# HLT-NUS Submission for 2019 NIST Multimedia Speaker Recognition Evaluation

Rohan Kumar Das, Ruijie Tao, Jichen Yang, Wei Rao, Cheng Yu and Haizhou Li Department of Electrical and Computer Engineering, National University of Singapore, Singapore E-mail: {rohankd, eleyji, haizhou.li}@nus.edu.sg, ruijie.tao@u.nus.edu

Abstract—This work describes the speaker verification system developed by Human Language Technology Laboratory, National University of Singapore (HLT-NUS) for 2019 NIST Multimedia Speaker Recognition Evaluation (SRE). The multimedia research has gained attention to a wide range of applications and speaker recognition is no exception to it. In contrast to the previous NIST SREs, the latest edition focuses on a multimedia track to recognize speakers with both audio and visual information. We developed separate systems for audio and visual inputs followed by a score level fusion of the systems from the two modalities to collectively use their information. The audio systems are based on x-vector based speaker embedding, whereas the face recognition systems are based on ResNet and InsightFace based face embeddings. With post evaluation studies and refinements, we obtain an equal error rate (EER) of 0.88% and an actual detection cost function (actDCF) of 0.026 on the evaluation set of 2019 NIST multimedia SRE corpus.

# I. INTRODUCTION

Speaker recognition is one of the widely studied areas in speech processing that has witnessed success over the last two decades [1], [2]. Such success has made possible to deploy the technologies from laboratory to real-world applications in the recent years [3]–[9]. The National Institutes of Standards and Technology (NIST) runs a speaker recognition evaluation (SRE<sup>1</sup>) to benchmark the robustness of various systems in adverse and realistic conditions [10]. Some of those adverse conditions included in the past few editions are short utterance, noisy environment, channel mismatch, language mismatch and multi-speaker test scenario [10]. The top performing teams in every editions show the scope towards addressing these issues by development of robust systems [11]–[14].

The latest edition 2019 NIST SRE runs on two tracks. The first track focuses on addressing the domain mismatch in terms of language, similar to that of the past two editions in 2016 and 2018 [15]. The only difference being removal of fixed set training condition and to observe the specific gain achieved by open set training condition. On the other hand, the second track sets out an entirely new direction to explore the multimedia speaker recognition for the first time as a part of SRE [16]. It aims to recognize a speaker's identify collectively with both audio and visual cues. This paper focuses on our participation to 2019 NIST multimedia SRE and the developed systems for the challenge with their associated results.

The research on processing multimedia for various applications have witnessed breakthrough in the recent years [17]. It has been successfully explored for applications such as speech separation [18], speaker diarization [19] and automatic speaker recognition [20]. While there are few works [21]–[24] exploring audio-visual multimedia information in the context of speaker recognition, they are not comparable to one another due to lack of benchmark data in the past. Further, the very recent explorations for benchmarking latest developments focus dominantly either on speaker recognition or face recognition, which is independent of multimedia processing. In this regard, the latest 2019 NIST multimedia SRE plays a pivotal role to have a common platform for evaluating performance of audiovisual speaker recognition systems.

Both speaker recognition and face recognition systems have evolved over the time to reach the current state-of-theart systems. In this work, we use x-vector based speaker embedding for speaker recognition studies [25]. A previous study [26] showed that mixed bandwidth systems help to achieve improved results for speaker recognition. Therefore, we resample the audio to build narrowband and wideband audio data based separate systems. Additionally, we focus on speech enhancement aspect as well to process the data from real-world scenario in the challenge. On the other hand, the visual system is based on ResNet-101 model for performing face recognition [27]. Post evaluation, we also developed more advanced InsightFace based system for face recognition studies [28]. Finally, a score level fusion is performed for the developed audio and visual systems to use audio-visual information for the 2019 NIST multimedia SRE collectively.

The rest of the paper is organized as follows. Section II describes the details of the audio systems developed for speaker recognition studies. In Section III, the visual systems developed for face recognition studies are described. The results and analysis for the submitted and post evaluation systems are reported in Section IV. Finally, the paper is concluded in Section V.

#### II. AUDIO SYSTEMS

This section describes the details of audio systems developed as a part of 2019 NIST multimedia SRE. The audio systems follow the x-vector based speaker modeling [25]. We focus on two different aspects while developing the audio systems. The first aspect deals with development of mixed

<sup>&</sup>lt;sup>1</sup>https://www.nist.gov/itl/iad/mig/speaker-recognition

bandwidth (narrowband and wideband) systems, whereas the second one focuses on applying speech enhancement to observe the impact on real-world data. We also consider a baseline audio system to compare to our audio systems. Next, we discuss the details of the developed audio systems for 2019 NIST multimedia SRE.

# A. Baseline

We use the standard Kaldi<sup>2</sup> system for having a baseline system for the audio track. The pre-trained x-vector model<sup>3</sup> is used to obtain a baseline system results. It is a system that uses wideband 16 kHz audio data for training the models and hence, we refer this system as x-vector<sub>16k</sub>-I in this work.

### B. Mixed Bandwidth Systems

Motivated by the gains achieved by mixed bandwidth systems in [26], we are interested in studying the impact of narrowband and wideband systems to use their different aspects. This also leverage us to use larger amount of corpora for data augmentation during the training of x-vector extractor than that is used for the baseline.

We develop two systems with narrowband and wideband audio data, respectively. Both systems are composed by three modules, namely, feature extraction, x-vector extraction and linear discriminant analysis (LDA) with probabilistic LDA (PLDA). The difference of these two systems is on the sampling rate of speech data. The former considers 8 kHz audio data, whereas the latter considers 16 kHz audio data. Therefore, we refer these two systems as x-vector<sub>8k</sub> and xvector<sub>16k</sub>-II, respectively. The details of experimental setups for the two systems are mentioned in the following.

1) Feature extraction: 23-dimensional mel frequency cepstral coefficient (MFCC) features are extracted from each utterance, which is followed by cepstral mean normalization with a window size of 3 seconds [29], [30]. An energy based voice activity detection (VAD) method is then used to remove the silence frames for training and development test data.

2) *x-vector extraction:* The setting of x-vector network architecture (e.g., the number of hidden units per layer) and the training process follows that in [31]. The summary of corpora used for training, development and data augmentation can be viewed from Table I. In addition, the specific details of the corpus used for training the extractor for narrowband and wideband systems are mentioned in Table II. We note that both narrowband and wideband systems follow the same network architecture to derive 512-dimensional x-vectors.

*3) LDA-PLDA:* After pre-processing on the input vectors, LDA is first applied to minimize the channel and distance variation and reduce the dimension of x-vectors to 150. Then, PLDA is utilized as the back-end classifier. The data used for LDA-PLDA training for both narrowband and wideband systems can be viewed from Table II.

 TABLE I

 Summary of speech corpora partitioned into training and development set.

Corpus	Usage
Previous NIST SREs	Training
Switchboard, VoxCeleb [32], [33]	
MUSAN Music [34] & Noise	Data Augmentation
RIRS Noise [35]	Data Augmentation
2019 NIST SRE Dev	Development

TABLE II			
DETAILS OF TRAINING SETS FOR X-VECTOR EXTRACTOR AND			
DA-PLDA MODEL IN NARROWBAND AND WIDEBAND SYSTEMS.			

System	Details			
	x-vector extractor			
	2,412,648 utterances from 11,577			
Narrowband	speakers in switchboard, previous SREs,			
	VoxCeleb1 and VoxCeleb2 corpora.			
Wideband	2,090,300 utterances from 7,185 speakers			
	in VoxCeleb1 and VoxCeleb2 corpora.			
LDA-PLDA				
Narrowband	176,081 utterances from 4,392 speakers			
	in previous SREs			
Wideband	200,000 utterances from 6,345 speakers			
	in VoxCeleb1 and VoxCeleb2 corpora			

# C. WPE System

I

The challenge data for 2019 NIST multimedia SRE are collected in uncontrolled environments and therefore, the audio segments are affected by background noise. Hence, we focused on speech enhancement aspect to improve the quality of input audio to the x-vector based system. The weighted prediction error (WPE) based dereverberation method is applied on the audio to perform speech enhancement [36]. It is a statistical method that uses delayed linear prediction model to find out the prediction error, which is then optimized to obtain the clean speech [36]. It is to be noted that we implement WPE only for the narrowband system for this study. The remaining system architecture follows the same as that of the x-vector<sub>8k</sub> system. We refer this WPE based system as x-vector<sub>8k</sub>-WPE in this work. Next, we discuss the visual systems developed for 2019 NIST multimedia SRE.

# **III. VISUAL SYSTEMS**

As discussed earlier, the 2019 NIST multimedia SRE involves both audio-visual cues for recognizing a person. Therefore, the enrollment video not only provides the target individual's voice segments, but also gives the target speaker's face information using the bounding boxes [16]. The visual system is expected to automatically determine whether the target person is present in another test video. In this work, we develop two different systems for face recognition that considers visual information to recognize the speakers. The first one is based on ResNet-101 architecture and serves as a baseline system for visual modality. On the other hand, the second one uses the latest InsightFace based face embedding framework, which we develop during post evaluation. Table III summarizes various corpora that we used for developing the

<sup>&</sup>lt;sup>2</sup>http://kaldi-asr.org/

<sup>&</sup>lt;sup>3</sup>http://kaldi-asr.org/models/m8



Fig. 1. Block diagram of the visual system for 2019 NIST multimedia SRE. In the enrollment videos, the baseline method crops the target faces based on the given face bounding boxes, whereas the post evaluation method detects faces in the specific frames.

TABLE III Summary of visual corpora partitioned into training and development set.

Corpus	Usage	System
MS-Celeb-1M & Asian-Celebrity [27]	Train Face Recognition	Baseline
WIDER FACE [37]	Train Face Detection	Post Evaluation
MS1MV2 [38]	Train Face Recognition	Post Evaluation
2019 NIST SRE Dev	Development	Both

visual systems. We discuss the details of the baseline and the post evaluation visual systems in the following subsections.

# A. Baseline: ResNet-101

In this subsection, we describe the baseline for the visual system that performs face recognition. It consists of a pipeline that receives raw images and produces the similarity scores between two segments. In the first stage, the specific faces in each enrollment video are cropped based on the given bounding boxes, then the frames in the whole test video at a rate of one frame per second are extracted. The face detection module takes the raw frames as input to detect the faces appeared. The detection contains a bounding box and five key points (eyes, nose and mouth), which is used to align the faces with similarity transformation. In the second stage, the aligned face images are converted to feature vectors with the recognition model. We use ResNet-101 trained on the cleaned version of MS-Celeb-1M database and Asian-Celebrity database as the recognition model [27]. In the third stage, we perform template matching operation to match faces between two segments. The pair-wise similarity scores are computed based on the extracted face embeddings, and the average of top-10 pairs is regarded as the final score. Fig. 1 shows the block diagram of the visual system with different stages involved to perform face recognition.

#### B. Post Evaluation: InsightFace

During post evaluation, we primarily focused on improving our visual system. The general pipeline of the system follows that of the baseline, but we used the latest methods at various stages for face recognition [39]. We highlight the key differences of the post evaluation visual systems in the following.

1) Face Detection: Face detection is performed with RetinaFace [40], which is the state-of-the-art framework in WIDER FACE hard test set [37]. The RetinaFace is a robust single-stage face detector, which added five facial landmarks as the extra supervision signal to improve the performance. We selected ResNet-50 model, which is trained on the WIDER FACE database for face detection. Followed by this, multi-task cascaded convolutional networks (MTCNN) [41] is applied to perform face alignment. In this study, for enrollment videos, we found the quality of some given bounding boxes are not so high due to the poses, lighting conditions and crop quality, hence, instead of using the given frames, 5 frames (given frames plus the  $\pm$  2 frames for each one) are selected as the specific frames to improve the performance. For these frames, we extracted and detected the faces, whose overlap rate with the given boxes is higher than the threshold.

2) Face Embedding: The face embeddings are extracted using InsightFace [28], which considers additive angular margin loss to obtain highly discriminative features for face recognition. This framework can find the maximum classification boundary in the angular space. We considered ResNet-100 model trained on the cleaned MS-Celeb-1M database (MS1MV2) [38] to obtain 512-dimensional face embeddings.

*3) Backends:* For each trial, the cosine similarity score between the extracted face embeddings in the enrollment video and the test video is computed. Instead of selecting the top-10 scores as that in the baseline method, the average of top 20% scores is taken here to get the final log-likelihood ratio (LLR) result using logistic regression.

TABLE IV					
PERFORMANCE OF AUDIO, VISUAL AND AUDIO-VISUAL (AV) SYSTEMS ON					
THE DEVELOPMENT SET OF 2019 NIST MULTIMEDIA SRE					

Modality	System	EER (%)	minDCF	actDCF
Audio	x-vector <sub>16k</sub> -I	12.96	0.525	0.548
	x-vector <sub>16k</sub> -II	12.38	0.463	0.838
	x-vector <sub>8k</sub>	12.04	0.467	0.632
	x-vector <sub>8k</sub> -WPE	11.42	0.462	0.638
	Fusion: Audio	12.04	0.459	0.466
Visual	ResNet-101	11.11	0.423	0.446
Challenge Submission				
Audio-visual	Fusion: AV	6.48	0.208	0.216
Post Evaluation				
Visual	InsightFace	5.65	0.318	0.342
Audio-visual	Fusion: AV	4.63	0.130	0.157

IV. RESULTS AND ANALYSIS

In this section, we report the results and analysis of different systems developed for 2019 NIST multimedia SRE. We used the scoring tool package given by the organizers to report the results in terms of equal error rate (EER), minimum detection cost function (minDCF) and actual detection cost function (actDCF) as per the evaluation plan [42]. It is noted that actDCF serves as the primary metric for the challenge. The score level fusion and calibration of various systems are performed using the Bosaris<sup>4</sup> toolkit [43] in this study.

Table IV shows the results for various audio and visual systems on the development set of 2019 NIST multimedia SRE. We note that scores belonging to the individual audio or visual systems are not calibrated. The scores are only calibrated for the fusion systems. It is observed that our mixed bandwidth and WPE based systems perform better than the baseline audio system. Further, the score level fusion of the audio systems improves the actDCF as well as minDCF compared to all the individual audio systems. This shows the gain achieved by complementary nature of information by these developed audio systems.

We then focus on the results of the visual systems. It is observed from Table IV that the visual system performs better than all the audio systems as well as fusion of all the audio systems, which shows its effectiveness for recognizing speakers. Further, we perform the fusion of all the four audio systems and the ResNet-101 based visual system, which was our primary submission to the audio-visual track of 2019 NIST multimedia SRE. We find the fused audio-visual system enhances the performance by a larger margin due to the complementary information for capturing speaker identity.

Post evaluation, we developed another visual system based on InsightFace framework, whose results are also reported in Table IV. This system performs even better than our baseline ResNet-101 based visual system that was used for audio-visual fusion during challenge submission. The gains achieved with this post evaluation system is due to the use of the stateof-the-art face detection as well as face recognition models. Further, the scoring strategy described in the previous section also helped to obtain an improved result. We then perform

TABLE V PERFORMANCE OF AUDIO, VISUAL AND AUDIO-VISUAL (AV) SYSTEMS ON THE EVALUATION SET OF 2019 NIST MULTIMEDIA SRE.

Modality	System	EER (%)	minDCF	actDCF
Audio	x-vector <sub>16k</sub> -I	7.79	0.407	0.610
	x-vector <sub>16k</sub> -II	7.30	0.317	0.558
	x-vector <sub>8k</sub>	8.14	0.414	0.548
	x-vector <sub>8k</sub> -WPE	7.02	0.306	0.529
	Fusion: Audio	6.46	0.314	0.321
Visual	ResNet-101	6.16	0.330	0.343
Challenge Submission				
Audio-visual	Fusion: AV	2.32	0.131	0.136
Post Evaluation				
Visual	InsightFace	1.55	0.049	0.085
Audio-visual	Fusion: AV	0.88	0.024	0.026

a fusion of this system with the four audio and ResNet-101 based visual systems. The resultant fused system obtained post evaluation significantly improves the performance of audio-visual speaker recognition.

Table V shows the results of various systems and their fusion on the evaluation set of 2019 NIST multimedia SRE. We observe that the performance of all the audio and visual systems are comparatively better than that on the development set. This points towards the fact that the development set is more challenging in nature. Further, the WPE based audio system emerges as the best system among the all audio system developed on the evaluation set. On the other hand, the InsightFace system outperforms the baseline visual system by a large margin. This also results in a significant gain in the audio-visual fusion system obtained with post evaluation study. We obtain an actDCF of 0.026 with the final audiovisual system in comparison to actDF of 0.136 audio-visual system submitted to the challenge. It is also noted that the final audio-visual system thus obtained performs comparable to the top two systems of 2019 NIST multimedia SRE.

#### V. CONCLUSION

This work reports the details of audio-visual systems developed by HLT-NUS for 2019 NIST multimedia SRE. The audio systems follow x-vector based framework with the focus on mixed bandwidth and speech enhancement. On the contrary, the visual systems are developed using ResNet-101 and InsightFace frameworks. A score level fusion of the developed audio and visual systems significantly improves the speaker recognition performance in comparison to systems based on individual modality. The future work will focus on associating temporal relation of audio-visual cues for performing multimedia speaker recognition.

#### VI. ACKNOWLEDGEMENT

This research work is supported by Programmatic Grant No. A1687b0033 from the Singapore Government's Research, Innovation and Enterprise 2020 plan (Advanced Manufacturing and Engineering domain), and Human-Robot Interaction Phase 1 (Grant No. 192 25 00054) from the National Research Foundation, Prime Minister's Office, Singapore under the National Robotics Programme.

<sup>&</sup>lt;sup>4</sup>https://sites.google.com/site/bosaristoolkit/

#### REFERENCES

- T. Kinnunen and H. Li, "An overview of text-independent speaker recognition: From features to supervectors," *Speech Communication*, vol. 52, pp. 12 – 40, 2010.
- [2] J. H. L. Hansen and T. Hasan, "Speaker recognition by machines and humans: A tutorial review," *IEEE Signal Processing Magazine*, vol. 32, no. 6, pp. 74–99, Nov 2015.
- [3] K. A. Lee, B. Ma, and H. Li, "Speaker verification makes its debut in smartphone," in *SLTC Newsletter*, February 2013.
- [4] K. A. Lee, A. Larcher, H. Thai, B. Ma, and H. Li, "Joint application of speech and speaker recognition for automation and security in smart home," in *Interspeech 2011*, 2011, pp. 3317–3318.
- [5] D. Chakrabarty, S. R. M. Prasanna, and R. K. Das, "Development and evaluation of online text-independent speaker verification system for remote person authentication," *International Journal of Speech Technology*, vol. 16, no. 1, pp. 75–88, 2013.
- [6] S. Dey, S. Barman, R. K. Bhukya, R. K. Das, Haris B. C., S. R. M. Prasanna, and R. Sinha, "Speech biometric based attendance system," in *National Conference on Communications*, 2014.
- [7] R. K. Das, S. Jelil, and S. R. M. Prasanna, "Development of multilevel speech based person authentication system," *Journal of Signal Processing Systems*, vol. 88, no. 3, pp. 259–271, 2017.
- [8] R. K. Das and S. R. M. Prasanna, "Investigating text-independent speaker verification from practically realizable system perspective," in *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC) 2018*, 2018, pp. 1483–1487.
- [9] S. Jelil, A. Shrivastava, R. K. Das, S. R. M. Prasanna, and R. Sinha, "SpeechMarker: A voice based multi-level attendance application," in *Interspeech 2019*, 2019, pp. 3665–3666.
- [10] C. S. Greenberg, L. P. Mason, S. O. Sadjadi, and D. A. Reynolds, "Two decades of speaker recognition evaluation at the National Institute of Standards and Technology," *Computer Speech & Language*, vol. 60, 2020.
- [11] J. Villalba, N. Chen, D. Snyder, D. Garcia-Romero, A. McCree, G. Sell, J. Borgstrom, F. Richardson, S. Shon, F. Grondin, R. Dehak, L. P. García-Perera, D. Povey, P. A. Torres-Carrasquillo, S. Khudanpur, and N. Dehak, "State-of-the-art speaker recognition for telephone and video speech: The JHU-MIT submission for NIST SRE18," in *Interspeech* 2019, 2019, pp. 1488–1492.
- [12] K. A. Lee, V. Hautamäki, T. H. Kinnunen, H. Yamamoto, K. Okabe, V. Vestman, J. Huang, G. Ding, H. Sun, A. Larcher, R. K. Das, H. Li, M. Rouvier, P.-M. Bousquet, W. Rao, Q. Wang, C. Zhang, F. Bahmaninezhad, H. Delgado, and M. Todisco, "14U submission to NIST SRE 2018: Leveraging from a decade of shared experiences," in *Interspeech 2019*, 2019, pp. 1497–1501.
- [13] P. A. Torres-Carrasquillo, F. Richardson, S. Nercessian, D. Sturim, W. Campbell, Y. Gwon, S. Vattam, N. Dehak, H. Mallidi, P. S. Nidadavolu, R. Li, and R. Dehak, "The MIT-LL, JHU and LRDE NIST 2016 speaker recognition evaluation system," in *Interspeech 2017*, 2017, pp. 1333–1337.
- [14] N. Kumar, R. K. Das, S. Jelil, D. B.K., H. Kashyap, K. S. R. Murty, S. Ganapathy, R. Sinha, and S. M. Prasanna, "IITG-Indigo system for nist 2016 SRE challenge," in *Interspeech 2017*, 2017, pp. 2859–2863.
- [15] S. O. Sadjadi, C. Greenberg, E. Singer, D. Reynolds, L. Mason, and J. Hernandez-Cordero, "The 2019 NIST speaker recognition evaluation CTS challenge," in *Speaker Odyssey 2020*, 2020, pp. 266–272.
  [16] S. O. Sadjadi, C. S. Greenberg, E. Singer, D. A. Reynolds, L. Mason, and
- [16] S. O. Sadjadi, C. S. Greenberg, E. Singer, D. A. Reynolds, L. Mason, and J. Hernandez-Cordero, "The 2019 NIST audio-visual speaker recognition evaluation," in *Speaker Odyssey 2020*, 2020, pp. 259–265.
- [17] T. Baltrušaitis, C. Ahuja, and L.-P. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 2, pp. 423–443, 2019.
- [18] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," ACM Transactions on Graphics, vol. 37, no. 4, pp. 112:1–112:11, 2018.
- [19] K. Hoover, S. Chaudhuri, C. Pantofaru, M. Slaney, and I. Sturdy, "Putting a face to the voice: Fusing audio and visual signals across a video to determine speakers," *CoRR*, vol. abs/1706.00079, 2017.
- [20] K. Noda, Y. Yamaguchi, K. Nakadai, H. G. Okuno, and T. Ogata, "Audio-visual speech recognition using deep learning," *Applied Intelligence*, vol. 42, no. 4, pp. 722–737, 2015.

- [21] B. Maison, C. Neti, and A. Senior, "Audio-visual speaker recognition for video broadcast news: some fusion techniques," in *IEEE Third Workshop* on Multimedia Signal Processing, 1999, pp. 161–167.
- [22] U. V. Chaudhari, G. N. Ramaswamy, G. Potamianos, and C. Neti, "Audio-visual speaker recognition using time-varying stream reliability prediction," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2003*, vol. 5, 2003, pp. V–712.
- [23] G. Chetty and M. Wagner, "Audio visual speaker verification based on hybrid fusion of cross modal features," in *Pattern Recognition and Machine Intelligence*, 2007, pp. 469–478.
- [24] E. D'Arca, N. M. Robertson, and J. R. Hopgood, "Robust indoor speaker recognition in a network of audio and video sensors," *Signal Processing*, vol. 129, pp. 137 – 149, 2016.
- [25] 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) 2018, 2018, pp. 5329–5333.
- [26] K. A. Lee, H. Yamamoto, K. Okabe, Q. Wang, L. Guo, T. Koshinaka, J. Zhang, and K. Shinoda, "NEC-TT system for mixed-bandwidth and multi-domain speaker recognition," *Computer Speech Language*, vol. 61, 2020.
- [27] J. Harvey, Adam. LaPlace. (2019) Megapixels: Origins, ethics, and privacy implications of publicly available face recognition image datasets. [Online]. Available: https://megapixels.cc/
- [28] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4690–4699.
- [29] S. Davis and P. Mermelstein, "Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 28, no. 4, pp. 357–366, Aug 1980.
- [30] S. Furui, "Cepstral analysis technique for automatic speaker verification," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 29, no. 2, pp. 254–272, Apr 1981.
- [31] D. Snyder, D. Garcia-Romero, G. Sell, A. McCree, D. Povey, and S. Khudanpur, "Speaker recognition for multi-speaker conversations using x-vectors," in *IEEE International Conference on Acoustics, Speech* and Signal Processing (ICASSP), 2019, pp. 5796–5800.
- [32] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: A large-scale speaker identification dataset," in *Interspeech 2017*, 2017, pp. 2616– 2620.
- [33] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Interspeech* 2018, 2018, pp. 1086–1090.
- [34] D. Snyder, G. Chen, and D. Povey, "MUSAN: A music, speech, and noise corpus," *CoRR*, vol. abs/1510.08484, 2015. [Online]. Available: http://arxiv.org/abs/1510.08484
- [35] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *IEEE International Conference on Acoustics, Speech* and Signal Processing (ICASSP) 2017, 2017, pp. 5220–5224.
- [36] T. Nakatani, T. Yoshioka, K. Kinoshita, M. Miyoshi, and B. Juang, "Speech dereverberation based on variance-normalized delayed linear prediction," *IEEE Transactions on Audio, Speech, and Language Pro*cessing, vol. 18, no. 7, pp. 1717–1731, Sep. 2010.
- [37] S. Yang, P. Luo, C.-C. Loy, and X. Tang, "WIDER FACE: A face detection benchmark," in *Proceedings of the IEEE conference on computer* vision and pattern recognition (CVPR), 2016, pp. 5525–5533.
- [38] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, "MS-Celeb-1M: A dataset and benchmark for large-scale face recognition," in *European Conference on Computer Vision (ECCV) 2016*, 2016, pp. 87–102.
- [39] R. Tao, R. K. Das, and H. Li, "Audio-visual speaker recognition with a cross-modal discriminative network," in *Interspeech 2020*, 2020.
- [40] J. Deng, J. Guo, Y. Zhou, J. Yu, I. Kotsia, and S. Zafeiriou, "Retinaface: Single-stage dense face localisation in the wild," *arXiv preprint* arXiv:1905.00641, 2019.
- [41] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.
- [42] "NIST 2019 Speaker Recognition Evaluation Plan", NIST, Aug 2019.
- [43] N. Brümmer and E. de Villiers, "The BOSARIS toolkit: Theory, algorithms and code for surviving the new DCF," *CoRR*, vol. abs/1304.2865, 2013.