

Comparative Analysis of Modality Fusion Approaches for Audio-Visual Person Identification and Verification

Aref Farhadipour* and Masoumeh Chapariniya and Teodora Vukovic and Volker Dellwo

Department of Computational Linguistics, University of Zurich, Switzerland

masoumeh.chapariniya@uzh.ch

teodora.vukovic2@uzh.ch

volker.dellwo@uzh.ch

Abstract

Multimodal learning involves integrating information from various modalities to enhance learning and comprehension. We compare three modality fusion strategies in person identification and verification by processing two modalities: voice and face. In this paper, a one-dimensional convolutional neural network is employed for x-vector extraction from voice, while the pre-trained VGGFace2 network and transfer learning are utilized for face modality. In addition, gammatonegram is used as speech representation in engagement with the Darknet19 pre-trained network. The proposed systems are evaluated using the K-fold cross-validation technique on the 118 speakers of the test set of the VoxCeleb2 dataset. The comparative evaluations are done for single-modality and three proposed multimodal strategies in equal situations. Results demonstrate that the feature fusion strategy of gammatonegram and facial features achieves the highest performance, with an accuracy of 98.37% in the person identification task. However, concatenating facial features with the x-vector reaches 0.62% for EER in verification tasks.

1 Introduction

Biometric modalities encompass distinct static physiological traits that remain consistent within the human body, like fingerprints, as well as dynamic behavioral traits that are unique characteristics displayed in response to interactions with the environment, such as gait. However, some modalities can be a combination of both static and dynamic, such as speech (Minaee et al., 2023). In the real world, the human brain simultaneously processes multiple modalities to recognize the identity

of each person. Although the exact mechanism of multimodal processing in the human brain remains unclear, the human mind can effortlessly identify individuals based on their faces and voices with minimal errors (Perrodin et al., 2015). Voice and face have garnered significant attention in the development of automatic identity recognition systems (Minaee et al., 2023; Farhadipour and Taghipour, 2023; Farhadipour and Veisi, 2024a). While both modalities can change over time or be vulnerable to spoof attacks, combining and analyzing mixed forms of these two modalities can potentially increase the uniqueness of biometric features for each person in identity recognition. In human-machine interaction, identity recognition plays a crucial role. It has been applied in various tasks such as access control, automatic monitoring of older individuals, rehabilitation programs for people with physical and mental disabilities, etc. Automatic identity recognition can be categorized into two main tasks: person verification and person identification. Person verification involves authenticating a claimed identity. This typically involves a two-class classification, comparing the claimed identity with unique specific and universal background models. In other words, the system determines whether to accept or reject the claimed identity. On the other hand, person identification refers to the process of identifying an identity in a multi-class classification scenario, where the system needs to determine the person's identity from a large pool of individuals who have been previously trained. While there has been significant research on developing multimodal person verification systems (Shah et al., 2023), there is still limited work done in the context of multimodal person identification.

In recent years, deep learning has emerged

Correspondence: aref.farhadipour@uzh.ch

as an effective approach to pattern recognition, allowing for the modeling of complex functions. This approach enables multimodal signal processing and can be implemented through different tasks. In this work, we utilized deep learning approaches in identity recognition tasks with three different strategies. The first strategy is sensor fusion, where raw information, such as sound and image data, are directly integrated into a classifier. This approach combines the information from different sources at the input level without front-end processing. The second strategy is feature fusion, which involves separately extracting low-dimensional features from each modality and then combining them to feed into a classifier. This approach allows for the extraction of modality-specific features before integration. The third strategy is score fusion, where separate classifiers are designed for each modality, and the final scores from each classifier are merged and fed into a decision-making system.

In this work, it is assumed that the person's voice and face information are accessible simultaneously. We trained two separate systems to learn and recognize each modality individually for person identification and verification. Furthermore, three multimodal strategies have been proposed, including all three discussed modes: sensor-level fusion, feature-level fusion, and score-level fusion in identification tasks. Finally, the best fusion mode that is explored in the identification task is utilized in the verification scenario.

To achieve these objectives, the pre-trained VGGFace2 (Cao et al., 2018) is utilized for the visual modality. However, for the voice modality, we proposed two methods consisting of gammatonegram representation (Farhadipour and Veisi, 2024b) and x-vector (Snyder et al., 2018). This approach results in the creation of the proposed FaceNet and VoiceNet models for each modality from the VoxCeleb2 dataset.

In the feature and score fusion modes, we mixed facial features with gammatonegram and x-vector separately, and a softmax layer was employed for multimodal learning by combining the information from both modalities. However, in score fusion mode, because we chose a two-dimensional space for mixing

the modalities, we just concatenated facial features with gammatonegram representation. The proposed systems are trained and evaluated using the test section of the VoxCeleb2 dataset, which consists of 118 speakers. A K-fold cross-validation approach is utilized to ensure robustness and reliability.

The rest of the paper is organized as follows. Section 2 presents an overview of related works. In Section 3, the proposed strategies for multimodal learning in person identification and verification are described in detail. Section 4 focuses on the evaluation setup. Experimental results are reported in section 5. Discussing the results and comparison with previous works are done in section 6. Finally, section 7 concludes the work and discusses some ideas for future trends.

2 Related Works

Multimodal learning systems in identity identification integrated different modalities, including the fusion of fingerprint and DNA (Ramana et al., 2022), face and gait (Prakash et al., 2023), and face, palmprint and iris (Aldjia and Leila, 2021). In integrating face and speech, the main works focused on speaker verification (Wang et al., 2022), and few works are accomplished in speaker identification. It should be noted that the existence of the annual VoxSRC challenge had a great impact on this tendency (Huh et al., 2023). In this part, we report some of the previous works in multimodal speaker identification.

In the paper (Chung et al., 2018), an EER of 4.42% was reported on the Voxceleb1 dataset in a speaker verification task. The same authors in another article (Nagrani et al., 2020) reported 2.95% as EER on the Voxceleb1 dataset. Authors in paper (Moufidi et al., 2023) utilized a residual neural network to encode depth videos, while a time delay neural network architecture was used to encode voice signals. To evaluate the performance, 1,000 random speakers from the VoxCeleb2 dataset were selected. In that work, the accuracy of the three systems based on voice, depth of the lip regions, and overall fusion reached 56.03%, 17.49%, and 64.11%, respectively.

Alam et al., (Alam et al., 2015) introduced a confidence-based score fusion framework

for audio-visual biometric identification. They proposed the confidence-ratio approach. The researchers utilized a database consisting of 43 individuals. The results indicated that, in the best case with clean data, an accuracy of 90.45% was achieved. Paper (Asadpour et al., 2011) employed a combination of audio and video features by implementing a hidden Markov model.

The study (Shah et al., 2023) suggested a two-branch network to extract facial and voice signal features, using a support vector machine to classify speakers based on single and multi-domain features. Their method got an accuracy of 97.2 % for Speaker identification performance on VoxCeleb1. In (Sari et al., 2021), a multi-view model with a shared classifier to map audio and video into a unified space was introduced. The unimodal and audio-visual fusion approaches achieved an equal error rate of 1.6% on the VoxCeleb2 dataset in the person verification.

Qian et al., (Qian et al., 2021) introduced three types of audio-visual deep neural networks: feature level, embedding level, and embedding level combination with joint learning. The VoxCeleb2 test trial list experiments demonstrated equal error rates of 5.08% and 2.89% for visual and audio modality systems, respectively. Tao et al., (Tao et al., 2020) employed an audio-visual cross-modal discrimination network for speaker recognition. The best setting of the system gained an accuracy of 86.12% on VoxCeleb2 for speaker identification in score-level fusion.

The paper (Stefanidi et al., 2020) introduced an approach to person identification using CNNs on the VoxCeleb1 audiovisual database. The outcomes achieved an 86.97% top-5 accuracy. Cai et al., (Cai et al., 2022) presented a self-supervised learning framework for speaker recognition, which combined clustering and deep representation learning. With multi-modal training data, their framework obtained an accuracy of 77.60% on the VoxCeleb2.

3 Multimodal Learning Strategies

This work uses two single-modality systems to analyze each modality before exploring multimodal learning strategies. For the first

system, VoiceNet, we examined two methods. The first one is the x-vector, which is a one-dimensional CNN with five convolutional layers, a statistical pooling layer (Wang et al., 2021), and three fully connected layers that focus on identity identification from voice modality. The architecture of this neural network is depicted in Figure 1. The second one is the gammatonegram, which represents the utterances as an image and fine-tunes the Darknet19 (Redmon and Farhadi, 2017) for learning the features of each speaker and classification task.

As shown in figure 1, the person's speech is extracted from the video file, and silence is removed using a Voice Activity Detector (VAD) (Giannakopoulos, 2009). The audio files are then transformed into MFCC (Mel-frequency cepstral coefficients) features, and the mismatch between training and test utterances is reduced using the CMVN (Cepstral Mean and Variance Normalization) technique (Prasad and Umesh, 2013) and serves as input for VoiceNet. We made another VoiceNet using gammatonegram separately. The final layer of VoiceNet consists of 118 neurons, representing the number of speakers. The outputs of VoiceNets present the performance of the person identification systems based on voice modality.

In addition to VoiceNet, the FaceNet is proposed for face recognition, as shown in Figure 1. FaceNet is built using a pre-trained CNN known as VGGFace2, which has been trained on more than 3.3 million face images (Cao et al., 2018). The transfer learning technique is employed, where the final layers of the network are replaced with new layers to learn the specific information of the new individuals.

The output of FaceNet represents the result of person identification using the face modality. To ensure comparability across different multimodal learning scenarios, an overall audio file and a single frame of the face are extracted from each video file. Since each video contains only one person, the 25th frame is extracted as the face modality.

Figure 1 illustrates the three proposed strategies for multimodal learning: sensor fusion, feature fusion, and score fusion. This study

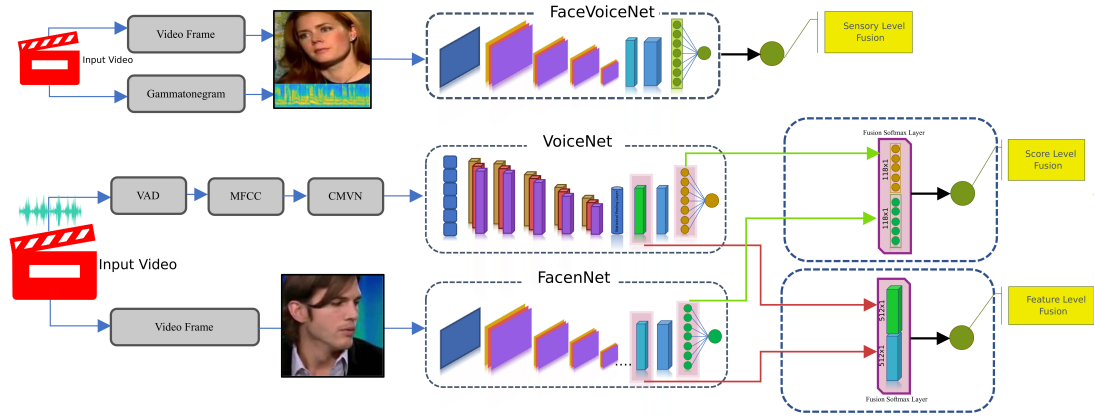


Figure 1: The block diagram of the proposed single-modality systems consists of VoiceNet, FaceNet, and three multimodal systems with different fusion strategies

uses FaceNet and VoiceNets as the foundation for feature and score fusion systems. However, We employ a separate network with mixed modalities as input for the sensor fusion mode. In the subsequent sections, we comprehensively explain each multimodal learning strategy.

3.1 Sensor Level Fusion

Since images are two-dimensional and sounds are one-dimensional signals in this multimodal learning strategy, we aim to combine them into a shared data space. We chose the gammatonegram visualization method (Farhadipour et al., 2014) to represent audio files. Sensor-level multimodal learning is performed on the VGGFace2 network and using transfer learning. As depicted in the sensor level fusion system in Figure 1, we create an integrated image to serve as input for the FaceVoiceNet. This input image incorporates both facial and voice information. The reasons for choosing voice as the modality to be mapped to a new space were influenced by two factors: the availability of a powerful pre-trained network for face recognition and the common practice of representing voice in the form of an image using the gammatonegram.

3.2 Feature Level Fusion

In the feature fusion mode, as depicted in Figure 1, a softmax layer is employed for multimodal learning. In this mode, VoiceNet is utilized to extract the x-vector, resulting in a feature vector of speakers with dimensions

of 512×1 . Similarly, from the FaceNet, the activation values of the last ReLU layer are extracted as the facial feature embedding, which is a 512×1 vector. These two vectors are concatenated and construct a bimodal vector with a size of 1024×1 . In another scenario, we replace gammatonegram and Darknet with the x-vector. In this situation, we extract a 118×1 vector from a middle layer of Darknet19 as a feature vector to concatenate with facial features and create a 630×1 dimension multimodal feature vector.

3.3 Score Level Fusion

The score fusion strategy aims to reach a function that accurately predicts the correct speaker ID by using the score vectors from the basis networks. In other words, our proposed system tries to estimate the f in eq. (1) to make D , which is the ground truth in the training process and the final decision in the testing phase.

$$D = f(x, y) \quad (1)$$

In this equation, x and y are the score vectors of VoiceNets and FaceNet, respectively. We used the softmax layer for modeling f based on its high modeling ability. Each input vector has 118 components, resulting in a bimodal vector with a dimension of 236×1 .

4 Evaluation Setup

In this work, the test part of the VoxCeleb2 dataset is used for evaluation. This part consists of 118 speakers and is extracted from YouTube videos in real-world conditions, including various types of noise such as laughter,

cross-talk, channel effects, music, and some other environmental sounds. We used this subset just to compare the performance of different strategies proposed in this work in a common situation. The test section of Vox-Celeb2 comprises a total of 4,911 unique video files and 36,237 utterances extracted from these videos (Chung et al., 2018). To evaluate the performance of the systems, the K-fold cross-validation technique is employed, with a value of $K=3$ based on a traditional strategy to use 70% of data for training and the rest for testing. This means that the evaluation is repeated three times, and in each iteration, two folds of the dataset are used for training, while the remaining fold is employed for testing. Under these conditions, approximately 24,000 bimodal utterances are used for training, while around 12,000 bimodal utterances are used for testing.

VoiceNet was trained in 8 epochs, with 128 batch sizes, and the learning rate was adjusted from $1e-3$ to $1e-5$. On the other hand, the VGGFace2 is trained in 15 epochs with batch size 32 and a constant learning rate of $1e-3$. For fine-tuning the Darknet19, we set the learning rate to $1e-4$ with 15 epochs. Cross-entropy was used as a loss function in the training process of CNNs and softmax layers. Moreover, stochastic gradient descent with momentum was utilized as the optimizer. The training process is stopped for multimodal learning using the softmax layer when the minimum gradient criteria are satisfied. In this work, it occurred around the 40th to 50th epoch. Two data augmentation techniques were used to avoid overfitting and making a general system: rotation operators within the range of $[-20, 20]$ degrees and vertical and horizontal translation by a distance of $[5, 5]$ pixels.

To evaluate proposed identification systems, there exist several metrics that can be utilized to evaluate the efficiency of a multi-class pattern recognition system. The system's effectiveness can be demonstrated from various perspectives by carefully selecting the appropriate metrics that enable comparisons with other works. The proposed evaluation parameters encompass precision, specificity, sensitivity, accuracy, and F1 score. The utiliza-

tion of a confusion matrix allows for a visual representation of these parameters (Lever, 2016). Verification systems have encountered two types of errors: false acceptance and false rejection. False acceptance refers to the claims that were accepted inaccurately. However, false rejection deals with the identity that is rejected incorrectly. Equational Error Rate (EER) is the optimum point at which these two errors are equal.

5 Experimental Results

Our main goal is to compare different strategies in modality fusion. In this part, we report experimental results in two scenarios. Person identification results are reported to compare the performance of different fusion modes in multi-class classification scenarios in equal situations. Based on the results of the identification task, the best fusion mode is evaluated in the verification scenario as a two-class classification besides single modality modes.

5.1 Person Identification

According to the results in Table 1, the performance of the proposed person identification systems in the single modality can be observed in face identification and two other speaker identification separated based on the x-vector and gammatonegram representation methodologies. The table includes the achievements for each fold separately and the average performance. Based on the results, it is evident that the system performs better in face single modality recognition than voice. The accuracy achieved by the FaceNet is 96.00%, while the VoiceNet reached an accuracy of 72.67% using x-vector and an accuracy of 61.64% based on gammatonegram. It seems that the x-vector could represent the speaker feature better than the gammatonegram in single modality mode.

Other parameters that provide more insight into the systems' performance can also be found in Table 1. These parameters present the ability of systems to accept correct utterances and reject incorrect utterances for each class. It is worth noting that the presence of babble noise and low sound quality significantly impact the performance of VoiceNets, resulting in its lower accuracy compared to FaceNet.

Metrics	Face Identification				x-Vector Speaker Identification				Gamma. Speaker Identification			
	Fold 1	Fold 2	Fold 3	Avg.	Fold 1	Fold 2	Fold 3	Avg.	Fold 1	Fold 2	Fold 3	Avg.
Precision	0.97	0.97	0.95	0.96	0.68	0.77	0.74	0.73	0.60	0.64	0.60	0.61
Sensitivity	0.97	0.97	0.95	0.96	0.68	0.77	0.74	0.73	0.60	0.64	0.60	0.61
Specificity	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
F-measure	0.97	0.97	0.95	0.96	0.68	0.77	0.74	0.73	0.60	0.64	0.60	0.61
Accuracy(%)	96.54	96.47	94.99	96.00	67.59	76.96	73.47	72.67	60.45	64.47	60.00	61.64

Table 1: Results of single-modality face identification and speaker identification systems with three different feature sets for three folds and presentation of average values

Metrics	Sensor Fusion				Score Fusion				Feature Fusion			
	Fold 1	Fold 2	Fold 3	Avg.	Fold 1	Fold 2	Fold 3	Avg.	Fold 1	Fold 2	Fold 3	Avg.
Precision	0.95	0.95	0.91	0.94	0.97	0.97	0.95	0.96	0.99	0.99	0.98	0.99
Sensitivity	0.95	0.95	0.91	0.94	0.97	0.97	0.95	0.96	0.99	0.99	0.98	0.99
Specificity	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
F-measure	0.95	0.95	0.91	0.94	0.97	0.97	0.95	0.96	0.99	0.99	0.98	0.99
Accuracy(%)	95.02	95.10	90.70	93.61	96.73	96.81	95.18	96.24	98.59	98.94	97.59	98.37

Table 2: Performance of proposed multimodal identification systems in three different fusion strategies based on gammatonegram and facial features

Metrics	Score Fusion				Feature Fusion			
	Fold 1	Fold 2	Fold 3	Avg.	Fold 1	Fold 2	Fold 3	Avg.
Precision	0.97	0.97	0.96	0.97	0.99	0.99	0.97	0.98
Sensitivity	0.97	0.97	0.96	0.97	0.99	0.99	0.97	0.98
Specificity	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
F-measure	0.97	0.97	0.96	0.97	0.99	0.99	0.97	0.98
Accuracy(%)	96.88	97.27	95.94	96.70	98.81	98.88	97.28	98.33

Table 3: Results of proposed multimodal identification systems in two different fusion strategies based on x-Vector and facial features

In addition, it can be seen that the modalities can complement each other while being independent of each other. Table 2 presents the results of three multimodal strategies, each based on different fusion methods. In this table, gammatonegram is used for voice and FaceNet features for face modality. The first part of the table illustrates the performance of the proposed multimodal learning system in sensor fusion mode. Despite achieving precision, sensitivity, specificity, F-measure, and accuracy percentage of 0.94, 0.94, 1, 0.94, and 93.61%, respectively, the multimodal system in sensor fusion mode underperforms compared to the FaceNet single modality system.

This suggests that gammatonegram as a speech presentation method can potentially confuse the sensor-level multimodal system, leading to a decrease in performance in comparison with the single modality system. Although the multimodal learning system with input-level fusion of image and speech data does not improve efficiency, it provides valuable insights to researchers. The findings indicate that the fusion of sensors for these two modalities may not significantly enhance performance, underscoring the need for careful consideration of fusion strategies and modality compatibility in multimodal learning systems.

Based on the investigation in Table 2, the

multimodal system in score fusion mode has been evaluated using the information from previous single-modality systems based on facial features and gammatonegram representation. Each network’s final softmax layer data is used as input for the aggregated softmax layer in multimodal learning. The average performance across three different folds shows that the system achieves better precision, sensitivity, specificity, F-measure, and accuracy scores than the single modality mode. Specifically, the system achieves 96.24% accuracy, which is 0.24% higher than the result obtained by the FaceNet single-modality system.

Table 2 reveals that the proposed softmax layer in feature fusion mode demonstrates high efficiency, achieving 98.37% accuracy, 0.99 precision, 0.99 sensitivity, 1 specificity, and 0.99 F-measure. These results signify a significant improvement in the accuracy of the person identification task, with a 2.37% increase compared to the performance of FaceNet in single modality mode. This underscores the effectiveness of the feature fusion strategy in enhancing the system’s accuracy. The additional source data from different modalities provide extra information for identification, as they express different aspects of the same class. For instance, in fold 1, the speaker identification has the lowest

Metrics	Face Verification				Speaker Verification			
	Fold 1	Fold 2	Fold 3	Avg.	Fold 1	Fold 2	Fold 3	Avg.
EER(%)	0.48	0.50	2.07	1.01	5.99	5.37	3.65	5.12

Table 4: Results of proposed person verification systems in two single modes

Metrics	Gammatonegram and FaceNet				x-Vector and FaceNet			
	Fold 1	Fold 2	Fold 3	Avg.	Fold 1	Fold 2	Fold 3	Avg.
EER(%)	0.44	0.41	1.61	0.82	0.37	0.36	1.15	0.62

Table 5: Results of proposed person verification systems in feature fusion modes based on two different speech features

performance compared to the other two folds, but the facial recognition for this fold is done with high accuracy. This demonstrates that modality fusion can compensate for the shortcomings of single modalities.

Based on Table 3, the proposed softmax layer for facial features and x-vector in score fusion mode reaches 96.7% accuracy on average, which is better than the single modality mode. However, the feature fusion mode showcases remarkable effectiveness, achieving an accuracy of 98.33%, precision and sensitivity scores of 0.98 each, a specificity of 1, and an F-measure of 0.98. These findings underscore the successful enhancement in accuracy for the person identification task by 2.33% compared to FaceNet as the best achievement in single modality mode.

The average achievements of all five modes show that FaceNet performed the best in single-modality scenarios. Despite VoiceNet’s lower efficiency, when combined with the face embedding feature using gammatonegram and x-vector, it improved the identification system’s performance in two scenarios: feature fusion and score fusion. Our best achievement was in the feature fusion mode, where we combined facial features with the gammatonegram.

5.2 Person Verification

In this section, we investigate the results of the proposed system for the person verification task. Based on the results obtained in person identification, the feature-level fusion has shown the best performance. Therefore, in the verification scenario, we only evaluate the fusion of features consisting of gammatonegram, x-vector, and FaceNet’s activations in multimodal mode, besides two typical single modality-based person verification. In this

feature fusion mode, the proposed person verification system utilizes Within-Class Covariance Normalization (WCCN) and Linear Discriminant Analysis (LDA) to reduce intra-class variation and decrease the dimension of the feature vector, respectively. The LDA output yields a vector with 150 components as an eigenvector from the original vector with 1,024 elements.

Furthermore, Gaussian Probabilistic Linear Discriminant Analysis (GPLDA) is employed for decision-making. It is trained with 20 iterations and uses vectors with a length of 150 components. These parameters remain fixed in both single and multimodal scenarios. The results of this evaluation can be found in Table 4.

In the single modality mode in verification tasks, both proposed FaceNet and VoiceNet are used as feature extractors, and their features are fed into GPLDA for decision-making in a two-class classification. The results indicate that the average EER in the speaker verification using voice modality is 5.12%, while in the face verification mode, it is 1.01%. As expected, the system based on face modality performs better due to the superior performance of the proposed face feature extraction. In multimodal mode, we designed two systems based on two different voice features.

In the fusion of x-vector and facial features, the system achieves an average EER of 0.62% in the feature fusion mode. However, in concatenating the gammatonegram feature vector with the facial feature vector, the system reaches 0.82% of EER. Similar to the speaker identification scenario, this demonstrates an improved performance compared to the single modality mode.

6 Discussion

visualization of the activity status of layers can also provide useful information for understanding CNN’s functionality. Usually, an appropriate cognition of what happens inside the CNN network gets little attention, and one of the approaches is to consider the CNN as a black box. However, there are helpful visualization methods to understand what is happening inside of a CNN. We depict a representation method called the

Locally-Interpretable Model-agnostic Explanation (LIME) technique (Ribeiro et al., 2016) to show the crucial parts of the image that play a more significant role in network decision-making.

This visualization was made using 2D-CNNs consisting of FaceNet in Figure 2 and gammatonegram-based speaker identifiers in Figure 3. These figures make it possible to understand which parts of the images are focused on by the network. It is also possible to understand whether the network is focused on the essential and discriminative parts of the image or not. The image segmentation colors show each part’s priority according to a color bar. By examining the map and its corresponding image, it can be seen that the network focused on the distinguishing parts of the image, which is the reason for the acceptable result for FaceNet and weak performance on speaker identification using the gammatonegram feature.

In this work, we tried to analyze different strategies for audio-visual modality fusion in the identity recognition task. However, it could be helpful to compare the results of the proposed systems with previous work to understand the scale of the metrics. The results of the present work depict that utilizing gammatonegram representation for voice modality and VGGFace2 pre-trained network for face modality can properly depict the identity information in the identification task. However, utilizing x-vector as a voice feature vector beside fine-tuned VGGFace2 in the verification tasks can present the best result.

7 Conclusion

In this work, we analyzed different strategies for audio-visual modality fusion in the identity recognition tasks. The results of the present work depict that utilizing gammatonegram representation for voice modality and VGGFace2 pre-trained network for face modality can properly depict the identity information.

This study developed two separate single-modality systems for voice and face in two identification and verification tasks. The FaceNet architecture utilized VGGFace2, while we had two VoiceNet, the first one




Figure 2: Crucial parts of face images for decision-making in the FaceNet




Figure 3: Important part of gamma-tonegram image from the viewpoint of Darknet19 network for classification

was trained from scratch to extract the x-vector, and the second one was based on gammatonegram representation and fine-tuning of pre-trained Darknet19. Based on these systems, multimodal learning was applied in three fusion modes, as described in the article. A softmax layer was the classifier in feature fusion and score fusion approaches in the person identification task. However, for person verification, LDA was used for dimension reduction, and GPLDA was utilized for decision-making. The evaluation was conducted on 118 speakers from the VoxCeleb2 dataset. The results demonstrated that combining speech and face modalities using multimodal learning outperformed the single-mode approach in both identification and verification tasks. Additionally, the feature fusion mode was found to be the most effective strategy for these two modalities.

Future studies could enhance this research by investigating more efficient speech features, such as deep belief networks in auto-encoder architecture, to be used as input for the proposed VoiceNet for x-vector extraction. Moreover, utilizing para-linguistic systems, such as gender recognition from face or voice, could offer valuable insights in a score fusion scenario.

References

- Mohammad Rafiqul Alam, Mohammed Bennamoun, Roberto Togneri, and Ferdous Sohel. 2015. A confidence-based late fusion framework for audio-visual biometric identification. *Pattern Recognition Letters*, 52:65–71.
- Boucetta Aldjia and Boussaad Leila. 2021. Sensor level fusion for multi-modal biometric identification using deep learning. In *2021 International Conference on Recent Advances in Mathematics and Informatics (ICRAMI)*, pages 1–5.
- Vahid Asadpour, Mohammad Mehdi Homayounpour, and Farzad Towhidkhah. 2011. Audio-visual speaker identification using dynamic facial movements and utterance phonetic content. *Applied Soft Computing*, 11(2):2083–2093.
- Danwei Cai, Weiqing Wang, and Ming Li. 2022. Incorporating visual information in audio based self-supervised speaker recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:1422–1435.
- Qiong Cao, Li Shen, Weidi Xie, Omkar M Parkhi, and Andrew Zisserman. 2018. Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 67–74.
- Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. 2018. Voxceleb2: Deep speaker recognition. *arXiv preprint arXiv:1806.05622*.
- Aref Farhadipour, Mohammad Asgari, and Mohammad Reza Hasanabadi. 2014. Gammatonegram based speaker identification. In *2014 4th International Conference on Computer and Knowledge Engineering (ICCKE)*, pages 52–55. IEEE.
- Aref Farhadipour and Pouya Taghipour. 2023. Facial emotion recognition under mask coverage using a data augmentation technique. In *2023 13th International Conference on Computer and Knowledge Engineering (ICCKE)*, pages 001–006. IEEE.
- Aref Farhadipour and Hadi Veisi. 2024a. Analysis of deep generative model impact on feature extraction and dimension reduction for short utterance text-independent speaker verification. *Circuits, Systems, and Signal Processing*, pages 1–18.
- Aref Farhadipour and Hadi Veisi. 2024b. Gammatonegram representation for end-to-end dysarthric speech processing tasks: speech recognition, speaker identification, and intelligibility assessment. *Iran Journal of Computer Science*, pages 1–14.
- Theodoros Giannakopoulos. 2009. A method for silence removal and segmentation of speech signals, implemented in matlab. *University of Athens, Athens*, 2.
- Jaesung Huh, Andrew Brown, Jee weon Jung, Joon Son Chung, Arsha Nagrani, Daniel Garcia-Romero, and Andrew Zisserman. 2023. Voxsrc 2022: The fourth voxceleb speaker recognition challenge. *arXiv preprint arXiv:2302.10248*.
- Jake Lever. 2016. Classification evaluation: It is important to understand both what a classification metric expresses and what it hides. *Nature methods*, 13(8):603–605.
- Shervin Minaee, Amirali Abdolrashidi, Hang Su, Mohammed Bennamoun, and David Zhang. 2023. Biometrics recognition using deep learning: A survey. *Artificial Intelligence Review*, pages 1–49.
- Abderrazzaq Moufidi, David Rousseau, and Pejman Rasti. 2023. Attention-based fusion of ultrashort voice utterances and depth videos for multimodal person identification. *Sensors*, 23(13):5890.
- Arsha Nagrani, Joon Son Chung, Weidi Xie, and Andrew Zisserman. 2020. Voxceleb: Large-scale speaker verification in the wild. *Computer Speech & Language*, 60:101027.
- Catherine Perrodin, Christoph Kayser, Taylor J Abel, Nikos K Logothetis, and Christopher I Petkov. 2015. Who is that? brain networks and mechanisms for identifying individuals. *Trends in cognitive sciences*, 19(12):783–796.
- Ashwin Prakash, Athira Nambiar, Alexandre Bernardino, and et al. 2023. Multimodal adaptive fusion of face and gait features using keyless attention based deep neural networks for human identification. *arXiv preprint arXiv:2303.13814*.
- N Vishnu Prasad and Srinivasan Umesh. 2013. Improved cepstral mean and variance normalization using bayesian framework. In *2013 IEEE workshop on automatic speech recognition and understanding*, pages 156–161.
- Yanmin Qian, Zhengyang Chen, and Shuai Wang. 2021. Audio-visual deep neural network for robust person verification. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:1079–1092.
- N. Venkata Ramana, S. Anu H Nair, and KP Sanal Kumar. 2022. Hybrid biometric based person identification using machine learning. *Journal of Pharmaceutical Negative Results*, pages 1605–1614.
- Joseph Redmon and Ali Farhadi. 2017. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271.

- Marco Tullio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Leda Sari, Kritika Singh, Jiatong Zhou, Lorenzo Torresani, Nayan Singhal, and Yatharth Saraf. 2021. A multi-view approach to audio-visual speaker verification. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6194–6198.
- Saqlain Hussain Shah, Muhammad Saad Saeed, Shah Nawaz, and Muhammad Haroon Yousaf. 2023. Speaker recognition in realistic scenario using multimodal data. In *2023 3rd International Conference on Artificial Intelligence (ICAI)*, pages 209–213.
- David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur. 2018. X-vectors: Robust dnn embeddings for speaker recognition. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5329–5333.
- Anton Stefanidi, Artem Topnikov, Gennadiy Tupitsin, and Andrey Priorov. 2020. Application of convolutional neural networks for multimodal identification task. In *2020 26th Conference of Open Innovations Association (FRUCT)*, pages 423–428.
- Ruijie Tao, Rohan Kumar Das, and Haizhou Li. 2020. Audio-visual speaker recognition with a cross-modal discriminative network. *arXiv preprint arXiv:2008.03894*.
- Rui Wang, Junyi Ao, Long Zhou, Shujie Liu, Zhihua Wei, Tom Ko, Qing Li, and Yu Zhang. 2022. Multi-view self-attention based transformer for speaker recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6732–6736.
- Shuai Wang, Yexin Yang, Yanmin Qian, and Kai Yu. 2021. Revisiting the statistics pooling layer in deep speaker embedding learning. In *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, pages 1–5.