

To what extent can ASV systems naturally defend against spoofing attacks?

Jee-weon Jung^{*1}, Xin Wang^{*2}, Nicholas Evans³, Shinji Watanabe¹, Hye-jin Shim¹, Hemlata Tak⁴,
Siddhant Arora¹, Junichi Yamagishi², Joon Son Chung⁵

¹Carnegie Mellon University, USA, ²National Institute of Informatics, Japan

³EURECOM, France, ⁴Pindrop, USA

⁵Korea Advanced Institute of Science and Technology, South Korea

jeeveonj@ieee.org, wangxin@nii.ac.jp

Abstract

The current automatic speaker verification (ASV) task involves making binary decisions on two types of trials: target and non-target. However, emerging advancements in speech generation technology pose significant threats to the reliability of ASV systems. This study investigates whether ASV effortlessly acquires robustness against spoofing attacks (i.e., zero-shot capability) by systematically exploring diverse ASV systems and spoofing attacks, ranging from traditional to cutting-edge techniques. Through extensive analyses conducted on eight distinct ASV systems and 29 spoofing attack systems, we demonstrate that the evolution of ASV inherently incorporates defense mechanisms against spoofing attacks. Nevertheless, our findings also underscore that the advancement of spoofing attacks far outpaces that of ASV systems, hence necessitating further research on spoofing-robust ASV methodologies.

Index Terms: speaker verification, anti-spoofing, speech synthesis, voice conversion, spoofing-aware speaker verification

1. Introduction

Automatic Speaker Verification (ASV) systems are biometric technologies designed to confirm whether a given speech sample belongs to a claimed enrolled speaker [1]. These systems evaluate performance through a protocol consisting of two types of trials: target trials, where the voice in the speech sample matches the enrolled speaker, and non-target trials, where it does not. As a crucial component of an authentication system, ASV must not only differentiate between these trials but also robustly defend against spoofing attacks (also known as presentation attacks) – efforts to deceive the system with artificially generated voices. Yet, remarkably, the integration of such critical defenses against spoofing attacks has been considered in only a fractional segment of the existing ASV systems [2, 3].

Meanwhile, the explosive advancements in Text-To-Speech synthesis (TTS) and Voice Conversion (VC) have heightened vulnerabilities to spoofing attacks. The accessibility of sophisticated TTS and VC systems has inadvertently equipped fraudsters with the means to generate speech data mimicking specific target individuals [4]. While only a subset of the latest deep-learning-based TTS and VC systems reached the perceptual quality of bonafide speech [5–7], many systems, even those employing basic linear statistical models, have successfully fooled state-of-the-art ASV systems at the time (e.g., the i-vector PLDA [8] and the x-vector [9] systems evaluated in ASVspoof 2015 [2] and 2019 challenges [10]).

In response, a specialized area of study has emerged aimed at enhancing ASV systems with integrated spoofing detection

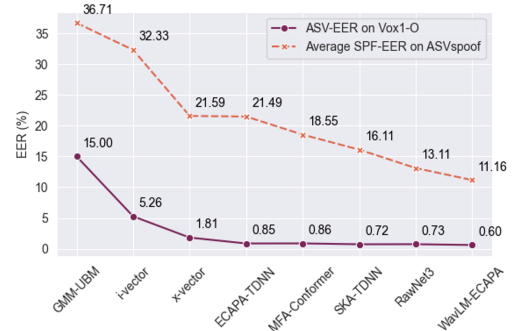


Figure 1: Average Spoof Equal Error Rates (SPF-EERs) on 29 different spoofing attacks, chronologically displayed using eight automatic speaker verification (ASV) systems. The SPF-EER adopts spoof trials in place of conventional non-target trials, where in a spoof trial, the test utterance is a system-generated voice of the target speaker. Conventional EERs (equivalent to SV-EER in [3]) of the ASV systems in the Vox1-O evaluation protocol are also reported as a reference.

capabilities, thus introducing a third type of trial: the spoof trial [2, 3, 11, 12]. Despite the addition, these extended systems, termed Spoofing-robust ASV (SASV), retain a binary classification function, authenticating target trials while rejecting all others. Initial efforts in SASV development combined separately developed ASV and countermeasure (CM) subsystems [2, 13–15].¹ More recent approaches have explored integrated solutions [11, 16, 17], utilizing the adaptive power of neural networks to assess both the identity of the speaker and the authenticity of the speech concurrently, aligning with the singular goal of accepting target trials.

While ASV systems are generally vulnerable to spoofing, not all attacks are equally effective. The ASVspoof 2019 [10] revealed that different attacks impact the ASV Equal Error Rate (EER) to varying extents. Further investigation [13] has uncovered instances where an ASV system can reject spoofing attacks, even when their training regimes did not explicitly account for such scenarios (i.e., *zero-shot* in terms of data and scenario). This finding indicates that ASV systems possess an inherent ability to reject spoof attempts, especially when the imitation of the target speaker’s characteristics falls short. This insight leads to two key questions: (1) Can future enhancements in ASV technology, focused exclusively on refining speaker discrimination as currently practiced, prove robust enough to counter both present and future spoofing challenges? (2) How effectively can current ASV systems, utilizing a zero-shot strategy, repel spoofing attacks, and which precise varieties of at-

^{*}Equal contribution.

¹The CM system does not require a target enrolled speaker. It makes a binary decision on whether the input is from humans or synthesized.

tacks are they adept at blocking?

To address the first question, we present a chronological analysis examining the correlation between advancements in ASV technology and their inherent defense against spoofing attacks. Our study spans a broad spectrum of ASV technologies, ranging from the Gaussian Mixture Model-Universal Background Model (GMM-UBM) [18] developed in the pre-deep learning era to the latest Self-Supervised Learning (SSL) representation-based Enhanced Deep Speaker Embedding (ECAPA-TDNN [19]). Illustrated in Figure 1 and detailed further in Section 5, our findings reveal a clear evolution in ASV systems’ natural defense mechanisms against spoofing attacks.² Notably, with conventional ASV error rates plummeting from 13.00% to 0.60%, the systems’ robustness to spoofing has increased fourfold, although these ASV systems were never explicitly designed to counter spoof scenarios. This demonstrates a remarkable *zero-shot* capability; the ASV systems effectively countered spoofing attacks without any prior training on such scenarios or data. However, it’s important to note that the current EER, above 10%, remains too high for practical application in real-world scenarios.

To address the second question, we conduct a parallel chronological analysis focused on the evolution of attack technologies. Our examination of 29 different TTS/VC systems reveals a consistent pattern; speech data spoofed using the latest algorithms proves significantly more adept at fooling ASV systems. This outcome reflects the continuous advancement in spoofing techniques.

Our analysis leads to two critical insights: (1) advancements in ASV are enhancing the robustness of existing spoofing attacks; (2) progress in TTS/VC is far outpacing that in ASV. Given this dynamic, we emphasize the urgent need for the ASV community to intensify efforts in crafting SASV systems.³ Whether through integrating ASV and CM subsystems or by developing single neural networks, the enhancement of spoofing robustness is paramount. Without it, ASV systems lack the necessary robustness to serve as reliable authentication mechanisms.

2. ASV systems

Given the limited space, we briefly outline the eight ASV systems under investigation in this work. RawNet3 [20] is the only system not relying on handcrafted acoustic features such as mel-spectrograms for input. Refer to the references for details.

GMM-UBM utilises a generative modeling approach for ASV [18,21,22]. It trains a GMM with the training set to form a Universal Background Model (UBM). The GMM-UBM system evaluates the log-likelihood ratio between the GMM adapted toward the target speaker and the UBM for each input.

i-vector. [8] captures speaker and channel variabilities in a unified low-dimensional space. It projects a GMM supervector, derived from a UBM, into this compact latent space using a total variability matrix.

x-vector. [9] leverages time-delay neural networks (TDNNs) for frame-level analysis, employing varying dilation rates. A statistics pooling layer then aggregates these frame embeddings

²This observation, while somewhat expected – since improved speaker discrimination inherently complicates successful spoofing attempts – represents the first documentation of a clear evolutionary trend.

³We envision a future where the term ASV inherently signifies SASV, recognizing spoofing-robustness as a fundamental requirement for any ASV system to effectively serve as an authentication mechanism.

into a singular utterance-level embedding through mean and standard deviation calculations, followed by two dense layers for further processing.

ECAPA-TDNN [19] employs Res2Net as the backbone processing block with a squeeze-excitation module for frame-level processing. Additional techniques, such as multi-layer feature aggregation, are adopted. Context- and channel-dependent statistics pooling is devised to aggregate frame embeddings into utterance-level embedding.

MFA-Conformer [23] combines CNN sub-sampling and Conformer encoder layers for frame embedding. An attentive statistics pooling layer combines frame embeddings into an utterance embedding. This process involves concatenating and normalizing frame embeddings from all encoder layer outputs through layer normalisation [24].⁴

SKA-TDNN. Building upon ECAPA-TDNN, SKA-TDNN [25] introduces a novel frame embedding processing mechanism (“msSKA block”) and a feature-enhancing module (“fcwSKA block”) employing selective kernel attention. This is followed by context- and channel-dependent pooling, batch normalization, and dense layers for utterance-level integration.

RawNet3 [20] is a CNN-based system optimized for processing raw waveforms directly, employing an analytic filterbank [26] layer as a substitute for traditional acoustic features. Its modeling approach for both frame and utterance levels parallels ECAPA-TDNN, with key adjustments for raw waveform input compatibility. These modifications include strategically using pooling layers to condense the sequence length.

WavLM-Large with ECAPA-TDNN. This system combines WavLM-Large [27], an SSL model with strong representations for several downstream tasks, including ASV, with ECAPA-TDNN. We freeze the WavLM and use its representation as a feature. ECAPA-TDNN is trained using the WavLM representations in place of conventional acoustic features.

3. Corpora

3.1. Training

For training our ASV systems, except for the GMM-UBM, we utilize the VoxCeleb 1 and 2 corpora [28,29], which feature celebrity utterances sourced from YouTube. These datasets are the preferred choice for recent ASV research. The combined training set encompasses over 2.7k hours of speech, derived from 1.2 million utterances across 7.3k speakers. The WavLM-Large [27] we use was pre-trained with external corpora. We train two GMM-UBMs, utilizing the training partitions of ASVspooF 2015 and ASVspooF 2019 logical access, respectively.

3.2. Evaluation

We select the ASVspooF 2015 [2] and 2019 logical access (LA) [10] corpora to assess our ASV systems’ performances when faced with spoofing attacks. Other corpora were excluded due to the lack of an official SASV evaluation protocol, which is crucial for our analysis.⁵ Table 1 describes the statistics

⁴In our implementation, we adopt the context- and channel-dependent statistics pooling from ECAPA-TDNN.

⁵Most corpora have evaluation protocols for assessing anti-spoofing systems where a binary decision of “bonafide” and “spooF” needs to be made. For the analysis in this work, where ASV systems are evaluated in a zero-shot manner with spooF attacks, we need an SASV evaluation protocol, where input is a pair of an enrollment speaker and a test

Table 1: Statistics of the ASVspoof 2015’s evaluation partition and 2019 LA’s development and evaluation partitions [2, 10].

Corpus	# Speaker	# Utterance	# Algorithms
ASVspoof 2015	46	193,404	10
ASVspoof 2019 LA	87	96,081	19

Table 2: Spoofing attacks in ASVspoof 2015 (S^*) and 2019 LA corpora (A^*). They are grouped based on acoustic and waveform models, and A18 to A10 are sorted based on when the system was proposed. DNN-based *acoustic* and *waveform* models are highlighted in the second and third columns, respectively.

Group	ID	Type	Acoustic model	Waveform model
1	A18	VC	i-vector + PLDA	LPC
	S5	VC	GMM	MLSA
	A06	VC	GMM	spectral filtering
	A19	VC	GMM	spectral filtering
	S2	VC	Linear reg.	STRAIGHT
	S1	VC	DTW	STRAIGHT
	S6	VC	GMM + GV	STRAIGHT
	S7	VC	GMM + GV	STRAIGHT
	S3	TTS	NLP + HMM	STRAIGHT
	S4	TTS	NLP + HMM	STRAIGHT
2	S8	VC	GMM-tensor	STRAIGHT
	S9	VC	DTW + Kernel reg.	STRAIGHT
	A05	VC	VAE	WORLD
	A17	VC	VAE	waveform filtering
	A13	TTS	TTS + VC(DNN)	waveform filtering
	A09	TTS	NLP + RNN	Vocaine
	A14	TTS	TTS + VC(DNN)	STRAIGHT
	A03	TTS	NLP + DNN	WORLD
	A02	TTS	NLP + HMM-DNN	WORLD
	A07	TTS	NLP + RNN-GAN	WORLD
A11	TTS	DNN(end2end)	Griffin-Lim	
3	A08	TTS	NLP + HMM-DNN	Dilated CNN
	A01	TTS	NLP + HMM-DNN	WaveNet
	A12	TTS	NLP + RNN	WaveNet
	A15	TTS	TTS + VC(DNN)	WaveNet
	A10	TTS	DNN(end2end)	WaveRNN
4	S10			
	A04	TTS	NLP + Unit-selection	Waveform concat.
	A16			

of the two corpora. Table 2 provides detailed descriptions of each spoofing attack algorithm, categorizing them as follows: “Group 1” excludes neural networks, “Group 2” incorporates neural networks solely in the acoustic model, “Group 3” utilizes neural networks in both acoustic and language models, and “Group 4” consists of non-parametric systems.

ASVspoof 2015. This corpus, pioneering in the realm of spoofing-robust ASV, includes an evaluation set with ten spoofing attacks derived from TTS and VC technologies, as detailed in Table 2 with IDs beginning with “S.”

ASVspoof 2019 logical access is the most widely used corpora for developing spoofing-robust ASV and anti-spoofing systems. Its development and evaluation partitions comprise 19 spoofing attacks. Details on the attacks are presented Table 2 with “ID” starting with “A.”

Vox1-O. Adopted as the benchmark protocol, this involves data from 40 speakers in the VoxCeleb1 test set. We adopt the protocol to show our ASV systems’ performance in Figure 1.

4. Experimental configurations

4.1. GMM-UBM and i-vector

GMM-UBM. We employ three GMM-UBM systems in this work. GMM-UBM performance in Figure 1 is from [28]. We trained two GMM-UBM systems using the MSR Identity toolkit [30]. These UBMs were trained on bonafide training

utterance. Each trial belongs to one of the three categories, “target”, “non-target”, and “spoof”.

data from ASVspoof 2015 and 2019 LA, excluding spoofed utterances. Each system was employed to assess the respective evaluation set.⁶

i-vector. We use the Kaldi toolkit [31] to train the i-vector system. Our implementation follows the original recipe with a modification: we omit the VoxCeleb2 test set from the training dataset to ensure consistency in training data across all ASV systems.

4.2. DNN-based systems

For the remaining six DNN-based systems, we leveraged publicly available pre-trained systems from ESPnet-SPK [32, 33]. Our methodology aimed to minimize hyper-parameter optimization, striving for consistency in configuration settings across these systems to mitigate the influence of irrelevant variables. Unless mentioned otherwise, the below configuration applies to all six systems. We set the batch size to 512 and trained the model for 40 epochs. Models use Adam optimizer [34] for training. A cosine annealing learning rate scheduler with warm-ups and restarts schedules the learning rate with the peak of $1e^{-3}$ and the minimum of $5e^{-6}$. We also augmented the data during training with Musan [35] noise and RIR [36] reverberations. 80-dimensional mel-spectrograms with a 25ms window and a 10ms shift size are used for the acoustic feature. Speaker embeddings have 192 dimensions. Training objective functions comprise AAM-softmax [37] with sub-center and inter top-k losses [38]. Details on hyper-parameters and training schemes are available at <https://github.com/espnet/espnet/blob/master/egs2/voxceleb/spk1/README.md>.

x-vector implementation follows the original architecture proposed in [9]. Its speaker embeddings are 512-dimensional. **MFA-Conformer** adopts six encoder blocks, each with 512-dimensional output and a CNN kernel size of 15. We apply a longer warm-up of 10k steps. **WavLM-Large with ECAPA-TDNN** uses the weighted summation of the WavLM-Large model’s layer-wise outputs.

4.3. Metrics

All outcomes are assessed using two evaluation metrics. The conventional EER for ASV involves target and non-target trials. The Spoofing Equal Error Rate (SPF-EER) is determined using target trials and spoofed non-target trials, evaluating the ASV system’s susceptibility when non-target trials are substituted with spoofing trials. Refer to [3, 39] for details.

5. Results

5.1. ASV performances

Figure 1 presents the EERs of the eight ASV systems on the Vox1-O dataset (dotted line) alongside their SPF-EER across 29 spoofing attacks in the ASVspoof 2015 and 2019 LA corpora. Performances on Vox1-O demonstrate the evolution of conventional ASV, while performances on ASVspoof highlight ASV systems’ zero-shot (i.e., effortless) capability against out-of-domain spoofing attacks.

5.2. Does ASV progression naturally enhance defense against spoofing attacks?

Analysis from Figure 1 demonstrates a significant improvement in ASV performance, with EER reducing from 36.71%

⁶This decision aimed to expose GMM-UBM systems to a lesser domain mismatch compared to the other seven ASV systems, attempting to compensate for GMM-UBM’s relatively lower performance. Despite this advantage, GMM-UBM remains the least effective.




Figure 2: Detailed analyses on chronologically sorted eight ASV systems. (left): different groups of spoofing attacks. (middle): TTS vs. VC attacks. (right) DNN vs. non-DNN-based spoofing attacks. Group 1: does not involve neural networks. Group 2: only the acoustic model is a neural network. Group 3: acoustic and waveform models both are neural networks. Group 4: non-parametric systems.




Figure 3: Analyses from the perspective of spoofing attacks. (left): averaged ASV results in terms of groups. (middle): five ASV systems' results in terms of groups. (right): averaged ASV results of all 29 spoofing attacks in terms of years.

Table 3: State-of-the-art SASV systems.

System	ASV-EER	SPF-EER
Jointly optimised [40]	0.11	0.17
Single neural network [11]	1.27	1.23

in the GMM-UBM system to 11.16% in the WavLM-Large with ECAPA-TDNN. Figure 2 offers an in-depth analysis of this progression. The left section of this figure indicates a general trend of increasing EER across Groups 1 to 4,⁷ suggesting that advancements in ASV technology do not inherently confer enhanced defenses against increasingly sophisticated spoofing attacks. The right section confirms this observation, showing that EERs for spoofing attacks utilizing neural network-based TTS and VC systems remain consistently more challenging than their counterparts. Intriguingly, the middle section reveals that VC attacks consistently yield lower EERs than TTS attacks, highlighting a distinct vulnerability pattern within ASV systems' response to different spoofing methodologies.

5.3. Insights from the spoofing attacks

Figure 3 offers a viewpoint from the perspective of spoofing attacks. The left-side plot shows a trend of increasing EERs chronologically across Groups 1 to 3. We also find that unit-selection-based non-parametric TTS systems, group 4, pose the most significant challenge. The middle plot examines the performance of the five individual ASV systems across different groups to identify any distinct trends. Interestingly, neural network-based ASV systems exhibit similar patterns of monotonic increase, whereas the i-vector system demonstrates generalized EER performances across various groups of attacks, albeit with generally poor performance. Finally, the right-side plot arranges the 29 attacks according to the systems' devised years, highlighting the consistent upward trajectory of EERs. This consistent increase across different timeframes

⁷Group 1 represents the earliest systems without neural networks, while Group 3 includes the most recent systems employing neural networks for both acoustic and waveform modeling.

underscores the growing necessity for research into spoofing-resistant ASV technologies, specifically SASV systems.

5.4. Discussion: potential of SASV research

In closing, we demonstrate the bright potential of SASV research through a snapshot of the current state-of-the-art SASV systems [11, 40] in Table 3. The results are benchmarked using the ASVspoof 2019 LA corpus in the way it has been utilized in the SASV 2022 challenge; they are SASV systems rejecting non-target and spoof trials, not CM systems making a binary decision of bonafide and spoof. It's important to note that the evaluations consider only 13 of the attacks, making a direct comparison with other results in this paper incompatible. Despite the 13 analyzed attacks being more complex than the 16 not examined, the results reveal that current SASV systems (1.23% SPF-EER) significantly outperform ASV systems (11.16% SPF-EER) in dismissing spoof trials while also accurately rejecting non-target trials. This outcome highlights the promising potential and evolving capability of SASV systems in the ongoing battle against spoofing threats.

6. Conclusion and future works

Our investigation centered on whether the evolution of contemporary ASV systems, which traditionally do not account for spoofing scenarios, inherently develops defenses against such attacks. Through comprehensive analysis, we have confirmed a gradual improvement in ASV systems' ability to reject spoofed inputs. However, this progress lags significantly behind the rapid advancements in speech-generation technologies. Our analyses call for an urgent push within the ASV community towards developing either integrated ASV and CM subsystems or adopting single, unified SASV neural network approaches. Given the pace of advancements in deep learning, we strongly advocate for a focused effort on the latter, anticipating it will yield significant breakthroughs in enhancing ASV's resilience against spoofing attacks.

7. Acknowledgement

Experiments of this work used the Bridges2 system at PSC and Delta system at NCSA through allocations CIS210014 and IR1120008P from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, #2138296. This study is partially supported by JST AIP Acceleration Research (JPMJCR24U3), JST PRESTO (JPMJPR23P9), Japan, and by the ANR Bruel project (ANR-22-CE39-0009), France.

8. References

- [1] J. P. Campbell, "Speaker recognition: A tutorial," *Proceedings of the IEEE*, vol. 85, no. 9, pp. 1437–1462, 1997.
- [2] Z. Wu, T. Kinnunen, N. Evans et al., "Asvspoof 2015: the first automatic speaker verification spoofing and countermeasures challenge," in *Proc. Interspeech*, 2015.
- [3] J.-w. Jung, H. Tak, H.-j. Shim et al., "Sasv 2022: The first spoofing-aware speaker verification challenge," in *Proc. Interspeech*, 2022.
- [4] "Mom warns of hoax using ai to clone daughter's voice," www.abcnews.go.com/GMA/Family/mom-warns-hoax-ai-clone-daughters-voice/story?id=98551351, Accessed: 2024-02-28.
- [5] J. Shen, R. Pang, R. J. Weiss et al., "Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions," in *Proc. ICASSP*, 2018, pp. 4779–4783.
- [6] Y. Zhao, W.-C. Huang, X. Tian et al., "Voice Conversion Challenge 2020 — Intra-lingual semi-parallel and cross-lingual voice conversion —," in *Proc. Joint Workshop for the Blizzard Challenge and Voice Conversion Challenge 2020*, 2020, pp. 80–98.
- [7] E. Cooper and J. Yamagishi, "How do Voices from Past Speech Synthesis Challenges Compare Today?," in *Proc. SSW*, 2021, pp. 183–188.
- [8] N. Dehak, P. J. Kenny, R. Dehak et al., "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.
- [9] D. Snyder, D. Garcia-Romero et al., "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. ICASSP*, 2018.
- [10] X. Wang, J. Yamagishi, M. Todisco et al., "Asvspoof 2019: A large-scale public database of synthesized, converted and replayed speech," *Computer Speech & Language*, vol. 64, pp. 101114, 2020.
- [11] S. H. Mun, H.-j. Shim, H. Tak et al., "Towards single integrated spoofing-aware speaker verification embeddings," in *Proc. Interspeech*, 2023.
- [12] N. Torgashov, R. Makarov, I. Yakovlev et al., "The ID R&D VoxCeleb speaker recognition challenge 2023 system description," *arXiv:2308.08294*, 2023.
- [13] W. Ge, H. Tak, M. Todisco and N. Evans, "Can spoofing countermeasure and speaker verification systems be jointly optimised?," in *Proc. ICASSP*, 2023, pp. 1–5.
- [14] X. Liu, M. Sahidullah, K. A. Lee and T. Kinnunen, "Generalizing speaker verification for spoof awareness in the embedding space," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [15] H.-J. Shim, H. Tak, X. Liu et al., "Baseline systems for the first spoofing-aware speaker verification challenge: Score and embedding fusion," in *Proc. Speaker Odyssey*, 2022.
- [16] H.-j. Shim, J.-w. Jung, J.-h. Kim and H.-j. Yu, "Integrated replay spoofing-aware text-independent speaker verification," *Applied Sciences*, vol. 10, no. 18, pp. 6292, 2020.
- [17] M. Todisco, H. Delgado, K. A. Lee et al., "Integrated presentation attack detection and automatic speaker verification: Common features and gaussian back-end fusion," in *Proc. Interspeech*, 2018.
- [18] D. A. Reynolds, T. F. Quatieri and R. B. Dunn, "Speaker verification using adapted gaussian mixture models," *Digital signal processing*, vol. 10, no. 1-3, pp. 19–41, 2000.
- [19] B. Desplanques, J. Thienpondt et al., "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Proc. Interspeech*, 2020.
- [20] J.-w. Jung, Y. Kim, H.-S. Heo et al., "Pushing the limits of raw waveform speaker recognition," in *Proc. Interspeech*, 2022.
- [21] A. Larcher, K. A. Lee, B. Ma and H. Li, "Text-dependent speaker verification: Classifiers, databases and rsr2015," *Speech Communication*, vol. 60, pp. 56–77, 2014.
- [22] T. Kinnunen, M. Sahidullah, I. Kukanov et al., "Utterance verification for text-dependent speaker recognition: a comparative assessment using the reddots corpus," in *Proc. Interspeech*, 2016.
- [23] Y. Zhang, Z. Lv, H. Wu et al., "MFA-Conformer: Multi-scale feature aggregation conformer for automatic speaker verification," in *Proc. Interspeech*, 2022.
- [24] J. L. Ba, J. R. Kiros and G. E. Hinton, "Layer normalization," in *NeurIPS - Deep Learning Symposium*, 2016.
- [25] S. H. Mun, J.-w. Jung et al., "Frequency and multi-scale selective kernel attention for speaker verification," in *Proc. SLT*, 2022.
- [26] M. Pariente, S. Cornell, A. Deleforge and E. Vincent, "Filterbank design for end-to-end speech separation," in *Proc. ICASSP*, 2020.
- [27] S. Chen, C. Wang et al., "WavLM: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, 2022.
- [28] A. Nagrani, J. S. Chung and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Proc. Interspeech*, 2017.
- [29] J. S. Chung, A. Nagrani and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Interspeech*, 2018.
- [30] S. O. Sadjadi, M. Slaney and L. Heck, "Msr identity toolbox," *Seattle, WA, USA: Microsoft*, 2013.
- [31] D. Povey, A. Ghoshal, G. Boulianne et al., "The Kaldi speech recognition toolkit," in *Proc. ASRU*, 2011.
- [32] S. Watanabe, T. Hori, S. Karita et al., "Espnet: End-to-end speech processing toolkit," in *Proc. Interspeech*, 2018.
- [33] J.-w. Jung, W. Zhang, J. Shi et al., "Espnet-spk: full pipeline speaker embedding toolkit with reproducible recipes, self-supervised front-ends, and off-the-shelf models," in *Proc. Interspeech*, 2024.
- [34] D. P. Kingma and J. Ba, "Adam: a method for stochastic optimization," in *Proc. ICLR*, 2015.
- [35] D. Snyder, G. Chen and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint:1510.08484*, 2015.
- [36] T. Ko, V. Peddinti, D. Povey et al., "A study on data augmentation of reverberant speech for robust speech recognition," in *Proc. ICASSP*, 2017.
- [37] J. Deng, J. Guo, N. Xue and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proc. CVPR*, 2019.
- [38] M. Zhao, Y. Ma, M. Liu and M. Xu, "The SpeakIn system for VoxCeleb speaker recognition challenge 2021," *arXiv:2109.01989*, 2021.
- [39] H.-j. Shim, J.-w. Jung, T. Kinnunen et al., "a-dcf: an architecture agnostic metric with application to spoofing-robust speaker verification," in *Proc. Speaker Odyssey*, 2024.
- [40] A. Alenin, N. Torgashov, A. Okhotnikov et al., "A subnetwork approach for spoofing aware speaker verification," in *Proc. Interspeech*, 2022.