International Journal of Sciences: Basic and Applied Research (IJSBAR)

ISSN 2307-4531
(Print & Online)

https://gssrr.org/index.php/JournalOfBasicAndApplied/index

Face Recognition Technologies: A Comprehensive Review

Mustafa Semih Sadak *

Piri Reis University, Tuzla, İstanbul, 34940, Türkiye
Email: msadak@pirireis.edu.tr

Abstract

It is hard to state how much face recognition has developed in the last decade. As a once-specialized type of biometric system, we see face recognition wherever we look: at the airport, through Instagram etc. This review will discuss that journey from its traditional systems to the deep learning systems that we see at the forefront of the field. From a detailed analysis of 127 academic publications from 2015 to 2024, what is evidently clear is that face recognition systems have undergone rapid algorithm development during this time period. From that literature, there are several emergent trends. Convolutional neural networks are still dominant, transformer-based architectures are a new development, and there is an uptick in work on the pressing issues of bias and privacy. Our analysis reflects that face recognition systems are more accurate but adept at modeling real-world situation variation in the terms of illumination, pose and the aging process. Additionally, the issue of demographic fairness is far from solved. Overall, the review provides an informative snapshot of the current state of face recognition and identifies what we observe as the most promising and important avenues towards future research.

Keywords: Face recognition; Deep learning; Biometrics; Computer vision; Neural networks; Privacy; Bias mitigation.

Received: 6/25/2025
Accepted: 8/9/2025
Published: 8/19/2025

** Corresponding author.*

1. Introduction

It would be easy to believe face recognition sprang up overnight. However, there is a long, tortuous journey from a niche research challenge to a pervasive technology that is now a part of our daily routine; moreover, it has involved a number of pivotal creeping breakthroughs. The growth of performance in face recognition systems is due in part to the confluence of three different factors: the availability of powerful computing, unprecedented quantities of training data, and, of course, the deep learning revolution that has transformed all of computer vision.

The tale of face recognition is itself just as fascinating. It has been around since the 1960s with Bledsoe's [1] early, semi-automated work, and has gone through at least three distinguishable paradigms. The 1970s and 1980s were dominated by geometric approaches using facial landmarks; the late 1980s and 1990s saw the advent of holistic approaches using Eigenfaces [2] and Fisherfaces [3]. In the 2000s, the field moved to local, keypoint-based representations - which leads us to the present era characterized by deep learning, ushered in by a number of important papers including DeepFace [4] and FaceNet [5], and indisputably the current state-of-the-art.

Our intent here is to contextualize the last decade, from 2015 to 2024. This time frame feels especially important, not only for the meteoric rise of deep learning rising to prominence, but also for a newly coined collective awareness of the ethical consequences of the technology. Our intention is to orient our fellow researchers through this complicated literature - an awareness of key contributions, knowledge gaps, and possible next steps.

2. Methodology

To better understand the field, we knew we had to be systematic. Our search process began with the normal academic databases - IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, and arXiv - to find papers published between January 2015 and December 2024. We ran through a list of search terms; some basic and some more specific. Included in the search terms were face recognition, deep learning, and face recognition deep learning Our initial search produced a staggering 1,847 papers. Once again, we had to be selective and we established a number of inclusion criteria: the work had to perform original research in face recognition algorithms, have experimental evaluation results on standard benchmark datasets, be peer-reviewed and be in English. Thus, we ruled out papers that could be viewed as review paper, position papers, and studies that only focused on face detection without a recognition component.

Upon deleting duplicates and filtering we arrived at 312 papers. The final, and arguably most subjective step was narrowing this group down to 127 papers that construct the reviews; we would favour papers that already had a significant number of citations or had cutting-edge proposals which we felt were genuinely new. Each final paper was then cross analyzed by its specific technical approach, experimental design, results, and stated limitations.

3. Traditional Approaches to Face Recognition

Before looking at current deep learning approaches it is essential to see where it all began with traditional face recognition methods. These traditional methods are not all obsolete, however, since they are still relevant to current research and certainly practical in a resource-limited situation.

3.1 Geometric Feature-Based Approaches

Geometric feature-based approaches for face recognition look at geometric features of the human face. These geometric features can be distances between the eyes, nose width, and mouth position, for example. Kanade [6] was a leader in developing these systems with an automated system that identified human faces by the 16 facial parameters. Although these systems seem intuitive, there was still a robustness issue when varying facial expressions and angles were presented.

3.2 Holistic Methods

The introduction of Principal Component Analysis (PCA) from Turk and Pentland [2], propelled face recognition into a new realm through the use of the Eigenfaces approach. In actuality, this kind of face recognition was concerned with dimensionality reduction where a face image was represented in a lower dimensional subspace. Belhumeur and his colleagues [3] continued the wonderful work using Fisherfaces, which added so-called Linear Discriminant Analysis (LDA) which incorporates between-class variability and maximizes between-class variability and minimizes within-class variability.

3.3 Local Feature Based Methods

The problem of holistic approaches is that they do not always handle partial occlusions and local variation in images, so researchers began to develop local feature-based methods. Ahonen and his colleagues [7] used Local Binary Patterns (LBP) for face recognition, which were robust to changes in illumination. The Scale-Invariant Feature Transform (SIFT) [8] and Speeded-Up Robust Features (SURF) [9] were also adapted for face recognition are also invariant to scale and rotation.

4. Deep Learning Revolution in Face Recognition

With the emergence of deep learning, face recognition research saw a dramatic shift. Various sources of large-scale datasets and computational resources had opened up the possibility for training deep neural networks capable of automatically learning hierarchical representations of faces.

4.1 Convolutional Neural Networks

The first deep-learning face recognition work which used Convolutional Neural Networks (CNNs) was DeepFace [4] by Taigman and his colleagues. It was reported to have achieved close to human performance on the Labeled Faces in the Wild (LFW) benchmark. Their network implementation consisted of a nine-layer deep neural network which was trained using four million facial images. This work emphasized the benefits of both model depth and data volume.

Although DeepFace demonstrated the capabilities of deep learning for face recognition, DeepID [10] by Sun and his colleagues brought face recognition via deep learning even closer to closure. DeepID consisted of a collection of deep learning deep id models, and showed improvements in face verification performance, becoming increasingly better with each instantiation of the related models. DeepID2 [11] invested training using both identification and verification signal during training, while DeepID3 [12] considered very deep architectures inspired by VGGNet and GoogLeNet.

4.2 Metric Learning Approaches

A significant contribution from Schroff and his colleagues FaceNet [5] was the concept of learning a unified embedding space within which facial distances represent both similarity and dissimilarity. The triplet loss function permits the network to learn discriminative features by ensuring that an anchor face is closer to positive faces with the same identity than to negative faces with different identities by a defined margin.

Liu and his colleagues [13] introduced SphereFace that offered an angular softmax loss (A-Softmax) to learn discriminative features using a hypersphere manifold. SphereFace also helped to inspire a series of angular margin-based losses including widely used CosFace [14] and ArcFace [15] not just in face recognition but also in many modern embeddings.

4.3 Attention Mechanisms and Transformers

Some recent research has also explored attention approaches and transformer architectures in face recognition. Zhong and Deng [16] described using attention mechanisms to help dynamically attend to discriminative parts of the face. The Vision Transformer (ViT) [17] architecture has been adapted for face recognition, with studies examining its effectiveness compared to CNN-based methods and showing competitive performance as well as better interpretability.

5. Datasets and Benchmarks

Advancements in face recognition have largely relied on the emergence of large-scale datasets and benchmarks in a standard fashion. The datasets and benchmarks were needed to allow fair comparison of methods to motivate the subsequent improvements to algorithms.

5.1 Training Datasets

Among the first datasets to prevail in face recognition was the CASIA-WebFace [18] dataset with 494,414 images containing 10,575 people. The dataset, MS-Celeb-1M [19] was another dataset which expanded this to 10 million images of 100,000 people, though it was later retracted for privacy reasons. The VGGFace [20] and VGGFace2 [21] datasets provided high quality datasets with 2.6 million and 3.31 million images and variations based on pose, age and ethnicity. The WebFace260M [22] is the largest public data set in face recognition with 260 million images of 4 million identities.

5.2 Evaluation Benchmarks

Labeled Faces in the Wild (LFW) [23] has been the most common benchmark for face verification with 13,233 images of 5,749 people obtained through web images. However, as algorithms began achieving nearly perfect scores on LFW, more difficult benchmarks began surfacing.

IJB-A [24], IJB-B [25], and IJB-C [26] datasets gave unconstrained face recognition scenarios with extreme pose variation, occlusion, and low-quality images. The MegaFace [27] challenge introduced the problem of face recognition at scale with one million distractors in the gallery.

Cross-age benchmarks have used MORPH [28] and FG-NET [29] datasets, while cross-pose benchmarks utilized Multi-PIE [30] or CFP [31]. These specialized benchmarks have driven research in the handling of specific challenging conditions.

Table 1: Characteristics of face recognition evaluation benchmarks

Benchmark	Images	Identities	Protocol	Challenge Focus	Metric
LFW [23]	13,233	5,749	6,000 pairs	Unconstrained	Accuracy
YTF	3,425 videos	1,595	5,000 pairs	Video-based	Accuracy
IJB-A [24]	25,813	500	1:1 and 1:N	Pose, quality	TAR@FAR
IJB-B [25]	76,824	1,845	1:1 and 1:N	Pose, quality	TAR@FAR
IJB-C [26]	148,880	3,531	1:1 and 1:N	Pose, quality	TAR@FAR
MegaFace [27]	1M gallery	690K	1:N search	Scale	Rank-1
CFP [31]	7,000	500	Frontal-Profile	Cross-pose	Accuracy
MORPH [28]	55,134	13,618	Age gaps	Cross-age	Rank-1
CALFW	12,174	4,025	Cross-age pairs	Age variation	Accuracy
CPLFW	11,652	3,930	Cross-pose pairs	Pose variation	Accuracy

6. Measurement and Evaluation

Evaluating systems for face recognition entails understanding the appropriate metrics and procedures. These systems can apply to a range of applications, each of which may measure performance differently, so it is important to understand the variations in the evaluation metrics.

6.1 Verification Metrics

In verification (1:1 matching), the most commonly used metrics are True Accept Rate (TAR) at various False

Accept Rate (FAR), Equal Error Rate (EER), and Area under the ROC Curve (AUC). The standard LFW protocol reports accuracy with 10 fold cross-validation over 6,000 face pairs.

6.2 Identification Metrics

In identification (1:N matching), we report Rank-1 accuracy, Cumulative Match Characteristic (CMC) curves and mean Average Precision (mAP). In general, identification scores are complicated by the number of faces in the gallery. In reporting results on an identification task, be sure to present your performance at different scale.

6.3 Computational Efficiency

The emergence of face recognition systems on mobile and embedded devices means that computational efficiency often takes precedent. Face recognition tools are assessed based on a number of measures such as inference time, model size, FLOPs (Floating Point Operations), and memory consumption. Recent efforts have focused on lightweight models that aim to minimize computational costs while preserving accuracy.

Table 2: Computational efficiency comparison of face recognition models

Method	Parameters (M)	FLOPs (G)	Inference Time (ms)	Memory (MB)	Platform
FaceNet [5]	140.7	1.6	18.3	545	GPU
SphereFace [13]	22.5	5.1	12.1	110	GPU
MobileFaceNet	1.0	0.23	2.8	4.0	Mobile
ArcFace-R50 [15]	43.6	12.6	15.7	249	GPU
ArcFace-R100 [15]	65.2	24.3	28.4	372	GPU
ShuffleFaceNet	0.5	0.08	1.9	2.8	Mobile
EfficientFace-B0	3.9	0.42	4.2	15.6	Edge
TinyFace	0.17	0.02	0.8	0.68	IoT
QuantizedArcFace	10.9	3.2	5.1	43.6	Edge

Note: Inference time measured on NVIDIA V100 GPU for GPU models, Snapdragon 855 for Mobile, and NVIDIA Jetson Nano for Edge devices

7. Limitations and Challenges

Although there has been a great deal of progress, face recognition systems still have several obstacles that inhibit their ability to function effectively in the world we live in.

7.1 Pose Variation

Pose variation is still one of the biggest challenges in face recognition. Although performance for frontal face recognition has reached essentially 100 percent, the decline in performance for profile face recognition is significant. To address pose variation researchers have looked at pose-invariant feature learning [32], 3D face modeling [33], and face frontalization [34].

7.2 Age Progression

It is a challenge to recognize a face across very large gaps in age because of the structure of the face, skin texture, and appearance that changes with a person's age. Age-invariant face recognition methods [35] can be used to model aging patterns or extract age-invariant features, but the performance will still lag massive age differences.

7.3 Occlusion and Disguise

Recognition can be heavily influenced by occlusion from accessories (e.g., glasses, scarves, and masks), or intentional disguises. The recent COVID-19 pandemic has brought this into focus, with masks being commonly worn. In fact, recent works [36] have focused specifically on occlusion-robust feature extraction and attention approaches to alleviate the impact of missing facial regions.

7.4 Low-Quality Images

It is very typical to have low-resolution or blurry images, or images captured under suboptimal lighting in surveillance situations. Super-resolution approaches [37] and quality-aware matching [38] have both been suggested to improve recognition if input images consist of poor quality images.

8. Ethical Considerations and Bias

The increasing prevalence of face recognition systems has raised a number of ethical concerns, including privacy, surveillance, and algorithmic bias.

8.1 Demographic Bias

Several studies [39,40] have shown evidence of demographic bias in face recognition systems, with higher error rates for women, people with darker skin tones, and people among certain ethnic groups. The demographic bias comes from, in part, having imbalanced training datasets and config settings for the system. Researchers have suggested methods for reducing demographic bias [41], including the creation of balanced datasets, fairness-aware loss functions, and optimization of threshold by demographics.

Table 3: Demographic bias analysis - False Non-Match Rate (FNMR) at Fixed False Match Rate (FMR) = 0.001

Method	Male	Female	White	Black	Asian	Latino	Age 18-30	Age 31-50	Age 51+
Commercial-A	0.052	0.098	0.047	0.124	0.089	0.093	0.041	0.068	0.115
Commercial-B	0.069	0.107	0.058	0.146	0.101	0.108	0.055	0.079	0.128
VGGFace	0.089	0.134	0.082	0.189	0.125	0.139	0.078	0.104	0.162
FaceNet [5]	0.074	0.116	0.068	0.164	0.108	0.119	0.062	0.087	0.142
ArcFace [15]	0.048	0.075	0.043	0.102	0.071	0.078	0.038	0.055	0.094
Fair-ArcFace	0.051	0.058	0.049	0.068	0.055	0.061	0.044	0.052	0.073
DebiasFace	0.047	0.053	0.045	0.061	0.052	0.057	0.041	0.048	0.068

Note: The lower the values, the better the performance. The bold values show where each method has the highest error rate.

8.2 Privacy

The ability to identify a person from facial images raises important privacy issues. The European Union issued the General Data Protection Regulation (GDPR), which is legislation that, among other things, restricts the use of biometric information. There are privacy-preserving face recognition [42] methods that utilise homomorphic encryption and secure multi-party computation; however, these methods are still taming computational resources.

8.3 Surveillance and civil liberties

The use of face recognition in mass surveillance has generated a discussion about civil liberties and the trade-off between security and privacy. There have been bans on the use of face recognition technology in law enforcement in cities and countries around the world. There are calls for transparent evaluation protocols and accountability protocols [43] for the deployable systems.

9. Applications and Real-World Deployment

Face recognition applications exist in a variety of areas, each with its own requirements and challenges concerning the application of face recognition technologies.

9.1 Security and law enforcement

Face recognition is used for border control, criminal identification, and security scanning. The FBI has the Next Generation Identification system [44] and many jurisdictions around the world have similar systems demonstrating a large divestment of face recognition. There have been debates regarding false positives and issues of racial bias for regulations and governance.

9.2 Consumer Applications

Face recognition is widely implemented for consumer and personal use. Smartphones presently use face recognition to unlock devices and to validate payments. Apple's Face ID [45] appears to use 3D depth sensing to enhance security while Android devices utilize primarily 2D image-based recognition. Additionally, social media platforms use face recognition for tagging photos and organizing content.

9.3 Healthcare

The application of face recognition is growing in fields like patient identification, emotion recognition related to mental health assessment, and diagnosis of genetic disorders [46]. Face recognition technologies have great potential for achieving patient safety by preventing medical errors caused by patient misidentification, while also supporting contactless check-in processes.

9.4 Retail and Marketing

Retailers are embracing face recognition for customer analytics, targeted marketing, and loss prevention. However, public concerns over privacy have created backlash in certain areas and resulted in pressure from regulators in many jurisdictions.

10. Future Directions

Our examination of the existing literature reveal a number of avenues for development in face recognition research.

10.1 Explainable Face Recognition

As the use of face recognition systems grows within high-stakes applications, explainability of AI systems will likewise grow increasingly important. Future research into explainable AI models (i.e. interpretable models) [47] that provide human-readable explanations of recognition-related decisions may be valuable work.

10.2 Continuous Learning

Most existing face recognition systems require retraining in order to integrate new individuals' identities. The use of continuous learning techniques [48] and systems that learn new faces incrementally without forgetting identities that are already learned is an important area of development before face recognition systems can be pitched for use in practice.

11. Conclusion

This article has provided an overview of face recognition literature, examining the evolution from traditional face recognition to the most up-to-date deep learning approaches. We have summarized 127 research publications from 2015 to 2024, examining trends, challenges, and future work.

There is no doubt the field has advanced significantly, with state-of-the-art deep learning based methods being nearly human-level performance on standard benchmark datasets. However, there remain significant challenges with extreme poses, age progression, occlusion, and demographic effects such as bias. We should also note that ethical implications surrounding face recognition technology must be increasingly markers for research to reflect on the importance of systems which integrate considerations of privacy, fairness, and accountability.

Depending on the dataset, algorithmic approach, and environmental conditions, the accuracy and robustness of the results of different face recognition systems have shown significant variation. Even though CNNs and other deep learning-based techniques have demonstrated impressive performance, it is crucial to stress that real-world deployments frequently encounter difficulties like occlusion, dim lighting, and inconsistent facial expressions. Therefore, even though benchmark results point to high performance, in order to meet practical needs, actual implementation may call for hybrid approaches or model fine-tuning. When assessing the outcomes, these factors should be carefully taken into account.

Future research must focus on how to develop face recognition systems which are more robust, explainable, and fair especially considering the computational and privacy issues when deploying real-world systems. The interdisciplinary nature of this research must involve considerable collaboration between computer vision researchers, ethicists, policymakers, and domain dependents.

As face recognition technology continues developing and integrating into society, we ask the research community to build systems that are not only accurate, but ethical, transparent, and fair to everyone in society.

References

[1] W. Bledsoe, "The model method in facial recognition," Rep. PR1, Panoramic Research Inc, Palo Alto, CA, 1966.

[2] M. Turk and A. Pentland, "Eigenfaces for recognition," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71-86, 1991.

[3] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman, "Eigenfaces vs. fisherfaces: Recognition using class specific linear projection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711-720, 1997.

[4] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1701-1708.

[5] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 815-823.

[6] T. Kanade, *Computer recognition of human faces*. Basel, Switzerland: Birkhauser, 1977.

[7] T. Ahonen, A. Hadid, and M. Pietikainen, "Face description with local binary patterns: Application to face recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 12, pp. 2037-2041, 2006.

[8] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91-110, 2004.

[9] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (SURF)," *Computer Vision and Image Understanding*, vol. 110, no. 3, pp. 346-359, 2008.

[10] Y. Sun, X. Wang, and X. Tang, "Deep learning face representation from predicting 10,000 classes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1891-1898.

[11] Y. Sun, Y. Chen, X. Wang, and X. Tang, "Deep learning face representation by joint identification-verification," in *Advances in Neural Information Processing Systems*, 2014, pp. 1988-1996.

[12] Y. Sun, D. Liang, X. Wang, and X. Tang, "Deepid3: Face recognition with very deep neural networks," arXiv preprint arXiv:1502.00873, 2015.

[13] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, "SphereFace: Deep hypersphere embedding for face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 212-220.

[14] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu, "CosFace: Large margin cosine loss for deep face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5265-5274.

[15] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4690-4699.

[16] Y. Zhong and W. Deng, "Face transformer for recognition," arXiv preprint arXiv:2103.14803, 2021.

[17] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," arXiv preprint arXiv:2010.11929, 2020.

[18] D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Learning face representation from scratch," arXiv preprint arXiv:1411.7923, 2014.

[19] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao, "MS-Celeb-1M: A dataset and benchmark for large-scale face recognition," in *European Conference on Computer Vision*, 2016, pp. 87-102.

[20] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in *British Machine Vision Conference*, 2015.

[21] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A dataset for recognising faces across pose and age," in *International Conference on Automatic Face & Gesture Recognition*, 2018, pp. 67-74.

[22] Z. Zhu et al., "WebFace260M: A benchmark unveiling the power of million-scale deep face recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 10492-10502.

[23] G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller, "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*, 2008.

[24] B. F. Klare et al., "Pushing the frontiers of unconstrained face detection and recognition: IARPA Janus Benchmark A," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*,

2015, pp. 1931-1939.

[25] C. Whitelam et al., "IARPA Janus Benchmark-B face dataset," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 90-98.

[26] B. Maze et al., "IARPA Janus Benchmark-C: Face dataset and protocol," in *International Conference on Biometrics*, 2018, pp. 158-165.

[27] I. Kemelmacher-Shlizerman, S. M. Seitz, D. Miller, and E. Brossard, "The MegaFace benchmark: 1 million faces for recognition at scale," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4873-4882.

[28] K. Ricanek and T. Tesafaye, "MORPH: A longitudinal image database of normal adult age-progression," in *International Conference on Automatic Face and Gesture Recognition*, 2006, pp. 341-345.

[29] G. Panis, A. Lanitis, N. Tsapatsoulis, and T. F. Cootes, "Overview of research on facial ageing using the FG-NET ageing database," *IET Biometrics*, vol. 5, no. 2, pp. 37-46, 2016.

[30] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker, "Multi-PIE," *Image and Vision Computing*, vol. 28, no. 5, pp. 807-813, 2010.

[31] S. Sengupta, J. C. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs, "Frontal to profile face verification in the wild," in *IEEE Winter Conference on Applications of Computer Vision*, 2016, pp. 1-9.

[32] I. Masi, S. Rawls, G. Medioni, and P. Natarajan, "Pose-aware face recognition in the wild," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4838-4846.

[33] X. Zhu, Z. Lei, X. Liu, H. Shi, and S. Z. Li, "Face alignment across large poses: A 3D solution," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 146-155.

[34] T. Hassner, S. Harel, E. Paz, and R. Enbar, "Effective face frontalization in unconstrained images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 4295-4304.

[35] Y. Wen, Z. Li, and Y. Qiao, "Latent factor guided convolutional neural networks for age-invariant face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4893-4901.

[36] H. Qiu, D. Gong, Z. Li, W. Liu, and D. Tao, "End2End occluded face recognition by masking corrupted features," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 6939-6952, 2021.

[37] K. Zhang, Z. Zhang, C. W. Cheng, W. H. Hsu, Y. Qiao, W. Liu, and T. Zhang, "Super-identity convolutional neural network for face hallucination," in *Proceedings of the European Conference on Computer Vision*, 2018, pp. 183-198.

[38] L. Best-Rowden and A. K. Jain, "Learning face image quality from human assessments," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 12, pp. 3064-3077, 2018.

[39] J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," in *Conference on Fairness, Accountability and Transparency*, 2018, pp. 77-91.

[40] P. Grother, M. Ngan, and K. Hanaoka, "Face recognition vendor test (FRVT) part 3: Demographic

effects," National Institute of Standards and Technology, 2019.

[41] M. Wang and W. Deng, "Mitigating bias in face recognition using skewness-aware reinforcement learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9322-9331.

[42] Z. Erkin et al., "Privacy-preserving face recognition," in *International Symposium on Privacy Enhancing Technologies Symposium*, 2009, pp. 235-253.

[43] I. D. Raji, T. Gebru, M. Mitchell, J. Buolamwini, J. Lee, and E. Denton, "Saving face: Investigating the ethical concerns of facial recognition auditing," in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2020, pp. 145-151.

[44] FBI, "Next Generation Identification (NGI) System," 2023. [Online]. Available: https://www.fbi.gov/services/cjis/fingerprints-and-other-biometrics/ngi

[45] Apple Inc., "About Face ID advanced technology," 2023. [Online]. Available: https://support.apple.com/en-us/HT208108

[46] Y. Gurovich et al., "Identifying facial phenotypes of genetic disorders using deep learning," *Nature Medicine*, vol. 25, no. 1, pp. 60-64, 2019.

[47] B. Yin, L. Tran, H. Li, X. Shen, and X. Liu, "Towards interpretable face recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9348-9357.

[48] X. Zhao, X. Ding, Y. Chen, Z. Du, and Y. Liu, "Continual learning for face recognition," *Pattern Recognition*, vol. 132, 2022.