

Facial Profile Biometrics: Domain Adaptation and Deep Learning Approaches

Malak Alamri¹²a, Sasan Mahmoodi²

¹*College of Computer and Information Science, Computer Science Department, Jouf University, Al Jouf, KSA*

²*School of Electronic and Computer Science, University of Southampton, Southampton, UK*

Mzalamri@ju.edu.sa, sm3@ecs.soton.ac.uk

Keywords: Facial Profile, Biometrics, Bilateral Symmetry, Domain Adaptation, Deep Learning.

Abstract: Previous studies indicate that human facial profiles are considered as a biometric modality and there is a bilateral symmetry in facial profile biometrics. This study examines the bilateral symmetry of the human face profiles and presents the analysis of facial profile images for recognition. A method from few-shot learning framework is proposed here to extract facial profile features. Based on domain adaptation and reverse validation, we introduce a technique known as reverse learning (RL) in this paper for the same side profiles to achieve a recognition rate of 85%. In addition, to investigate bilateral symmetry, our reverse learning model is trained and validated on the left side face profiles to measure the cross recognition of 71% for right side face profiles. Also in this paper, we assume that the right face profiles are unlabelled, and we therefore apply our reverse learning method to include the right face profiles in the validation stage to improve the performance of our algorithm for opposite side recognition. Our numerical experiments indicate an accuracy of 84.5% for cross recognition which, to the best of our knowledge, demonstrates higher performance than the state-of-the-art methods for datasets with similar number of subjects. Our algorithm based on few-shot learning can achieve high accuracies for a dataset characterized with as low as four samples per group.

1 INTRODUCTION

The increased global use of biometrics technology has resulted in an increase in security threats. Companies and government agencies have several difficulties recognising authorized users. Tokens, such as identification cards and passwords, are the most common security system authentication methods. However, they are becoming increasingly unreliable due to duplication and the limited capacity of human memory (Abdelwhab and Viriri, 2018). Thus, the use of biometrics, such as facial recognition and fingerprints, is growing in popularity worldwide. Face recognition is a very important security measure that works by identifying possible suspects in videos or surveillance frames. A video may only record a small section of the face, making correct identification impossible in some cases (see Figure 1).

It should be highlighted that the majority of biometrics literature focuses on the front face. Despite its significance, the facial profile has received less attention. It is common knowledge that facial recogni-



Figure 1: The Surveillance Camera Footage Released in Link with a Stabbing at Block 1 Club in 2019 (Everett, 2019).

tion systems may rely heavily on the ability to compare and match the features of two face images. However, a profile view of a face may disclose certain aspects of the face's structure that are not visible in a frontal view (Alamri and Mahmoodi, 2022). Combining the matching results derived from the frontal and profile views of a face can assist in reducing instances of incorrect identification. In the past decade, a variety of algorithms for automatically recognising individuals based on their facial profiles have been introduced. Several facial profile identification techniques are commonly used, including Scale-space fil-

^a <https://orcid.org/0000-0002-3484-0994>

tering (Lipoščak and Loncaric, 1999), Dynamic time warping (Bhanu and Zhou, 2004), Attributed string (Gao and Leung, 2002) and Hidden Markov model (HMM) (Wallhoff and Rigoll, 2001; Wallhoff et al., 2001). A facial profile biometric system enhances the effectiveness of multiple facial identification systems and makes recognition more realistic.

There are several deep learning-based methods that have explicitly investigated facial profile recognition and proposed methods for extracting discriminative features from facial profiles. Sengupta et al. illustrated how several algorithms performed when using a restricted protocol and how each one degraded from frontal–frontal to frontal–profile (Sengupta et al., 2016). In their study, the frontal–frontal and frontal–profile experiment achieved classification accuracy of 96.40% and 84.91%, respectively, using a deep features-based method. Moreover, in recent years, generative models, such as generative adversarial network (GAN)-based methods, have been widely used to synthesise the frontal view from the profile view in order to improve facial profile recognition systems (Zhao et al., 2018; Li et al., 2019; Yin et al., 2020). In addition, deep-learning-based methods have also demonstrated high levels of performance. Facial recognition has been greatly improved by deep learning techniques, which are trained on a large-scale dataset and demonstrate high-level recognition rates under challenging conditions (Parkhi et al., 2015; Deng et al., 2019; Meng et al., 2021). However, it is important to note that all deep learning methods require a large number of samples per class during training to achieve acceptable performance levels.

Our contributions in this paper are as follows: 1) higher recognition rate than the literature is achieved on large datasets containing more than 200 subjects, 2) a technique in the framework of few-shot learning is proposed here to extract features from facial profile images even with as low as four samples per class, 3) our method is based on domain adaptation (Wang et al., 2020), and reverse validation (Morvant et al., 2011; Zhong et al., 2010) to guarantee a high accuracy to improve recognition performance for bilateral symmetry by minimizing the conditional distributions between the training and validation data (Zhong et al., 2010), 4) the ability to identify mirror symmetric images is used generally since mirrored images were captured in different conditions than the first set of images, 5) facial profile appears largely to be bilaterally symmetric. We are the first to use few-shot learning (FSL) to examine the bilateral symmetry of a facial profile. Additionally, we consider challenging facial profile images with various planar

poses, while measuring the performance by training pre-processing facial profile dataset and applying face alignment to register all facial profiles for fair comparisons. The work given here represents the most recent state-of-the-art methods in the recognition of facial profiles and their bilateral symmetry.

The second section of this paper discusses the image preprocessing phases, including face detection, face alignment. In addition, the third section explains the feature extraction stage using two few-shot learning techniques. Section four describes the experimental design, recognition performance, and subsequent discussions. Moreover, Section five summarises the conclusions drawn from the previous sections and highlights the implications of the study.

2 FACIAL PROFILE AND LANDMARKS DETECTION

To begin with, our focus was on detecting the bounding box of a facial profile in all of the images in the dataset and then identifying the important landmarks within the bounding box as described in (Alamri, 2024). To achieve this, we used a pre-trained histogram of oriented gradients (HOG) with a linear support vector machine (SVM) object detector, as introduced in (Dalal and Triggs, 2005).

For landmarks detection, the ensemble of regression trees (ERT) (Kazemi and Sullivan, 2014) method is used to directly estimate the positions of facial landmarks by utilizing a sparse subset of pixel intensities. The majority of face alignment algorithms rely on a face detection bounding box to initialise the shape. Thus, to detect facial landmarks, the face must first be extracted from the image. This extracted region of interest (ROI) is then used to obtain the landmarks. Shape predictors, also known as landmark predictors, are used to predict key (x,y)-coordinates of a given 'shape'. Following the algorithm proposed by Kazemi and Sullivan (Tzimiropoulos and Pantic, 2013), shape predictors are used to locate individual facial structures such as the eyes, eyebrows, nose, lips/mouth, and jawline. They require two inputs: the greyscale version of the image and a rectangle object containing the coordinates of the face area.

The alignment and transformation framework was based on the relevant literature (Zeng and Yi, 2011; Walker et al., 1991). These transformations involved estimating a combination of rotation, translation, and scale that mapped the key points from one set to another on a template.




Figure 2: Convolutional neural network architecture.

3 EXTRACTING VISUAL FEATURES FROM FACIAL PROFILE IMAGES

Deep neural networks are widely used in computer vision for extracting features. In this section, we propose a structure inspired by previous research (Wang et al., 2020; Parnami and Lee, 2022; Garcia and Bruna, 2017) to leverage few-shot learning based on the ResNet50 neural network (He et al., 2016) and ArcFace (Deng et al., 2019). A typical CNN consists of feature extraction and classification components. During training, the model learns the unique facial features and generates feature embeddings through the feature extraction process. In the FSL strategy, after training is completed, it becomes possible to skip the classification step and generate feature embeddings for each face image. Figure 2 illustrates the CNN architecture and the key stages used for feature extraction.

- ResNet50 (He et al., 2016), short for residual network, is a specific type of CNN known for its ability to maintain a lower error rate, even in deeper networks. Our network accepts input images of size 244×244 pixels and computes 7×7 grid feature maps in the last layer before the fully connected network. The pre-trained model was optimised on the ImageNet dataset, which consisted of multiple classes. Due to the difference between the datasets used to train the networks and our target data, we chose to exclude the output of the last fully connected layer. Our numerical experiments demonstrated that the learned feature space efficiently represented human faces using these $7 \times 7 \times 2048$ feature maps, resulting in 49 2048-dimensional vectors.

We then applied Sequential Floating Forward Se-

lection (SFFS) (Shirbani and Soltanian Zadeh, 2013) and principal component analysis (PCA) (Jolliffe, 2002) algorithms as the next step to reduce the dimensionality of the feature space which resulted in achieving the best performance. After reducing the dimensionality, we used the feature vectors to train a KNearest Neighbour (KNN) recognition method with $k = 4$ identified as the optimal value to achieve the best performance.

- ArcFace (Deng et al., 2019) is an innovative deep face recognition algorithm proposed by Jiankang Deng et al. It is considered a discriminative model. The authors proposed the additive angular margin loss function, which has proven to be highly effective in obtaining discriminative features for facial recognition. It consistently outperforms other state-of-the-art loss functions. ArcFace addresses the challenge of effectively learning discriminative face features by incorporating an angular margin that pushes the learned features of different classes apart in the angular space. By enhancing the separability between classes, ArcFace significantly improves the model's ability to distinguish between similar faces. This model is trained using the LFW dataset (Huang et al., 2008) to obtain pre-trained weights. The LFW dataset consists of 13,233 facial photos collected from the web. The ArcFace model expects inputs of size 112×112 and returns 512-dimensional vector representations. Importantly, ArcFace simplifies the process by compressing the extracted features to only 512 components, thereby eliminating the need for PCA.

We selected ResNet50 (Montero et al., 2022) as the backbone for all of the network architectures tested in the ArcFace repository, due to its optimal balance between accuracy and parameter count. In this context, the base model was ResNet50, and we utilised ArcFace as the loss function. Generally, the term backbone refers to the feature-extracting network responsible for processing input data into a specific feature representation. The backbone plays a crucial role in extracting and encoding features from the input data, capturing both low-level and high-level features. After feature extraction, we calculated the ArcFace loss and used it to update the network parameters through backpropagation. We chose ArcFace (Deng et al., 2019) as the baseline for two reasons: it employs a SoftMax-loss-based methodology, eliminating the need for an exhaustive training data preparation stage, and it has been demonstrated to yield the best results for the original face

recognition task.

4 FACIAL PROFILE RECOGNITION USING DOMAIN ADAPTATION AND REVERSE VALIDATION METHODOLOGY

Domain adaptation, a subfield of machine learning, aims to train a model on a source dataset and achieve a high level of accuracy on a target dataset that differs significantly from the source (Farahani et al., 2021). Conversely, if some samples in the target dataset lack labels, reverse validation (Morvant et al., 2011; Wang et al., 2020) is employed to identify the best model that minimises the difference between the conditional distributions of the source and target datasets. In this study, we propose a method grounded in domain adaptation and reverse validation, assuming that some facial profile image samples lack labels and exhibit distinct conditional distributions from the source (training) dataset. Inspired by (Ganin et al., 2016), we introduced an algorithm called reverse learning (RL) based on domain adaptation and reverse validation. Unlike conventional prediction methods, our approach employed a two-step prediction process. We use SFFS to select the most optimal features from the distributions of the training and validation data. In our RL algorithm, the two-step prediction helped identify the best model with optimised hyper-parameters and features between the training and validation datasets, addressing issues associated with covariate shift assumptions (David et al., 2010).




Figure 3: Training process of our reverse learning (RL) algorithm: ① training the model on the training set, ② training the model on the validation set.

In the first step, the process began with a simple training procedure using SFFS, As shown in Figure 3. In the second step of the prediction process, after



Figure 4: Sample facial profile images from the XM2VTSDB dataset.

making predictions for the validation set, we trained on the validation set, made predictions for the training set, and aimed to minimise errors using the SFFS algorithm. In the second prediction step, the validation set served as the training set. The second prediction was finally used to quantify the error on the training set, named the reverse learning metric. Finally, we evaluated the performance of our system by presenting unseen test data from the target dataset.

4.1 Face Profile Dataset

We assessed our method using the XM2VTSDB database, a research resource established and maintained by the University of Surrey (see (Messer et al., 1999) for details). This database is an extended version of the M2VTS database, as it comprises more video recordings of each subject during each session compared to M2VTS. Participants from the XM2VTSDB database attended four sessions and the database was developed over a significant period, enabling a wide range of appearance variations of the individual subjects (see Figure 4).

4.2 Performance Analysis

In this section, we introduce a domain adaptation approach that is employed to assess the effect of this strategy on the recognition accuracy of facial profiles from opposite sides of the face. We detected a noticeable shift in data within the same feature space when evaluating the model fitted on the training set. This discrepancy arose from the limited number of samples per class and the variations in subjects' appearances across the four images, as well as between the training and test sets. To address this challenge, we employed a technique that aimed to identify the gap between the training, validation, and test sets.

We conducted a series of experiments on side-view face recognition to assess the effectiveness of the RL method compared to traditional approaches. We then employed a similar technique in two distinct settings:

4.2.1 Experiment for same view facial profile (left side):

In this setting, we employed the leave-one-out cross-validation (LOOCV) method to measure the recognition rate. In addition, we opted to use KNN as a simple classifier. The model was trained and evaluated from the perspective of the left side, with the test data also originating from the left side. We evaluated the performance of this experiment by analysing four samples from the left side view of each individual in the dataset. For this experiment, 80% of the data on the same side was allocated for training and validation, while the remaining 20% served as the unseen testing set. By utilising ResNet50 features, we achieved an average recognition accuracy of 85% when training KNN with 80% of the samples in the dataset (traditional strategy). Interestingly, we observed an increase in accuracy to 91% when employing the RL strategy. However, when using ArcFace features, the average recognition accuracy dropped to 71%. Notably, this accuracy remained consistent even with the application of the RL approach.




Figure 5: Recognition via CMC performance for ResNet50 features.

4.2.2 Experiment for mirror bilateral symmetry view (left and right side):

In this case, the facial profile side images were initially mirrored from right to left, following the way in which previous processes had been applied to detect facial profiles. We evaluated the performance of this experiment by analysing eight samples from the left and right side view of each individual in the dataset. In this experiment, we trained and validated on the left side and measured the recognition rate on the right side as unseen test data. Our numerical experiments revealed that utilising RL can enhance recognition accuracy to 75%, compared to 70% with the traditional method, based on ResNet50 features.

Additionally, the model was trained on left-side facial profiles and validated on unlabelled right-side profiles. Testing was then conducted using unseen data from the right side. By using right facial profiles as unlabelled data in the validation process, a recognition rate of 82% was achieved. This marked a significant improvement of 12% compared to traditional methods. Interestingly, when ArcFace features were utilised, the accuracy remained consistent across all three techniques, at 56%. However, it is worth noting that the RL approach did not lead to any enhancement in accuracy.

One of the most important techniques for evaluating model performance is cumulative matching characteristics (CMC) curves (DeCann and Ross, 2013). The CMC curve shows the recognition at different ranks, indicating the probability of finding an accurate match at a particular rank. Figures 5 and 6 present the recognition performance of facial profiles. In Figure 5, which used ResNet50 features, the accuracy for same-side recognition improved from 85% at rank-1 to 96% at rank-10 when employing RL. Moreover, the accuracy for BS increased from 71% at rank-1 to 91% at rank-10 when using only right-side images as test data. However, when right-side images were also included in the validation stage, the accuracy rose from 82% at rank-1 to 96% at rank-10.

In Figure 6, which used ArcFace features, the accuracy for RL on the same side improved from 71% at rank-1 to 87% at rank-10. Additionally, the accuracy for BS using only right-side images as test data improved from 56% at rank-1 to 81% at rank-10. Additionally, the accuracy for BS using only right-side images as test data improved from 56% at rank-1 to 81% at rank-10. Similarly, when right-side images were also used in the validation stage, the accuracy increased from 56% at rank-1 to 83% at rank-10. The inferior performance of ArcFace compared to ResNet50 can be attributed to the geometric nature of its loss function, which does not effectively optimise the model for this specific task.

Table 1: Recognition rates with and without domain adaptation based on RL algorithm; (BS) bilateral symmetry, (L) left side, (R) right side, (SD) standard deviation.

View	Method	Dataset			Model					
		Training	Validation	Test	ResNet50	ArcFace	Accuracy	SD	Accuracy	SD
Same side	Traditional	L	L	L	85%	0.016	71%	0.012		
	RL	L	L	L	91%	0.014	71%	0.011		
BS	Traditional	L	L	R	70%	0.021	56%	0.009		
	RL	L	L	R	75%	0.011	56%	0.008		
	RL	L	R	R	82%	0.042	56%	0.041		




Figure 6: Recognition via CMC performance for ArcFace features.

mize for the best results. In particular, ArcFace’s loss function can be limiting, especially in cases such as facial profile. Our numerical experiment results are summarised in Table 1 and reveal the following two findings:

- 1) Using ResNet50 features, the recognition rate for the same side exceeded 85% with the traditional method and improved by 6% with the RL method. For the opposite side, the recognition rate using the traditional method is 70%, which increased to 75% with the RL method when the left side was

included in both the validation and training sets. Notably, the accuracy significantly improved by 12% when right facial profiles were included in the validation process.

- 2) When employing ArcFace features, the recognition rates for both the same and opposite sides remained unchanged, and the model performance did not alter even after introducing the RL method.

Our numerical experiments utilising ResNet50 indicated that the RL method proposed here outperformed the state-of-the-art for datasets containing more than 200 subjects. The dataset utilised in this study differs from the one used in previous work. To clarify, our dataset presented greater challenges due to the following reasons: 1) it comprised 230 subjects, which is substantially larger than all datasets considered in (Ding et al., 2013; Bhanu and Zhou, 2004; Xu and Mu, 2007), 2) some of the facial profiles in our dataset were occluded, yet our method still achieved a high recognition rate of 91% using our RL method, surpassing all methods presented in (Ding et al., 2013; Bhanu and Zhou, 2004; Xu and Mu, 2007), and 3) in the method proposed here, we assumed that the validation set was unlabelled. Such an assumption rendered our dataset more challenging than those in the literature.

In our study of 230 subjects, we achieved a cross-recognition rate of 82%, which is the highest among our results. This performance can be compared to the work presented in (Toygar et al., 2018), where only 46 subjects with left and right profile images were used. The best cross-recognition rate reported in (Toygar et al., 2018) was 81.52%, when the model was trained on the left side and tested on the right side. However, in our study, facial profiles with a larger number of subjects achieved an 82% cross-recognition rate. This indicates that facial profiles contain sufficient information to be considered an independent and important biometric modality. Table 2 presents the recognition rates of facial profiles from various methods found in the literature.

Table 2: Recognition rates of facial profiles in the literature with various methods.

Publication	Dataset	Method	Accuracy
Same side view			
(Ding et al., 2013)	44	DWT + RF	92.50% 90%
(Bhanu and Zhou, 2004)	30	DWT	
(Xu and Mu, 2007)	38	PCA	77.63%
Other side view			
(Toygar et al., 2018)	46	BSIF + LPQ + LBP	81.52%

Table 1 presents the results which demonstrate that ResNet50 performed better than ArcFace in both same-side and cross-recognition experiments. While ArcFace demonstrated strong performance when dealing with a slightly angled view of faces and excelled in large-scale face identification tasks, its performance was lower when applied to facial profiles.

5 CONCLUSION

This study introduced a method for facial profile recognition that combines few-shot learning, domain adaptation, and reverse validation techniques. We employed a similarity registration technique to ensure the precise alignment of all facial profile images. By utilising two pre-trained CNN models, ResNet50 and ArcFace, we implemented few-shot learning. Among these models, ResNet50 demonstrated superior performance compared to ArcFace. Specifically, our RL algorithm, which utilised ResNet50, outperformed the traditional training methods discussed in this study. The results obtained from our RL method reveal significant improvements in classification rates. Specifically, the recognition rate for same side (left side) reached 91%, surpassing the state-of-the-art performance achieved with datasets of similar sizes.

Additionally, promising outcomes were observed in cross-recognition, with a rate of 82% achieved when right-side images were used in the validation stage. Furthermore, a recognition rate of 75% was attained when left-side images were employed in validation, with right-side images used for testing. Numerical experiments indicated a notable 7% enhancement in cross-recognition accuracy when right-side faces were included in the validation process. Therefore, we can conclude that depending on the applica-

tion, recognition is viable even when facial profiles are the sole biometric modality available, including scenarios involving bilateral symmetry. Finally, we obtained promising results using only four samples per subject. Alternatively, leveraging a learning transformation or training a neural network from scratch could have been considered.

REFERENCES

Abdelwhab, A. and Viriri, S. (2018). A survey on soft biometrics for human identification. *Machine Learning and Biometrics*, 37.

Alamri, M. (2024). *Facial profile recognition using comparative soft biometrics*. PhD thesis, University of Southampton.

Alamri, M. and Mahmoodi, S. (2022). Face profile biometric enhanced by eyewitness testimonies. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 1127–1133. IEEE.

Bhanu, B. and Zhou, X. (2004). Face recognition from face profile using dynamic time warping. In *Proceedings of the 17th International Conference on Pattern Recognition, 2004. ICPR 2004.*, volume 4, pages 499–502. IEEE.

Dalal, N. and Triggs, B. (2005). Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, volume 1, pages 886–893. Ieee.

David, S. B., Lu, T., Luu, T., and Pál, D. (2010). Impossibility theorems for domain adaptation. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 129–136. JMLR Workshop and Conference Proceedings.

DeCann, B. and Ross, A. (2013). Relating roc and cmc curves via the biometric menagerie. In *2013 IEEE Sixth International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–8. IEEE.

Deng, J., Guo, J., Xue, N., and Zafeiriou, S. (2019). Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4690–4699.

Ding, S., Zhai, Q., Zheng, Y. F., and Xuan, D. (2013). Side-view face authentication based on wavelet and random forest with subsets. In *2013 IEEE International Conference on Intelligence and Security Informatics*, pages 76–81. IEEE.

Everett, A. (17-4-2019). Warrington guardian-new cctv images after man stabbed in face in block 1 club.

Farahani, A., Vogheli, S., Rasheed, K., and Arabnia, H. R. (2021). A brief review of domain adaptation. *Advances in data science and information engineering: proceedings from IC DATA 2020 and IKE 2020*, pages 877–894.

Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. (2016). Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030.

Gao, Y. and Leung, M. K. (2002). Human face profile recognition using attributed string. *Pattern Recognition*, 35(2):353–360.

Garcia, V. and Bruna, J. (2017). Few-shot learning with graph neural networks. *arXiv preprint arXiv:1711.04043*.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.

Huang, G. B., Mattar, M., Berg, T., and Learned-Miller, E. (2008). Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*.

Jolliffe, I. T. (2002). *Principal component analysis for special types of data*. Springer.

Kazemi, V. and Sullivan, J. (2014). One millisecond face alignment with an ensemble of regression trees. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1867–1874.

Li, P., Wu, X., Hu, Y., He, R., and Sun, Z. (2019). M2fpa: A multi-yaw multi-pitch high-quality dataset and benchmark for facial pose analysis. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10043–10051.

Lipoščak, Z. and Loncaric, S. (1999). A scale-space approach to face recognition from profiles. In *Computer Analysis of Images and Patterns: 8th International Conference, CAIP'99 Ljubljana, Slovenia, September 1–3, 1999 Proceedings 8*, pages 243–250. Springer.

Meng, Q., Zhao, S., Huang, Z., and Zhou, F. (2021). Magface: A universal representation for face recognition and quality assessment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14225–14234.

Messer, K., Matas, J., Kittler, J., Luettin, J., and Maitre, G. (1999). Xm2vtsdb: The extended m2vts database. In *Second international conference on audio and video-based biometric person authentication*, volume 964, pages 965–966. Citeseer.

Montero, D., Nieto, M., Leskovsky, P., and Aginako, N. (2022). Boosting masked face recognition with multi-task arcface. In *2022 16th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*, pages 184–189. IEEE.

Morvant, E., Habrard, A., and Ayache, S. (2011). Sparse domain adaptation in projection spaces based on good similarity functions. In *2011 IEEE 11th International Conference on Data Mining*, pages 457–466. IEEE.

Parkhi, O. M., Vedaldi, A., and Zisserman, A. (2015). Deep face recognition.

Parnami, A. and Lee, M. (2022). Learning from few examples: A summary of approaches to few-shot learning. *arXiv preprint arXiv:2203.04291*.

Sengupta, S., Chen, J.-C., Castillo, C., Patel, V. M., Chellappa, R., and Jacobs, D. W. (2016). Frontal to profile face verification in the wild. In *2016 IEEE winter conference on applications of computer vision (WACV)*, pages 1–9. IEEE.

Shirbani, F. and Soltanian Zadeh, H. (2013). Fast sffs-based algorithm for feature selection in biomedical datasets. *AUT Journal of Electrical Engineering*, 45(2):43–56.

Toygar, Ö., Alqaralleh, E., and Afaneh, A. (2018). Symmetric ear and profile face fusion for identical twins and non-twins recognition. *Signal, Image and Video Processing*, 12:1157–1164.

Tzimiropoulos, G. and Pantic, M. (2013). Optimization problems for fast aam fitting in-the-wild. In *Proceedings of the IEEE international conference on computer vision*, pages 593–600.

Walker, M. W., Shao, L., and Volz, R. A. (1991). Estimating 3-d location parameters using dual number quaternions. *CVGIP: image understanding*, 54(3):358–367.

Wallhoff, F., Muller, S., and Rigoll, G. (2001). Recognition of face profiles from the mugshot database using a hybrid connectionist/hmm approach. In *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*, volume 3, pages 1489–1492. IEEE.

Wallhoff, F. and Rigoll, G. (2001). A novel hybrid face profile recognition system using the feret and mugshot databases. In *Proceedings 2001 International Conference on Image Processing (Cat. No. 01CH37205)*, volume 1, pages 1014–1017. IEEE.

Wang, Y., Yao, Q., Kwok, J. T., and Ni, L. M. (2020). Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34.

Xu, X. and Mu, Z. (2007). Feature fusion method based on kcca for ear and profile face based multimodal recognition. In *2007 IEEE international conference on automation and logistics*, pages 620–623. IEEE.

Yin, Y., Jiang, S., Robinson, J. P., and Fu, Y. (2020). Dual-attention gan for large-pose face frontalization. In *2020 15th IEEE international conference on automatic face and gesture recognition (FG 2020)*, pages 249–256. IEEE.

Zeng, H. and Yi, Q. (2011). Quaternion-based iterative solution of three-dimensional coordinate transformation problem. *J. Comput.*, 6(7):1361–1368.

Zhao, J., Cheng, Y., Xu, Y., Xiong, L., Li, J., Zhao, F., Jayashree, K., Pranata, S., Shen, S., Xing, J., et al. (2018). Towards pose invariant face recognition in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2207–2216.

Zhong, E., Fan, W., Yang, Q., Verschueren, O., and Ren, J. (2010). Cross validation framework to choose amongst models and datasets for transfer learning. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2010, Barcelona, Spain, September 20-24, 2010, Proceedings, Part III 21*, pages 547–562. Springer.