C-P Map: A Novel Evaluation Toolkit for Speaker Verification

Lantian Li, Di Wang, Wenqiang Du, Dong Wang

Center for Speech and Language Technologies, Tsinghua University

{lilt,wangdi,duwq}@cslt.org, wangdong99@mails.tsinghua.edu.cn

Abstract

Evaluation trials are used to probe performance of automatic speaker verification (ASV) systems. In spite of the clear importance and impact, evaluation trials have not been seriously treated in research and engineering practice. This paper firstly presents a theoretical analysis on evaluation trials and highlights potential bias with the most popular cross-pairing approach used in trials design. To interpret and settle this problem, we define the concept of trial config and C-P map derived from it. The C-P map measures the performance of an ASV system on various trial configs in a 2-dimensional map. On the map, each location represents a particular trial config and its corresponding color represents the system performance. Experiments conducted on representative ASV systems show that the proposed C-P map offers a powerful evaluation toolkit for ASV performance analysis and comparison. The source code for C-P map has been release at https://gitlab.com/csltstu/sunine.

1. Introduction

Automatic speaker verification (ASV) aims to testify the claimed speaker identity of a speech signal [1, 2, 3, 4, 5]. After decades of research, current ASV systems have achieved fairly good performance, at least in benchmark evaluations. For example, with a bunch of state-of-the-art architectures/techniques, such as ResNet-based topology [6, 7, 8], self-attentive pooling [9], angular margin loss [10, 11] and score normalization & calibration [12, 13], researchers have reported equal error rates (EER) less than 1.0% on the VoxCeleb dataset, partly due to VoxSRC [14, 15], one of the most popular speaker verification challenges. Similar low EERs were also reported on SITW [16], another famous benchmark dataset.

In spite of the impressive EER results in benchmark evaluations, the genuine performance of modern ASV systems is dubious in real deployment applications. Practitioners often observe substantial performance gap between benchmark test and experience of end users, which we call *benchmark-deployment gap*.

To interpret and settle the benchmark-deployment gap, numerous studies have been carried out. Most of the research focuses on the *data theme*, which hypothesizes that the performance gap is largely attributed to acoustic mismatch. Based on this belief, various benchmark datasets were elaborately designed to simulate real-life acoustic conditions. For example, HI-MIA [17] was designed for near-far field mismatch, NIST SRE [18, 19] involved long-short mismatch and channel mismatch. Interestingly, even with the acoustic mismatch, researchers found that good performance can still be attained on benchmark tests. A possible reason is that these elaborately designed benchmarks, although involve some types of acoustic mismatch, yet reflect the true complexity of real-life conditions. Recently, a more challenging dataset CN-Celeb [20, 21] was released, with an intentional design to involve more real-life complexity, in particular multi-genre and cross-genre phenomena. Experiments on CN-Celeb demonstrated much higher EER results than on other benchmarks. Solving real-life data complexity remains a hot and challenging topic.

Besides the data complexity, there is another issue that seems equally important but less attended: the bias on evaluation trials. It is well known that trials are used to probe and measure performance of ASV systems. Therefore, if the trials are not appropriately designed, the performance of the target system cannot be correctly measured. Unfortunately, in most existing benchmark tests, the trials are designed simply based on cross-pairing. This trivial design leads to trials that might be quite different from those encountered in real-life scenarios. In particular, cross-pairing tends to produce a large proportion of *easy* trials, leading to over optimistic performance estimation. We argue that the *trial bias* problem is a primary source for the benchmark-deployment gap.

In this paper, we firstly conduct a theoretical analysis on evaluation trials and highlight the potential trial bias problem. Our analysis shows that an ASV system may exhibit very different performance when probed by different settings of trials, and therefore a prudent performance evaluation should involve multiple and diverse trial settings. We call each setting of trials a *trial config*, and the performance with various trial configs form a 2-dimensional map, named *Config-Performance (C-P) map*. We will show that C-P map is a powerful tool for system analysis, tuning and comparison. For instance, the benchmarkdeployment gap can be interpreted as the performance discrepancy between two different trial configs on the C-P map.

The rest of the paper is organized as follows. Section 2 presents a theoretical analysis on evaluation trials and discuss the trial bias problem, and Section 3 presents the concept of trial config and C-P map derived from it. Section 4 presents the experimental results with C-P map of several representative ASV systems, and Section 5 concludes the paper.

2. Trials and bias on trials

In speaker verification, the performance of a system is evaluated by a set of trials. Each trial is an individual test case, generally involving an enrollment prototype and a test speech. With modern embedding-based systems (e.g., i-vector [22] or xvector [23] systems), both the enrollment prototype and the test speech are represented by embedding vectors, and the verification task is formulated as testifying whether the two embedding vectors belong to the same speaker or not.

Trials are used to measure performance of ASV models and

This work was supported by the National Natural Science Foundation of China (NSFC) under Grants No.62171250, and also the Tsinghua-SPD Bank Joint Research Center for Digital Finance Technologies. Dong Wang is the corresponding author.

systems, and the role is like a prober. Obviously, if the prober is not well designed, the performance measurement cannot be reliable. Unfortunately, the importance of trials have not widely recognized, and the trivial *cross-pairing* approach remains the most popular for trial design. Specifically, suppose there are N speakers in the test set, each having K utterance, **full** crosspairing selects any two utterances (or the corresponding embedding vectors) as a trial and labels it as positive if the two utterances are from the same speaker or negative if they are from different speakers. This leads to NK(K-1) positive trials and $N(N-1)K^2$ negative trials. A variant of this full cross-pairing is **enrollment-fixed** cross-pairing, which fixes the enrollment speech, and selects only test speech by cross-pairing. This leads to NK positive trials and N(N-1)K negative trials, where K is the number of test utterances per speaker.

A notable problem of the cross-pairing approach is that it produces a large proportion of *easy trials*: trials that are easy to make decisions. This is particularly the case for negative trials. Unfortunately, the trials encountered in real-life applications might be quite different from those created by the crosspairing, and usually more challenging. For instance, real situations tend to involve users with similar accent and utterances recorded in similar acoustic environment, making the negative trials more challenging. In addition, the number of positive trials could be much larger than negative trials, so the probability of meeting easy trials is low. Finally, if a user finds a failure, he/she tends to try it over and over, leading to more hard trials. All the above issues suggest that there is a *trial bias* between the situations of benchmark test and real-life deployment. We argue that trial bias is an important source for the benchmarkdeployment gap mentioned in the previous section. Fig. 1 illustrates the phenomenon of trial bias.



Figure 1: Illustration of trial bias, where the axis represents the scores produced by an ASV system. The top figure (a) shows the scores of trials created by cross-pairing, and the bottom figure (b) shows the scores of trials encountered in a real application. The red star represents positive trials and the blue circle represents negative trials. Note that the class-conditional distributions of the scores are different in the two scenarios, which reflects the bias on trials.

It should be highlighted that from the perspective of performance evaluation, the bias is only related to the score distributions of the two trial classes (positive and negative). This means (1) only the *scores* matter, not any other properties of the test speech, such as acoustic environment, speaker traits, etc. (2) only the *distributions* matter, not any other quantities such as the number of speakers and the utterances per speaker. These two points are important and will be used to define the concept of trial config and C-P map.

3. Trial config and C-P map

3.1. Concept of C-P map

To gain a deeper understanding for the bias on trials, we introduce the concept of trial config. Given a set of enrollment/test utterances, a trial config is defined as a subset of trials selected to test against the target system. The full cross-pairing is the largest trial config and involves all the possible trials; the enrollment-fixed cross-pairing is another trial config where utterances are separated into an enrollment set and a test set, and only cross-set pairing is permitted. Clearly, for a particular ASV system, performance with different trial configs are different, reflecting performance of the target system under different deployment situations. By collecting all the trial configs and the corresponding performance, we can evaluate the target system in a more thorough way. This idea can be implemented as a config-performance (C-P) map. In this map, the x-axis corresponds to subsets of positive trials and the y-axis corresponds to subsets of negative trials, so each location (x, y) on the map corresponds to a particular trial config. Let the color at (x, y)represents the performance measurement, we obtain a C-P map.

A general C-P map is hard to read unless the neighboring trial configs form a spatial structure. A natural structure can be obtained by sorting the trial configs according to their 'hardness'. The hardness can be defined by human, by agreement of several system, or by a system itself. As an example, Fig. 2 shows a C-P map where the trial configs are sorted by the private hardness of an i-vector system. Specifically, we sort the trials according to the scores assigned by an ASV system, and then construct trial configs by selecting trials from the ordered list. For the target trials sets (x-axis), we gradually select trials with higher scores from left to right, and for the non-target trials sets (y-axis), we gradually select trials with lower scores from bottom to up. The color in the map represents the EER values corresponding to each trial config. This leads to an ordered C-P map, where the trial configs on the left and bottom are harder than those on the right and up.



Figure 2: Illustration of config-performance (C-P) map. The green star and the red start illustrates the trial configs applied in benchmark test and real-life applications respectively.

Although the hardness is defined privately by an ASV system, the C-P map still reveals lots of information. For instance, the large proportion of the high-performance area (right and up) suggests that there are many easy trials in the full cross-pairing trial config, and the low-performance area close to the origin point implies that the system performance is not perfect. Moreover, the benchmark-deployment gap we mentioned in previous sections can be clearly explained with the C-P map. First notice that the benchmark test and real-life deployment apply two different trial configs, one is created with cross-pairing design and the other is determined by the application scenario. These two configs are illustrated in Fig. 2 by the green star (right-up corner) and the red star (on the bottom left) respectively. The position of the red star reflects the fact that in real applications, negative trials could be more challenging, and positive trials tend to be more than negative trials in number. It is clear that the two trial configs lead to quite different EER results, which is precisely the benchmark-deployment gap.

If the order of the trial configs are fixed, the C-P map is more useful. For instance, it is possible to select several ASV baseline systems to evaluate each trial, and use the averaged score to sort the trials and construct the ordered trial configs. Using these ordered trial configs to draw C-P maps for different systems, we can get detailed comparison among systems. This offers a powerful tool that can be used to confirm effective techniques or identify spurious innovations.

3.2. Validity of C-P map

A critical question arises here is: are the performance measurements shown at different locations on the C-P map comparable, if they are based on different sets of trials? If the answer is NO, then the C-P map is not very useful, as the patterns shown in the map will be meaningless.

To answer this question, firstly notice that the performance metrics we consider in this study, i.e, EER and minDCF, are fully determined by the score distributions of the positive and negative trial classes. Take EER as an example, if we use p(c)and q(c) to denote the score distributions of positive and negative trials respectively, then it is easy to compute the EER threshold θ by solving the following equation:

$$\int_{-\infty}^{\theta} p(c) \mathrm{d}c = \int_{\theta}^{+\infty} q(c) \mathrm{d}c \tag{1}$$

where the left quantity represents the false rejection (FR) rate and the right quantity represents the false alarm (FA) rate, when the decision threshold is set to θ . According to the monotonic property of cumulative distribution functions (CDFs), the solution for θ is unique. Once θ is determined, the EER can computed as the left or right quantity of Eq. (1). Fig. 3 shows an example of the computation when the score distributions of both the positive and negative classes are Gaussian.



Figure 3: Illustration of EER computation when the score distributions of the negative and positive trials are both Gaussian. Suppose the variances of the two Gaussians are both 1.0, and the means are 0.0 and 3.0 respectively. According to the symmetry, it is easy to know the decision threshold is 1.5, and looking the CDF table of a normal Gaussian shows that EER is 6.68%.

This means that no matter which trials are used in the test set, only if they can approximate the score distributions of the positive and negative classes, the EER/minDCF results will approach to the true values. For instance, if one increases the number of test speakers and the number of test utterances per speaker, the number of negative trials will be increased much faster than the number of positive trials, but the EER/minDCF will be the same, supposing that the new speakers and utterances are sampled from the same underlying distribution as the original test set.

Since the evaluation metrics are determined by distributions of scores of trials rather than trials themselves, it is obvious to conclude that the performance values in the C-P map are comparable, only if every trial config involves sufficient trials. Note that the condition of the above statement means that on the C-P map, the origin point and its vicinity are not well defined and should be interpreted with caution.

4. Experiments

4.1. Data

VoxCeleb [24, 25] was used in our experiments. It is a largescale audio-visual speaker dataset collected by the University of Oxford, UK¹. Specifically, the development set of *VoxCeleb2* was used to train the i-vector and x-vector systems, which contains 5,994 speakers in total and entirely disjoints from the Vox-Celeb1 and SITW datasets. Trials of the cleaned *VoxCeleb1-O* was used for tuning and the cleaned *VoxCeleb1-E* was used for evaluation.

4.2. Basic systems

In our experiment, we firstly follow the VoxCeleb recipe of the Kaldi toolkit [26] to build our i-vector [22] and x-vector [23] baselines². These basic recipes may not achieve the best performance on a particular dataset, but have been demonstrated to be highly competitive and generalizable by many researchers with their own data and model settings. Moreover, using these recipes allow others to reproduce our results easily. No data augmentation is used.

- i-vector: The i-vector model was built following the Kaldi VoxCeleb/v1 recipe. The acoustic features comprise 23-dimensional MFCCs plus the log energy, augmented by the first- and second-order derivatives, resulting in a 72-dimensional feature vector. Moreover, cepstral mean normalization (CMN) is employed to normalize the channel effect, and an energy-based voice active detection (VAD) is used to remove silence segments. The UBM consists of 2,048 Gaussian components, and the dimensionality of the i-vector space is 400. For the back-end model, LDA is firstly used to reduce dimensionality to 200, and then PLDA [27] is employed to score the trials.
- x-vector: The x-vector model was created following the Kaldi VoxCeleb/v2 recipe. The acoustic features are 30-dimensional MFCCs. The DNN architecture adopts TDNN topology which involves 5 time-delay (TD) layers to learn frame-level deep speaker features, and a statistic pooling (TSP) layer is used to accumulate the frame-level features to utterance-level statistics, including the mean and standard deviation. After the pooling

¹http://www.robots.ox.ac.uk/~vgg/data/voxceleb/

²https://github.com/kaldi-asr/kaldi/tree/master/egs/voxceleb/

System	Front-End	Back-End	VoxCeleb1-O EER(%) minDCF		VoxCeleb1-E EER(%) minDCF	
1	GMM i-vector	PLDA	5.819	0.5189	5.872	0.5038
2	TDNN + TSP + Softmax	PLDA	4.558	0.4882	4.290	0.4343
3	TDNN + TSP + AM-Softmax	Cosine	3.430	0.3370	3.389	0.3619
4	ResNet34 + TSP + AM-Softmax	Cosine	1.633	0.1770	1.688	0.1900
5	ResNet34 + TSP + AAM-Softmax	Cosine	1.803	0.1961	1.747	0.1946
6	ResNet34 + ASP + AM-Softmax	Cosine	1.521	0.1642	1.504	0.1669

Table 1: EER(%) and minDCF results with the modern ASV systems on VoxCeleb1 evaluation trials.

layer, 2 fully-connection (FC) layers are used as the classifier, for which the outputs correspond to the number of speakers in the training set. Once trained, the 512dimensional activations of the penultimate layer are read out as an x-vector. The back-end model is the same as in the i-vector system.

4.3. More powerful systems

We also constructed more powerful x-vector systems to reproduce the performance of the SOTA speaker recognition techniques. A bunch of state-of-the-art architectures/techniques, including ResNet34 topology [6], attentive statistics pooling (ASP) strategy [28] and angular margin-based training objectives (AM-Softmax [10] and AAM-Softmax [11]) are employed in our experiments. Specifically, the acoustic features are 80-dimensional Fbanks. The DNN architecture adopts the ResNet34 topology for frame-level feature extraction. The TSP and ASP strategies are used to construct utterance-level representations. These representations are then transformed by a fully-connected layer to generate logits and are fed to a softmax layer to generate posterior probabilities over speakers. The training objective employs AM-Softmax and AAM-Softmax, and the margin factor is set to 0.2 and the scale factor is set to 30. Once trained, the 256-dimensional activations of the last fully-connection layer are read out as an x-vector. The simple cosine distance is used to score the trials. No data augmentation is used. The source code has been published online to help readers reproduce our systems³.

The EER and minDCF (p-target is 0.01) results of these ASV systems are presented in Table 1. It can be observed that modern ASV systems can obtain reasonable performance, even with the basic settings.

4.4. C-P map

In this section, we use the C-P map as a tool to analyze performance of the modern ASV systems on the cleaned *VoxCeleb1*-*E* dataset. In order to define a unified arrangement for the trial configs in the C-P map, the scores from the two basic systems (i-vector and x-vector) are averaged and used to sort the positive and negative trials respectively. The C-P maps are then constructed based on the sorted trials, following the process presented in Section 3.1.

Fig. 4 and Fig. 5 show the C-P maps of 6 systems corresponding to Table 1 with EER and minDCF metrics respectively. Firstly, it can be seen that all these systems can obtain satisfactory performance on the trial configs in the right-up corner. However, as the trial configs approach to the origin point in the left-bottom corner, the performance gap of different systems becomes significant. This indicates that these systems are good at recognizing easy trials but are different on tackling hard trials. Secondly, comparing these C-P maps, we observe that with the development of speaker recognition techniques (from System 1 to System 6), the performance improvement is steadily obtained. Among all these technique advances, neural topology change from TDNN to ResNet34 seems the most significant.

4.5. Delta C-P map

This section uses C-P map to compare two systems on various trial configs. To make the comparison more clear, the relative change ratio (RCR) at each location (x, y) on two C-P maps is computed as follows:

$$\operatorname{RCR}(x,y) = \frac{\operatorname{CP}_{ref}(x,y) - \operatorname{CP}_{test}(x,y)}{\operatorname{CP}_{ref}(x,y)}, \qquad (2)$$

where $CP_i(x, y)$ is the metric value (e.g., EER, minDCF) at the location (x, y) on the C-P map of system *i*. CP_{ref} represents the C-P map of the reference system and CP_{test} represents the C-P map of the test system.

Once RCR(x, y) is computed at all the locations, we can construct a delta C-P map, to compare two systems in an allround way. Obviously, for the trial config at location (x, y)in the delta C-P map, if its RCR is positive, it means the test system beats the reference system. In contrast, if RCR(x, y)is negative, the test system loses. To avoid numerical uncertainty, we define a small tolerance value ϵ (10⁻⁵ in this study), and regard the comparison is tied if $|RCR(x, y)| < \epsilon$. We can therefore compute the proportion of the trial configs on the delta C-P map with three situations (win: tie: lose), which provides a thorough comparison between two systems. Experiments are conducted to compare some representative speaker recognition techniques by delta C-P map, from the perspectives of model, topology, training objective and pooling strategy. All these results demonstrate that delta C-P map is a powerful tool to compare performance of systems in a thorough way.

4.5.1. i-vector vs. x-vector

Firstly, the delta C-P map between the basic i-vector system and x-vector system is presented, as shown in Fig. 6. The C-P map of the i-vector system (System 1) corresponds to $CP_{ref}(\cdot)$ and the C-P map of the x-vector system (System 2) corresponds to $CP_{test}(\cdot)$. It can be observed that x-vector wins the game on almost all the trial configs. This indicates that the discriminative

³https://gitlab.com/csltstu/sunine



Figure 4: The C-P maps of 6 systems tested on VoxCeleb-E trials with EER metric.



Figure 5: The C-P maps of 6 systems tested on VoxCeleb-E trials with minDCF metric.

model is superior to the probabilistic model at least under our experimental condition.

AM-Softmax overwhelmingly outperforms the standard Softmax, even though with a simpler back-end model.

4.5.2. Softmax vs. AM-Softmax

Secondly, we compared two training objectives Softmax and AM-Softmax by delta C-P map, as shown in Fig. 7. The C-P map of the Softmax system (System 2) corresponds to $CP_{ref}(\cdot)$ and the C-P map of the AM-Softmax system (System 3) corresponds to $CP_{test}(\cdot)$. It can be observed that the margin-based

4.5.3. TDNN vs. ResNet34

Thirdly, we compare two popular neural topologies, TDNN and ResNet34 by delta C-P map, as shown in Fig. 8. The C-P map of the TDNN system (System 3) corresponds to $CP_{ref}(\cdot)$ and the C-P map of the ResNet34 system (System 4) corresponds to $CP_{test}(\cdot)$. It is clear to see that the ResNet34 topology is



Figure 6: The delta C-P maps of x-vector against i-vector with EER/minDCF metrics.



Figure 7: The delta C-P maps of AM-Softmax against Softmax with EER/minDCF metrics.







Figure 9: The delta C-P maps of AAM-Softmax against AM-Softmax with *EER/minDCF* metrics.



Figure 10: The delta C-P maps of ASP against TSP with EER/minDCF metrics.



Figure 11: The roadmap of speaker recognition techniques measured by C-P map and delta C-P map.

superior to TDNN by a large margin, demonstrating the success of ResNet34 in speaker recognition.

4.5.4. AM-Softmax vs. AAM-Softmax

Fourthly, two margin-based training objectives AM-Softmax and AAM-Softmax are compared by delta C-P map, as shown in Fig. 9. The C-P map of the AM-Softmax (System 4) corresponds to $CP_{ref}(\cdot)$ and the C-P map of AAM-Softmax (System 5) corresponds to $CP_{test}(\cdot)$. It can be observed that the performance gap between the two objectives are quite marginal, even though AM-Softmax shows a bit advantage than AAM-Softmax on most trial configs.

4.5.5. TSP vs. ASP

Finally, two pooling strategies TSP and ASP are compared by delta C-P map, as shown in Fig. 10. The C-P map of the TSP system (System 4) corresponds to $CP_{ref}(\cdot)$ and the C-P map of the ASP system (System 6) corresponds to $CP_{test}(\cdot)$. It can be found that ASP outperforms TSP on the whole, demonstrating its advantage.

4.6. Roadmap

This section combines all these C-P maps and delta C-P maps in Section 4.4 and Section 4.5, and constructs a roadmap that illustrates the development of speaker recognition techniques in recent years. By this roadmap, it is clear to see which technique is effective and which innovation is revolutionary. This further demonstrates that the proposed C-P map is a very valuable tool for technique analysis and system comparison.

5. Conclusions

This paper is inspired by the ubiquitous benchmark-deployment discrepancy. Our hypothesis is that this problem is largely attributed to the different trials encountered in different situations. To formulate the significant impact of trials on performance measurement, we propose the concept of C-P map which represents performance of a speaker verification system on various trials configs in a 2-dimensional map. We empirically show that C-P map is a novel evaluation toolkit for ASV system analysis and comparison, and can be used to confirm really effective techniques. As for the future work, more comprehensive analysis will be conducted to understand the behavior of different models at different locations on the C-P map. Besides, other recently proposed techniques will be analyzed by C-P map to verify their effectiveness.

6. References

- Joseph P Campbell, "Speaker recognition: A tutorial," *Proceedings of the IEEE*, vol. 85, no. 9, pp. 1437–1462, 1997.
- [2] Douglas A Reynolds, "An overview of automatic speaker recognition technology," in 2002 IEEE international conference on acoustics, speech, and signal processing. IEEE, 2002, vol. 4, pp. IV–4072.
- [3] Tomi Kinnunen and Haizhou Li, "An overview of textindependent speaker recognition: From features to supervectors," *Speech communication*, vol. 52, no. 1, pp. 12– 40, 2010.
- [4] John HL Hansen and Taufiq Hasan, "Speaker recognition by machines and humans: A tutorial review," *IEEE Signal* processing magazine, vol. 32, no. 6, pp. 74–99, 2015.
- [5] Zhongxin Bai and Xiao-Lei Zhang, "Speaker recognition based on deep learning: An overview," *Neural Networks*, vol. 140, pp. 65–99, 2021.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [7] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck, "ECAPA-TDNN: Emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *INTERSPEECH*, 2020, pp. 3830–3834.
- [8] Tianyan Zhou, Yong Zhao, and Jian Wu, "ResNeXt and Res2Net structures for speaker verification," in 2021 IEEE Spoken Language Technology Workshop (SLT). IEEE, 2021, pp. 301–307.
- [9] Yingke Zhu, Tom Ko, David Snyder, Brian Mak, and Daniel Povey, "Self-attentive speaker embeddings for text-independent speaker verification.," in *INTER-SPEECH*, 2018, vol. 2018, pp. 3573–3577.
- [10] Feng Wang, Jian Cheng, Weiyang Liu, and Haijun Liu, "Additive margin softmax for face verification," *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.
- [11] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690–4699.
- [12] Patrick Kenny, "Bayesian speaker verification with heavytailed priors.," in *Odyssey*, 2010, vol. 14.
- [13] Pavel Matejka, Ondrej Novotný, Oldrich Plchot, Lukas Burget, Mireia Diez Sánchez, and Jan Cernocký, "Analysis of score normalization in multilingual speaker recognition.," in *INTERSPEECH*, 2017, pp. 1567–1571.
- [14] Joon Son Chung, Arsha Nagrani, Ernesto Coto, Weidi Xie, Mitchell McLaren, Douglas A Reynolds, and Andrew Zisserman, "VoxSRC 2019: The first Voxceleb speaker recognition challenge," arXiv preprint arXiv:1912.02522, 2019.
- [15] Arsha Nagrani, Joon Son Chung, Jaesung Huh, Andrew Brown, Ernesto Coto, Weidi Xie, Mitchell McLaren, Douglas A Reynolds, and Andrew Zisserman, "VoxSRC 2020: The second Voxceleb speaker recognition challenge," arXiv preprint arXiv:2012.06867, 2020.

- [16] Mitchell McLaren, Luciana Ferrer, Diego Castan, and Aaron Lawson, "The speakers in the wild (SITW) speaker recognition database.," in *INTERSPEECH*, 2016, pp. 818–822.
- [17] Xiaoyi Qin, Hui Bu, and Ming Li, "Hi-mia: A farfield text-dependent speaker verification database and the baselines," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing* (*ICASSP*). IEEE, 2020, pp. 7609–7613.
- [18] Alvin F Martin and Craig S Greenberg, "NIST 2008 speaker recognition evaluation: Performance across telephone and room microphone channels," in *Tenth Annual Conference of the International Speech Communication Association*, 2009.
- [19] Seyed Omid Sadjadi, Timothée Kheyrkhah, Audrey Tong, Craig S Greenberg, Douglas A Reynolds, Elliot Singer, Lisa P Mason, Jaime Hernandez-Cordero, et al., "The 2016 NIST speaker recognition evaluation.," in *INTER-SPEECH*, 2017, pp. 1353–1357.
- [20] Yue Fan, JW Kang, LT Li, KC Li, HL Chen, ST Cheng, PY Zhang, ZY Zhou, YQ Cai, and Dong Wang, "CN-Celeb: a challenging chinese speaker recognition dataset," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7604–7608.
- [21] Lantian Li, Ruiqi Liu, Jiawen Kang, Yue Fan, Hao Cui, Yunqi Cai, Ravichander Vipperla, Thomas Fang Zheng, and Dong Wang, "CN-Celeb: multi-genre speaker recognition," *Speech Communication*, 2022.
- [22] Najim Dehak, Patrick J Kenny, Réda Dehak, Pierre Dumouchel, and Pierre Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio*, *Speech, and Language Processing*, vol. 19, no. 4, pp. 788– 798, 2011.
- [23] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *IEEE International Conference* on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2018, pp. 5329–5333.
- [24] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, "Voxceleb: A large-scale speaker identification dataset," in *INTERSPEECH*, 2017, pp. 2616–2620.
- [25] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman, "Voxceleb2: Deep speaker recognition," in *INTER-SPEECH*, 2018, pp. 1086–1090.
- [26] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al., "The Kaldi speech recognition toolkit," in *IEEE 2011* workshop on automatic speech recognition and understanding. IEEE Signal Processing Society, 2011, number EPFL-CONF-192584.
- [27] Sergey Ioffe, "Probabilistic linear discriminant analysis," in *European Conference on Computer Vision (ECCV)*. Springer, 2006, pp. 531–542.
- [28] Koji Okabe, Takafumi Koshinaka, and Koichi Shinoda, "Attentive statistics pooling for deep speaker embedding," in *INTERSPEECH*, 2018, pp. 2252–2256.