

REPLAY AND SYNTHETIC SPEECH DETECTION WITH RES2NET ARCHITECTURE

Xu Li¹, Na Li², Chao Weng², Xunying Liu¹, Dan Su², Dong Yu², Helen Meng¹

¹The Chinese University of Hong Kong

²Tencent AI Lab, Tencent

{xuli, xyliu, hmmeng}@se.cuhk.edu.hk, lina011779@126.com, {cweng, dansu, dyu}@tencent.com

ABSTRACT

Existing approaches for replay and synthetic speech detection still lack generalizability to unseen spoofing attacks. This work proposes to leverage a novel model structure, so-called Res2Net, to improve the anti-spoofing countermeasure’s generalizability. Res2Net mainly modifies the ResNet block to enable multiple feature scales. Specifically, it splits the feature maps within one block into multiple channel groups and designs a residual-like connection across different channel groups. Such connection increases the possible receptive fields, resulting in multiple feature scales. This multiple scaling mechanism significantly improves the countermeasure’s generalizability to unseen spoofing attacks. It also decreases the model size compared to ResNet-based models. Experimental results show that the Res2Net model consistently outperforms ResNet34 and ResNet50 by a large margin in both physical access (PA) and logical access (LA) of the ASVspoof 2019 corpus. Moreover, integration with the squeeze-and-excitation (SE) block can further enhance performance. For feature engineering, we investigate the generalizability of Res2Net combined with different acoustic features, and observe that the constant-Q transform (CQT) achieves the most promising performance in both PA and LA scenarios. Our best single system outperforms other state-of-the-art single systems in both PA and LA of the ASVspoof 2019 corpus.

Index Terms— ASV anti-spoofing, replay detection, synthetic speech detection, Res2Net, multi-scale feature

1. INTRODUCTION

Automatic speaker verification (ASV) systems aim at confirming a claimed speaker identity against a spoken utterance, which has been widely applied into commercial devices and authorization tools. However, it is also broadly noticed that malicious attacks can easily degrade a well-developed ASV system, and such attacks may be classified into impersonation [1], replay [1], voice conversion (VC) [2], text-to-speech [3] synthesis (TTS) and the recently emerged adversarial attacks [4, 5].

Over the past decade, the speaker verification community has held several ASVspoof challenges [6–8] to develop countermeasures mainly against replay, speech synthesis and voice conversion attacks. ASVspoof 2019 [8] is the latest one that includes all previous attacks within two sub-challenges: physical access (PA) and logical access (LA). PA considers spoofing attacks from replay while LA refers to attacks generated with TTS and VC techniques.

A challenging issue for constructing reliable countermeasures is the defense against unseen attacks. This is fully considered in ASVspoof 2019 [8], where the evaluation partition consists mostly

of unseen spoofing attacks, generated with replay configurations and spoofing algorithms which differ from those in the training and development partitions. Countermeasures that perform well in the development set are still very likely to fail in the evaluation set, due to the lack of generalizability to unseen spoofing attacks [8]. This issue challenges the generalizability of developed countermeasures in face of unseen attack [9, 10]. To construct generalized countermeasures, existing efforts may be divided into three categories: feature engineering [11, 12], system modeling [13–15] and loss criteria [16]. Among them, system modeling is most essential for data-driven countermeasures development with little human effort, and is the focus of this paper.

Many existing neural network architectures have been applied to designing powerful countermeasures against spoofing attacks, such as light convolutional neural network (LCNN) [17, 18], residual neural network (ResNet) [14, 15] and their variations [11, 16]. These models exhibit a strong ability in time and frequency domain modeling and achieve promising performance to capture the spoofing cues. However, when applied to unseen spoofing attacks, the performance of state-of-the-art (SOTA) systems is still limited [8]. ResNet and LCNN based systems mainly have two dimensions to control the model’s capacity, i.e. width and depth. But simply increasing the width and depth is not efficient for improving the model’s capacity. Specifically, since anti-spoofing countermeasures development requires high generalizability to unseen spoofing attacks, only increasing the width and depth can easily lead to over-fitting, due to the huge amount of parameters.

In this work, we make the first attempt to leverage a novel Res2Net [19] model architecture in anti-spoofing systems, motivated by their promising performance on vision tasks [19–21]. Res2Net focuses on the revision of the ResNet block to enable multiple feature scales. It splits the feature maps within one block into multiple channel groups and designs a residual-like connection across different channel groups. Such residual-like connection increases the possible receptive fields, resulting in multiple feature scales. This additional multi-scale reception improves the system’s capacity and helps the system perform better when generalized to unseen spoofing attacks. Meanwhile, it also decreases the model size compared to traditional ResNet-based models. Our experiments verified a significant improvement of Res2Net50 over ResNet34 and ResNet50 in detecting both replay and synthetic speech audios. Moreover, integration with the squeeze-and-excitation (SE) block [22] can further enhance performance. For feature engineering, we investigate the generalizability of the Res2Net model that incorporates with different acoustic features, and observe that the constant-Q transform (CQT) achieves the most promising results in both PA and LA subsets of ASVspoof 2019. Our best single model outperforms the other SOTA single systems in both physical and logical scenarios.

This work was done when Xu Li was an intern at Tencent AI Lab.

Table 1. The overall model architectures of ResNet34, ResNet50, Res2Net50 and SE-Res2Net50. The type of a residual block and the number of channels are specified inside the brackets, while the repeat times of each block on one stage are specified outside the brackets. “2-d fc” denotes a fully connected layer with 2 output units.

Stage	ResNet34	ResNet50	Res2Net50	SE-Res2Net50
Conv1	conv2d, 7×7 , 16, $stride = 2$ max pool, 3×3 , $stride = 2$		[conv2d, 3×3 , 16, $stride = 1$] $\times 3$	
Conv2	[Basic BLK, 16] $\times 3$	[Bottleneck BLK, 16] $\times 3$	[Res2Net BLK, 16] $\times 3$	[SE-Res2Net BLK, 16] $\times 3$
Conv3	[Basic BLK, 32] $\times 4$	[Bottleneck BLK, 32] $\times 4$	[Res2Net BLK, 32] $\times 4$	[SE-Res2Net BLK, 32] $\times 4$
Conv4	[Basic BLK, 64] $\times 6$	[Bottleneck BLK, 64] $\times 6$	[Res2Net BLK, 64] $\times 6$	[SE-Res2Net BLK, 64] $\times 6$
Conv5	[Basic BLK, 128] $\times 3$	[Bottleneck BLK, 128] $\times 3$	[Res2Net BLK, 128] $\times 3$	[SE-Res2Net BLK, 128] $\times 3$
global average pool, 2-d fc, softmax				

The contributions of this work include: 1) Leveraging the Res2Net model architecture into anti-spoofing and verifying its significant improvements over ResNet34 and ResNet50 models; 2) Investigating the generalizability of the Res2Net model incorporated with different acoustic features in both PA and LA scenarios; 3) Developing a single model that outperforms other SOTA single systems for both PA and LA in ASVspoof 2019.

The rest of this paper is organized as follows: Section 2 illustrates the Res2Net block and its integration with the SE block, followed by the overall model architectures. Experimental setup and results are discussed in Section 3 and 4, respectively. We conclude this work in Section 5.

2. APPROACH

2.1. Res2Net block

The Res2Net [19] architecture aims at improving multi-scale representation by increasing the number of available receptive fields. This is achieved by connecting smaller filter groups within one block in a hierarchical residual-like style. The comparison among the basic block [23], bottleneck block [23] and Res2Net block is illustrated in Fig. 1. The Res2Net block is modified from the bottleneck block. After a 1×1 convolution, it evenly splits the input feature maps by the channel dimension into s subsets, denoted by x_i , where $i \in \{1, 2, \dots, s\}$. Except for x_1 , each x_i is processed by a 3×3 convolutional filter $K_i(\cdot)$. Starting from $i = 3$, x_i is added with the output of K_{i-1} before being fed into $K_i(\cdot)$. This process can be formulated as Eq. 1:

$$y_i = \begin{cases} x_i, & i = 1 \\ K_i(x_i), & i = 2 \\ K_i(x_i + y_{i-1}), & 2 < i \leq s \end{cases} \quad (1)$$

where s is defined as the scale dimension [19], indicating the number of partitions applied to split feature maps. This hierarchical residual-like connection enables multiple-size of receptive fields within one block, resulting in multiple feature scales. Finally, it concatenates all splits and passes them through a 1×1 convolution filter to maintain the channel size of this residual block.

Note that the hierarchical residual-like connection between filter groups also decreases the amount of model parameters. We assume that the feature maps after the first 1×1 convolution is $X \in \mathbb{R}^{I \times H_{in} \times W_{in}}$, and the one before the last 1×1 convolution is $Y \in \mathbb{R}^{O \times H_{out} \times W_{out}}$, where H_{in} , W_{in} , H_{out} and W_{out} are feature dimensions, and I and O are the number of channels for X and Y , respectively. In the bottleneck block, it transforms X to Y by using a filter $W \in \mathbb{R}^{I \times O \times 3 \times 3}$ with a parameter size of $9 \times I \times O$. While in the Res2Net block, it splits X into s partitions with each

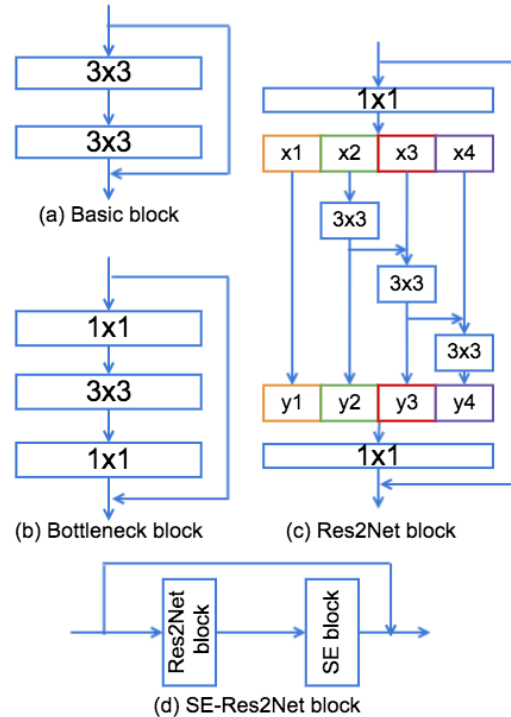


Fig. 1. The illustration of the basic block, bottleneck block, Res2Net block (scale dimension $s = 4$, the box in each color represents the feature maps within a channel group), and SE-Res2Net block.

$x_i \in \mathbb{R}^{\frac{I}{s} \times H_{in} \times W_{in}}$, for $i \in \{1, 2, \dots, s\}$. Except for x_1 , it transforms each x_i into $y_i \in \mathbb{R}^{\frac{O}{s} \times H_{out} \times W_{out}}$ by using a filter $K_i \in \mathbb{R}^{\frac{I}{s} \times \frac{O}{s} \times 3 \times 3}$, resulting in a total parameter size of $\frac{9 \times I \times O \times (s-1)}{s^2}$, which is smaller than the parameter size of W .

2.2. Integration with the squeeze-and-excitation block

The squeeze-and-excitation (SE) block [22] adaptively re-calibrates channel-wise feature responses by explicitly modeling the inter-dependencies between channels. This inter-dependencies modeling assigns different impact weights to channels, which improves the model’s capacity to focus on channel information that is most related with spoofing cues. Motivated by this, we stack the Res2Net and SE blocks together to form the SE-Res2Net block, as shown in Fig. 1 (d). Our experimental results demonstrated that this integration can

Table 2. Results on the ASVspoo 2019 physical and logical access in terms of EER (%) and t-DCF of different network architectures. (The input features for PA and LA are Spec and LFCC, respectively.)

System	# params	Physical Access				Logical Access			
		Dev Set		Eval Set		Dev Set		Eval Set	
		EER (%)	t-DCF	EER (%)	t-DCF	EER (%)	t-DCF	EER (%)	t-DCF
ResNet34	1.33M	0.83	0.022	1.46	0.041	0.39	0.013	5.75	0.131
SE-ResNet34	1.34M	0.57	0.015	1.32	0.037	0.35	0.011	4.69	0.103
ResNet50	1.05M	0.91	0.024	1.59	0.043	0.94	0.028	6.44	0.146
SE-ResNet50	1.09M	0.70	0.020	1.37	0.038	0.47	0.008	5.06	0.109
Res2Net50	0.88M	0.45	0.012	0.91	0.026	0.36	0.011	4.55	0.099
SE-Res2Net50	0.92M	0.52	0.012	0.74	0.021	0.23	0.005	2.87	0.079
Stat-SE-Res2Net50	0.96M	0.35	0.001	1.00	0.027	0.20	0.004	2.86	0.068

bring further improvement for both PA and LA scenarios.

2.3. Overall model architecture

This work evaluates and compares the performances of ResNet34, ResNet50, Res2Net50 and SE-Res2Net50, and an overview of their architectures are shown in Table 1. Notice that the Conv1 Stages of Res2Net50 and SE-Res2Net50 are slightly different from that of ResNet34 and ResNet50. The reason is that we adopt the updated “v1b” version of Res2Net provided by [19], which has been experimentally demonstrated to be more effective. The remaining parts for all models are identical, except for the block type. ResNet34, ResNet50, Res2Net50 and SE-Res2Net50 adopt the basic, bottleneck, Res2Net and SE-Res2Net blocks, respectively. Based on our experiments, we decrease the block expansion in ResNet50, Res2Net50 and SE-Res2Net50 from 4 to 2, to reduce the number of model parameters and prevent over-fitting issues. The scale dimension s in Res2Net-based models is experimentally set as 4.

3. EXPERIMENTAL SETUP

Dataset: All experiments are conducted under the ASVspoo 2019 challenge, which contains two subset evaluations: PA and LA. The detailed description of these two subsets are shown in Table 3. For each subset, our models are trained on the training partition, and the development partition is used for model selection. The evaluation partition consists mostly of unseen spoofing attacks, generated with either replay configurations or spoofing algorithms which differ from those in the training and development partitions. Systems are evaluated by the tandem detection cost function (t-DCF) [8] and equal error rate (EER) [8]. The log-probability of the bonafide class is adopted as the score for t-DCF and EER computation.

Feature extraction: Three acoustic features are evaluated in our experiments: log power magnitude spectrogram (Spec), linear frequency cepstral coefficients (LFCC) and constant-Q transform (CQT). The Spec is extracted with a “Hanning” window having size of 25ms and step-size of 10ms, and 512 FFT points are applied. The LFCC exactly follows the official baseline provided in the ASVspoo 2019 [8], extracted with 20ms window length, 512 FFT points and 20 filters with their delta and double delta coefficients, making 60-dimensional feature vectors. The CQT is extracted with 16ms step size, Hanning window, 9 octaves with 48 bins per octave. All features are truncated along the time axis to reserve exactly 400 frames. The feature less than 400 frames would be extended by repeating their contents.

Training strategy: The binary cross entropy is used as the loss function to train all models. Adam [24] is adopted as the optimizer

Table 3. Summary of the ASVspoo2019 corpus

Partition	Physical Access		Logical Access	
	#Bonafide	#Spoofed	#Bonafide	#Spoofed
Train	5,400	48,600	2,580	22,800
Dev.	5,400	24,300	2,548	22,296
Eval.	18,090	116,640	7,355	63,882

with $\beta_1 = 0.9$, $\beta_2 = 0.98$ and weight decay 10^{-9} . The learning rate scheduler increases the learning rate linearly for the first 1000 warm-up steps and then decreases it proportionally to the inverse square root of the step number [25]. All models are trained with 20 epochs, and the model with lowest EER on development set is chosen to be evaluated. Our codes have been made open-source¹.

4. RESULTS

4.1. Effectiveness of the Res2Net architecture

This section evaluates the effectiveness of the proposed Res2Net-based architectures for detecting spoofing samples. We leverage the results of ResNet34-based and ResNet50-based models for comparison, as shown in Table 2. The input features for PA and LA evaluation are Spec and LFCC, respectively. Upon comparing ResNet34 and ResNet50, we observe that ResNet34 outperforms ResNet50 in all conditions, which indicates that simply increasing the depth is not efficient for generalizability enhancement. After involving Res2Net50 into comparison, we observe that it significantly outperforms both ResNet34 and ResNet50. Specifically, Res2Net50 respectively outperforms ResNet34 and ResNet50 by a relative EER reduction of 37.7% and 42.8% on the PA evaluation set, and 20.9% and 29.3% on the LA evaluation set. Similar gains are also observed under the t-DCF metric. We also observe that when compared to ResNet34 and ResNet50, the model size of Res2Net50 is relatively reduced by 33.8% and 16.2%, respectively, which verifies the efficiency of the Res2Net architecture in detecting spoofing attacks.

Furthermore, integration with the SE block further improves the performance for all model architectures. For Res2Net50, the SE block achieves a relative EER reduction of 18.7% and 36.9% for PA and LA evaluation set, respectively. For horizontal comparisons, SE-Res2Net50 outperforms SE-ResNet34 and SE-ResNet50 by a relative EER reduction of 43.9% and 46.0% respectively on the PA evaluation set, and 38.8% and 43.3% respectively on the LA evaluation set. We also observe similar gains under the t-DCF metric.

A further attempt replaces the average pooling (AvgPool) layer in SE-Res2Net50 by the statistics pooling (StatPool) layer, the per-

¹<https://github.com/lixucuhk/ASV-anti-spoofing-with-Res2Net>

Table 4. Results on the ASVspoof 2019 physical access in terms of EER (%) and t-DCF of SE-Res2Net50 with different input features.

Features	Dev Set		Eval Set	
	EER (%)	t-DCF	EER (%)	t-DCF
Spec	0.519	0.0120	0.741	0.0207
LFCC	0.833	0.0229	1.465	0.0434
CQT	0.329	0.0086	0.459	0.0116
fusion	0.096	0.0028	0.287	0.0075
ASVspoof 2019 Baseline [8]				
CQCC-GMM	9.87	0.1953	11.04	0.2454
LFCC-GMM	11.96	0.2554	13.54	0.3017

Table 5. Results on the ASVspoof 2019 logical access in terms of EER (%) and t-DCF of SE-Res2Net50 with different input features.

Features	Dev Set		Eval Set	
	EER (%)	t-DCF	EER (%)	t-DCF
Spec	0	0	8.783	0.2237
LFCC	0.228	0.0051	2.869	0.0786
CQT	0.432	0.0143	2.502	0.0743
fusion	0	0	1.892	0.0452
ASVspoof 2019 Baseline [8]				
CQCC-GMM	0.43	0.0123	9.57	0.2366
LFCC-GMM	2.71	0.0663	8.09	0.2116

formance of which is shown at the bottom in Table 2. We observe that the StatPool layer improve the performance on development set for both PA and LA scenarios, but less improvement is observed on the LA evaluation set and performance decline is observed on the PA evaluation set. This suggests that the StatPool layer can lead to over-fitting issues for ASV anti-spoofing tasks.

4.2. Feature engineering

This section evaluates different acoustic features incorporated with SE-Res2Net50. The results of PA and LA scenarios are shown in Table 4 and Table 5, respectively. We have different observations in PA and LA scenarios. For instance, the Spec performs relatively well in the PA scenario, while the performance is much worse compared to the other two features in the LA scenario. Actually, we also observe limited generalization in the LA scenario for the Spec feature integrated with SE-ResNet34 and SE-ResNet50. One possible reason is that LA attacks are synthetic speech where the phase information is fake and estimated during the synthesis process. Hence, most spoofing cues for LA attacks exist in the phase information of the audio. While the Spec features we adopted are the log magnitudes of the spectrum, which contain no phase information, resulting in poor generalization for LA attacks. The LFCC is more suitable in detecting LA attacks than PA attacks. Interestingly, the CQT performs the best in both scenarios. Moreover, we also conduct experiments on other features, such as filter-banks and constant-Q cepstral coefficients (CQCC). But we observe similar or worse performance than the three features above.

When compared to the baseline systems [8] provided by ASVspoof2019, the CQT-based system achieves a relative EER reduction over 95.8% in the PA evaluation set, and a relative EER reduction over 69.1% in the LA evaluation set. To take advantage of the complementarity across these three features, we designed a fusion system operated by simply averaging the output scores of three systems. It achieves an EER of 0.287% and t-DCF of 0.0075

Table 6. Performance comparison of our proposed system to some known state-of-the-art single systems on the ASVspoof 2019 PA and LA evaluation set. Some systems reported results on only one access, and we denote the absent results as “-” in the table.

System	Physical Access		Logical Access	
	EER (%)	t-DCF	EER (%)	t-DCF
Spec+ResNet+CE [13]	3.81	0.0994	9.68	0.2741
MFCC+ResNet+CE [13]	-	-	9.33	0.2042
CQCC+ResNet+CE [13]	4.43	0.1070	7.69	0.2166
Spec+ResNet+CE [15]	1.29	0.036	11.75	0.216
Joint-gram+ResNet+CE [14]	1.23	0.0305	-	-
GD-gram+ResNet+CE [14]	1.08	0.0282	-	-
LFCC+LCNN+A-softmax [17]	4.60	0.1053	5.06	0.1000
FFT+LCNN+A-softmax [17]	-	-	4.53	0.1028
CQT+LCNN+A-softmax [17]	1.23	0.0295	-	-
FG-CQT+LCNN+CE [18]	-	-	4.07	0.102
Spec+LCGRNN+GKDE-Softmax [16]	1.06	0.0222	3.77	0.0842
Spec+LCGRNN+GKDE-Triplet [16]	0.92	0.0198	3.03	0.0776
MGD+ResNeWt+CE [11]	2.15	0.0465	-	-
CQTMGD+ResNeWt+CE [11]	0.94	0.0250	-	-
Fbanks&CQT+ResNeWt+CE [11]	0.52	0.0134	-	-
Ours: CQT+SE-Res2Net50+CE	0.459	0.0116	2.502	0.0743

in the PA evaluation set, and an EER of 1.892% and t-DCF of 0.0452 in the LA evaluation set.

4.3. Comparison with the state-of-the-art single systems

This section compares our best single system, i.e. SE-Res2Net50 incorporated with the CQT, with some of the reported SOTA single systems on the ASVspoof2019 PA and LA evaluation sets. Some of the top single systems are shown in Table 6, according to our best knowledge. The systems are denoted by a name consisting of input features, system architecture and loss criteria.

We observe that only a few systems achieve either an EER below 1.0% on the PA evaluation set, or an EER below 4.0% on the LA evaluation set, and rarely do systems have promising performance on both, which indicates the challenge in detecting unseen spoofing attacks. Most of the well-performing systems explore powerful model architectures and loss criteria. For PA attacks, the system in [11] has very promising performance and achieved the top single system in the PA scenario of the ASVspoof2019 competition. This is outperformed by our system at a relative EER reduction of 11.5%, depicting the effectiveness of our proposed system. For LA attacks, our system also outperforms other SOTA systems by a large margin.

5. CONCLUSION

This work proposes to leverage the Res2Net architecture into ASV anti-spoofing systems. Res2Net revises the ResNet block to enable multiple feature scales, which effectively improves the countermeasure’s generalizability to unseen spoofing attacks. Experimental results show significant improvement of Res2Net50 over the ResNet34 and ResNet50 systems in both PA and LA scenarios. Moreover, we incorporated different acoustic features into Res2Net50, and found that the CQT achieves the most promising results for both PA and LA attacks. Our best single model outperforms other SOTA single systems in both PA and LA scenarios. Future work will explore loss criteria with better generalization ability to deal with ASV anti-spoofing tasks.

6. ACKNOWLEDGEMENT

This work is partially supported by HKSAR Government’s Research Grants Council General Research Fund (Project No. 14208718).

7. REFERENCES

- [1] Z. Wu, N. Evans, T. Kinnunen, J. Yamagishi, F. Alegre, and H. Li, "Spoofing and countermeasures for speaker verification: A survey," *speech communication*, vol. 66, pp. 130–153, 2015.
- [2] T. Kinnunen, Z. Wu, K. Lee, F. Sedlak, E. Chng, and H. Li, "Vulnerability of speaker verification systems against voice conversion spoofing attacks: The case of telephone speech," in *ICASSP*. IEEE, 2012, pp. 4401–4404.
- [3] V. Shchemelinin and K. Simonchik, "Examining vulnerability of voice verification systems to spoofing attacks by means of a TTS system," in *ICSC*. Springer, 2013, pp. 132–137.
- [4] R. K. Das, X. Tian, T. Kinnunen, and H. Li, "The attacker's perspective on automatic speaker verification: An overview," *arXiv preprint arXiv:2004.08849*, 2020.
- [5] X. Li, J. Zhong, X. Wu, J. Yu, X. Liu, and H. Meng, "Adversarial attacks on GMM i-vector based speaker verification systems," in *ICASSP*. IEEE, 2020, pp. 6579–6583.
- [6] Z. Wu, T. Kinnunen, N. Evans, J. Yamagishi, C. Hanilçi, M. Sahidullah, and A. Sizov, "Asvspoof 2015: the first automatic speaker verification spoofing and countermeasures challenge," in *Interspeech*, 2015.
- [7] T. Kinnunen, M. Sahidullah, H. Delgado, M. Todisco, N. Evans, J. Yamagishi, and K. A. Lee, "The asvspoof 2017 challenge: Assessing the limits of replay spoofing attack detection," in *Interspeech*, 2017.
- [8] 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," in *Interspeech*, 2019.
- [9] R. K. Das, J. Yang, and H. Li, "Assessing the scope of generalized countermeasures for anti-spoofing," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6589–6593.
- [10] X. Li, N. Li, J. Zhong, X. Wu, X. Liu, D. Su, D. Yu, and H. Meng, "Investigating robustness of adversarial samples detection for automatic speaker verification," *arXiv preprint arXiv:2006.06186*, 2020.
- [11] X. Cheng, M. Xu, and T. F. Zheng, "Replay detection using cqt-based modified group delay feature and resnewt network in asvspoof 2019," in *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2019, pp. 540–545.
- [12] Z. Wu, E. S. Chng, and H. Li, "Detecting converted speech and natural speech for anti-spoofing attack in speaker recognition," in *Interspeech*, 2012.
- [13] M. Alzantot, Z. Wang, and M. B. Srivastava, "Deep residual neural networks for audio spoofing detection," *arXiv preprint arXiv:1907.00501*, 2019.
- [14] W. Cai, H. Wu, D. Cai, and M. Li, "The DKU replay detection system for the asvspoof 2019 challenge: On data augmentation, feature representation, classification, and fusion," *arXiv preprint arXiv:1907.02663*, 2019.
- [15] C.-I. Lai, N. Chen, J. Villalba, and N. Dehak, "Assert: Anti-spoofing with squeeze-excitation and residual networks," *arXiv preprint arXiv:1904.01120*, 2019.
- [16] A. Gomez-Alanis, J. A. Gonzalez-Lopez, and A. M. Peinado, "A kernel density estimation based loss function and its application to asv-spoofing detection," *IEEE Access*, vol. 8, pp. 108530–108543, 2020.
- [17] G. Lavrentyeva, S. Novoselov, A. Tseren, M. Volkova, A. Gorlanov, and A. Kozlov, "Stc antispoofing systems for the asvspoof2019 challenge," *arXiv preprint arXiv:1904.05576*, 2019.
- [18] Z. Wu, R. K. Das, J. Yang, and H. Li, "Light convolutional neural network with feature genuinization for detection of synthetic speech attacks," *arXiv preprint arXiv:2009.09637*, 2020.
- [19] S. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. H. Torr, "Res2net: A new multi-scale backbone architecture," *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [20] Z. Cao and H. J. Lee, "Improved res2net model for person re-identification," in *2019 IEEE First International Conference on Cognitive Machine Intelligence (CogMI)*. IEEE, 2019, pp. 235–240.
- [21] W. Zhou, Y. Chen, C. Liu, and L. Yu, "Gfnet: Gate fusion network with res2net for detecting salient objects in rgb-d images," *IEEE Signal Processing Letters*, 2020.
- [22] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 2015.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.