



## BUT Systems and Analyses for the ASVspoof 5 Challenge

Johan Rohdin<sup>1,2</sup>, Lin Zhang<sup>1</sup>, Oldřich Plchot<sup>1</sup>, Vojtěch Staněk<sup>1</sup>, David Mihola<sup>1</sup>, Junyi Peng<sup>1</sup>,  
Themis Stafylakis<sup>2</sup>, Dmitriy Beveraki<sup>2</sup>, Anna Silnova<sup>1</sup>, Jan Brukner<sup>1</sup>, Lukáš Burget<sup>1</sup>

<sup>1</sup>Brno University of Technology, <sup>2</sup>Omilia Conversational Intelligence  
{rohdin, qzhang, iplchot}@fit.vutbr.cz, tstafylakis@omilia.com

### Abstract

This paper describes the BUT submitted systems for the ASVspoof 5 challenge, along with analyses. For the conventional deepfake detection task, we use ResNet18 and self-supervised models for the closed and open conditions, respectively. In addition, we analyze and visualize different combinations of speaker information and spoofing information as label schemes for training. For spoofing-robust automatic speaker verification (SASV), we introduce effective priors and propose using logistic regression to jointly train affine transformations of the countermeasure scores and the automatic speaker verification scores in such a way that the SASV LLR is optimized.

### 1. Introduction

Automatic speaker verification (ASV) systems are widely used to verify the identity of speakers. However, ASV systems are also vulnerable to spoofing attacks [1, 2, 3]. Although the main purpose of generative models is to facilitate people’s lives, not to attack biometric models or present false information, the advancement of generative models poses an increased threat to biometric systems and society. Thus, it is desirable to explore anti-spoofing systems (also known as countermeasures - CM or presentation attack detection - PAD) to detect and prevent spoofing attacks. To encourage researchers to work on this important task in the speech processing field, the ASVspoof [4, 5, 6, 7] challenge has been held since 2015. So far, three different spoofing scenarios have been discussed in previous years: (1) physical access (PA) for replay attacks, (2) logical access (LA) for spoofed speech generated by text-to-speech (TTS) synthesis and/or voice conversion (VC) attacks, and (3) deepfake (DF) for strongly compressed LA attacks.

This year’s ASVspoof 5 [8] involves two tracks. Track 1, like in previous years, involves conventional deepfake detection that discriminates bona fide speech from spoofed speech. For the main discussed spoofing scenario, ASVspoof 5 combined LA and DF based on advanced TTS/VC systems and introduced adversarial attacks (specifically focusing on Malafide [9] and its upgraded version Malacopula [10]) to the unseen evaluation set. Track 2 in ASVspoof 5 merged with the spoofing-robust automatic speaker verification (SASV) [11] 2022 challenge, with a newly defined task-agnostic metric, a-DCF [12]. Conveniently, newly proposed score fusion based on a non-linear combination of CM and ASV log-likelihood ratios (LLRs) after score calibration [13] and a single integrated system based on SKA-TDNN [14] are provided as baselines. The challenge also defines two conditions in each track - close and open - depending on whether it is allowed to use external data.

For track 1 close condition, we followed top-ranked teams from previous years and utilized ResNet18 as our submitted system. Furthermore, we explored the influence of training with

speaker labels in combination with spoofed/bona fide labels. As for the open condition, given the promising performance of SSL models for spoof detection [15, 16, 17], we compared different SSL models as front-end. We utilized our previously proposed Multi-head Factorized Attentive Pooling (MHFA) [18] to efficiently aggregate information from transformer layers through an attention mechanism, which showed superior results compared to a simple pooling layer.

For track 2, there are two common approaches: (1) score/embedding fusion [19, 20] between two independent ASV and CM, or (2) a single integrating model [21, 22] that optimizes ASV and CM simultaneously. Among these, score fusion is more widely used as it is more intuitive and makes the model decision more explainable. Most existing score fusion studies focus on simple score summation. However, in order to take a decision that minimizes the expected cost for a trial, a non-linear combination of the CM and ASV is necessary [13]. In this work, we provide a more general treatment of the SASV scoring problem. We derive the general SASV LLR and show that the optimal SASV decision for any choice of cost parameters can be taken from the SASV LLR where the priors have been replaced by *effective priors*, which are obtained by absorbing the costs of various incorrect decisions into the original priors. We then use these results to jointly optimize the calibration of the CM and ASV scores to provide an accurate SASV LLR.

### 2. System for track 1 deepfake detection

#### 2.1. Task and database

Track 1 of ASVspoof 5 focuses on the stand-alone speech deepfake detection task, distinguishing bona fide samples from spoofed ones, just like in previous competitions [4, 5, 6, 7].

The data of this challenge were collected during ASVspoof 5 Phase 1. The overall dataset is based on the Multilingual Librispeech (MLS) dataset (English-language subset) [23], and synthetic data are collected from community volunteers. Eight, eight, and sixteen spoofing attacks are considered for training, development, and evaluation sets, respectively. There are 400 and 785 speakers involved in the training and development set respectively. To measure the performance of deepfake detection, the minimum detection cost function (minDCF), the cost of log-likelihood ratio ( $C_{lr}$ ) [24], and the equal error rate (EER) are considered in track 1 of the challenge. More details can be found in the summary paper of the challenge [8].

#### 2.2. ResNet18 for the closed condition

##### 2.2.1. Details of the system

For the closed condition, we chose ResNet18 [25] as our system with MUSAN [26] noise subset and room impulse responses for

data augmentation, as it is the most used system by top-ranked teams in the previous years [27]. We use 80-dimensional Mel-filterbank with a window length of 25 ms and a frame shift of 10 ms. After extracting embedding from ResNet18, we use the temporal statistics pooling layer, a linear layer with 256 units, a ReLU activation function, a batch normalization layer, and another linear layer with softmax activation for calculating cross-entropy loss with K-class classification. The Likelihoods for bonafide/spoof were computed by summing the likelihoods over speakers and spoof types where applicable, after which the LLR were computed. This approach may not be ideal but due to time constraints we did not explore alternative approaches. In the next subsection, we will analyze different K based on whether and how we utilize speaker information for classification.

### 2.2.2. Comparison on different labeling schemes

To explore whether speaker information could help deepfake detection in the ASVspoof 5 challenge, we analyzed different label schemes considering different speaker identity and bona fide/spoof classes in this subsection. This is motivated by conflicting conclusions in existing studies. Some studies propose that simultaneously optimizing speaker classification and deepfake detection would enhance the robustness of deepfake detection [28]. Whereas others claim that reducing speaker variability would be beneficial for deepfake detection [29]. We examined five types of labeling schemes based on ResNet18 introduced in the previous subsection. Three (*spk-binspf*, *spk-mulspf*, *spk-onespf*) of these schemes include speaker identity, while the other two (*mulspf*, *binspf*) focus on bona fide/spoof(s) classification, as shown below. The numbers of classes K for each are annotated, given there are four hundred speakers, eight different spoofing methods (A01 ~ A08), and one bona fide in the training set.

- *spk-binspf* (K = 800): The label is a combination of speaker ID and bona fide/spoof,
- *spk-mulspf* (K = 3600): The label is a combination of speaker ID and bona fide/A01/.../A08,
- *spk-onespf* (K = 401): The label is speaker ID in case of bona fide, else “spoof,”
- *mulspf* (K = 9): The label is bona fide/A01/.../A08,
- *binspf* (K = 2): The label is bona fide/spoof, the same as in common deepfake detection.

Results using the above five label schemes on the development set of track 1 are shown in Table 1, and visualization of their embedding spaces by UMAP [30] are shown in Figure 1.

First, we focus on the training set regarding the seen scenario to discuss three questions:

(1) *Should we consider speaker information for spoofed speech?* Figure 1 (a) to (c) show the models considering speaker ID. In the (a) *spk-binspf* and (b) *spk-mulspf*, which treat spoofing methods for each speaker as independent classes, we observe that although some bona fide samples are clustered in the center, most samples are mixed and difficult to distinguish. In (c) *spk-onespf*, after we integrate all spoofed samples as a single spoof class without considering “spoofed” speaker information, the spoofed speech begins to distinguish from bona fide speech. This model achieves the lowest EER and promising minDCF compared with other models. This hints that it is challenging to train the model when assigning speaker ID to the spoofed speech. This could be due to the reduced number of samples for each class or the confusion introduced by speaker information for spoofed speech.

Table 1: Results of ResNet18 with different label schemes on the development set of track 1.

ID	Model	minDCF	EER (%)	$C_{lr}$	actDCF
1	<i>spk-binspf</i>	0.2401	13.640	2.6596	0.9349
2	<i>spk-mulspf</i>	0.2708	13.818	0.7985	0.6129
3	<i>spk-onespf</i>	0.1624	11.891	3.1482	0.7794
4	<i>mulspf</i>	0.1811	12.456	5.4365	0.9978
5	<i>binspf</i>	0.1374	12.156	2.6360	0.6720

(2) *Should we consider speaker information for bona fide speech?* Visualization on comparing Figure 1 (b) vs. (d), and (c) vs. (e) shows that when we remove speaker information and focus on deepfake detection as in (d) and (e), the samples are more clustered according to their bona fide/spoof labels. This is understandable, as the models in (b) and (c) contain speaker information, and the differences between speaker characteristics might be larger than the differences between spoofed and bona fide samples, which makes it difficult to learn how to distinguish bona fide from spoof.

(3) *Should we integrate different spoofing methods?* Comparing (d) and (e), we can observe that integrating different spoofing methods as a single spoof class helps the model distinguish spoofed samples from bona fide. Given that we didn’t treat different spoofing methods independently, it is understandable that A01 to A06 are mixed. Notably, A07 and A08 are still well distinguished from others even though they are trained under the same label. Similar observations can be found in (c-d). This is acceptable, as A07 is generated by FastPitch [31] and A08 is generated by VITS [32], which are different compared to the other six spoofing methods based on GlowTTS and GradTTS.

Next, we move to the development set, which contains unseen spoofing methods and disjoint speakers. Across all five label schemes in Figure 1 (a-e), we observe that A11 (Tacotron 2 [33]), A13 (StarGANv2-VC [34]) and A14 (YourTTS [35]) are well distinguished in all label schemes compared with other spoofing methods. This could be because A11, A13, and A14 utilized similar or the same components with spoofing methods that the model encountered during training. For example, Tacotron2 technology applied in A11 is also utilized by A07 (FastPitch [31]) to estimate the duration of the input symbols. YourTTS used in A14 is built based upon VITS [32] that is applied by A08 of the training set. This shows that the model is still limited to the seen scenarios, and its performance is restricted by the degree of mismatch from the training set. Meanwhile, A12 (In-house unit-selection) [8] is difficult to distinguish from the bona fide, which is understandable as unit-selection selecting segments from bona fide utterances to consist the desired spoofed speech, and it is unseen during training. More robust technology for detecting such unit-selection attacks is worth exploring in future work. Additionally, exploring adaptation methods to handle mismatched scenarios more effectively is essential for future research.

Finally, we submit the system 5 (ResNet18-*binspf*) that uses bona fide/spoof classes as our final system for the track 1 - close condition. It achieves minDCF=0.5809, actDCF=0.8537,  $C_{lr}$ = 4.0994, and EER=23.34% in the evaluation set.

### 2.3. Pretrained SSL with MHFA for the open condition

SSL models have attracted attention in the deepfake speech detection area due to their promising performance [15, 16, 17]. Therefore, we used pretrained SSL models as our CM system

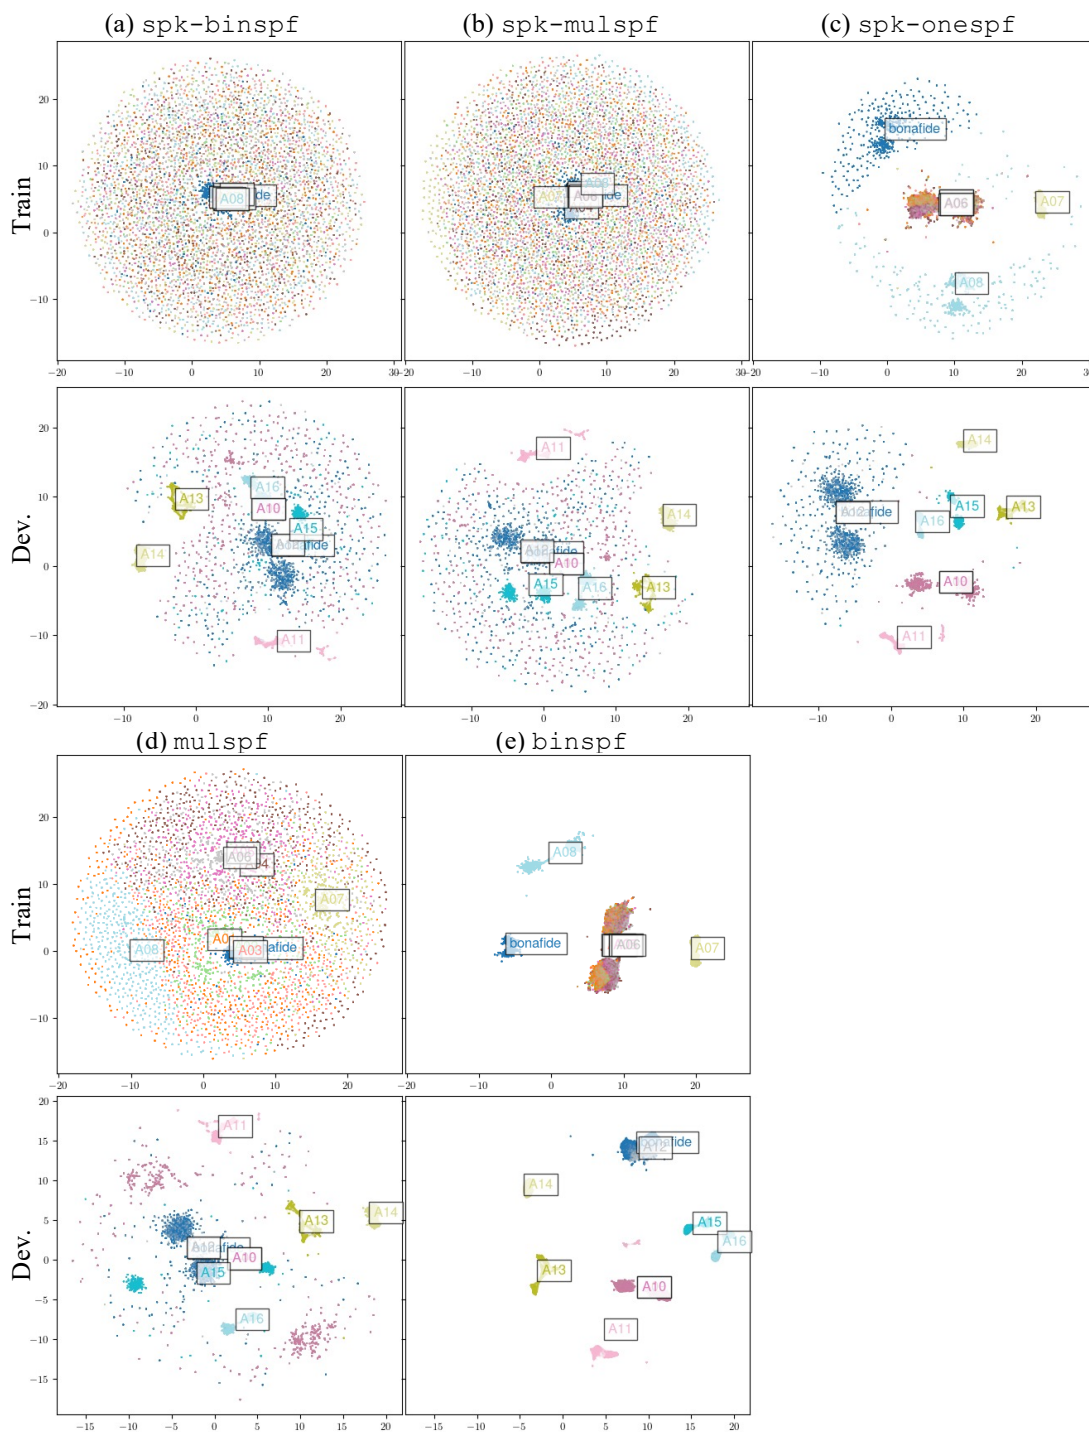


Figure 1: *Embedding space of different label schemes.*

submitted to track 1-open.

Specifically, we compared different pretrained SSL models including Wav2vec2 model [36], WavLM model [37], Hubert [38], and data2vec [39] as shown in Table 2. All SSL models are in their Base version given the prohibition using of LibriLight in ASVspoof challenge [8]. In addition, when we aggregated hidden features extracted from transformer layers of SSL models, we compared learnable weighted sum [40] with average pooling

(AP) vs. MHFA [18]. MHFA is our recently proposed pooling method that utilizes two sets of normalized layer-wise weights to generate attention maps and compressed features. We followed the configuration from our previous paper [18] with the number of heads set as 64. During training, only the parameters of MHFA while keeping SSL models frozen. No data augmentation was applied in these experiments. Results are shown in Table 2.

Table 2: Performances of pretrained SSL models (Base versions) on the development set of Track 1-open condition.

ID	Model	minDCF	EER (%)	$C_{lr}$	actDCF
1	Wav2vec2 + AP	0.1312	5.094	0.2924	0.1674
2	Wav2vec2 + MHFA	0.0848	3.300	0.5876	0.1097
3	HuBERT + MHFA	0.2497	11.164	0.8306	1.0000
4	WavLM + MHFA	0.1400	6.881	0.8268	1.0000
5	Data2vec + MHFA	0.2231	8.400	0.8067	0.2724
6	Fusion 2 + 4	0.0763	2.974	0.87861	1.0000
7	Fusion 6 + ResNet18	0.0693	3.048	0.90159	1.0000

Comparing systems 1 and 2, MHFA outperformed simple AP. Comparing 2 ~ 5, Wav2vec2 and WavLM achieve better performance. Thus, we submitted the system 7 – equal weight averaging of max-min normalized prediction scores from systems 2, 4 from Table 2 and system 5 from Table 1. The fused system 7 achieves minDCF=0.2573, actDCF=1.0000,  $C_{lr}$ =0.9955, EER=9.28% in the evaluation set.

### 3. System for track 2

#### 3.1. Task and database

Track 2 of ASVspoof 5 involves a spoofing-robust automatic speaker verification (SASV) task [11]. This is a newly introduced track that has emerged in recent years, aiming to integrate ASV and CM systems and accepting the speech only if it is spoken by the bona fide target speaker. The training and development sets are the same as in track 1 with the additional well-known Voxceleb2 [41] database from the speaker verification field. Voxceleb2 provided 5,994 speakers for training. For measuring performance, different from the SASV 2022 [11], which utilizes SASV-EER as the main metric, this year’s challenge uses the newly introduced min a-DCF [12] as the primary metric, with min t-DCF [42] and t-EER [43] as the supplement metrics. More details can be found in the summary paper of the challenge [8].

#### 3.2. ASV systems

We based our ASV systems on the ResNet architecture by following the exact recipes with the Voxceleb dataset from the Wespeaker toolkit<sup>1</sup> [44] while omitting speech and music parts from the MUSAN data that were not allowed to be used in the challenge. We experimented on ResNet34, ResNet101, and ResNet221 with Additive Angular Margin (AAM) softmax loss. We compared models w/ and w/o large margin fine-tuning. As enrollment embeddings, we used the average embedding from multiple enrollment utterances. We also analyzed normalizing extracted embeddings by subtracting the mean of the ASVspoof5 train set/voxceleb2 dev set. For scoring, we used simple cosine similarity. The results are shown in Table 3.

#### 3.3. Spoofing-robust automatic speaker verification system

Our SASV system is based on combining the CM LLR and the ASV LLR into a SASV LLR, i.e., the LLR for the hypotheses

- $\mathcal{H}_A$  (Accept hypothesis): The speech is bona fide and from the target speaker.
- $\mathcal{H}_R$  (Reject hypothesis):  $\mathcal{H}_A$  is not true.

<sup>1</sup><https://github.com/wenet-e2e/wespeaker>

Table 3: EERs (%) of ASV systems on the development set of Track 2.

ID	Model	Large Margin	Mean Norm.	EER (%)
1	ResNet34	✓	-	5.935
2	ResNet34	✓	ASVspoof5 train	5.605
3	ResNet34	✓	Voxceleb2 dev.	5.952
4	ResNet34	-	ASVspoof5 train	5.464
5	ResNet101	-	ASVspoof5 train	5.233
6	ResNet221	-	ASVspoof5 train	5.101

Since for binary decision problems, the optimal decision can be taken from the LLR, this is the optimal score for a SASV system. Under some assumptions, the SASV LLR can be computed from the LLR of the CM system, the LLR of the ASV system, and the priors of the cost parameters. The formula for the resulting SASV LLR has been provided before, e.g., in [13] and [45]. Here, we present a slightly more general form of the LLR and introduce the concept of *effective priors* for the SASV task, inspired by this concept in speaker verification [46]. The use of effective priors shows how to make optimal decisions in the scenarios where the different types of false accept have different costs, which is not the case in this evaluation<sup>2</sup> but may be useful in other scenarios and, more importantly for this evaluation, enables us to do fusion/calibration using logistic regression.

The details of our approach are provided in the following subsections. We denote the speech  $X$  and the properties of the speech as follows:

- $B$ : speech is bona fide
- $S$ : speech is spoofed
- $T$ : speech is from the target speaker
- $N$ : speech is not from the target speaker

where  $B$  and  $S$  are disjoint and  $T$  and  $N$  are disjoint. Note that  $\mathcal{H}_A = (B, T)$  and  $\mathcal{H}_R$  is the union of  $(B, N)$ ,  $(S, T)$  and  $(S, N)$ <sup>3</sup>.

##### 3.3.1. Optimal scores and decisions

Given some speech data,  $X$ , the expected cost of rejecting a trial is  $C_{\text{miss}}P(\mathcal{H}_A|X)$ , where  $C_{\text{miss}}$  is the cost of false rejection. If the cost of incorrectly accepting a spoofed utterance,  $C_{\text{fa,spoof}}$ , and incorrectly accepting an impostor,  $C_{\text{fa,imp}}$ <sup>4</sup>, is the same (as is the case in ASVspoof 5), the expected cost of accepting a trial is  $C_{\text{fa}}P(\mathcal{H}_R|X)$ , where  $C_{\text{fa}} = C_{\text{fa,spoof}} = C_{\text{fa,imp}}$ . The problem is then a standard binary decision theory problem for which the optimal decision is to accept if (see e.g. the BOSARIS toolkit manual [46])

$$\frac{C_{\text{miss}}P(\mathcal{H}_A|X)}{C_{\text{fa}}P(\mathcal{H}_R|X)} = \frac{C_{\text{miss}}P(X|\mathcal{H}_A)P(\mathcal{H}_A)}{C_{\text{fa}}P(X|\mathcal{H}_R)P(\mathcal{H}_R)} > 1$$

$$\Leftrightarrow \log \frac{P(X|\mathcal{H}_A)}{P(X|\mathcal{H}_R)} > \log \frac{C_{\text{fa}}}{C_{\text{miss}}} - \log \frac{P(\mathcal{H}_A)}{P(\mathcal{H}_R)}. \quad (1)$$

Although the evaluation plan did not explicitly ask participants to provide LLRs for track 2, the above shows that using LLRs

<sup>2</sup>Both false acceptance of an impostor (ASV false accept) and false acceptance of spoofed speech (CM false accept) have cost 10 in this evaluation.

<sup>3</sup>When we refer to spoofed speech from a target/non-target speaker, we mean spoofed speech with simulated characteristics similar to those of the true target/non-target speaker.

<sup>4</sup>Note that in the ASVspoof 5 evaluation plan, this is denoted  $C_{\text{fa}}$ .

and an appropriate threshold leads to optimal decisions.<sup>5</sup>

### 3.3.2. LLR

The LLR is given by

$$\begin{aligned}
& \log \frac{P(X|\mathcal{H}_A)}{P(X|\mathcal{H}_R)} \\
&= \log \frac{P(X|B, T)}{P(X|B, N)p_{\text{BN}} + P(X|S, T)p_{\text{ST}} + P(X|S, N)p_{\text{SN}}} \\
&= -\log \frac{P(X|B, N)p_{\text{BN}} + P(X|S, T)p_{\text{ST}} + P(X|S, N)p_{\text{SN}}}{P(X|B, T)} \\
&= -\log \left( \frac{P(X|B, N)}{P(X|B, T)}p_{\text{BN}} + \frac{P(X|S, T)}{P(X|B, T)}p_{\text{ST}} \right. \\
&\quad \left. + \frac{P(X|S, N)}{P(X|B, T)}p_{\text{SN}} \right), \tag{2}
\end{aligned}$$

where

- $p_{\text{BN}} = P(B, N|\mathcal{H}_R)$
- $p_{\text{ST}} = P(S, T|\mathcal{H}_R)$
- $p_{\text{SN}} = P(S, N|\mathcal{H}_R)$

while  $p_{\text{BT}} = P(B, T|\mathcal{H}_A) = 1$  is kept implicit. The first term after the final equal sign in Eq. (2) is just the inverted LLR for ASV on *bona fide* speech and the second term is just the inverted *speaker-dependent* LLR for a CM. Our CM systems are speaker-independent, which is incorrect according to the above formula but hopefully a good enough approximation. Note that the LLR depends on conditional priors,  $p_{\text{BN}}$ ,  $p_{\text{ST}}$  and  $p_{\text{SN}}$ . This is common to LLRs that are composed of several LLRs for simpler *subevents*, see [47] and [48] for two other examples. In many SASV scenarios (including the challenge if we understand correctly), we do not expect spoofing of non-target speakers, i.e.,  $p_{\text{SN}} = 0$ .

### 3.3.3. Effective priors

Analogously to common practice in ASV [46], it simplifies matters to convert the cost parameters into an *equivalent* set of costs parameters<sup>6</sup> where  $C_{\text{miss}} = C_{\text{fa, spoof}} = C_{\text{fa, imp}} = C_{\text{fa, spoof, imp}} = 1$  but where the priors are replaced with the *effective priors*. With general costs, we shall accept the trial if

$$\frac{P(X|B, T)P(B, T)C_{\text{miss}}}{P(X|B, N)P(B, N)C_{\text{fa, imp}} + P(X|S, T)P(S, T)C_{\text{fa, spoof}} + P(X|S, N)P(S, N)C_{\text{fa, spoof, imp}}} > 1. \tag{3}$$

By dividing the numerator and denominator by

$$\begin{aligned}
Z &= P(B, T)C_{\text{miss}} + P(B, N)C_{\text{fa, imp}} \\
&\quad + P(S, T)C_{\text{fa, spoof}} + P(S, N)C_{\text{fa, spoof, imp}}, \tag{4}
\end{aligned}$$

we obtain

$$\frac{P(X|B, T)P'(B, T)}{P(X|B, N)P'(B, N) + P(X|S, T)P'(S, T) + P(X|S, N)P'(S, N)} > 1, \tag{5}$$

<sup>5</sup>Strictly speaking, we do not need to provide the threshold, and the LLRs could be subjected to any monotonically rising function and still be optimal for the challenge metric since it does not care about calibration.

<sup>6</sup>Contrary to Section 3.3.1, we here also include spoofed impostors. The cost of false acceptance of such trials is denoted  $C_{\text{fa, spoof, imp}}$ .

where, e.g.,

$$P'(B, N) = C_{\text{fa, imp}}P(B, N)/Z, \tag{6}$$

and the other *effective priors*,  $P'(B, T)$ ,  $P'(S, T)$  and  $P'(S, N)$  are defined similarly. This means that the decision that is optimal according to Ineq. (5) is also optimal according to Eq. (3) and vice versa. Thus, we can work with  $C_{\text{miss}} = C_{\text{fa, spoof}} = C_{\text{fa, imp}} = C_{\text{fa, spoof, imp}} = 1$  and the original priors replaced by the effective priors,  $P'(\cdot, \cdot)$ , when taking decisions such as calibration and fusion models. In this way, we can handle situations when originally, e.g.,  $C_{\text{fa, spoof}} \neq C_{\text{fa, imp}}$ . In the case when  $C_{\text{fa, spoof}} = C_{\text{fa, imp}}$ , working with effective priors is still helpful for training, e.g., calibration models (see the next subsection).

### 3.3.4. Logistic regression based calibration/fusion

Let  $p_{\text{SN}} = 0$  and denote

$$\text{llr}_{\text{cm}}(X) = \log \frac{P(X|B, T)}{P(X|S, T)} \tag{7}$$

and

$$\text{llr}_{\text{asv}}(X) = \log \frac{P(X|B, T)}{P(X|B, N)}, \tag{8}$$

then

$$\begin{aligned}
\text{llr}_{\text{sasv}}(X) &= \log \frac{P(X|\mathcal{H}_A)}{P(X|\mathcal{H}_R)} \\
&= -\log \left( p'_{\text{BN}}e^{-\text{llr}_{\text{cm}}(X)} + p'_{\text{ST}}e^{-\text{llr}_{\text{asv}}(X)} \right) \tag{9}
\end{aligned}$$

We note that although the primary metric, min a-DCF, is calibration insensitive, i.e., it does not care whether the SASV LLR is calibrated, proper calibration of the CM and ASV LLR is still important for min a-DCF due to the complex relation between them and the SASV LLR given by Eq. (9). Calibrating the raw CM and ASV LLRs (denoted  $\text{llr}_{\text{cm}}^{\text{raw}}$  and  $\text{llr}_{\text{asv}}^{\text{raw}}$ ) with affine transformations, we obtain the *corrected* SASV LLR

$$\begin{aligned}
&\text{llr}_{\text{sasv}}(X, \tilde{a}_0, \tilde{a}_1, \tilde{c}_0, \tilde{c}_1) \\
&= -\log \left( p'_{\text{BN}}e^{-\tilde{c}_1\text{llr}_{\text{cm}}^{\text{raw}}(X) - \tilde{c}_0} + p'_{\text{ST}}e^{-\tilde{a}_1\text{llr}_{\text{asv}}^{\text{raw}}(X) - \tilde{a}_0} \right). \tag{10}
\end{aligned}$$

We then learn the calibration parameters  $\tilde{a}_0$ ,  $\tilde{a}_1$ ,  $\tilde{c}_0$  and  $\tilde{c}_1$  jointly with logistic regression with the three classes,  $(B, T)$ ,  $(B, N)$  and  $(S, T)$ , being weighted according to their effective priors, i.e.,

$$\begin{aligned}
&\tilde{a}_0, \tilde{a}_1, \tilde{c}_0, \tilde{c}_1 = \\
&\arg \min_{\substack{\tilde{a}_0, \tilde{a}_1, \\ \tilde{c}_0, \tilde{c}_1}} \sum_D \frac{P'(D)}{N_D} \sum_{i=1}^{N_D} L(X, S_D, a_0, a_1, c_0, c_1) \tag{11}
\end{aligned}$$

where

$$\begin{aligned}
&L(X, S_D, a_0, a_1, c_0, c_1) \\
&= \log \left( 1 + e^{-S_D(\text{llr}_{\text{sasv}}(X, a_0, a_1, c_0, c_1) + \tau)} \right), \tag{12}
\end{aligned}$$

$D = \{(B, T), (B, N), (S, T)\}$ ,  $N_D$  is the number of trials for class  $D$ ,

$$S_D = \begin{cases} 1 & \text{for } (B, T) \\ -1 & \text{for } (B, N) \text{ and } (S, T) \end{cases} \tag{13}$$

Table 4: *Impact of calibration for the SASV LLR of track 2 in terms of min a-DCF on the ASVspoof5 development set. No calibration refers to combining the CM and ASV LLR according Eq. (9), i.e., without any calibration and Calibration refers to combining the CM and ASV LLR according Eq. (10) with the calibration parameters optimized according to Eq. (12).*

No calibration	Calibration
0.17874	0.16854

Table 5: *Results of various systems on the closed condition of track 2 on the ASVspoof5 development set. The number under the CM system heading refers to the ID in Table 1 and the number under the ASV system heading refers to the ID in Table 3.*

CM system	ASV system	min a-DCF	min t-DCF	t-EER (%)
1	2	0.16854	0.35234	8.196
5	2	0.12696	0.20858	6.569
5	5	0.12529	0.20858	5.977
5	6	0.12527	0.20924	6.026

Table 6: *Results of various systems on the open condition of Track 2 on the ASVspoof5 development set. The number under the CM system heading refers to the ID in Table 2 and the number under the ASV system heading refers to the ID in Table 3.*

CM system	ASV system	min a-DCF	min t-DCF	t-EER (%)
2	6	0.04761	0.18886	2.514
7	6	0.07287	0.15573	2.026

and

$$\tau = \log \frac{P'(B, T)}{P'(B, N) + P'(S, T)}. \quad (14)$$

Logistic regression on LLR is a standard approach for calibration and fusion in speaker verification [49] which encourages good calibration as well as discrimination.

### 3.3.5. Experiments

We implemented the logistic regression calibration in Pytorch [50]. For optimization, we used its L-BFGS [51] optimizer with default settings. The result using `spk-binspf` of Track one as CM LLR and system 2 of Table 1 the ASV LLR are presented in Table 4. We can see that the proposed calibration improves the min a-DCF with around 1% absolute. Due to time constraints, we have not evaluated the effect of calibration for systems other than `spk-binspf`, nor have we compared it with alternative approaches for combining the CM and ASV LLR. A brief discussion of some of the conceptual aspects is provided in the next subsection while further experimental evaluations and analysis should be part of future work.

The results for the closed condition of track 2 are shown in Table 5. Consistent with the results for track 1, we can see that CM system 5 outperforms CM system 1. The ASV system has a very minor impact on min a-DCF but a larger impact on t-EER. For the closed condition of track 2, we submitted the system in the last row. It achieves min a-DCF=0.389, min t-DCF=0.778, t-EER=20.850% in the evaluation set.

The results for the open condition of track 2 are shown in Table 6. Contrary to track 1, the fusion of several CM systems did not improve in the primary metric for track 2. Due to limited time, we did not explore this further and submitted the system in

the first row. It achieves min a-DCF=0.180, min t-DCF=0.543, t-EER=8.390% in the evaluation set.

### 3.3.6. Discussion

Equation (9) is the same as [13] and [45]. In those papers, it was suggested to tune  $p'_{BN}$  and  $p'_{ST}$  with a grid search. In addition, discriminative calibration of the CM and ASV LLR was done individually before combining them to form the SASV LLR in [13]. In [45], both the ASV LLR and the CM LLR were estimated by a generative (Gaussian) fusion of the raw CM and raw ASV scores. However we keep  $p'_{BN}$  and  $p'_{ST}$  as specified by the cost parameters (the effective versions) and instead jointly learn affine transformations of  $llr_{cm}$  and  $llr_{asv}$  that optimizes the SASV LLR on the left side of Eq. (9). A few points can be made:

- Most calibration methods rarely produce scores that are well-calibrated at all operating points. Joint optimization of the CM and ASV calibration should calibrate these LLRs to be optimal for the operating point of the SASV task. This speaks in favor of our proposed method.
- Tuning  $p'_{BN}$  and  $p'_{ST}$  corresponds to adjusting the offsets of the CM and ASV LLR. This is less powerful than affine transformations. However, individual precalibration of the CM and ASV LLR followed by tuning of  $p'_{BN}$  and  $p'_{ST}$  for the SASV as in [13] and [45] task could be sufficient.
- Tuning parameters with a grid search as in [13] and [45] allows optimizing the performance of DCF at one specific operating point. Since the DCF of one operating point is not a continuous function, it cannot be optimized by gradient-based methods such as L-BFGS.
- Logistic regression corresponds to optimizing the calibration for a wide range of operating points instead of the one specified by DCF [24]. This could make the score less optimized for the specific operating point but, on the other hand, reduce the risk of overfitting to this specific operating point.

The pros and cons of different calibration/fusion methods need to be analyzed in future work.

## 4. Conclusion

This paper described BUT systems and analyses for the ASVspoof 5 challenge. For track 1, we constructed ResNet18 with analyses on different speaker and spoofing label schemes for the close condition and pretrained SSL model with MHFA for the open condition. For track 2, we defined SASV LLR in a more general form with an introduced concept of effective priors. Introducing effective priors enables optimal decision for the SASV task regardless of cost parameters. It also enables calibrating SASV LLRs as well as evaluating the quality of such LLRs with calibration sensitive metrics.

## 5. Acknowledgements

This work was partly supported by the European Union’s Horizon Europe grant agreement No. 101135916 “ELOQUENCE,” and by the Czech Ministry of Interior project No. VB02000060 “NABOSO.” We acknowledge VSB – Technical University of Ostrava, IT4Innovations National Supercomputing Center, Czech Republic, for awarding this project access to the LUMI supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CSC (Finland) and the LUMI consortium through the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (grant ID: 90254).

## 6. References

- [1] Nicholas Evans, Tomi Kinnunen, and Junichi Yamagishi, “Spoofing and countermeasures for automatic speaker verification,” in *Proc. Interspeech*, 2013, pp. 925–929.
- [2] Tim Ring, “Europol: the ai hacker threat to biometrics,” *Biometric Technology Today*, vol. 2021, no. 2, pp. 9 – 11, 2021.
- [3] Anton Firc and Kamil Malinka, “The dawn of a text-dependent society: deepfakes as a threat to speech verification systems,” in *Proc. ACM/SIGAPP SAC*, 2022, p. 1646–1655.
- [4] Zhizheng Wu, Tomi Kinnunen, Nicholas Evans, Junichi Yamagishi, Cemal Hanilçi, Md Sahidullah, and Aleksandr Sizov, “ASVspooF 2015: the First Automatic Speaker Verification Spoofing and Countermeasures Challenge,” in *Proc. Interspeech*, 2015, pp. 2037–2041.
- [5] 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, pp. 2–6.
- [6] Andreas Nautsch, Xin Wang, Nicholas Evans, Tomi H. Kinnunen, Ville Vestman, Massimiliano Todisco, Héctor Delgado, Md Sahidullah, Junichi Yamagishi, and Kong Aik Lee, “ASVspooF 2019: Spoofing countermeasures for the detection of synthesized, converted and replayed speech,” *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, no. 2, pp. 252–265, 2021.
- [7] Junichi Yamagishi, Xin Wang, Massimiliano Todisco, Md Sahidullah, Jose Patino, Andreas Nautsch, Xuechen Liu, Kong Aik Lee, Tomi Kinnunen, Nicholas Evans, and Héctor Delgado, “ASVspooF 2021: accelerating progress in spoofed and deepfake speech detection,” in *Proc. ASVspooF Workshop*, 2021, pp. 47–54.
- [8] Xin Wang, Héctor Delgado, Hemlata Tak, Jee-weon Jung, Hye-jin Shim, Massimiliano Todisco, Ivan Kukanov, Xuechen Liu, Md Sahidullah, Tomi Kinnunen, Nicholas Evans, Kong Aik Lee, and Junichi Yamagishi, “ASVspooF 5: Crowdsourced speech data, deepfakes, and adversarial attacks at scale,” in *Proc. ASVspooF Workshop 2024 (accepted)*.
- [9] Michele Panariello, Wanying Ge, Hemlata Tak, Massimiliano Todisco, and Nicholas Evans, “Malafide: a novel adversarial convolutive noise attack against deepfake and spoofing detection systems,” in *Proc. Interspeech*, 2023, pp. 2868–2872.
- [10] Massimiliano Todisco, Michele Panariello, Xin Wang, Hector Delgado, Kong-Aik Lee, and Nicholas Evans, “Malacopula: Adversarial automatic speaker verification attacks using a neural-based generalised hammerstein model,” in *Proc. ASVspooF Workshop 2024 (accepted)*.
- [11] Jee weon Jung, Hemlata Tak, Hye jin Shim, Hee-Soo Heo, Bong-Jin Lee, Soo-Whan Chung, Ha-Jin Yu, Nicholas Evans, and Tomi Kinnunen, “SASV 2022: The First Spoofing-Aware Speaker Verification Challenge,” in *Proc. Interspeech*, 2022, pp. 2893–2897.
- [12] Hye jin Shim, Jee weon Jung, Tomi Kinnunen, Nicholas Evans, Jean-François Bonastre, and Itshak Lapidot, “a-DCF: an architecture agnostic metric with application to spoofing-robust speaker verification,” in *Proc. Odyssey*, 2024, pp. 158–164.
- [13] Xin Wang, Tomi Kinnunen, Lee Kong Aik, Paul-Gauthier Noé, and Junichi Yamagishi, “Revisiting and improving scoring fusion for spoofing-aware speaker verification using compositional data analysis,” in *Proc. Interspeech*, 2024, p. (accepted).
- [14] Sung Hwan Mun, Hye jin Shim, Hemlata Tak, Xin Wang, Xuechen Liu, Md Sahidullah, Myeonghun Jeong, Min Hyun Han, Massimiliano Todisco, Kong Aik Lee, Junichi Yamagishi, Nicholas Evans, Tomi Kinnunen, Nam Soo Kim, and Jee weon Jung, “Towards Single Integrated Spoofing-aware Speaker Verification Embeddings,” in *Proc. Interspeech*, 2023, pp. 3989–3993.
- [15] Xin Wang and Junichi Yamagishi, “Investigating Self-Supervised Front Ends for Speech Spoofing Countermeasures,” in *Proc. Odyssey*, 2022, pp. 100–106.
- [16] Hemlata Tak, Massimiliano Todisco, Xin Wang, Jee weon Jung, Junichi Yamagishi, and Nicholas Evans, “Automatic Speaker Verification Spoofing and Deepfake Detection Using Wav2vec 2.0 and Data Augmentation,” in *Proc. Odyssey*, 2022, pp. 112–119.
- [17] Piotr Kawa, Marcin Plata, Michał Czuba, Piotr Szymański, and Piotr Syga, “Improved DeepFake Detection Using Whisper Features,” in *Proc. Interspeech*, 2023, pp. 4009–4013.
- [18] Junyi Peng, Oldřich Plhot, Themos Stafylakis, Ladislav Mošner, Lukáš Burget, and Jan Černocký, “An attention-based backend allowing efficient fine-tuning of transformer models for speaker verification,” in *Proc. SLT*, 2023, pp. 555–562.
- [19] Hye jin Shim, Hemlata Tak, Xuechen Liu, et al., “Baseline Systems for the First Spoofing-Aware Speaker Verification Challenge: Score and Embedding Fusion,” in *Proc. Odyssey*, 2022, pp. 330–337.
- [20] Xingming Wang, Xiaoyi Qin, Yikang Wang, Yunfei Xu, and Ming Li, “The DKU-OPPO System for the 2022 Spoofing-Aware Speaker Verification Challenge,” in *Proc. Interspeech*, 2022, pp. 4396–4400.
- [21] Chang Zeng, Lin Zhang, Meng Liu, and Junichi Yamagishi, “Spoofing-Aware Attention based ASV Backend with Multiple Enrollment Utterances and a Sampling Strategy for the SASV Challenge 2022,” in *Proc. Interspeech*, 2022, pp. 2883–2887.
- [22] Alexander Alenin, Nikita Torgashov, Anton Okhotnikov, Rostislav Makarov, and Ivan Yakovlev, “A Subnetwork Approach for Spoofing Aware Speaker Verification,” in *Proc. Interspeech*, 2022, pp. 2888–2892.
- [23] Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert, “MLS: A Large-Scale Multilingual Dataset for Speech Research,” in *Proc. Interspeech*, 2020, pp. 2757–2761.
- [24] Niko Brümmer and Johan du Preez, “Application-independent evaluation of speaker detection,” *Computer Speech & Language*, vol. 20, no. 2, pp. 230–275, 2006.
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016, pp. 770–778.

- [26] David Snyder, Guoguo Chen, and Daniel Povey, “MUSAN: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.
- [27] Xuechen Liu, Xin Wang, Md Sahidullah, Jose Patino, Héctor Delgado, Tomi Kinnunen, Massimiliano Todisco, Junichi Yamagishi, Nicholas Evans, Andreas Nautsch, et al., “ASVspoof 2021: Towards spoofed and deepfake speech detection in the wild,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 2507–2522, 2023.
- [28] Yichuan Mo and Shilin Wang, “Multi-task learning improves synthetic speech detection,” in *Proc. ICASSP*, 2022, pp. 6392–6396.
- [29] Gajan Suthokumar, Vidhyasaharan Sethu, Kaavya Sriskandaraja, and Eliathamby Ambikairajah, “Adversarial multi-task learning for speaker normalization in replay detection,” in *Proc. ICASSP*, 2020, pp. 6609–6613.
- [30] Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger, “UMAP: Uniform manifold approximation and projection,” *The Journal of Open Source Software*, vol. 3, no. 29, pp. 861, 2018.
- [31] Adrian Łańcucki, “Fastpitch: Parallel text-to-speech with pitch prediction,” in *Proc. ICASSP*, 2021, pp. 6588–6592.
- [32] Jaehyeon Kim, Jungil Kong, and Juhee Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *Proc. ICML*, 2021, pp. 5530–5540.
- [33] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerrv-Ryan, et al., “Natural tts synthesis by conditioning wavenet on mel spectrogram predictions,” in *Proc. ICASSP*, 2018, pp. 4779–4783.
- [34] Yinghao Aaron Li, Ali Zare, and Nima Mesgarani, “StarGANv2-VC: A Diverse, Unsupervised, Non-Parallel Framework for Natural-Sounding Voice Conversion,” in *Proc. Interspeech*, 2021, pp. 1349–1353.
- [35] Edresson Casanova, Julian Weber, Christopher D Shulby, Arnaldo Candido Junior, Eren Gölgel, and Moacir A Ponti, “YourTTS: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone,” in *Proc. ICML*, 2022, pp. 2709–2720.
- [36] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: a framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, 2020, p. 12.
- [37] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, and 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, pp. 1505–1518, Oct. 2022.
- [38] 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*, vol. 29, pp. 3451–3460, 2021.
- [39] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli, “Data2vec: A general framework for self-supervised learning in speech, vision and language,” in *Proc. ICML*, 2022, pp. 1298–1312.
- [40] Shu wen Yang, Po-Han Chi, and et. al., “SUPERB: Speech Processing Universal PERFORMANCE Benchmark,” in *Proc. Interspeech*, 2021, pp. 1194–1198.
- [41] Joon Son Chung, Arsha Nagrani, and Andrew Senior, “VoxCeleb2: Deep Speaker Recognition,” in *Proc. Interspeech*, 2018, pp. 1086–1090.
- [42] Tomi Kinnunen, Héctor Delgado, Nicholas Evans, Kong Aik Lee, Ville Vestman, Andreas Nautsch, Massimiliano Todisco, Xin Wang, Md Sahidullah, Junichi Yamagishi, et al., “Tandem assessment of spoofing countermeasures and automatic speaker verification: Fundamentals,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2195–2210, 2020.
- [43] Tomi H Kinnunen, Kong Aik Lee, Hemlata Tak, Nicholas Evans, and Andreas Nautsch, “t-EER: Parameter-free tandem evaluation of countermeasures and biometric comparators,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [44] Hongji Wang, Chengdong Liang, Shuai Wang, Zhengyang Chen, Binbin Zhang, Xu Xiang, Yanlei Deng, and Yanmin Qian, “Wespeaker: A research and production oriented speaker embedding learning toolkit,” in *Proc. ICASSP*. IEEE, 2023, pp. 1–5.
- [45] Massimiliano Todisco, Héctor Delgado, Kong Aik Lee, Md Sahidullah, Nicholas Evans, Tomi Kinnunen, and Junichi Yamagishi, “Integrated Presentation Attack Detection and Automatic Speaker Verification: Common Features and Gaussian Back-end Fusion,” in *Proc. Interspeech*, 2018, pp. 77–81.
- [46] Niko Brümmer and Edward de Villiers, “The BOSARIS Toolkit: Theory, algorithms and code for surviving the new DCF,” *arXiv preprint arXiv:1304.2865*, 2013.
- [47] Niko Brümmer, “LLR transformation for SRE’12,” *Agnitio Research*, 2012.
- [48] Yosef Solewicz, Noa Cohen, Johan Rohdin, Srikanth Madikeri, and Jan Honza Černecký, “Speaker Recognition on Mono-Channel Telephony Recordings,” in *Proc. Odyssey*, 2022, pp. 193–199.
- [49] Niko Brummer, Lukas Burget, Jan Cernocky, Ondrej Glembek, Frantisek Grezl, Martin Karafiat, David A. van Leeuwen, Pavel Matejka, Petr Schwarz, and Albert Strasheim, “Fusion of heterogeneous speaker recognition systems in the STBU submission for the NIST speaker recognition evaluation 2006,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 7, pp. 2072–2084, 2007.
- [50] Adam Paszke, Sam Gross, et al., “Pytorch: An imperative style, high-performance deep learning library,” in *Proc. NeurIPS*, 2019, pp. 8024–8035.
- [51] Dong C Liu and Jorge Nocedal, “On the limited memory bfgs method for large scale optimization,” *Mathematical programming*, vol. 45, no. 1, pp. 503–528, 1989.