of RawNet3 included SGDR scheduling [25] with a restart after each epoch. The checkpoint of the highest validation accuracy was selected for later tests on the full In-The-Wild dataset. We present our results using Equal Error Rate (EER) metric as a fraction. EER is commonly used in DF and spoofing problems. To ensure reproducibility, we ran each process with a fixed randomness seed. Each experiment was run on a single NVIDIA TITAN RTX GPU (24GB VRAM).

## 4. Benchmarks

### 4.1. Baseline comparison

Our baseline comparison covered LCNN, MesoNet, SpecRNet and RawNet3 models. We tested front-ends of LFCC, MFCC and using Whisper’s encoder. The encoder was not optimized (the weights were frozen), i.e., we used it purely as a feature extractor and based solely on its pre-trained features.

Note that the results presented in Tab. 1, similarly to those reported in [23], deviate from the low errors typically reported in the DeepFake detection literature. This phenomenon is caused by training the models on an ‘artificial’ dataset created in a controlled manner (ASVspoof), whereas the evaluation is done on real-world samples from In-The-Wild dataset [23]. The distribution of the artifacts in both sets differs significantly. As shown, the models do not have sufficient generalization capabilities, and when verified on recordings of a substantially different nature, the detection capabilities deteriorate significantly. Additionally note that, both, in the case of [23] that was trained on LA subset of ASVspoof 2019 [26], and in this paper (trained on DF subset of ASVspoof 2021), some of the models perform worse than random guessing (EER=0.5). These results do not undermine the models in the traditional setup. In fact,

Table 1: The comparison of EER scores on In-The-Wild dataset. Using Whisper’s encoder as a front-end contributes to the significant enhancement in case of SpecRNet and LCNN networks.

Model	Front-end	EER
SpecRNet	LFCC	0.5184
SpecRNet	MFCC	0.6897
SpecRNet	Whisper	0.3644
LCNN	LFCC	0.7756
LCNN	MFCC	0.6762
LCNN	Whisper	0.3567
MesoNet	LFCC	0.5451
MesoNet	MFCC	0.3132
MesoNet	Whisper	0.3856
RawNet3	-	0.5199

LCNN with LFCC front-end and SpecRNet with MFCC, i.e., two architectures that achieved the worst results on In-The-Wild dataset, scored a satisfying EERs of 0.0149 and 0.0218 during validation on ASVspoof 2021. Notably, in the case of all detection models, LFCC and MFCC front-ends tend to provide features that were well-suited in the case of ASVspoof, and when trained and verified on the data from the same source have high efficacy [27], yet do not occur regularly in the DeepFakes from the In-The-Wild dataset. We discuss the nature of the extracted features in more depth in Sect. 5. Interestingly, smaller architectures – SpecRNet and MesoNet, seem to generalize better and provide higher efficacy than LCNN. In fact, MesoNet (MFCC-variant) achieved the lowest EER. These results are similar to the ones reported in [23], where the model achieved the best results among the spectrogram-based architectures. One of the reasons may be the lower number of parameters, which results in a lesser degree of ‘adjusting’ towards the artifacts specific to the ASVspoof datasets, thus, higher generalization capabilities. Using Whisper-based features helps with generalizing. In the case of SpecRNet, we achieve a 29.71% improvement in comparison with LFCC and 47.17% in comparison with MFCC. In the case of LCNN, we improve both by 54% and 47.25%. Following the intuition that additional information may improve the detection, we investigate the synergy of the feature extractors in Sect. 4.2. Moreover, to improve the results even further and to address the worse results in the case of MesoNet, we decided to unfreeze the model (Sect. 4.3).

**Constant Q-cepstral coefficients** One of the popular spectrogram-based front-end used in speech and audio signal processing is Constant Q-cepstral coefficients (CQCC) [28]. Works like [8, 9] tested these features for spoofing detectors trained on ASVspoof 2021 (LA). To provide an extensive investigation of different front-ends, we used the CQCC with the LCNN model. However, when training the architecture with the same parameters as other feature extractors, the results on ASVspoof 2021 DF were unsatisfactory, achieving only around 60% accuracy on train and test datasets. Consequently, we did not proceed with training other models using CQCC or evaluating them on the In-The-Wild dataset.

#### 4.2. Concatenated front-ends

Works like [21] showed that using the concatenation of multiple front-ends could increase the detectors’ effectiveness. The discussed pipeline did not differ from the one in Sect. 4. We considered spectrogram-based models and used them with a con-

catenation of the classical front-ends and Whisper’s encoder.

Table 2: The EER values of the models using concatenated front-end features. The only enhancement is visible in case of two models, whereas for other models it had a negative impact.

Model	Front-end	EER
SpecRNet	Whisper + LFCC	0.3485
SpecRNet	Whisper + MFCC	0.4116
LCNN	Whisper + LFCC	0.6270
LCNN	Whisper + MFCC	0.6117
MesoNet	Whisper + LFCC	0.8029
MesoNet	Whisper + MFCC	0.3822

Table 3: Evaluation of the fine-tuned Whisper feature extraction (as a sole front-end and in concatenation, similarly as in Sect. 4.2), on DeepFake In-The-Wild dataset. Note that ‘EER (frozen)’ refers to the case, where Whisper’s encoder was not fine-tuned (cf. Tab. 1, Tab. 2) and ‘EER (tuned)’ to the results of fine-tuned extractor.

Model	Front-end	EER (frozen)	EER (tuned)
SpecRNet	Whisper + LFCC	0.3485	0.3795
SpecRNet	Whisper + MFCC	0.4116	0.3769
SpecRNet	Whisper	0.3644	0.3338
LCNN	Whisper + LFCC	0.6270	0.6270
LCNN	Whisper + MFCC	0.6117	0.5899
LCNN	Whisper	0.3567	0.3290
MesoNet	Whisper + LFCC	0.8029	0.5526
MesoNet	Whisper + MFCC	0.3822	0.2672
MesoNet	Whisper	0.3856	0.3362

The comparison between the results of concatenated features (Tab. 2) and a single feature extractor (cf. Tab. 1) shows that LCNN and SpecRNet models based on cepstral front-ends improve when trained with Whisper features. Detection enhances by up to 40.32% for SpecRNet and up to 19.15% for LCNN. This suggests some positive synergy between the features and that additional knowledge is gained. However, the synergy does not guarantee that the results of joint-features detection outperform the Whisper-based. This may be caused by ‘covering’ some of the important Whisper features by the spectrogram-based front-ends. We investigate this issue in Sect. 5. Conversely, one may notice a negative synergy between the Whisper-extracted features and the others (not a substantial one in the case of MFCC and significant in the case of LFCC). One may assume that the antagonistic effect is due to the architecture details and ‘compression’ of the provided information, which results in performing the classification based on the ‘noised’ information rather than an enhanced set of features. In Sect. 5, we show how the extracted features may be presented and discuss the issue further.

#### 4.3. Whisper fine-tuning

The following benchmark concerned models using Whisper’s encoder. This time, however, we did not treat the encoder strictly as a front-end algorithm but rather fine-tuned it to the problem of DF detection. The intuition was the following — while the features produced by Whisper tend to provide better performance than the typically used front-ends (cf. Tab. 1), this model was trained for a different purpose (ASR), and on the biased (in the sense of DF) data. We assume that the encoder fine-




Figure 1: Saliency maps of front-ends of a bona fide sample for the MesoNet.




Figure 2: Gradient calculated on a spoofed signal for the SpecRNet. Note that the y-axis ranges differ.

tuned to a specific task, in our case DeepFake detection, might yield even better results. For this purpose, we fine-tuned models using the Whisper front-end and evaluated the results with a fine-tuned version of the feature extractor. After the initial training presented in Sect. 4.1, we trained models for additional 5 epochs. This time, however, we unfroze Whisper layers and performed fine-tuning with a learning rate of  $10^{-6}$ .

One may note that for all architectures (with a notable exception of SpecRNet with Whisper and LFCC features, where we got results worse by less than 9%), the fine-tuning provided an improvement. Notably, unfrozen Whisper features allowed us to improve even the previous best result — MesoNet with MFCC features – by 14.69%. The best model we obtained, MesoNet with fine-tuned Whisper+MFCC, scored an EER of 0.2672. This surpasses the state-of-the-art results reported in [23], where authors obtained 0.3394 EER evaluating RawNet2 [5] on DeepFake In-The-Wild after training model on 4s samples from ASVspoof 2019 [26]. Our results indicate that unfreezing the model and using Whisper extracted features may improve the results of detecting DeepFakes from a significantly different distribution than the set the model was trained on, which would address the vital issue of generalization.

## 5. Features comparison

In order to check if the different front-ends indeed generate different features, we analyze which parts of the input data most significantly affect the detection results. Additionally, we compare two architectures that achieved the highest results — SpecRNet containing a recurrent layer (GRU) and MesoNet, which primarily consists of convolutional and max-pooling layers. We use a technique known from adversarial attacks [29] that relies on calculating the gradient on the input data.

We observed that the choice of the model’s architecture has great importance in processing front-ends. As the MesoNet consists of 4 max-pooling layers, the final linear layers of

the model receive *max-pooled* information from spatially distributed blocks of size equal to  $32 \times 32$  (see Fig. 1). In turn, information in SpecRNet is processed using the GRU, and the decision is mainly impacted by the end part of the signal (Fig. 2a). Having that said, the models utilizing Whisper features often rely on some characteristics extracted from one or more narrow signal slices (see two peaks around 220k and 360k in Fig. 2b). We suppose that Whisper works well with recurrent NN because it extracts *prominent* attributes that do not tend to be hidden, passing through the recurrent sequence. As we have not assessed any mechanism for spatially-independent processing of the combination of front-ends, our findings suggest a negative impact of Whisper features on other front-ends. Specifically, models that solely rely on MFCC or LFCC for decision-making tend to prioritize the lower band of the front-end and disregard delta and double-delta features (Fig. 1a and 1b). Conversely, the 2D Whisper features show the most significant impact of the full band, with a noticeable focus on the middle (Fig. 1c). Fine-tuning Whisper features obtain further performance improvement; surprisingly, in this case, the importance of delta and double-delta features increases (Fig. 1d). We suppose that using the Whisper front-end effectively captures the speech features, and including deltas could enhance the results by providing additional coefficients to describe the spectrum dynamic.

## 6. Summary

In this paper, we show that using Whisper [12] as a feature extractor in DeepFake detection may improve the efficacy of the detection architectures, particularly in the case of evaluation on samples with significantly different distribution than the training set. We trained the models on ASVspoof 2021 DF and evaluated them on the DeepFakes In-The-Wild dataset. Using fine-tuned Whisper as a sole feature extractor, we achieved EER of  $0.33 \pm 0.01$  for all the investigated architectures, which is better than the recent results reported in [23]. Moreover, using Whisper and MFCC joint features for MesoNet, we gained EER as low as 0.2672, showing that even *tiny.en* version of Whisper significantly helps in the generalization of DF detection scoring the state-of-the-art results on the In-The-Wild dataset. In future work, we would like to investigate more Whisper models, following the intuition that larger models used as feature extractors may improve the generalization. Additionally, we are interested in exploring the combinations of front-ends concerning architectures of various models.

## 7. Acknowledgements

This work has been partially funded by Department of Artificial Intelligence, Wrocław University of Science and Technology.

## 8. References

- [1] M. Masood, M. Nawaz, K. M. Malik, A. Javed, A. Irtaza, and H. Malik, "Deepfakes generation and detection: State-of-the-art, open challenges, countermeasures, and way forward," *Applied Intelligence*, pp. 1–53, 2022.
- [2] X. Wu, R. He, Z. Sun, and T. Tan, "A light cnn for deep face representation with noisy labels," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 11, pp. 2884–2896, 2018.
- [3] P. Kawa, M. Plata, and P. Syga, "Specrnet: Towards faster and more accessible audio deepfake detection," in *2022 IEEE 21st International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. IEEE, 2022.
- [4] D. Afchar, V. Nozick, J. Yamagishi, and I. Echizen, "Mesonet: a compact facial video forgery detection network," in *2018 IEEE International Workshop on Information Forensics and Security (WIFS)*, 2018, pp. 1–7.
- [5] H. Tak, J. Patino, M. Todisco, A. Nautsch, N. Evans, and A. Larcher, "End-to-end anti-spoofing with rawnet2," *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Jun 2021. [Online]. Available: <http://dx.doi.org/10.1109/ICASSP39728.2021.9414234>
- [6] J.-w. Jung, H.-S. Heo, H. Tak, H.-j. Shim, J. S. Chung, B.-J. Lee, H.-J. Yu, and N. Evans, "Aasist: Audio anti-spoofing using integrated spectro-temporal graph attention networks," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6367–6371.
- [7] H. Tak, M. Todisco, X. Wang, J. weon Jung, J. Yamagishi, and N. Evans, "Automatic Speaker Verification Spoofing and Deepfake Detection Using Wav2vec 2.0 and Data Augmentation," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2022)*, 2022, pp. 112–119.
- [8] T. Kinnunen, H. Delgado, N. Evans, K. A. Lee, V. Vestman, A. Nautsch, M. Todisco, X. Wang, M. Sahidullah, J. Yamagishi, and D. A. Reynolds, "Tandem assessment of spoofing countermeasures and automatic speaker verification: Fundamentals," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2195–2210, 2020.
- [9] X. Wang and J. Yamagishi, "Investigating Self-Supervised Front Ends for Speech Spoofing Countermeasures," in *Proc. The Speaker and Language Recognition Workshop (Odyssey 2022)*, 2022, pp. 100–106.
- [10] M. Sahidullah and G. Saha, "Design, analysis and experimental evaluation of block based transformation in mfcc computation for speaker recognition," *Speech Communication*, vol. 54, no. 4, pp. 543–565, 2012. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167639311001622>
- [11] S. Davis and P. Mermelstein, "Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 28, no. 4, pp. 357–366, 1980.
- [12] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," 2022. [Online]. Available: <https://arxiv.org/abs/2212.04356>
- [13] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [14] D. Hendrycks and K. Gimpel, "Gaussian error linear units (gelus)," *arXiv preprint arXiv:1606.08415*, 2016.
- [15] R. Child, S. Gray, A. Radford, and I. Sutskever, "Generating long sequences with sparse transformers," *arXiv preprint arXiv:1904.10509*, 2019.
- [16] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [17] C. Wang, J. Yi, J. Tao, H. Sun, X. Chen, Z. Tian, H. Ma, C. Fan, and R. Fu, "Fully automated end-to-end fake audio detection," in *Proceedings of the 1st International Workshop on Deepfake Detection for Audio Multimedia*, ser. DDAM '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 27–33. [Online]. Available: <https://doi.org/10.1145/3552466.3556530>
- [18] J. M. Martín-Doñas, I. G. Torre, A. Álvarez, and J. Arellano, "The vicomtech spoofing-aware biometric system for the sasv challenge," *arXiv preprint arXiv:2204.01399*, 2022.
- [19] J. weon Jung, Y. Kim, H.-S. Heo, B.-J. Lee, Y. Kwon, and J. S. Chung, "Pushing the limits of raw waveform speaker recognition," in *Proc. Interspeech 2022*, 2022, pp. 2228–2232.
- [20] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, "Self-normalizing neural networks," *Advances in neural information processing systems*, vol. 30, 2017.
- [21] P. Kawa, M. Plata, and P. Syga, "Attack agnostic dataset: Towards generalization and stabilization of audio deepfake detection," in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 4023–4027. [Online]. Available: <https://doi.org/10.21437/Interspeech.2022-10078>
- [22] J. Yamagishi, X. Wang, M. Todisco, M. Sahidullah, J. Patino, A. Nautsch, X. Liu, K. A. Lee, T. Kinnunen, N. Evans *et al.*, "Asvspoof 2021: accelerating progress in spoofed and deepfake speech detection," *arXiv preprint arXiv:2109.00537*, 2021.
- [23] N. M. Müller, P. Czempin, F. Dieckmann, A. Froghyar, and K. Böttinger, "Does audio deepfake detection generalize?" in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 2783–2787. [Online]. Available: <https://doi.org/10.21437/Interspeech.2022-10078>
- [24] J. Frank and L. Schönherr, "WaveFake: A Data Set to Facilitate Audio Deepfake Detection," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021.
- [25] I. Loshchilov and F. Hutter, "SGDR: stochastic gradient descent with warm restarts," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [26] M. Todisco, X. Wang, V. Vestman, M. Sahidullah, H. Delgado, A. Nautsch, J. Yamagishi, N. Evans, T. Kinnunen, and K. A. Lee, "Asvspoof 2019: Future horizons in spoofed and fake audio detection," *arXiv preprint arXiv:1904.05441*, 2019.
- [27] X. Liu, X. Wang, M. Sahidullah, J. Patino, H. Delgado, T. Kinnunen, M. Todisco, J. Yamagishi, N. Evans, A. Nautsch *et al.*, "Asvspoof 2021: Towards spoofed and deepfake speech detection in the wild," *arXiv preprint arXiv:2210.02437*, 2022.
- [28] M. Todisco, H. Delgado, and N. Evans, "Constant q cepstral coefficients: A spoofing countermeasure for automatic speaker verification," *Computer Speech & Language*, vol. 45, pp. 516–535, 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0885230816303114>
- [29] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=rJzIBfZAB>