



Automatic Quality Assessment for Audio-Visual Verification Systems. The *LOVe* submission to NIST SRE Challenge 2019

Grigory Antipov, Nicolas Gengembre, Olivier Le Blouch, Gaël Le Lan

Orange Labs, 4 rue du Clos Courtel, Cesson Sévigné, France

{grigory.antipov,nicolas.gengembre,olivier.leblouch,gael.lelan}@orange.com

Abstract

Fusion of scores is a cornerstone of multimodal biometric systems composed of independent unimodal parts. In this work, we focus on quality-dependent fusion for speaker-face verification. To this end, we propose a universal model which can be trained for automatic quality assessment of both face and speaker modalities. This model estimates the quality of representations produced by unimodal systems which are then used to enhance the score-level fusion of speaker and face verification modules. We demonstrate the improvements brought by this quality-dependent fusion on the recent NIST SRE19 Audio-Visual Challenge dataset.

1. Introduction

Accurate person recognition in videos is a holy grail of biometrics. Evaluation campaigns and datasets were organized to tackle this problem for applications such as indexation of TV content [1] or secure bimodal authentication [2]. That task had been rarely addressed for uncontrolled data, due to the lack of a challenging bimodal dataset, until the last Speaker Recognition Evaluation (SRE) organized by National Institute of Standards and Technology (NIST). The 2019's edition proposes an Audio-Visual challenge [3], for which participants are allowed to use both speech and image signals to detect persons of interest in social media format videos.

Multimodal biometric fusion can be performed at various levels: representation, score or decision [4]. Due to the variety of representations and classifiers depending on each biometrics, it is usually easier to fuse at the score level. Score-level fusion strategies usually involve beforehand normalization making scores to have similar statistics [5]. Rapidly, the problem of degraded input data for some biometric modalities led to integrate quality information into the fusion strategy [6], where fingerprint images qualities were annotated by human experts.

Kimura et al. [7] employed quality features for score-level fusion of face, gait and height biometrics, using spatial and temporal resolution as quality estimates. Logistic regression was used for fusion, and the authors used a Gaussian process regression with a nonlinear kernel function to predict optimal fusion weights from qualities. More generally, some previous works [8, 9] proposed quality estimates based on the hand-crafted *expert features* extracted from the input signal. On the contrary, in this work, we propose a trainable model to automatically estimate a quality based on a face or speaker feature representation. More precisely, our contributions are as follows: (1) we propose a universal model which is separately trained for quality estimation of speaker and face modalities, (2) we design 2 independent unimodal systems for speaker and face verification, respectively, and (3) we apply our quality assessment model to improve the fusion of the unimodal verification systems on the NIST SRE19 Audio-Visual challenge dataset.

The rest of the paper is organized as following. Section 2 is dedicated to related works. Section 3 details our quality assessment system and the corresponding quality-based fusion. Sections 4 and 5 present face and speaker recognition pipelines, respectively. Section 6 experimentally evaluates our systems, and Section 7 concludes the work.

2. Related Works

2.1. Face and Speaker Embeddings

Contemporary face and speaker verification systems have similar structure. Indeed, in both cases, Neural Networks (NNs) are trained to extract feature representations (commonly referred to as *embeddings*) of the respective input modalities. For face verification, a large panoply of approaches has been proposed using Convolutional Neural Networks (CNNs) [10, 11, 12, 13, 14], the current state-of-the-art (SOTA) being obtained by training with CosFace [13] and ArcFace [14] losses. Similarly, the SOTA embeddings for speaker recognition are also neural-based and known as *x-vectors* [15].

However, the main difference between face and speaker verification is in the way these embeddings are compared to compute the similarity between the respective face images or speech recordings. The two cases are detailed below.

At the dawn of the CNN-based face verification various methods have been explored to similarity computation, such as weighted χ^2 , Siamese networks [10] or Joint Bayesian [11]. As of today, the prevailing approach is a direct cosine in the face embeddings space. Such approach is indeed effective when the input face images allow extraction of discriminative embeddings, but the recent work of Shi and Jain [16] showed that it falls short when the face embeddings are extracted from the low-quality or occluded face images.

Contrary to face verification embeddings, the so-called *x-vectors* are almost never used without a Probabilistic Linear Discriminant Analysis [17] (PLDA) scoring backend to deal with intra- and inter-speakers variabilities. End-to-end approaches enabling cosine similarities of embeddings remain quite marginal, mostly because of language adaptation, but the issue is more and more investigated [18]. In this work, we experiment with a pipeline inspired by Snyder [19] and enriched by two complementary scoring backends: standard PLDA and direct cosine similarities.

2.2. Automatic Quality Assessment

While to the best of our knowledge no previous work proposed a dedicated system for automatic quality estimation of various modality representations, some studies are closely related to our main contribution. More precisely, the already mentioned work of Shi and Jain [16] showed that a pair of noisy face embeddings can be very close to each other in the embeddings space

while being extracted from faces of different persons which leads to false positive errors. Therefore, Shi and Jain argued that for face verification on challenging datasets, it is essential to estimate the *uncertainty* of the embedding and propose to do it via probabilistic modeling. The authors reported that their model’s convergence is unstable and highly dependent on the hyper-parameters, but their approach illustrates that quality estimation, which could be seen as a form of uncertainty, can be directly derived from the embeddings. In the similar spirit, a very recent work of Silnova et al. [20] successfully employed probabilistic uncertainty estimation for speaker diarization.

Otherwise, the aggregation models [21, 22] estimate relative weights for a set of input embeddings, to compute a weighted average. However, these weights cannot be considered as absolute embedding qualities, as they are estimated w.r.t. a given set of embeddings. In this work, we propose a simple model, inspired by the aggregation model of Xie and Zisserman [22], which estimates the quality of embeddings independently. We apply this model for face and speaker modalities.

3. Quality and Score-Level Fusion

3.1. Quality Assessment

When fusing scores obtained from multiple modalities, one expects that the fusion algorithm learns the strengths and weaknesses of the unimodal algorithms. For example, in an audio-visual verification task, such as the NIST SRE19 challenge, an ideal fusion approach would pay more attention to one of the two unimodal scores depending on the confidence of the audio and video systems (*e.g.* it should pay more attention to the audio score, if the video signal is too dark and indistinguishable, and the other way around if the audio is recorded in a noisy environment). Importantly, the unimodal scores themselves are often not sufficient to estimate the confidence of the respective predictions (indeed, as discussed in Section 2, a pair of blurry faces is often wrongly predicted as positive with a high verification score). Therefore, in this Section, we propose a separate module for estimating the confidence of the unimodal systems.

Our quality assessment module is presented in Figure 1. This module is a neural network designed on the top of the (priorly trained) speaker / face embeddings extractor. As depicted in Figure 1-(a), at inference time, the embeddings (denoted e_i) are the inputs of the quality assessment module which predicts a single real-valued q_i (*i.e.* the *quality*) in the range $[0., 1.]$ by applying a sigmoid.

During training (*cf.* Figure 1-(b)), the quality module is optimized to find a weighted average of a set of M embeddings ($M = 3$ in the figure) belonging to the same person. More precisely, the weights q_i (*i.e.* the qualities) of each embedding reflect to what extent the embedding participates in the linear combination:

$$\bar{e} = \frac{\sum_{i=1\dots M} q_i e_i}{\sum_{i=1\dots M} q_i} \quad (1)$$

The resulting averaged embedding \bar{e} is used for identity classification (the optimization is performed with the ArcFace loss [14]). During this training process, the quality assessment module learns to assign larger weights to embeddings that are easier to classify, and lower weights to hardly classifiable ones. Therefore, the weights learned this way implicitly represent the qualities of the embeddings.

3.2. Quality-Based Score-Level Fusion

The presented quality assessment module is used for the score-level fusion of audio and video scores in the following way. For a single embedding, each of the 2 modalities generates 3 scalars, namely: the verification score s , the quality of the enrollment embedding q_e , and of the test one q_t . Therefore, each *trial* produces 6 scalars, namely: the 2 scores s_{spk} and s_{face} and the corresponding quality estimates $\mathbf{q}_{spk} = (q_e, q_t)_{spk}$ and $\mathbf{q}_{face} = (q_e, q_t)_{face}$. These 6 values are fused via the quality-based score-level fusion with the Cllr-logistic regression [23, 24], as shown in the following equation:

$$LLR_{spk+face} = \sum_{i \in \{spk, face\}} a_i s_i + b_i q_{e_i} + c_i q_{t_i} + d \quad (2)$$

where a_i, b_i, c_i and d are the learned regression parameters.

Finally, it is important to notice, that the embedding qualities can be used not only for the inter-modality fusion as presented above, but also for fusion of scores inside a single modality. Indeed, as it is common for challenges and as presented in Sections 4 and 5, we propose several embedding extractors for each modality. Therefore, the intra-modality verification scores can also be fused taking into account the qualities of the respective embeddings (*cf.* Section 6).

4. Face Verification Pipeline

In this Section, we present the visual part of our audio-visual system based on the faces detected in the video frames. More precisely, video face verification is performed in 4 steps: Firstly, (1) the faces in video frames are detected and aligned. Then (2) they are passed through an embedding extraction CNN. After that, (3) the resulting face embeddings issued from various faces of a video should be combined (aggregated) taking into account the fact that they can potentially belong to different persons in the same video. Finally, (4) the similarity between a trial of videos is estimated by computing the distance between the respective aggregated embeddings.

4.1. Extraction of Face Embeddings

We use SOTA standard approaches for steps (1) and (2). Thus, face detection is performed at 1 frame per second using the well-known MTCNN model [25] for detection and DLIB [26] for alignment. For face embeddings extraction, we have trained 4 CNNs, namely ResNet-50 [27], PyramidNet [28], ArcFace-50 and ArcFace-100 [14] using CosFace and ArcFace losses. All 4 face embeddings extraction CNNs have been trained for the identity classification on a subset of the mixture of two public datasets, namely: MsCeleb1M [29] and MegaFace [30], containing about 7M photos of 100K classes (identities). PyramidNet (the 101-layers architecture has been chosen) and ResNet-50 have been trained with the CosFace loss, while ArcFace-50 and ArcFace-100 have been trained with the ArcFace loss.

4.2. Aggregation of Face Embeddings

Naturally, face embeddings extracted from different frames of a video are aggregated differently depending on whether the video is used for enrollment or test. Indeed, enrollment videos have several annotated *key* faces where the person to be enrolled is manually annotated with bounding boxes. Therefore, the aggregation model serves to construct a single aggregated key embedding from the embeddings extracted from key faces. For

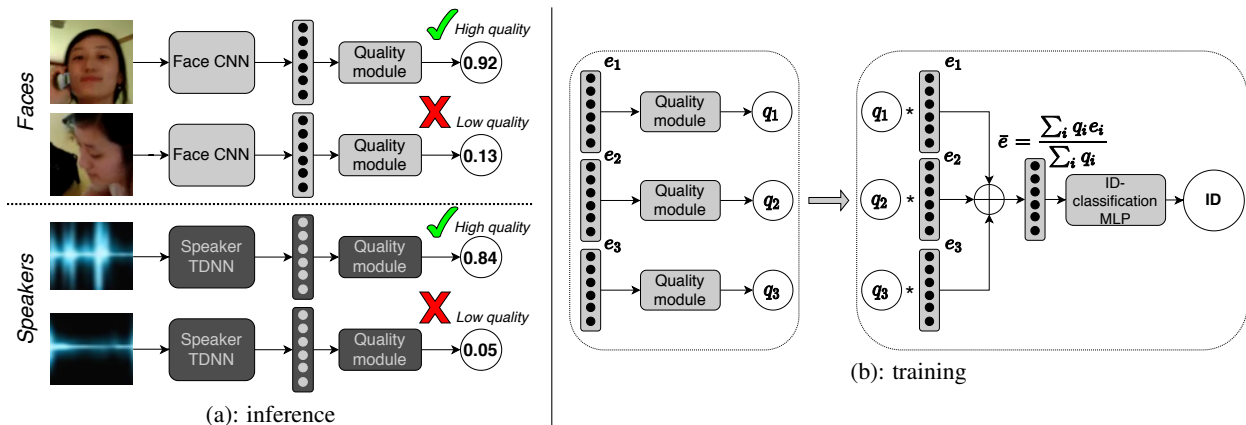


Figure 1: Proposed quality assessment modules for face and speaker embeddings. At inference time (a), the quality modules take the respective embeddings and output their real-valued qualities. A (face / speaker) quality module is trained (b) to optimize a weighted average embedding for further identity classification.

test videos, there is no prior on number of identities in a video. Thus, we firstly cluster (agglomerative hierarchical clustering algorithm is employed) embeddings extracted in a video, and then aggregate them inside each obtained cluster.

We do not simply average face embeddings to aggregate them, but we rather train a separate 2-layer Transformer-based model for this sake. This model learns how to downplay the influence of poor quality embeddings and how to pay more attention to discriminative embeddings in order to minimize the classification loss. The Transformer aggregation model is trained after the face embeddings extraction CNN (the latter being fixed during the training of the former) for the identity classification objective on the VGGFace2 dataset [31]. As a result, the usage of the Transformer-based aggregation instead of the trivial average aggregation improves the score of our face verification pipeline by more than 20% according to the challenge’s metric.

4.3. Calculation of Similarity between Face Embeddings

In order to compute a similarity between a trial of enrollment and test videos $sim(enrol, test)$ (which is the ultimate step of our face verification pipeline), we calculate the similarities between the aggregated key enrollment embedding $\bar{x}_{enrol}^{(K)}$ and the aggregated embeddings of each test cluster $\bar{x}_{test}^{(c_i)}$, and take the maximum: $sim(enrol, test) = \max_i \{sim(\bar{x}_{enrol}^{(K)}, \bar{x}_{test}^{(c_i)})\}$.

The most common way of calculating a similarity between a pair of face embeddings is to simply use a cosine similarity: $sim(\cdot, \cdot) \equiv \cos(\cdot, \cdot)$. Nevertheless, we have empirically observed that normalizing the cosine similarities with Adaptive-Symmetric NORMalization (*as-norm*) [32] allows to largely avoid the false positive verification errors (which are the most costly ones according to the challenges’ metrics). Therefore, we employ *as-norm* for score normalization in our final pipeline (the IJC-B dataset [33] containing about 3500 identities is used to calculate the reference cohort).

5. Speaker Verification Pipeline

This Section describes the global speaker verification pipeline. Firstly, (1) our two speaker identification systems are defined, followed by (2) the processing of test recordings by diarization and cluster selection, then (3) the experimental details.

5.1. Speaker identification

The two speaker identification systems in our pipeline are all based on a same *x-vector* topology. The first one respects the full Kaldi [34] process described in [19]. The second one is an end-to-end Pytorch [35] variant where the standard LLR is replaced by cosine similarity. This is achieved by replacing the top layers by an additional embedding layer, and refining the entire network with the ArcFace loss [14].

5.2. Diarization and cluster selection

For test recordings, potentially containing several different speakers, an automatic *diarization* step is mandatory to segment test recordings and extract speaker homogeneous clusters. That is where the procedure defined in [19] is applied to our process. On test recordings, embeddings are extracted from 1.5 second segments with a 0.75 second overlap. LLRs or cosine similarities are then computed between all pairs of embeddings. After an agglomerative hierarchical clustering, we take the union of each of the individual diarization results in a set of ways to partition a recording that has at most K speakers. Finally, the computation of the trial score, or *cluster selection* step, consists in taking the maximum of all scores obtained by an *enrollment embedding* versus all the *cluster embeddings* of a test recording (similarly to the face verification pipeline).

5.3. Experimental details

The speaker identification systems are trained on the VoxCeleb dataset [36] with the Kaldi framework and Pytorch. In order to increase the diversity of the acoustic conditions in the VoxCeleb dataset, a 8-fold augmentation strategy is used that adds seven corrupted copies of the original recordings to the training list. The recordings are corrupted by either digitally adding noise (i.e., babble, general noise, music), convolving with simulated room impulse response (RIR), high-pass and low-pass filtering, GSM-AMR and pitch/tempo shifting.

For speech parameterization, we extract 30-dimensional MFCCs (including c0) from 25 ms frames every 10 ms using a 30-channel mel-scale filterbank spanning the frequency range 20 Hz – 7600 Hz. Before dropping the non-speech frames using a proprietary neural phonetic-based speech activity detection, a short-time cepstral mean subtraction is applied over a 3-second

Table 1: Audiovisual fusion results on SRE19-DEV and SRE19-EVAL. The columns q_{spk} and q_{face} indicate whether the respective qualities are used for intra- / inter-modality fusion (cf. Section 3.2). “LR” = Logistic Regression (Equation 2). “EER” = Equal Error Rate. Row 4 reports our official Challenge primary result, based on an earlier version of the visual system, not presented in the paper.

#	Fusion Rule	Quality q_{spk}		Quality q_{face}		SRE19-DEV			SRE19-EVAL		
		intra	inter	intra	inter	EER	minC	actC	EER	minC	actC
1	Audio-only					8.33	0.220	0.246	2.65	0.138	0.148
2	Visual-only					8.33	0.179	0.186	2.21	0.066	0.069
3	Sum					3.99	0.094	0.122	0.66	0.031	0.043
4	LR (official)					3.96	0.087	0.101	0.66	0.039	0.042
5	LR					3.99	0.094	0.101	0.66	0.032	0.034
6	LR		✓			3.70	0.097	0.104	0.60	0.032	0.035
7	LR				✓	2.78	0.094	0.097	0.88	0.027	0.027
8	LR		✓		✓	2.78	0.094	0.095	0.66	0.025	0.029
9	LR	✓	✓	✓	✓	2.78	0.088	0.091	0.69	0.021	0.022

sliding window. Our experiments led us to a full speaker verification pipeline resulting in a fusion of six versions of a similar process. They differ from the speaker identification systems used for the *diarization* and *cluster selection* steps by exploring complementary backend configurations. For the sake of simplicity, these variants are not detailed in this paper. The final submission for our audio system is a fusion of these six variants scores computed by Cllr-logistic regression [23], with a target prior of 0.05, which parameters were trained on SRE19-DEV.

6. Fusion results

In this Section, we detail the score-level fusion results on both development and evaluation sets of challenge dataset (SRE19-DEV and SRE19-EVAL), for various fusion strategies, with or without quality estimates. Results are reported in Table 1 in terms of EER , min_C and act_C , as specified in [3]¹. The official SRE19 scoring toolkit was used to compute these metrics.

6.1. Challenge Metrics

A predefined detection cost function $C(\theta)$ (Equation 3), depending on the decision threshold θ , allows to compute two metrics: min_C and act_C (the latter is the calibrated one).

$$\begin{cases} C(\theta) = P_{FN}(\theta) + \frac{1-P_{tar}}{P_{tar}} \times P_{FP}(\theta) \\ min_C = \min_{\theta}(C) \\ act_C = C\left(\theta = \log\left(\frac{1-P_{tar}}{P_{tar}}\right)\right) \end{cases} \quad (3)$$

$P_{FP}(\theta)$ (resp. $P_{FN}(\theta)$) is the false positive (resp. false negative) probability given the threshold θ . The target prior P_{tar} is set by the organizers at 0.05. The act_C corresponds to the cost when θ is Bayes-optimal [37]. It is the reference metric of the challenge, which simulates performances of the system in operational conditions. This metric is much more sensitive to false positive errors than to false negative ones.

The NIST SRE19 Audio-Visual task is threshold-dependent. Therefore calibration is required. Before fusion, all scores are calibrated on SRE19-DEV (for which labels are available), using the *pybosaris* implementation² of Cllr-logistic regression [23] with a target prior of 0.05. As mentioned in Section 3.2, Cllr-logistic regression is also used for fusion (LR in table 1), with the same target prior.

¹According to the NIST SRE19 Challenge rules, we can only report our proper challenge results, but not the ones of other competitors.

²<https://gitlab.eurecom.fr/nausch/pybosaris>

6.2. Fusion: Score Only - baseline

The first 3 rows of Table 1 present the final speaker (row 1) and face (row 2) recognition performances and those obtained when applying the sum rule on speaker and face LLRs [38] (row 3). These first results show that the audio and visual systems are relatively well calibrated with a gap between min_C and act_C of 0.026 in the worst case (speaker system on SRE19-DEV). It can also be noted that the task of person verification is harder on SRE19-DEV than on SRE19-EVAL, since better performances are obtained on SRE19-EVAL overall. The visual system gives better performances than the audio one with a min_C of 0.179 vs. 0.220 on SRE19-DEV (resp. 0.066 vs. 0.138 on SRE19-EVAL). Applying a simple sum rule shows both modalities are complimentary since the min_C drops to 0.094 (resp. 0.031), while the EER drops from 8.33% to 3.99% (resp. from (2.65%, 2.21%) for both modalities to 0.66%). Replacing the sum rule by the Cllr-logistic regression (LR in the Table, row 5) does not change EER and min_C results but improves act_C performances, as this formula calibrates the fused LLRs.

6.3. Quality-based fusion

Rows 6-9 of Table 1 present quality-based fusion performances. Qualities are included in the Cllr-logistic regression, as explained in Section 3.2.

Results show that incorporating quality estimates into the logistic regression improves all evaluation metrics. Contrastive experiments are proposed (rows 6 & 7), where the quality of only one modality is provided. This shows that the proposed quality estimation works better for face embeddings than for speaker embeddings. However, incorporating both quality estimates is complementary (row 8). Finally, row 9 of Table 1 presents an experiment where quality estimates are also used for the fusion of within-modality system variants : this approach gives the best results overall.

7. Conclusion

We proposed a universal automatic quality estimation module integrated in a multimodal system for audio-visual person verification. We firstly presented our face and speaker verification pipelines. Then we demonstrated that our quality estimation module improves the fusion of the audio-visual scores both on the intra- and inter-modality levels. This was validated on the NIST SRE19 Audio-Visual Challenge.

8. References

- [1] O. Galibert and J. Kahn, "The first official repere evaluation," in *First Workshop on Speech, Language and Audio in Multimedia*, 2013.
- [2] K. Messer, J. Matas, J. Kittler, J. Luetin, and G. Maitre, "Xm2vtsdb: The extended m2vts database," in *Second international conference on audio and video-based biometric person authentication*, vol. 964, 1999, pp. 965–966.
- [3] 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 *Odyssey - The Speaker and Language Recognition Workshop*, 2020.
- [4] A. Ross and A. Jain, "Information fusion in biometrics," *Pattern recognition letters*, vol. 24, no. 13, pp. 2115–2125, 2003.
- [5] A. Jain, K. Nandakumar, and A. Ross, "Score normalization in multimodal biometric systems," *Pattern recognition*, vol. 38, no. 12, pp. 2270–2285, 2005.
- [6] J. Fierrez-Aguilar, J. Ortega-Garcia, J. Gonzalez-Rodriguez, and J. Bigun, "Discriminative multimodal biometric authentication based on quality measures," *Pattern recognition*, vol. 38, no. 5, pp. 777–779, 2005.
- [7] T. Kimura, Y. Makihara, D. Muramatsu, and Y. Yagi, "Quality-dependent score-level fusion of face, gait, and the height biometrics," *Information and Media Technologies*, vol. 9, no. 3, pp. 346–350, 2014.
- [8] L. Ferrer, L. Burget, O. Plchot, and N. Scheffer, "A unified approach for audio characterization and its application to speaker recognition," in *Odyssey - The Speaker and Language Recognition Workshop*, 2012.
- [9] N. Brummer, L. Burget, P. Kenny, P. Matejka, E. De Villiers, M. Karafiat, M. Kockmann, O. Glembek, O. Plchot, D. Baum *et al.*, "Abc system description for nist sre 2010," *Proc. NIST 2010 Speaker Recognition Evaluation*, pp. 1–20, 2010.
- [10] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *Proceedings of CVPR*, Columbus, USA, 2014.
- [11] Y. Sun, Y. Chen, X. Wang, and X. Tang, "Deep learning face representation by joint identification-verification," in *Proceedings of NeurIPS*, Montreal, Canada, 2014.
- [12] Y. Wen, K. Zhang, Z. Li, and Y. Qiao, "A discriminative feature learning approach for deep face recognition," in *Proceedings of ECCV*, Amsterdam, Netherlands, 2016.
- [13] F. Wang, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face verification," *IEEE Signal Processing Letters*, vol. 25, no. 7, pp. 926–930, 2018.
- [14] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of CVPR*, Long Beach, USA, 2019.
- [15] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [16] Y. Shi, A. K. Jain, and N. D. Kalka, "Probabilistic face embeddings," in *Proceedings of ICCV*, Seoul, Korea, 2019.
- [17] S. Ioffe, "Probabilistic linear discriminant analysis," in *European Conference on Computer Vision*. Springer, 2006, pp. 531–542.
- [18] G. Bhattacharya, J. Monteiro, J. Alam, and P. Kenny, "Generative adversarial speaker embedding networks for domain robust end-to-end speaker verification," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6226–6230.
- [19] D. Snyder, D. Garcia-Romero, G. Sell, A. McCree, D. Povey, and S. Khudanpur, "Speaker recognition for multi-speaker conversations using x-vectors," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5796–5800.
- [20] A. Silnova, N. Brümmer, J. Rohdin, T. Stafylakis, and L. Burget, "A unified approach for audio characterization and its application to speaker recognition," in *Odyssey - The Speaker and Language Recognition Workshop*, 2020.
- [21] J. Yang, P. Ren, D. Zhang, D. Chen, F. Wen, H. Li, and G. Hua, "Neural aggregation network for video face recognition," in *Proceedings of CVPR*, Honolulu, USA, 2017.
- [22] W. Xie and A. Zisserman, "Multicolumn networks for face recognition," in *Proceedings of BMVC*, Newcastle, UK, 2018.
- [23] N. Brummer, L. Burget, J. Cernocky, O. Glembek, F. Grezl, M. Karafiat, D. A. Van Leeuwen, P. Matejka, P. Schwarz, and A. 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.
- [24] J. Kittler, N. Poh, O. Fatukasi, K. Messer, K. Kryszczuk, J. Richiardi, and A. Drygajlo, "Quality dependent fusion of intramodal and multimodal biometric experts," in *Biometric Technology for Human Identification IV*, vol. 6539. International Society for Optics and Photonics, 2007, p. 653903.
- [25] 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.
- [26] D. E. King, "Dlib-ml: A machine learning toolkit," *Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of CVPR*, Las Vegas, USA, 2016.
- [28] D. Han, J. Kim, and J. Kim, "Deep pyramidal residual networks," in *Proceedings of CVPR*, Honolulu, USA, 2017.
- [29] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, "Ms-celeb-1m: A dataset and benchmark for large-scale face recognition," in *Proceedings of ECCV*, Amsterdam, Netherlands, 2016.
- [30] I. Kemelmacher-Shlizerman, S. M. Seitz, D. Miller, and E. Brossard, "The megaface benchmark: 1 million faces for recognition at scale," in *Proceedings of CVPR*, Las Vegas, USA, 2016.
- [31] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "Vg-gface2: A dataset for recognising faces across pose and age," in *Proceedings of FG*, Xi'an, China, 2018.
- [32] S. Cumani, P. D. Batsu, D. Colibro, C. Vair, P. Laface, and V. Vasilakakis, "Comparison of speaker recognition approaches for real applications," in *Proceedings of INTERSPEECH*, Florence, Italy, 2011.
- [33] B. Maze, J. Adams, J. A. Duncan, N. Kalka, T. Miller, C. Otto, A. K. Jain, W. T. Niggel, J. Anderson, J. Cheney *et al.*, "Iarpa janus benchmark-c: Face dataset and protocol," in *Proceedings of ICB*, Gold Coast, Australia, 2018.
- [34] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldi speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. CONF. IEEE Signal Processing Society, 2011.
- [35] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimeshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, 2019, pp. 8024–8035.
- [36] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: a large-scale speaker identification dataset," in *Proceedings of INTERSPEECH*, Stockholm, Sweden, 2017.
- [37] M. H. De Groot, "Optimal statistical decisions," 1970.
- [38] J. Kittler, M. Hatef, R. P. Duin, and J. Matas, "On combining classifiers," *IEEE transactions on pattern analysis and machine intelligence*, vol. 20, no. 3, pp. 226–239, 1998.