Fully Automated End-to-End Fake Audio Detection

Chenglong Wang
University of Science and Technology
of China, Hefei, China
NLPR, Institute of Automation,
Chinese Academy of Sciences
Beijing, China
chenglong.wang@nlpr.ia.ac.cn

Jiangyan Yi
NLPR, Institute of Automation,
Chinese Academy of Sciences, China
School of Artificial Intelligence,
University of Chinese Academy of
Sciences
Beijing, China
jiangyan.yi@nlpr.ia.ac.cn

Jianhua Tao
NLPR, Institute of Automation,
Chinese Academy of Sciences, China
School of Artificial Intelligence,
University of Chinese Academy of
Sciences, China
CAS Center for Excellence in Brain
Science and Intelligence Technology
Beijing, China

Haiyang Sun
NLPR, Institute of Automation,
Chinese Academy of Sciences, China
School of Artificial Intelligence,
University of Chinese Academy of
Sciences
Beijing, China

Xun Chen
University of Science and Technology
of China
Hefei, China

Zhengkun Tian
NLPR, Institute of Automation,
Chinese Academy of Sciences, China
School of Artificial Intelligence,
University of Chinese Academy of
Sciences
Beijing, China

Haoxin Ma
NLPR, Institute of Automation,
Chinese Academy of Sciences, China
School of Artificial Intelligence,
University of Chinese Academy of
Sciences
Beijing, China

Cunhang Fan
School of Computer Science and
Technology, Anhui University
Hefei, China

Ruibo Fu
NLPR, Institute of Automation,
Chinese Academy of Sciences, China
School of Artificial Intelligence,
University of Chinese Academy of
Sciences
Beijing, China

ABSTRACT

The existing fake audio detection systems often rely on expert experience to design the acoustic features or manually design the hyperparameters of the network structure. However, artificial adjustment of the parameters can have a relatively obvious influence on the results. It is almost impossible to manually set the best set of parameters. Therefore this paper proposes a fully automated end-to-end fake audio detection method. We first use wav2vec pre-trained model to obtain a high-level representation of the speech. Furthermore, for the network structure, we use a modified version of the differentiable architecture search (DARTS) named light-DARTS. It learns deep speech representations while automatically learning and optimizing complex neural structures consisting of convolutional operations and residual blocks. The experimental results on the ASVspoof 2019 LA dataset show that our proposed system achieves an equal error rate (EER) of 1.08%, which outperforms the state-of-the-art single system.

CCS CONCEPTS

• **Security and privacy** → **Privacy protections**.

KEYWORDS

fake audio detection, ASVspoof, wav2vec, DARTS, self-supervised

1 INTRODUCTION

In the last few years, the development of speech synthesis technology [1–4] is increasing rapidly driven by deep learning. These models can generate high quality audio that is comparable to the human voice. Although the technology itself has no negative attributes, it is easy to be misused, such as faking audio to commit fraud and spread public opinion. Therefore, an increasing number of work [5–10] have been focused on detecting the fake audio recently.

In order to improve the performance of fake audio detection systems, recent works have focused on two aspects: improving the acoustic features of audio and designing effective classification models. It is particularly important to select acoustic features that

can effectively distinguish fake audios from real audios. Todisco et al. [11] apply the constant Q-transform instead of the short-time Fourier transform to process speech signals, which outperforms the Mel frequency cepstrum coefficients(MFCC). Sahidullah et al. [12] replace the Mel scale filters with linear filters and propose linear frequency cepstrum coefficients(LFCC), making it more focused on high frequency band features compared to MFCC. With the development of unsupervised pre-training, Yang et al. [13] first use the wav2vec pre-trained model as the feature extractor which obtains more robust acoustic feature. At the latest ADD2022 challenge [10], Donas et al. [14] also propose to use an improved version of wav2vec 2.0 as a feature extractor and win the first place on Track 1. The input to wav2vec feature extractor is the raw waveform, which learns speech information from a large amount of unlabeled speech that may contain information useful for fake audio detection.

Another approach is to design an effective classification model that learns a distinguished representation of real and fake audio. Gaussian mixture model (GMM) is the traditional classification model. With the development of deep learning, the performance of convolutional neural networks (CNN) [15, 16] is better than GMM. For example, light convolution neural network (LCNN) [17] with max feature map (MFM) activation function can not only separate noise signals from information signals by competitive learning, but also play a role in feature selection. The residual network (ResNet) [18] proposes a residual module that addresses the problem of network degradation. Compared with hand-designed approaches, differentiable architecture search (DARTS) [19] makes great achievements in the design of deep neural networks to automate the manual process of architecture design. Ge et al. [20] first successfully use a Partially-Connected DARTS approach to the fake audio detection tasks. Although Ge et al. [20] implement automation in the network structure, they choose LFCC and fast Fourier transform (FFT) in the feature level, which still requires manual parameter setting. For example, LFCC features need to set the window length, fft transform length, filter set coefficients and so on. However, artificial adjustment of the parameters can have a relatively obvious influence on the results. It is almost impossible to manually set the best set of parameters. Therefore, it's necessary to design a fully automated fake audio detection system.

To address these problems, we propose a fully automated end-to-end fake audio detection. More generally, we use pre-trained wav2vec model as feature extractor instead of traditional acoustic features. Inspired by light-CNN [21], we propose a light-DARTS based on DARTS to learn deep speech representations automatically. As shown in Figure 1 (d), we apply the max feature map (MFM) module to DARTS, which plays the role of feature selection. The main contributions of this study can be summarized as follows:

- To our best knowledge, this is the first work to propose light-DARTS for fake audio detection task.
- We propose a a fully automated end-to-end fake audio detection method. We use wav2vec features as inputs for ligth-DARTS.

The experimental results on the ASVspoof 2019 LA dataset show that our proposed method can acquire the EER of 1.08%. This result demonstrates the effectiveness of our proposed method.

The rest of this paper is organized as follows: Section 2 presents the related work. Section 3 illustrates the our proposed method. Experiments, results and discussions are reported in Section 4 and 5, respectively. Finally, the paper is concluded in Section 6.

2 RELATED WORK

2.1 DARTS

Compared with hand-designed approaches, DARTS [3] has made great achievements in the design of deep neural networks to automate the manual process of architecture design. Liu et al. [19] proposed the DARTS algorithm to use softmax to serialize the discrete search space for gradient-based structural search of neural networks. Instead of fixing the execution of a single operation at a specific layer, the authors define the structure search as a convex optimization problem by selecting the optimal operation from a set of operations and then alternating the operational parameters of the network and the structural parameters corresponding to the operations by gradient descent. Specifically, the learnable parameters of the operations in the network are optimized on the training dataset and the structural parameters corresponding to the different operations are trained on the validation dataset. Ultimately, a sparse network structure is obtained by solving this convex optimization problem to select the most appropriate operation for each layer.

### 2.2 Self-supervised Pretraining Models

Self-supervised pre-training becomes a hot issue in recent years and has been applied in the field of natural language processing and computer vision already. For speech processing, more and more self-supervised speech models are proposed, such as wav2vec [22, 23], HuBERT [24], wav2vec 2.0 which allowing to learn high-level representations of the speech signal. The above methods have been successfully applied in the fields of speech recognition, speaker recognition and emotion recognition [25, 26]. As for fake audio detection, only a few works have explored this similar approach [14, 27, 28]. Motivated by [27], we propose to use a pre-trained wav2vec style model as a feature extractor instead of traditional acoustic features, which solves exactly the problem of acoustic features requiring manual parameter setting. Furthermore, the features extracted from models pretrained on large-scale datasets help reduce overfitting, thereby improving reliability and domain robustness.

3 OUR PROPOSED SYSTEM

As shown in Figure 1 (a), the fully automated end-to-end fake audio detection consists of two modules: (1) The feature extraction module: It consists of a wav2vec feature extractor, which obtains a high-level representation of the speech. (2) The Light-DARTS module: We apply the max feature map (MFM) module to DARTS, which plays the role of feature selection.

3.1 Wav2vec Feature Extractor

The input to wav2vec [23] feature extractor is the raw waveform, which learns speech information from a large amount of unlabeled speech that may contain information useful for fake audio detection. In addition, pre-trained models trained from a large amount of

(a) Overall framework of Fully Automated End-to-End Fake Audio Detection

(b) Wav2vec feature extractor   (c) Wav2vec 2.0 feature extractor   (d) L-DARTS module
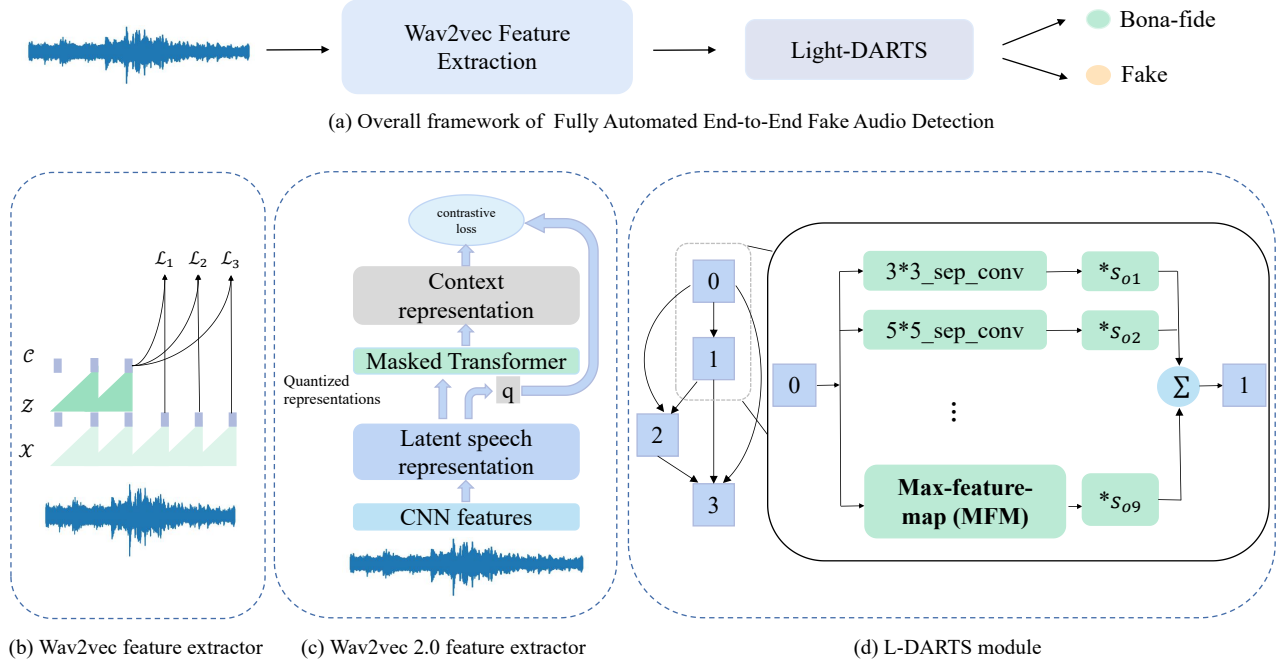
Figure 1: (a) Overall framework of the fully automated fake audio detection system. The system consists of the feature extraction module and the light-DARTS module. (b) Illustration of the wav2vec feature exctraction module. (c) Illustration of the wav2vec 2.0 feature exctraction module. (d) Illustration of the light-DARTS module. We apply the max feature map (MFM) module to DARTS.

unlabeled information can yield more generalized generic speech features, which may also be useful for distinguishing real from fake speech. Therefore, wav2vec is used as a feature extractor in this paper.

3.1.1 Pretrained wav2vec. As shown in Figure 1 (b), the wav2vec model contains two convolutional neural networks, an encoder network that maps the original input audio signal to a hidden space, and a context network that combines multiple time step outputs of the encoder network. The encoder generates a representation z_i for each time step i, while the context network combines the multiple encoder time step outputs into a new representation c_i for each time step i.

3.1.2 Pretrained wav2vec 2.0. As shown in Figure 1 (c), the wav2vec 2.0 [22] model includes three stages. Firstly, the raw speech is send into a local encoder which contains several convolutional layers (CNN). The stride of the local encoder is 20ms and the receptive field is 25ms. The local encoder embeddings are then fed as input to the context encoder, which consists of several transformer modules. In the open source two pre-trained models, the base model contains 12 transformer blocks with 8 attention heads per block, while the large model contains 24 transformer blocks with 16 attention heads per block. At the third stage, the local encoder representation is input to the quantization module, which consists of 2 codebooks, each with 320 entries. The linear mapping then converts the local encoder representation into logits. Given the logits, Gumbel-Softmax [29]

is applied to a sample of each codebook. The selected codes are concatenated and the resulting vectors are linearly transformed to obtain a quantized representation of the local encoder output, which is used in the objective function. The high-level embeddings is represented as:

$$\hat{X}_{embedding}(t) = W2V(X(t)) \in \mathbb{R}^{T*F} \tag{1}$$

where $W2V(\cdot)$ denotes the pretrained wav2vec model or wav2vec2.0 model, which aims to obtain high-level representations $\hat{X}_{embedding}(t)$ of raw speech $X(t)$.

3.2 Light-DARTS

Inspired by light-CNN [21], we propose a light-DARTS based on DARTS to learn representations of wav2vec features automatically. Search space, using a cell-based approach for spatial definition, finds a computational cell as the building block of the final architecture, and the learned cells can be stacked to form a convolutional network. By training out one small cell to form a large network. Suppose the node is represented as $x^{(i)}$, then each directed edge (i, j) will be associated with a node $x^{(i)}$ with some operation $o^{(i,j)}$ associated with it. To construct a continuous search space and generate an architecture, the predecessors and edges must be computed, i.e., a continuous relaxation optimization must be performed. First, the intermediate nodes based on the predecessor nodes are calculated

as follows.

$$x^{(j)} = \sum_{i<j} o^{(i,j)}(x^{(i)}) \tag{2}$$

where o denotes a set of candidate operations (e.g., convolution, max-pool, zero), and each operation denotes a function $o(\cdot)$ to be applied to the node $x^{(i)}$.

Secondly, in order to make the search space continuous, we relax the classification choice of a particular operation to $softmax$, among all possible operations, by the following formula.

$$\bar{o}^{(i,j)}(x) = \sum_{o \in O} \frac{\exp(\alpha_o^{(i,j)})}{\sum_{o' \in O} \exp(\alpha_{o'}^{(i,j)})} o(x) \tag{3}$$

where the operational mixing weights of a pair of nodes (i, j) are parameterized by a vector of dimension $|O|$ parametrized. The task of architectural search will be reduced to learning a set of continuous variables $\alpha = \alpha^{(i,j)}$. Finally replacing each mixed operation with the most probable operation $\bar{o}^{(i,j)}$, to obtain a discrete architecture. The formula is as follows.

$$o^{(i,j)} = \arg\max_{o \in O} \alpha_o^{(i,j)} \tag{4}$$

Then, light-DARTS jointly learns operation parameters α and weight parameters w (e.g., convolution kernels) using the gradient descent method by solving a bilevel optimization problem, where operation parameters α are upper-level variables and weight parameters w are lower-level variables:

$$\min_{\alpha} \mathcal{L}_{val}(\omega^*(\alpha), \alpha)$$
$$\omega^*(\alpha) = \arg\min_{\omega} \mathcal{L}_{train}(\omega, \alpha) \tag{5}$$

$\mathcal{L}_{train}$ and $\mathcal{L}_{val}$ denote the training and validation losses.

The key to handling edges is the combination of node connections and activation functions into a matrix, where each element represents the weight of the connection and activation function. When searching, light-DARTS traverses all nodes and finds the average weight of the relevant connections, thus turning the search space into a continuous space. The objective function becomes differentiable, and both the structure weights and the network weights are optimized. At the end of the search, the connection and activation function with the highest weight are selected to form the final network.

Furthermore, we introduce the MFM module in the architecture search space, which plays the role of feature selection. The MFM operation which combines two feature maps and outputs element-wise maximum one as shown in Fig. 2 can be written as:

$$h(x) = \max(x_1, x_2) \tag{6}$$

By using MFM, we select and retain 50% of the information from the input feature maps through element-level maximum operations across feature channels.

We define all possible candidate operations in a neural cell in the search space O, containing the following nine operations.

- 3*3 separable convolutions
- 5*5 separable convolutions
- 3*3 dilated convolutions
- 5*5 dilated convolutions
- 3*3 average pooling
- 3*3 max pooling
- skip connnection
- zero
- max feature map

4 EXPERIMENTS

4.1 Dataset

4.1.1 ASVspoof 2019 challenge database. Our focus is on the interference of speech synthesis and speech conversion with real speech, so we use the ASVspoof2019 logical access (LA) dataset [8]. The LA data set contains three subsets: training set, development set and evaluation set.

4.1.2 ASVspoof 2021 challenge database. We focused on the deep fake (DF) tracks for the ASVSpoof2021 challenge [7], where the training and development sets are the same as the ASVspoof 2019 LA database, but the evaluation set differs. The DF evaluation set shows audio coding and compression artifacts which includes approximately 600K of audio processing with various commercial audio codecs.

4.1.3 ADD 2022 challenge database. We focused on the low-quality fake audio detection(LF) track for the ADD2022 challenge. The train and development sets contain clean speech based on the multi-speaker Mandarin speech corpus AISHELL-3 [19]. The training set contains 27084 audio and the development set contains 28324 audio. The evaluation set contains a variety of ambient noise and background music.

4.2 Evaluation Metrics

In this work, in order to evaluate the results of different fake audio detection systems, the equal error rate (EER) is used as the evaluation metrics. The EER is the operating point where the false rejection rate (FRR) and false acceptance rate (FAR) are equal.

4.3 Experimental Setup

The wav2vec pretrained model variant "wav2vec large", which we use as a pretrained feature extractor using additional linear transformations and a larger context network, is trained on 960 hours Librispeech [30]. The model's downsampling factor is 160. Thus, there is a 512-dimensional vector for every 10 ms of speech. To form batches, the 400 time frames are fixed by truncating or concatenating. Therefore, the shape of wav2vec feature is $400 * 512$. For wav2vec 2.0 feature extractor, we have adopted two versions: 'wav2vec 2.0 base' and 'wav2vec 2.0 large'. The base model contains 768 model dimension, 12 transformer blocks, inner dimension (FFN) 3,072 and 8 attention heads. The large model contains 24 transformer blocks with model dimension 1,024, inner dimension 4,096 and 16 attention heads. Like the wav2vec features, we also performed truncating or concatenating in the temporal dimension. Therefore, the shape of wav2vec 2.0 base and large are $400 * 768$ and $400 * 1024$, respectively.

In the search stage, we followed the same experimental setup as the original DARTS [19] study. The neural cells were divided into two categories, namely normal cells and reduction cells. The number of cells N is 8, contains 6 normal cells and 2 reduction cells. The reduction cells were inserted into the 1/3 and 2/3 locations of the entire network. At each step, the model is trained for 50 epochs by Adam optimiser with the learning rate 0.0001.

Table 1: The results of EER(%) for our proposed different systems on ASVspoof2019 LA.

Feature	Network Architecture	Dev	Eval
LFCC	LCNN	0.13	4.75
Wav2vec	LCNN	0.03	3.51
Wav2vec 2.0-base	LCNN	0.02	3.32
Wav2vec 2.0-large	LCNN	0.14	3.56
LFCC	DARTS	0.01	4.82
Wav2vec	DARTS	0.02	2.18
Wav2vec 2.0-base	DARTS	0.01	2.23
Wav2vec 2.0-large	DARTS	0.05	1.97
LFCC	light-DARTS	0.05	4.35
Wav2vec	light-DARTS	0.06	1.51
Wav2vec 2.0-base	light-DARTS	0.01	1.19
Wav2vec 2.0-large	light-DARTS	0.02	**1.08**

Table 2: The results of EER(%) for our proposed different systems on ASVspoof2021 DF. For the characteristics of this test set, we also evaluated narrow-band (NB) FIR filters using a data enhancement approach, following a similar procedure as in [21]. Where T23 refers to the first place team of the ASVspoof2021 DF track.

Feature	Network Architecture	DF
LFCC	light-DARTS	16.34
Wav2vec	light-DARTS	9.21
Wav2vec 2.0-base	light-DARTS	8.16
Wav2vec 2.0-large	light-DARTS	**7.86**
T23	best result in ASVspoof2021 [7]	15.64

Table 3: The results of EER(%) for our proposed different systems on ADD2022 Track 1. For the characteristics of this test set, we also evaluated narrow-band (NB) FIR filters using a data enhancement approach, following a similar procedure as in [21]. Where A01 refers to the first place team of the ADD2021 Track 1.

Feature	Network Architecture	Track 1
LFCC	light-DARTS	24.18
Wav2vec	light-DARTS	22.52
Wav2vec 2.0-base	light-DARTS	21.23
Wav2vec 2.0-large	light-DARTS	**20.11**
A01	best result in ADD 2022 [10]	21.71

5 RESULTS AND DISCUSSION

5.1 Results on ASVspoof2019

Table 1 shows the results of our proposed different systems. We can find the following observations. Firstly, the wav2vec features perform better than the LFCC features, one of the accepted baselines. More specifically, the EER of w2v features is 3.51% while the EER of LFCC features is 4.75% with both fixed LCNN as back-end

Table 4: Performance comparison between the proposed model and other models on the evaluation set of ASVspoof2019 LA.

| Metheds	Dev	Eval
W2V-XLSR-LGF [28]	-	1.28
FFT-L-SENet [31]	-	1.14
Attention-based CNN [32]	0.16	1.87
RAWNET2 [33]	-	1.12
LFCC-PC-DARTS [20]	0.002	4.87
W2V-Siamese [13]	0.004	1.15
W2V2.0-light-DARTS(ours)	0.02	**1.08**

classifiers. The reason is that the features extracted from models pre-trained on large-scale datasets help reduce overfitting, thereby improving reliability and domain robustness. In addition, when the pre-trained model is replaced with wav2vec 2.0-base, the result is further improved to 3.32%. A possible explanation for this might be that compared to wav2vec, wav2vec 2.0 employs the transformer architectures to model the dependencies between feature sequences, resulting in a more discriminative representation of speech. Interestingly, the EER of wav2vec 2.0-large is observed to 3.56%, even slightly higher than wav2vec. This inconsistency may be due to the fact that the speech representation obtained from the pre-trained model of wav2vec 2.0-large is not adapted to the LCNN structure. These results indicate that the wav2vec pre-trained model is effective as a feature extractor, they may contain information useful for fake audio detection.

Secondly, for the network structure, the performances of DARTS is better than LCNN. For example, we control wav2vec feature remains consistent. We can see that the EER of LCNN is higher than that of DARTS, which decreased from 3.51% to 2.18%. A similar conclusion is reached for wav2vec 2.0-base and wav2vec 2.0-large. This result may be explained by the fact that the advantage of darts over LCNN is that it automatically tunes the parameters of the network structure, thus eliminating the need for manual parameter setting. A surprising aspect of the data is in the result of LFCC feature. When we control LFCC feature remains consistent, we find that the EER of the DARTS is slightly higher than that of LCNN. This discrepancy could be attributed to parameters that yield optimal results on the validation set do not necessarily yield optimal results on the test set. These results indicate that DARTS whose architecture and parameters are learned automatically is effective for fake audio detection task.

Thirdly, the results are further improved when the network structure is replaced with light-DARTS. When we control for LFCC, wav2vec, wav2vec 2.0-base, and wav2vec 2.0-large separately to remain consistent, the results all show a significant decrease in EER. More specifically, using the wav2vec 2.0 large feature, we could achieve EER of 1.08% for evaluation with the best checkpoint. A possible explanation for this might be that the MFM we used in light-DARTS not only separates the noisy signal from the informative signal, but also plays the role of feature selection between the two feature maps. Thus, it is effective for fake audio detection.
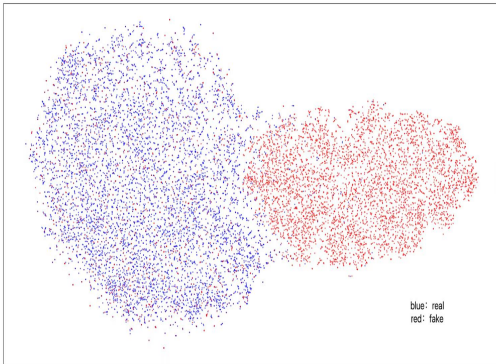
Figure 2: A t-SNE visual of LFCC-LCNN system. Blue dots represents real speech and red dots represents fake speech.
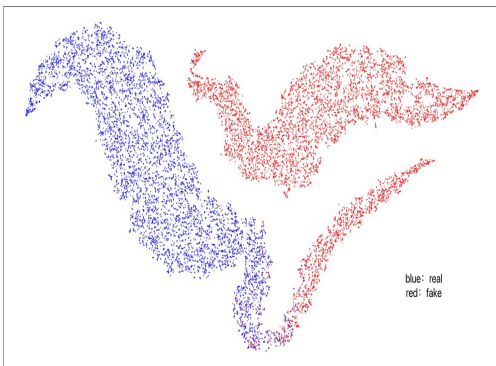
Figure 3: A t-SNE visual of wav2vec 2.0-large + light-DARTS system. Blue dots represents real speech and red dots represents fake speech.

Finally, our proposed fully automated end-to-end system is effective for fake audio detection task. As shown in the last three row of Table 1, wav2vec features combined with light-DARTS improves the performance of detecting fake audio while eliminating the need for manual setting of parameters. More specifically, the wav2vec 2.0 large feature combined with light-DARTS achieve the EER of 1.08%, which achieves best performance of single system.

5.2 Results on ASVspoof2021

To verify the generalizability of our proposed method, we also tested it on ASVspoof2021 DF dataset. The results of the correlational analysis are presented in Table 2. The overall conclusion remains consistent with section 5.1. More specifically, the "wav2vec 2.0 large + light-DARTS" reduces the EER to 7.86% on DF. In addition, our proposed method compares far better with ASVspoof2021 DF track

first place. The above results demonstrate the generalizability of our proposed method to different datasets.

5.3 Results on ADD 2022

In order to further validate the effectiveness of the proposed method on Chinese dataset, we have conducted relevant experiments on ADD 2022 Track 1. Table 3 provides the experimental data on ADD 2022 Track 1. The results show that the "wav2vec 2.0-base + light-DARTS" and "wav2vec 2.0-base + light-DARTS" achieved EERs of 21.23% and 20.11%, respectively, on the test set, both better than the first place in the competition results. The above results show that the proposed method works on the Chinese dataset as well.

5.4 Comparison with Other Systems

In order to evaluate the performance of our proposed method, we also compared the proposed system with other state-of-the-art systems on the evaluation set of the ASVspoof 2019 LA database. Recent published novel systems are compared, such as sub-band models [31], raw waveform based models [33] and frequency attention model [32]. Table 4 shows the performance comparison with other systems on the evaluation set of the ASVspoof 2019 LA database. From Table 4 we can find that our proposed fully automated end-to-end system achieves an EER of 1.08%, outperforms the second-ranked system (EER 1.12%) among all known systems. These results verify that the proposed method is quite effective for fake audio detection.

5.5 Visualization Analysis

To visualize the effectiveness of the proposed method, we also used t-SNE [34] to visualize LFCC-LCNN and wav2vec 2.0-large + light-DARTS, respectively. Both models are trained on the asvspoof2019 LA dataset and take the penultimate layer of the network. As shown in Figure 2 and Figure 3, we can see that the real and fake speech of LFCC-LCNN are not clearly distinguished, and there are many red dots embedded inside the blue dots. The real and fake speech of wav2vec 2.0-large + light-DARTS, on the other hand, is clearly separated, and only a little mixed at the boundary. The above visualization results are consistent with our experimental results.

6 CONCLUSION

The existing fake audio detection systems often rely on expert experience to design the acoustic features or manually design the hyperparameters of the network structure. However, artificial adjustment of the parameters can have a relatively obvious influence on the results. It is almost impossible to manually set the best set of parameters. In order to reduce the impact of manually set parameters on the stability of the fake audio detection system, this paper designs a fully automated end-to-end fake audio detection method. For features, we use wav2vec pre-trained model to obtain a high-level representation of the speech. As for classification model, we propose a light-DARTS to learn deep speech representations automatically. The results show that our proposed method achieves better detection results while not requiring manual parameter setting. The EER of our proposed method is relatively reduced by 77.26% compared to the baseline. In the future, we will further explore the application of DARTS on fake audio detection.

ACKNOWLEDGMENTS

REFERENCES

[1] Tao Wang, Jianhua Tao, Ruibo Fu, Jiangyan Yi, Zhengqi Wen, and Chunyu Qiang. Bi-level speaker supervision for one-shot speech synthesis. In *INTERSPEECH*, pages 3989–3993, 2020.

[2] Yuxuan Wang, R.J. Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, Quoc Le, Yannis Agiomyrgiannakis, Rob Clark, and Rif A. Saurous. Tacotron: Towards End-to-End Speech Synthesis. In *Proc. Interspeech 2017*, pages 4006–4010, 2017.

[3] Jose Sotelo, Soroush Mehri, Kundan Kumar, João Felipe Santos, Kyle Kastner, Aaron C Courville, and Yoshua Bengio. Char2wav: End-to-end speech synthesis. In *ICLR (Workshop)*, 2017.

[4] Jean-Marc Valin and Jan Skoglund. Lpcnet: Improving neural speech synthesis through linear prediction. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5891–5895. IEEE, 2019.

[5] Jiangyan Yi, Ye Bai, Jianhua Tao, Haoxin Ma, Zhengkun Tian, Chenglong Wang, Tao Wang, and Ruibo Fu. Half-Truth: A Partially Fake Audio Detection Dataset. In *Proc. Interspeech 2021*, pages 1654–1658, 2021.

[6] Haoxin Ma, Jiangyan Yi, Jianhua Tao, Ye Bai, Zhengkun Tian, and Chenglong Wang. Continual Learning for Fake Audio Detection. In *Proc. Interspeech 2021*, pages 886–890, 2021.

[7] Junichi Yamagishi, Xin Wang, Massimiliano Todisco, Md Sahidullah, Jose Patino, Andreas Nautsch, Xuechen Liu, Kong Aik Lee, Tomi Kinnunen, Nicholas Evans, et al. Asvspoof 2021: accelerating progress in spoofed and deepfake speech detection. *arXiv preprint arXiv:2109.00537*, 2021.

[8] Massimiliano Todisco, Xin Wang, Ville Vestman, Md. Sahidullah, Héctor Delgado, Andreas Nautsch, Junichi Yamagishi, Nicholas Evans, Tomi H. Kinnunen, and Kong Aik Lee. ASVspoof 2019: Future Horizons in Spoofed and Fake Audio Detection. In *Proc. Interspeech 2019*, pages 1008–1012, 2019.

[9] Tomi Kinnunen, Md. Sahidullah, Héctor Delgado, Massimiliano Todisco, Nicholas Evans, Junichi Yamagishi, and Kong Aik Lee. The ASVspoof 2017 Challenge: Assessing the Limits of Replay Spoofing Attack Detection. In *Proc. Interspeech 2017*, pages 2–6, 2017.

[10] Jiangyan Yi, Ruibo Fu, Jianhua Tao, Shuai Nie, Haoxin Ma, Chenglong Wang, Tao Wang, Zhengkun Tian, Ye Bai, Cunhang Fan, et al. Add 2022: the first audio deep synthesis detection challenge. *arXiv preprint arXiv:2202.08433*, 2022.

[11] Massimiliano Todisco, Héctor Delgado, and Nicholas Evans. Constant q cepstral coefficients: A spoofing countermeasure for automatic speaker verification. *Computer Speech & Language*, 45:516–535, 2017.

[12] Md. Sahidullah, Tomi Kinnunen, and Cemal Hanilçi. A comparison of features for synthetic speech detection. In *Proc. Interspeech 2015*, pages 2087–2091, 2015.

[13] Yang Xie, Zhenchuan Zhang, and Yingchun Yang. Siamese network with wav2vec feature for spoofing speech detection. In *Proc. Interspeech*, pages 4269–4273, 2021.

[14] Juan M Martín-Doñas and Aitor Álvarez. The vicomtech audio deepfake detection system based on wav2vec2 for the 2022 add challenge. *arXiv preprint arXiv:2203.01573*, 2022.

[15] Cheng-I Lai, Nanxin Chen, Jesús Villalba, and Najim Dehak. ASSERT: Anti-Spoofing with Squeeze-Excitation and Residual Networks. In *Proc. Interspeech 2019*, pages 1013–1017, 2019.

[16] Moustafa Alzantot, Ziqi Wang, and Mani B. Srivastava. Deep Residual Neural Networks for Audio Spoofing Detection. In *Proc. Interspeech 2019*, pages 1078–1082, 2019.

[17] Galina Lavrentyeva, Sergey Novoselov, Egor Malykh, Alexander Kozlov, Oleg Kudashev, and Vadim Shchemelinin. Audio replay attack detection with deep learning frameworks. In *Interspeech*, pages 82–86, 2017.

[18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[19] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: differentiable architecture search. *CoRR*, abs/1806.09055, 2018.

[20] Wanying Ge, Michele Panariello, Jose Patino, Massimiliano Todisco, and Nicholas Evans. Partially-connected differentiable architecture search for deepfake and spoofing detection. *arXiv preprint arXiv:2104.03123*, 2021.

[21] Xiang Wu, Ran He, Zhenan Sun, and Tieniu Tan. A light cnn for deep face representation with noisy labels. *IEEE Transactions on Information Forensics and Security*, 13(11):2884–2896, 2018.

[22] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33:12449–12460, 2020.

[23] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. wav2vec: Unsupervised Pre-Training for Speech Recognition. In *Proc. Interspeech 2019*, pages 3465–3469, 2019.

[24] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460, 2021.

[25] Zhiyun Fan, Meng Li, Shiyu Zhou, and Bo Xu. Exploring wav2vec 2.0 on speaker verification and language identification. *arXiv preprint arXiv:2012.06185*, 2020.

[26] Leonardo Pepino, Pablo Riera, and Luciana Ferrer. Emotion recognition from speech using wav2vec 2.0 embeddings. *arXiv preprint arXiv:2104.03502*, 2021.

[27] Zhiqiang Lv, Shanshan Zhang, Kai Tang, and Pengfei Hu. Fake audio detection based on unsupervised pretraining models. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 9231–9235. IEEE, 2022.

[28] Xin Wang and Junichi Yamagishi. Investigating self-supervised front ends for speech spoofing countermeasures. *arXiv preprint arXiv:2111.07725*, 2021.

[29] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *stat*, 1050:5, 2017.

[30] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5206–5210. IEEE, 2015.

[31] Yuxiang Zhang, Wenchao Wang, and Pengyuan Zhang. The Effect of Silence and Dual-Band Fusion in Anti-Spoofing System. In *Proc. Interspeech 2021*, pages 4279–4283, 2021.

[32] Hefei Ling, Leichao Huang, Junrui Huang, Baiyan Zhang, and Ping Li. Attention-based convolutional neural network for asv spoofing detection. *Proc. Interspeech 2021*, pages 4289–4293, 2021.

[33] Hemlata Tak, Jose Patino, Massimiliano Todisco, Andreas Nautsch, Nicholas Evans, and Anthony Larcher. 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, 2021.

[34] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.