

# SASV 2022: The First Spoofing-Aware Speaker Verification Challenge

Jee-weon Jung<sup>\*,1</sup>, Hemlata Tak<sup>\*,2</sup>, Hye-jin Shim<sup>3</sup>, Hee-Soo Heo<sup>1</sup>, Bong-Jin Lee<sup>1</sup>,  
Soo-Whan Chung<sup>1</sup>, Ha-Jin Yu<sup>3</sup>, Nicholas Evans<sup>2</sup>, and Tomi Kinnunen<sup>4</sup>

<sup>1</sup>Naver Corporation, South Korea,

<sup>2</sup>EURECOM, Sophia Antipolis, France,

<sup>3</sup>School of Computer Science, University of Seoul, South Korea,

<sup>4</sup>University of Eastern Finland, Finland

sasv.challenge@gmail.com

## Abstract

The first spoofing-aware speaker verification (SASV) challenge aims to integrate research efforts in speaker verification and anti-spoofing. We extend the speaker verification scenario by introducing *spoofed* trials to the usual set of target and impostor trials. In contrast to the established ASVspoof challenge where the focus is upon separate, independently optimised spoofing detection and speaker verification sub-systems, SASV targets the development of integrated and jointly optimised solutions. Pre-trained spoofing detection and speaker verification models are provided as open source and are used in two baseline SASV solutions. Both models and baselines are freely available to participants and can be used to develop back-end fusion approaches or end-to-end solutions. Using the provided common evaluation protocol, 23 teams submitted SASV solutions. When assessed with target, bona fide non-target and spoofed non-target trials, the top-performing system reduces the equal error rate of a conventional speaker verification system from 23.83% to 0.13%. SASV challenge results are a testament to the reliability of today's state-of-the-art approaches to spoofing detection and speaker verification.

**Index Terms:** spoofing-aware speaker verification, audio spoofing detection, anti-spoofing, speaker verification

## 1. Introduction

Automatic speaker verification (ASV) systems determine whether or not an input utterance contains speech uttered by a given, known speaker. As one of the most efficient, convenient, natural and non-intrusive biometric characteristics, ASV has found widespread application, most notably, in telephony-based scenarios. Reliability is crucial and must be maintained not only in the face of target and impostor trials or utterances, but also so-called spoofed utterances, namely manipulated, synthesised or specially crafted utterances designed to deceive or manipulate the ASV system.

The resistance of ASV systems to spoofing attacks has been studied only relatively recently and within the context of the ASVspoof initiative and associated challenge series [1]. ASVspoof tackles the threat of spoofing attacks using countermeasures (CMs), separate detection sub-systems in the form of binary classifiers designed to distinguish between bona fide and spoofed inputs. While there are other approaches [2, 3], CMs are usually combined with the ASV system in the form of a gate, the role of which is to detect and reject spoofed utterances such that they are never treated by the ASV sub-system.

While the combination of CM and ASV sub-systems have the potential to improve security through increased robustness to spoofing attacks (rejection of spoofed trials), there is also the potential for degraded usability (rejection of bona fide target trials). An integrated approach to assessment is hence desirable and should gauge the impact of spoofing and CMs upon the ASV system. Such a strategy implies that neither CMs nor ASV systems should be assessed independently from each other. Ultimately, there is but a single goal – reliable ASV.

The minimum tandem detection cost function (min t-DCF) [4] metric has been developed with this vision and is one approach to the assessment of integrated CM and ASV sub-systems. Given the potential dependence between CM and ASV sub-systems, and given the integrated approach to evaluation, it seems logical that the sub-systems themselves should also be jointly developed and optimised. This is not the case with ASVspoof, for which CMs are developed by challenge participants, while the ASV system is designed by the organisers. Herein lies the original goal of the new Spoofing-Aware Speaker Verification (SASV) Challenge.<sup>1</sup>

SASV aims to promote the study of jointly optimised or fused CM+ASV solutions in addition to single, integrated solutions [5–13]. The former is a more flexible and incremental approach to SASV. Fusion can be applied using current CM and ASV sub-systems and can hence exploit future advances in both. With the addition of a new fusion model, however, complexity is increased, even if perhaps only modestly. The fusion of two separate sub-systems is also somewhat at odds from the spirit of solutions to the *single* task of reliable ASV.

Single, integrated solutions, envisaged on a longer-term horizon, represent more of step change and a more substantive SASV solution; they demand the design of entirely new solutions to the single problem of reliable ASV. The development of single, integrated solutions is expected to be more challenging; it calls for the learning of a new latent space for the representation of both speaker identity and artefacts related to utterance authenticity (spoofing). No matter what the solution strategy, the primary objective is progress in reliable ASV, where reliability implies resistance to both spoofed as well as bona fide non-target inputs.

In building upon long-established best practice in addition to the momentum generated through ASVspoof, the new SASV Challenge is built upon a common evaluation protocol and benchmarking framework, metrics, open source baseline solutions and use of publicly available ASVspoof data, specifically the 2019 Logical Access (LA) database [14]. The two different solution strategies are described in Section 2. SASV

<sup>\*</sup>These authors contributed equally to this work.

<sup>1</sup><https://sasv-challenge.github.io>

datasets and new metrics are introduced in Sections 3. Baseline sub-systems and solutions are described in Section 4. SASV Challenge results and solution strategies are presented in Sections 5 and 6. Conclusions are presented in Section 7.

## 2. Related Works

In this section, we describe prior, related work in two approaches to SASV: (i) back-end fusion; (ii) single, integrated models [15–22]. Among the first category are decision-level, score-level, and embedding-level approaches to back-end fusion. Decision-level fusion includes various cascade approaches. Approaches to score-level fusion can be either parameter-free (e.g., score-sum ensemble) or parameter-driven (e.g., Gaussian mixture model) where both utilise separate scores from ASV and CM sub-systems [18]. Embedding-level fusion can also be achieved using a model operating upon embeddings that lie in different latent spaces [21–23]. Prior work includes Gomez-Alanis et al. [22] and Shim et al. [21] which both propose deep neural network (DNN)-based models to jointly optimise ASV and CM embeddings and hence produce single SASV scores and decisions. Kanervisto et al. [23] reports a tandem solution to jointly optimise ASV and CM systems using reinforcement learning.

Single, integrated models have also been proposed. These models classify utterances into target and non-target classes where the latter comprises both traditional bona fide as well as spoofed non-target trials [3, 19, 20, 24]. Sizov et al. [3] propose a two-stage PLDA approach to the joint optimisation of speaker and synthesis (spoofing) channel variations in an i-vector space. First, a PLDA model is trained using only embeddings extracted from bona fide speech. Then, a synthesis (spoofing) channel subspace is trained using only embeddings extracted from spoofed speech. Zhao et al. [24] propose an integrated spoofing-robust automatic speaker verification (SR-ASV) system. It uses a multi-task learning framework with max feature map activation [25] and residual convolutional blocks to extract discriminative embeddings and scores from task-specific, CM and ASV layers. None of the work described above was performed using common databases and protocols.

## 3. Databases, protocols and metrics

Described in this section are the two publicly available databases used for the SASV challenge. The VoxCeleb2 [26] database is used for the training of ASV sub-systems (see Section 4.1). The ASVspoof 2019 LA database [14] is used for the training of CM sub-systems (see Section 4.2).

### 3.1. VoxCeleb2

The VoxCeleb2 database<sup>2</sup> was collected by crawling online videos of celebrity interviews. The database is extracted from 150,480 unique videos with an average individual utterance length of 7.8 seconds. The development partition of the VoxCeleb2 database is used for the training of ASV sub-systems. It contains over 2,000 hours of data corresponding to 5,994 speakers (61% male). The absence of spoofed utterances necessitates use of the VoxCeleb2 database in conjunction with additional databases containing spoofed utterances.

<sup>2</sup><https://www.robots.ox.ac.uk/~vgg/data/voxceleb/vox2.html>

Table 1: *Trial types used for estimation of each EER. “+” indicates the positive class and “-” indicates the negative class. Each trial involves enrolment utterance(s) and a test utterance. The enrolment utterance(s) is bona fide (i.e. genuine) whereas test utterance belongs to either of the three types.*

	Target	Bona fide non-target	Spoofed non-target
SV-EER	+	-	
SPF-EER	+		-
SASV-EER	+	-	-

### 3.2. ASVspoof 2019

The ASVspoof 2019 LA database [14] is generated from the VCTK<sup>3</sup> source database [27], which includes speech data captured from 107 speakers (46 male, 61 female). It consists of disjoint train, development, and evaluation partitions. Each partition contains both bona fide and spoofed utterances. The latter are generated using 19 state-of-the-art VC, TTS and hybrid TTS-VC attack algorithms (6 for the train and development partitions, 13 for the evaluation partition). Containing 25,380 utterances corresponding to 20 different speakers and both bona fide and spoofed utterances, the ASVspoof database can be used for the training of CM sub-systems but can also be used in conjunction with the VoxCeleb2 database for SASV research.

### 3.3. SASV protocols

While the SASV protocols<sup>4</sup> exploit ASVspoof 2019 LA data, they are different to those used by participants of the ASVspoof challenge; they are not CM protocols and are, instead, ASV protocols or, more specifically, SASV protocols. The latter involve three types of trial: (i) target, bona fide trials uttered by the same speaker as enrolment utterance(s); (ii) bona fide non-target trials uttered by a different speaker as enrolment utterance(s); (iii) spoofed non-target trials containing speech which is either synthesised or converted in order to resemble the voice of the same speaker as enrolment utterance(s). Disjoint protocols are provided to challenge participants to support their development and evaluation of SASV solutions.

### 3.4. Metrics

SASV performance is assessed using the classical EER (SASV-EER) as the primary metric. In keeping with [17, 18, 21], the task hence remains one of binary classification: target vs non-target. We define *non-target* as the set of bona fide non-target (impostor) trials and spoofed non-target trials. Without adequate countermeasures, both can cause increases in the false acceptance rate. As depicted in Table 1, two additional EER estimates serve as secondary metrics. The speaker verification EER (SV-EER) involves combinations of target trials and bona fide non-target trials whereas the spoofing EER (SPF-EER) involves combinations of target trials and spoofed non-target trials. The SV-EER and SPF-EER are estimated using different subsets of the full protocol used in estimating the SASV-EER. As such, they reflect the reliability of the model under different, extreme conditions in which there are either no spoofed non-target trials or no bona fide non-target trials.

<sup>3</sup><http://dx.doi.org/10.7488/ds/1994>

<sup>4</sup>[https://github.com/sasv-challenge/SASVC2022\\_Baseline/tree/main/protocols](https://github.com/sasv-challenge/SASVC2022_Baseline/tree/main/protocols)

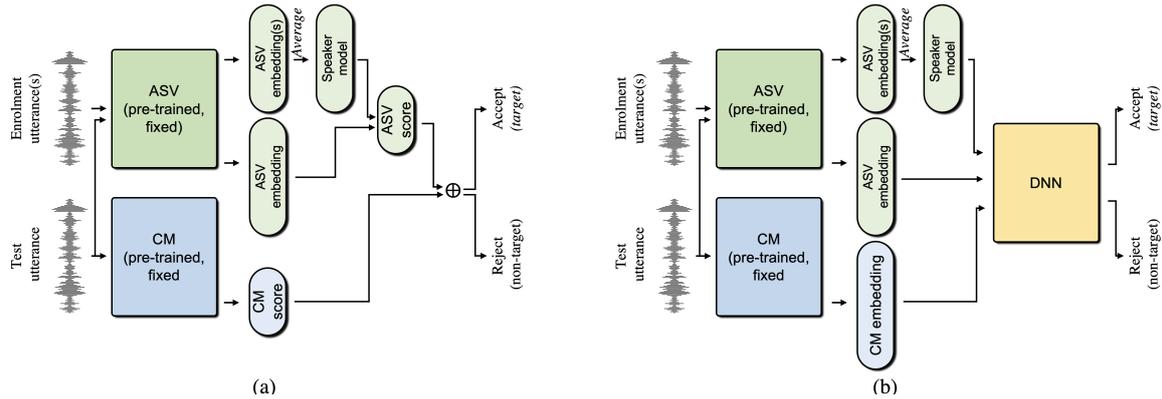


Figure 1: Illustration of the two baseline solutions of the SASV challenge. Both solutions utilise a pre-trained and fixed DNN-based ASV and CM sub-systems. (a): Baseline1 is a score-sum fusion where the cosine similarity score of ASV and CM output score are summed. (b): Baseline2 is a DNN-based fusion which involves speaker and spoofing embeddings.

## 4. SASV baseline solutions

There are two baseline SASV systems. Both comprise standalone CM and ASV sub-systems.

### 4.1. ASV and CM sub-systems

ECAPA-TDNN is an efficient, robust state-of-the-art speaker embedding extractor [28] consisting of a ResNet backbone architecture [29] with squeeze-excitation (SE) modules [30], a channel and context dependent statistics pooling layer and multi-layer feature aggregation to aggregate frame-level embeddings into utterance-level embeddings. Inputs are 80-dimensional mel-filterbank features. After aggregating frame-level representations, ASV embeddings are extracted using an affine transform with a fully-connected (FC) layer. Data augmentation techniques are applied using the room impulse response database [31] and additive noise recordings from the MUSAN database [32]. The network is trained using the recipe in [33] and a reproducible open source implementation.<sup>5</sup> Further details are available in [28].

AASIST [34] is an end-to-end state-of-the-art spoofing countermeasure system. It is based upon a RawNet2-based encoder [35, 36] and a spectro-temporal graph attention network (RawGAT-ST) [37], novel heterogeneous graph attention layers and max graph operations to integrate temporal and spectral representations. The output is generated using a readout operation and a hidden FC layer with two neurons. CM embeddings of 160 dimensions are extracted prior to the FC output layer. AASIST is also available as open source.<sup>6</sup> Further details are available in [34].

### 4.2. Baselines

Open source baselines include: i) **B1**: Score-sum fusion; and ii) **B2**: DNN back-end fusion.<sup>7</sup> Both utilise pre-trained (ECAPA-TDNN<sup>5</sup> and AASIST<sup>6</sup>) ASV and CM models. As illustrated in Figure 1-(a), baseline **B1** is a score-level back-end fusion method that combines ASV and CM sub-system outputs through score addition. Since it requires neither additional training nor fine-tuning, back-end processing is parameter-free. ASV scores are the cosine similarity between enrolment and test

<sup>5</sup><https://github.com/TaoRuijie/ECAPATDNN>

<sup>6</sup><https://github.com/clovaai/aasist>

<sup>7</sup>[https://github.com/sasv-challenge/SASVC2022\\_Baseline](https://github.com/sasv-challenge/SASVC2022_Baseline)

Table 2: SASV 2022 Challenge results in evaluation SASV EER (%). B1 and B2 correspond to the challenge baselines.

#	Team ID	EER	#	Team ID	EER
1	IDVoice	0.13	14	VTCC	1.86
2	DKU_OPPO	0.21	15	DeepASV	2.48
3	Hyu	0.28	16	SHELEZYAKA	2.77
4	DoubleRoc	0.37	17	xmuspeech	2.89
5	FlySpeech	0.56	18	HCCL	4.30
5	IRLAB	0.56	19	magnum	4.48
7	VicomSpeech	0.84	20	CAU_KU	4.95
8	CUHK_NTU	0.89	21	Tandem	6.22
9	MARG	1.15		B2	6.54
10	NIL_TJU	1.19	22	DHU	12.48
11	UR_AIR	1.34		B1	19.31
12	clips	1.36		ASV sub-system	23.83
13	orange_Lium	1.48	23	Souvik	24.32
	B1-v2	1.71			

utterances. A softmax non-linearity is optionally applied to the CM scores, which are otherwise unbounded, to normalise the scores within a (0,1) range. The version of B1 with the softmax non-linearity is referred to as B1-v2.

As illustrated in Figure 1-(b), baseline **B2** utilises an embedding-level fusion whereby a DNN operates upon a pair of speaker embeddings extracted from enrolment and test utterances, and a CM embedding extracted only from the test utterance. The DNN is a vanilla multi-layer perceptron comprising three FC layers with leaky ReLU non-linear activation functions after each layer. The output layer consists of two neurons which correspond to the target and non-target (both bona fide and spoofed) classes. The model is trained using the training partition of the ASVspool 2019 LA database. The larger-scale VoxCeleb2 database cannot be used for training since it contains only bona fide utterances. Full details regarding challenge baselines are available in [38].

## 5. Challenge results

From among 53 registrations, we received 23 submissions. SASV-EER results for each submission and the baselines for the evaluation partition are presented in Table 2.<sup>8</sup> Without any

<sup>8</sup> Full results are available at [https://sasv-challenge.github.io/challenge\\_results](https://sasv-challenge.github.io/challenge_results)

countermeasure, the ECAPA-TDNN ASV sub-system gives an SASV-EER of 23.83%. Among the baseline solutions and with an SASV-EER of 1.71%, baseline B1-v2 gives the best performance. With v2 having been developed after the release of the evaluation plan, in the remainder of this paper we emphasise v1 of the B1 baseline which gives an SASV-EER of 19.31%, a result which serves to show the importance of score normalisation prior to fusion. B2 gives an SASV-EERs of 6.54%.

From among the 23 submissions, 21 outperformed the B1 and B2 baselines. The submissions of 8 teams show SASV-EERs of below 1%, a substantial reduction from that of the B1 baseline and is even below the SV-EER of 1.63% (assessment without spoofed trials) of the ECAPA-TDNN ASV system [38]. These results are extremely encouraging and demonstrate the synergistic merit of combined ASV and CM systems. The top-performing system achieves an SASV-EER of 0.13% which corresponds to a 92% relative reduction compared to the SV-EER of the ECAPA-TDNN ASV system.

Detection error trade-off (DET) curves for all 23 submissions in addition to the baseline are plotted in Figure 2. Highlighted green, brown, red and blue profiles correspond to the top-performing submission and the B2, B1 baselines and ECAPA-TDNN ASV system respectively. The plots show substantial differences in the trade-off between false rejections and false acceptances and point toward the potential for tuning combined ASV and CM solutions for operating points other than the SASV-EER. Even so, the DET profile for the top-performing submission is below that of every other system. The elbow that is visible in many profiles stems from assessment being performed using both bona fide and spoofed non-target trials. Even if SASV-EER results show higher correlation with SV-EER than with SPF-EER results, the SASV-EER is still a composite of the SV-EER and the SPF-EER.

## 6. Solution strategies

All three systems with the lowest SASV-EER employed similar strategies [39–41]. First, they are all fusions of independent ASV and CM sub-systems. Second, they are all ensembles of multiple ASV+CM systems. Third, each of the three systems gives similarly low SV-EERs and SPF-EERs indicating that they are insensitive to the proportion of spoofed non-target trials. SV-EER and SPF-EER results for each system, for which we have no space to report here, are available online.<sup>8</sup>

There are differences too. The IDVoice [39] system fuses ASV and CM sub-systems at the score-level. The DKU\_OPPO system performs fusion at the decision-level, whereas embedding level fusion is employed in the Hyu system. Among their different systems, IDVoice use four different scores: ASV cosine similarity with and without score normalisation; CM cosine similarity; the score output of an end-to-end CM. Quality measurement functions [42] are also used to normalise the score for each trial according to various meta information such as the duration of speech. The DKU\_OPPO [40] system uses a modified cascade framework which combines decisions produced by the ASV sub-system with scores produced by the CM. Team HYU [41] focused on developing a latent space in which both speaker identity and spoofing artefacts can be captured using ASV and CM sub-system embeddings. A DNN with condition layers is used to map ASV and CM embeddings into a single SASV embedding [43]. There are further differences in the training strategies which are described in the participants’ system descriptions and associated articles.<sup>8</sup>

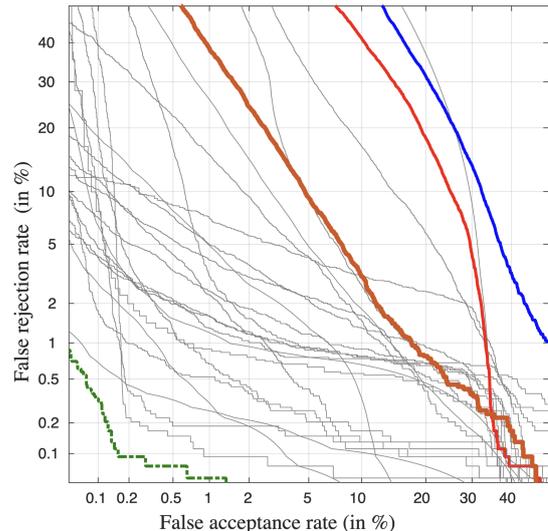


Figure 2: DET curves of challenge submissions on the evaluation protocol. green: best performing system, brown: B2, red: B1 (without softmax), blue: ECAPA-TDNN ASV system.

## 7. Discussion and Conclusions

The SASV Challenge was formed to promote the study of jointly optimised or fused and single solutions to reliable ASV. Reliability in this sense implies resistance to both bona fide and spoofed non-target trials. We received submissions from 23 participating teams who all designed competing solutions using a common training and development protocol. Results show substantial improvements over two Challenge baselines. At 0.13%, the lowest equal error rate for a spoofing-aware speaker verification system is even well below the equal error rate of a conventional, state-of-the-art speaker verification system assessed in the absence of spoofed trials. This is an especially encouraging result which not only indicates the reliability of today’s speaker verification and spoofing countermeasure technologies, but also the synergistic merit in their combination. There is also evidence that such impressive levels of reliability are consistent no matter what fraction of non-target trials that are spoofed; similarly low SV-EER and SPF-EER results show that performance is insensitive to the spoofing prior. We hope that these findings help to belay concerns of vulnerabilities to spoofing.

Looking to the future, we acknowledge some limitations of the current database. All top-performing systems remain ensemble combinations of independent countermeasure and speaker verification sub-systems. We anticipate future advances in joint optimisation stemming from the development of truly integrated, single system approaches. Research in this direction will likely depend on the availability of larger databases containing not only spoofing attacks generated with diverse attack algorithms, but also data collected from *many* more speakers. We hope that these developments may improve yet further on what are already extremely encouraging results.

## 8. Acknowledgements

We thank members of the 23 teams who submitted results to this, inaugural edition. We also thank advisory committee members, Héctor Delgado, Kong Aik Lee, Massimiliano Todisco, Md Sahidullah, and Xuechen Liu, for their help and support.

## 9. References

- [1] J. Yamagishi, X. Wang et al., “ASVspoof2021: Accelerating progress in spoofed and deep fake speech detection,” in *Proc. ASVspoof 2021 Workshop (Interspeech satellite)*, 2021.
- [2] P. L. De Leon, M. Pucher, J. Yamagishi et al., “Evaluation of speaker verification security and detection of hmm-based synthetic speech,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 8, pp. 2280–2290, 2012.
- [3] A. Sizov, E. Khoury, T. Kinnunen et al., “Joint speaker verification and antispoofing in the i-vector space,” *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 4, 2015.
- [4] T. Kinnunen, H. Delgado et al., “Tandem assessment of spoofing countermeasures and automatic speaker verification: Fundamentals,” *IEEE/ACM TASLP*, vol. 28, 2020.
- [5] J.-w. Jung, H. Tak, H.-j. Shim et al., “SASV Challenge 2022: A Spoofing Aware Speaker Verification Challenge Evaluation Plan,” *arXiv preprint:2201.10283*, 2022.
- [6] Y. Zhang, G. Zhu et al., “A Probabilistic Fusion Framework for Spoofing Aware Speaker Verification,” *arXiv preprint:2202.05253*, 2022.
- [7] P. Zhang, P. Hu and X. Zhang, “Norm-constrained score-level ensemble for spoofing aware speaker verification,” in *Proc. Interspeech*, 2022.
- [8] Z. Teng, Q. Fu et al., “SA-SASV: An End-to-End Spoof-Aggregated Spoofing-Aware Speaker Verification System,” in *Proc. Interspeech*, 2022.
- [9] W. Kang, J. Alam and A. Fathan, “End-to-end framework for spoof-aware speaker verification,” in *Proc. Interspeech*, 2022.
- [10] X. Liu, M. Sahidullah and T. Kinnunen, “Spoofing-Aware Speaker Verification with Unsupervised Domain Adaptation,” *arXiv preprint:2203.10992*, 2022.
- [11] L. Zhang, Y. Li, H. Zhao and L. Xie, “Backend Ensemble for Speaker Verification and Spoofing Countermeasure,” in *Proc. Interspeech*, 2022.
- [12] H. Wu, L. Meng, J. Kang et al., “Spoofing-Aware Speaker Verification by Multi-Level Fusion,” in *Proc. Interspeech*, 2022.
- [13] Z. Chang, L. Zhang et al., “Spoofing-Aware Attention based ASV Back-end with Multiple Enrollment Utterances and a Sampling Strategy for the SASV Challenge 2022,” in *Proc. Interspeech*, 2022.
- [14] X. Wang, J. Yamagishi et al., “ASVspoof 2019: a large-scale public database of synthesized, converted and replayed speech,” *Computer Speech & Language (CSL)*, vol. 64, 2020, 101114.
- [15] F. Alegre, A. Amehraye and N. Evans, “Spoofing countermeasures to protect automatic speaker verification from voice conversion,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, IEEE, 2013, pp. 3068–3072.
- [16] E. Khoury, T. Kinnunen, A. Sizov et al., “Introducing i-vectors for joint anti-spoofing and speaker verification,” in *Proc. Interspeech*. Citeseer, 2014.
- [17] M. Sahidullah, H. Delgado, M. Todisco et al., “Integrated spoofing countermeasures and automatic speaker verification: An evaluation on ASVspoof 2015,” in *Proc. Interspeech*, 2016.
- [18] 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.
- [19] J. Li, M. Sun and X. Zhang, “Multi-task learning of deep neural networks for joint automatic speaker verification and spoofing detection,” in *Proc. APSIPA*, 2019.
- [20] J. Li, M. Sun, X. Zhang and Y. Wang, “Joint decision of anti-spoofing and automatic speaker verification by multi-task learning with contrastive loss,” *IEEE Access*, 2020.
- [21] H.-j. Shim, J.-w. Jung, J.-h. Kim and H.-j. Yu, “Integrated replay spoofing-aware text-independent speaker verification,” *Applied Sciences*, 2020.
- [22] A. Gomez-Alanis, J. A. Gonzalez-Lopez et al., “On Joint Optimization of Automatic Speaker Verification and Anti-Spoofing in the Embedding Space,” *IEEE Transactions on Information Forensics and Security*, vol. 16, 2021.
- [23] A. Kanervisto, V. Hautamäki, T. Kinnunen and J. Yamagishi, “Optimizing Tandem Speaker Verification and Anti-Spoofing Systems,” *IEEE/ACM Transactions on Audio Speech and Language Processing (TASLP)*, vol. 30, pp. 477–488, 2022.
- [24] Y. Zhao, R. Togneri and V. Sreeram, “Multi-task learning-based spoofing-robust automatic speaker verification system,” *Circuits, Systems, and Signal Processing*, pp. 1–22, 2022.
- [25] X. Wu, R. He, Z. Sun and T. Tan, “A light cnn for deep face representation with noisy labels,” *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 11, 2018.
- [26] J. S. Chung, A. Nagrani and A. Zisserman, “Voxceleb2: Deep speaker recognition,” in *Proc. Interspeech*, 2018.
- [27] J. Yamagishi, C. Veaux, K. MacDonald et al., “CSTR VCTK corpus: English multi-speaker corpus for cstr voice cloning toolkit,” 2019, <http://dx.doi.org/10.7488/ds/1994>.
- [28] B. Desplanques, J. Thienpondt and K. Demuynck, “ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification,” in *Proc. Interspeech*, 2020.
- [29] S. Gao, M.-M. Cheng, K. Zhao et al., “Res2net: A new multi-scale backbone architecture,” *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- [30] J. Hu, L. Shen and G. Sun, “Squeeze-and-excitation networks,” in *Proc. CVPR*, 2018.
- [31] T. Ko, V. Peddinti, D. Povey et al., “A study on data augmentation of reverberant speech for robust speech recognition,” in *Proc. ICASSP*, 2017.
- [32] D. Snyder, G. Chen and D. Povey, “Musan: A music, speech, and noise corpus,” *arXiv preprint:1510.08484*, 2015.
- [33] R. K. Das, R. Tao and H. Li, “HLT-NUS SUBMISSION FOR 2020 NIST Conversational Telephone Speech SRE,” *arXiv preprint:2111.06671*, 2021.
- [34] J.-w. Jung, H.-S. Heo, H. Tak et al., “AASIST: Audio Anti-Spoofing using Integrated Spectro-Temporal Graph Attention Networks,” in *Proc. ICASSP*, 2022.
- [35] J.-w. Jung, S.-b. Kim, H.-j. Shim et al., “Improved rawnet with filter-wise rescaling for text-independent speaker verification using raw waveforms,” in *Proc. Interspeech*, 2020, pp. 1496–1500.
- [36] H. Tak, J. Patino, M. Todisco et al., “End-to-end anti-spoofing with Rawnet2,” in *Proc. ICASSP*, 2021.
- [37] H. Tak, J.-w. Jung, J. Patino et al., “End-to-end spectro-temporal graph attention networks for speaker verification anti-spoofing and speech deepfake detection,” in *Proc. ASVspoof 2021 Workshop (Interspeech satellite)*, 2021.
- [38] H. 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 Workshop*, 2022.
- [39] A. Alenin, N. Torgashov, A. Okhotnikov et al., “A Subnetwork Approach for Spoofing Aware Speaker Verification,” in *Proc. Interspeech*, 2022.
- [40] X. Wang, X. Qin, Y. Wang et al., “The DKU-OPPO System for the 2022 Spoofing-Aware Speaker Verification challenge,” in *Proc. Interspeech*, 2022.
- [41] J.-H. Choi, J.-Y. Yang, Y. Jeoung and J.-H. Chang, “HYU submission for the SASV challenge 2022: Reforming speaker embeddings with spoofing-aware conditioning,” in *Proc. Interspeech*, 2022.
- [42] A. Alenin, A. Okhotnikov, R. Makarov et al., “The ID R&D System Description for Short-duration Speaker Verification Challenge 2021,” in *Proc. Interspeech*, 2021, pp. 2297–2301.
- [43] E. Perez, F. Strub, H. De Vries et al., “Film: Visual reasoning with a general conditioning layer,” in *Proc. AAAI*, 2018.