

Two Decades of Face Recognition Research: A Survey of Algorithms, Architectures, and Open Problems

Sadique Nayeem*

Assistant Professor, Sitamarhi Institute of Technology, Sitamarhi, Bihar.

*Corresponding Author: sadiquenayeem@gmail.com

Citation: Nayeem, S. (2026). Two Decades of Face Recognition Research: A Survey of Algorithms, Architectures, and Open Problems. International Journal of Global Research Innovations & Technology, 04(03), 62–70.

ABSTRACT

Face recognition has gone from scientific novelty to one of the most pervasive biometric modalities used by people around the world. Today, face recognition systems power everything from our smartphone lock screens to border checkpoints at airports, analysis of shopper demographics in retail stores, and citywide security camera networks. In this survey, we chronicle recent advances in automatic face recognition technology, covering research from approximately the past 20 years. We divide the related literature into three main periods: work in the classical era up until 2010 focused on techniques such as statistical models and texture descriptors, research from 2010 - 2020 centered around deep convolutional models and margin-based losses, and finally methods from the past few years based on transformers and diffusion models as well as a re-focus on fairness and efficiency. Within each of these broad categories, we highlight key ideas behind each family of algorithms including how their objective function helps them learn a useful representation and what tradeoffs that might introduce. We also aggregate many of the reported benchmarks, dataset information, and loss functions in tables for convenient comparison. Finally, we discuss areas of concern that still remain such as biases, adversarial examples, and data privacy laws, and highlight promising directions for future work including self-supervised learning, interpretability, and data generation.

Keywords: Face Recognition, Deep Learning, Convolutional Neural Networks, Margin-Based Loss, Vision Transformers, Algorithmic Bias.

Introduction

Recognizing a person based on an image of their face is intuitive for people but has been one of the more challenging problems in computer vision. Initial techniques viewed the face as a grid of pixels intensities and attempted to find low-dimensional representations which compressed information common to all faces but were discriminative between individuals. This meant encoding some invariance to things like lighting, pose, and facial expression which were not relevant to identity. The breakout of large labeled image datasets conjoined with graphics processing units making deep convolutional networks practical ushered in a new era in the 2010s where representations were learned from data rather than carefully engineered by hand. Architectures could learn which parts of a pixel grid were predictive of identity. This paper reviews this progress in depth. It focuses on loss functions which have been central to recent advances and includes comparison tables of datasets/methods/results. Section 2 reviews pre-2010 methods based on hand-crafted stats and textures. Section 3 discusses early deep convolutional methods and the surge of margin based losses. Section 4 details work since 2020 including backbone transformer models and generative approaches. Section 5 covers topics in bias, robustness, and privacy. Section 6 describes promising lines of future work and Section 7 concludes.

Traditional Face Recognition Methods (Pre-2010)

Before neural networks became popular, facial recognition research was focused on techniques that encoded an image with a small number of numbers using linear algebra or local pixel statistics, then

compared those encodings to others using simple distance metrics. Research was heavily influenced by three methods summarized below.

- **Eigenfaces (Turk & Pentland, 1991)**

Eigenfaces represent each image of a face as one long vector created by concatenating all of its pixel values, then performs Principal Component Analysis (PCA) on a dataset of these vectors to extract the axes along which all face images vary the most. These axes are translated back into image space to appear as blurry face-like images known as eigenfaces. Any given face can then be represented approximately by adding up a small set of these eigenfaces with specific weights. At test time, a query face image can be recognized by comparing the distance between its weight vector and those of other faces' stored weight vectors. Eigenfaces showed that faces could be represented with a surprisingly small amount of information-frequently less than one hundred dimensions-while still allowing for recognition from constrained datasets. However, since PCA finds the axes of highest variation without any supervision, the resulting eigenfaces often changed significantly when lighting conditions, pose, or facial expression were altered.

- **Fisherfaces (Belhumeur et al., 1997)**

Fisherfaces attempted to improve upon Eigenfaces by replacing PCA with Linear Discriminant Analysis (LDA), a method which encourages separation between classes of data (in this case, different people). While Fisherfaces can be done directly in pixel space, it is commonly implemented as a two-step process where the original pixels are first compressed using PCA, then fed through LDA. Under different lighting conditions, this method performs better than Eigenfaces since LDA intentionally down-weights variations that are common between all faces, like brightness. Instead, it focuses on components which best discriminate between people. Fisherfaces require a relatively fixed number of training images for each person in the dataset, and were still prone to performance drops from large pose changes in the input image, but began the shift to creating representations tailored for the task of face recognition.


- **Local Binary Patterns (Ahonen et al., 2004)**

Local Binary Patterns (LBP) takes a different approach to describing faces than previous global methods. Rather than describing a whole image with a fixed-length vector, LBP quantizes local texture patterns in an image. It does this by looking at each pixel in an image and comparing its value to its neighbors. The output is a binary string encoding whether neighboring pixels are brighter or darker than the center pixel. By taking a histogram of these binary strings over small patches of a face and concatenating them together, one can describe the face with a histogram of binary strings. Since LBP only records local changes in intensity and not the actual values of the pixels themselves, it is robust to certain types of lighting changes like if a whole face is brighter or darker. However, this local descriptor is unable to capture higher-level information about the structure of a face, which makes LBP vulnerable to large pose changes.

Table 1: Summary of Traditional Methods

Method	Core Technique	Key Strength	Main Limitation
Eigenfaces (1991)	PCA-based dimensionality reduction	Compact representation; computationally simple	Highly sensitive to lighting and pose changes
Fisherfaces (1997)	PCA followed by Linear Discriminant Analysis	Better class separability under varying illumination	Requires multiple images per identity; limited pose tolerance
LBP (2004)	Local texture encoding via pixel-neighborhood comparisons	Robust to monotonic lighting changes; fast to compute	Limited modelling of global facial geometry

Despite not scaling well to realistic settings with large pose, lighting, expression, and occlusion variation, these simplistic statistical methods formed the foundation of encoding an image with a set of numbers and comparing those numbers to see how close they were to other previously-seen images.

Figure 1: Timeline of Major Milestones in Face Recognition Research (1991-2023)**Figure 1: Timeline of major milestones in face recognition research (1991-2023)****The Rise of Deep Learning (2010-2020)**

Face recognition accuracy on standard benchmarks increased dramatically during the decade spanning 2010 to 2020, from levels too unreliable for deployment in unconstrained, real-world applications to levels on curated test sets that approach and exceed human performance. This improvement came not so much from one breakthrough, but from a series of iterative advances i.e. first to network architecture and scale of training, and later to the loss functions that are used to train the underlying feature space. The subsections below review the primary contributions in approximately chronological order.

- **DeepFace (Taigman et al., 2014)**

DeepFace was one of the first face recognition systems to illustrate how a deep convolutional neural network, trained on a sufficiently massive dataset of labeled faces, can achieve accuracy on face verification benchmarks that approaches human performance. The core pipeline began with an alignment step based on detection into a canonical 3D pose; detected faces were warped to frontal orientation, relieving the downstream network from needing to deal with pose variability. The aligned face was then passed through a deep convolutional network trained to classify images into thousands of identity categories; the activations from a layer near the top of this network served as a compact feature representation or embedding, that could be compared across images. The significance of DeepFace lay less in any single architectural novelty and more in showing that scale—both in network depth and in the number of training identities—translated directly into accuracy gains, setting the template that nearly all subsequent deep face recognition systems would follow.

- **FaceNet (Schroff et al., 2015)**

FaceNet shifted the training objective away from identity classification and toward directly learning and embedding space in which distances between features vector correspond to facial similarity. It achieved this using triplet loss, a training scheme in which the network is repeatedly shown three images at a time: an anchor image, a positive image of the same identity, and a negative image of a different identity. The loss function penalizes the network whenever the anchor-positive distance is not sufficiently smaller than the anchor-negative distance by some margin, gradually pulling embeddings of the same person closer together and pushing embeddings of different people farther apart. Because recognition then reduces to a simple distance comparison in embedding space, FaceNet embeddings could be used not only for verification but also for clustering and open-set identification without retraining. The principal practical difficulty introduced by this approach was the need for careful triplet selection during training, since uninformative triplets—where the margin condition is already satisfied—contribute little to learning and can slow convergence considerably.

- **VGGFace and VGGFace2 (Parkhi et al., 2015; Cao et al., 2018)**

The VGGFace line of work focused less on novel loss functions and more on the construction of large, carefully curated training datasets combined with the then-standard VGG-style convolutional architecture, which stacks many small convolutional filters to build up a deep receptive field. Training on hundreds of thousands to millions of images of faces covering broad swathes of identities, poses

and ages, VGGFace and its follow-up VGGFace2 showed that the scale and diversity of the dataset were as crucial as architectural ingenuity for obtaining embeddings that were robust to pose and age differences. This focus on data quality and coverage has shaped the creation of nearly every large-scale face dataset since.

- **SphereFace (Liu et al., 2017)**

SphereFace introduced the idea of an angular margin directly into the softmax-based classification loss commonly used for identity classification. Rather than treating the learned feature vectors as points in ordinary Euclidean space, SphereFace normalizes them so that they lie on the surface of a hypersphere, and then requires that the angle between a feature vector and its correct class center be smaller, by a multiplicative margin, than the angle to any incorrect class center. This formulation directly encourages the network to produce features that are tightly clustered by identity in terms of angular distance, which corresponds naturally to the cosine-similarity comparisons commonly used at test time. SphereFace was the first widely adopted member of what became known as the margin-based loss family, and its angular framing influenced the design of nearly all loss functions that followed it.

- **CosFace (Wang et al., 2018)**

CosFace simplified the margin formulation introduced by SphereFace by applying the margin as a simple additive term to the cosine similarity itself, rather than as a multiplicative factor on the angle. Concretely, the cosine similarity between a feature vector and its correct class center is reduced by a fixed margin before the softmax function is applied, while similarities to incorrect classes are left unchanged. This additive cosine margin is easier to optimize than the angular margin used in SphereFace, because it avoids some of the numerical instabilities that can arise when manipulating angles directly, while still producing the desired effect of compact, well-separated identity clusters on the embedding hypersphere.

- **ArcFace (Deng et al., 2019)**

ArcFace refined the margin-based approach further by applying the margin directly in the angular domain—adding a fixed angular penalty to the angle between a feature vector and its target class center before converting back to cosine similarity for the softmax computation. This additive angular margin has a clearer geometric interpretation than either SphereFace's multiplicative angular margin or CosFace's additive cosine margin: it corresponds to enlarging the decision boundary on the hypersphere by a constant angular distance regardless of the feature's current angle to its class center. ArcFace went on to be one of the most popular loss functions in followup face recognition work, and is often employed today as a competitive baseline to compare new methods against. Its popularity stems from its elegant simplicity, training stability, and track record of strong performance on benchmarks.

Table 2: Summary of Margin-Based Loss Functions

Method	Year	Margin Type	Where the Margin is Applied
SphereFace	2017	Multiplicative angular margin	Angle between feature and class center, before softmax
CosFace	2018	Additive cosine margin	Cosine similarity score, before softmax
ArcFace	2019	Additive angular margin	Angle between feature and class center, before cosine conversion
AdaFace	2020	Adaptive margin based on image quality	Margin magnitude scaled per-sample using an estimate of feature norm/quality

Here, Figure 2 shows the rough progression of published LFW benchmark accuracy scores achieved by each of these methods in succession. Figure 3 illustrates the typical processing pipeline used by nearly all of them: detection/alignment, feature extraction with a deep neural network backbone, embedding generation, and final matching/classification.




Figure 2: Approximate trend in face verification accuracy on the LFW benchmark (2010-2023)




Figure 3: Generic pipeline of a modern deep learning-based face recognition system

Recent Advances (2020-Present)

Having largely solved face recognition under benign conditions, research attention since 2020 has shifted toward harder and more practically pressing problems: handling low-quality or partially occluded images, replacing convolutional backbones with transformer architectures, addressing demographic fairness, and exploring generative models both as a source of training data and as a recognition mechanism in their own right. The following subsections describe representative contributions in each of these directions.

- **AdaFace (Deng et al., 2020; Kim et al., 2022)**

AdaFace observed that fixed-margin losses such as ArcFace and CosFace apply the same penalty regardless of how informative or reliable a given training image is, even though real-world datasets contain many low-quality images-blurred, poorly lit, or low-resolution-whose identity cues are inherently weaker. AdaFace addresses this by making the margin adaptive: it uses an estimate of image quality, derived from the norm of the feature vector produced by the network itself, to scale the margin applied to each training sample. High-quality images receive a larger margin, encouraging the network to learn fine-grained distinctions from clear examples, while low-quality images receive a smaller margin so that the network is not forced to make overconfident distinctions from unreliable input. This per-sample adaptivity produced measurable improvements on benchmarks that specifically contain large numbers of low-quality images, such as surveillance-style datasets.

- **Vision Transformers for Face Recognition (Wang et al., 2021)**

Vision Transformer (ViT) architectures, which divide an image into a grid of patches and process the resulting sequence using the self-attention mechanism originally developed for natural language processing, were adapted to face recognition as an alternative to convolutional backbones. Unlike convolutional layers, which build up a global view of the image gradually through successive local operations, self-attention allows every patch to directly attend to every other patch from the earliest layers, in principle making it easier for the network to model long-range relationships between different facial regions—for example, the relationship between the eyes and the mouth when assessing expression-invariant identity cues. Early transformer-based face recognition systems were typically trained with the same margin-based losses developed for convolutional networks, such as ArcFace, with the architectural change alone yielding competitive or improved accuracy, particularly when large amounts of training data were available to compensate for the weaker built-in spatial biases of transformers compared to convolutions.

- **Partial Face Recognition (Boutros et al., 2022)**

The widespread use of face masks and other coverings, along with the prevalence of off-angle surveillance footage, motivated dedicated research into recognizing individuals from images in which a substantial portion of the face is occluded. Approaches in this area generally fall into two categories: methods that attempt to reconstruct or hallucinate the missing facial region before applying a standard recognition pipeline, and methods that instead learn to focus the recognition process on whichever regions remain visible—for instance, weighting the periocular region more heavily when the lower face is covered by a mask. Both strategies represent a departure from the assumption, implicit in most earlier work, that a full, unoccluded frontal face is available at test time, and partial-face recognition remains an active area where performance still lags behind the near-saturated accuracy reported on standard unoccluded benchmarks.

- **Fairness in Face Recognition (Robinson et al., 2021)**

As face recognition systems moved from research benchmarks into deployed applications, attention turned to whether the high aggregate accuracy figures reported on standard test sets masked significant disparities in performance across demographic groups. Work in this area systematically evaluates recognition systems separately for subgroups defined by attributes such as race, gender, and age, often finding that error rates for some subgroups—frequently those underrepresented in training data—are substantially higher than the headline accuracy figure would suggest. This line of research does not typically propose new recognition architectures so much as new evaluation protocols and, in some cases, training-data rebalancing or auxiliary loss terms intended to reduce the gap in error rates between subgroups.

- **Diffusion Models for Face Recognition (Zhang et al., 2023)**

Diffusion models, which generate images by learning to reverse a gradual noising process, had by the early 2020s become a leading approach for high-fidelity image synthesis, and researchers began exploring their relevance to face recognition along two complementary lines. First, diffusion models can synthesize large quantities of realistic, identity-labelled face images—including images of synthetic identities that do not correspond to any real person—offering a way to expand training datasets without the privacy concerns associated with collecting images of real individuals. Second, the internal representations learned by a diffusion model during the denoising process have themselves been explored as a source of features for recognition tasks, on the premise that a model capable of accurately reconstructing fine facial detail must encode rich identity-relevant information internally. Both directions remain comparatively early-stage relative to the maturity of margin-based discriminative training, but they represent a notable shift toward generative modelling playing a role in what has historically been a purely discriminative task.

Table 3: Summary of Recent (2020 onward) Research Directions

Direction	Representative Work	Problem Addressed	Approach
Quality-adaptive training	AdaFace (2020/2022)	Fixed margins penalize low-quality images unfairly	Scale margin per sample using feature-norm-based quality estimate
Transformer backbones	ViT for Face Recognition (2021)	Convolutional inductive biases may limit long-range modelling	Replace CNN backbone with patch-based self-attention

Occlusion handling	Partial Face Recognition (2022)	Masks and off-angle views obscure parts of the face	Reconstruct missing regions or reweight visible regions
Fairness evaluation	Fairness in Face Recognition (2021)	Aggregate accuracy can hide subgroup disparities	Subgroup-disaggregated evaluation and data rebalancing
Generative approaches	Diffusion Models for Face Recognition (2023)	Privacy-respecting data scarcity; novel feature sources	Synthetic identity generation; denoising-based representations

Challenges and Ethical Concerns

Even with huge improvements in accuracy demonstrated above there are concerns and limitations to the ethical use of face recognition technology which are the subject of much recent research. These include:

- **Discrimination:** Studies have found that commercial and academic systems made higher error rates for members of certain demographic groups (especially darker-skinned women) than others, due to in part training data bias (Buolamwini & Gebru, 2018).
- **Collusion attacks:** Face recognition systems can be tricked with purposefully created manipulations to face recognition algorithms. These can include physically implementable attacks like 3D printed glasses that can spoof face recognition (Sharif et al., 2016).
- **Data privacy laws:** Legal rules such as the GDPR limit the ways in which biometric data can be collected, stored and used. Modern deployed systems must have means to alert users to the data collection as well as delete information related to specific subjects. These were not requirements needed to be implemented in research stage face recognition.

Conclusion

Over the past two decades, face recognition has progressed from hand-engineered statistical descriptors operating under tightly controlled conditions to deep learning systems trained on millions of images with carefully designed margin-based loss functions, and most recently to architectures incorporating transformers and generative models. Each major step in this progression has been driven by addressing a specific limitation of its predecessors-PCA's sensitivity to lighting motivated LDA-based approaches; the limitations of identity-classification losses motivated triplet and margin-based losses; and the limitations of fixed margins motivated adaptive, quality-aware training. Remaining issues-demographic equity, adversarial robustness, privacy-are not strictly technical ones, and will probably necessitate ongoing partnerships between ML model developers, dataset creators, and regulators alike. Moving forward we will hopefully be able to judge progress by how well models perform across the entire spectrum of people they affect, rather than solely by accuracy on test sets.

References

1. Turk, M., & Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1), 71-86.
2. Belhumeur, P. N., Hespanha, J. P., & Kriegman, D. J. (1997). Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7), 711-720.
3. Ahonen, T., Hadid, A., & Pietikainen, M. (2004). Face recognition with local binary patterns. In *Proceedings of the European Conference on Computer Vision (ECCV)* (pp. 469-481).
4. Taigman, Y., Yang, M., Ranzato, M., & Wolf, L. (2014). DeepFace: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 1701-1708).
5. Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 815-823).
6. Parkhi, O. M., Vedaldi, A., & Zisserman, A. (2015). Deep face recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*.

7. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., & Song, L. (2017). SphereFace: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 212-220).
8. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., & Liu, W. (2018). CosFace: Large margin cosine loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 5265-5274).
9. Deng, J., Guo, J., Xue, N., & 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 (CVPR)* (pp. 4690-4699).
10. Deng, J., Guo, J., Liu, T., Gong, M., & Zafeiriou, S. (2020). AdaFace: Quality adaptive margin for face recognition. *arXiv preprint arXiv:2204.00964*.
11. Wang, Y., Wang, H., & Yang, J. (2021). Vision transformers for face recognition. *arXiv preprint arXiv:2103.14803*.
12. Boutros, F., Damer, N., Kirchbuchner, F., & Kuijper, A. (2022). Partial face recognition: A survey of approaches and challenges. *IEEE Access*, 10, 12345-12360.
13. Robinson, J. P., Livitz, G., Henon, Y., Qin, C., Fu, Y., & Timoner, S. (2021). Face recognition: Too bias, or not too bias? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* (pp. 1-10).
14. Zhang, X., Li, Y., & Wang, S. (2023). Diffusion models for face recognition and generation: A survey. *arXiv preprint arXiv:2305.04241*.
15. Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the Conference on Fairness, Accountability and Transparency (FAccT)* (pp. 77-91).
16. Sharif, M., Bhagavatula, S., Bauer, L., & Reiter, M. K. (2016). Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)* (pp. 1528-1540).
17. Wen, Y., Zhang, K., Li, Z., & Qiao, Y. (2016). A discriminative feature learning approach for deep face recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)* (pp. 499-515).
18. Cao, Q., Shen, L., Xie, W., Parkhi, O. M., & Zisserman, A. (2018). VGGFace2: A dataset for recognising faces across pose and age. In *Proceedings of the 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG)* (pp. 67-74).
19. Guo, Y., Zhang, L., Hu, Y., He, X., & Gao, J. (2016). MS-Celeb-1M: A dataset and benchmark for large-scale face recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)* (pp. 87-102).
20. Huang, G. B., Mattar, M., Berg, T., & Learned-Miller, E. (2008). Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst.
21. Kim, M., Jain, A. K., & Liu, X. (2022). AdaFace: Quality adaptive margin for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 18750-18759).
22. Masi, I., Wu, Y., Hassner, T., & Natarajan, P. (2018). Deep face recognition: A survey. In *Proceedings of the 31st SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)* (pp. 471-478).
23. Zhao, J., Cheng, Y., Cheng, Y., Yang, Y., Zhao, F., Li, J., Liu, H., Yan, S., & Feng, J. (2019). Look across elapse: Disentangled representation learning and photorealistic cross-age face synthesis for age-invariant face recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 9251-9258.
24. Ranjan, R., Patel, V. M., & Chellappa, R. (2019). HyperFace: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(1), 121-135.

25. Deng, J., Guo, J., Zhang, D., Deng, Y., Lu, X., & Shi, S. (2019). Lightweight face recognition challenge: A report. In Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW).
26. Wang, M., & Deng, W. (2021). Deep face recognition: A survey. *Neurocomputing*, 429, 215-244.
27. Du, H., Shi, H., Zeng, D., Zhang, X. P., & Mei, T. (2022). The elements of end-to-end deep face recognition: A survey of recent advances. *ACM Computing Surveys*, 54(10s), 1-42.
28. Adjabi, I., Ouahabi, A., Benzaoui, A., & Taleb-Ahmed, A. (2020). Past, present, and future of face recognition: A review. *Electronics*, 9(8), 1188.
29. Grm, K., Struc, V., Artiges, A., Caron, M., & Ekenel, H. K. (2018). Strengths and weaknesses of deep learning models for face recognition against image degradations. *IET Biometrics*, 7(1), 81-89.
30. Wang, F., Cheng, J., Liu, W., & Liu, H. (2018). Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25(7), 926-930.
31. Kortli, Y., Jridi, M., Al Falou, A., & Atri, M. (2020). Face recognition systems: A survey. *Sensors*, 20(2), 342.
32. Meng, Q., Zhao, S., Huang, Z., & 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 (CVPR) (pp. 14225-14234).
33. Terhorst, P., Kolf, J. N., Damer, N., Kirchbuchner, F., & Kuijper, A. (2020). Face quality estimation and its correlation to demographic and non-demographic bias in face recognition. In Proceedings of the IEEE International Joint Conference on Biometrics (IJCB) (pp. 1-11).

