IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

A COMPREHENSIVE SURVEY: DEEP LEARNING FOR FACE RECOGNITION

[1]Sujal Sunil Sawardekar, [2]Pranjal Suhas Sawant, [3]Prachi Bajirao Patil, [4]Mr. Laxman Subbu Naik
[123]Student, [4]Assitant Professor
[1234]Department of Computer Engineering,
[1234]Rajendra Mane College of Engineering & Technology, Ambav, Maharashtra, India

Abstract: The evolution of face recognition (FR) as a biometric technology has been dramatically improved by advances made with deep learning. This meta-analysis survey provides a comprehensive overview of 19 peer-reviewed research articles that can be categorized into 4 themes; A) foundational deep learning architecture(s) and evaluation, B) challenges of occlusion robustness (for example: recognition of individuals wearing a mask or sunglasses), C) a lightweight model design for deployment on an edge device, and D) methods of performing adversarial attacks on FR systems. Within each theme, this survey includes an identification of the key contributions made to the published literature, an analysis comparing different approaches, and an identification of current gaps in FR research. The results indicate that while improvements made through deep learning have elevated FR accuracy to a level that is comparable to that of humans on benchmark datasets, significant challenges remain with the application of FR systems in real-world settings, including simultaneous occlusion, resource-constrained hardware, and adversarial manipulation of the model and/or data. Future research directions will focus on developing unified multi-challenged frameworks, improving demographic fairness, and creating lightweight architectures that are aware of adversarial attacks.

Keywords: Face Recognition, Deep Learning, Convolutional Neural Network (CNN), Vision Transformer (ViT), Occlusion, MobileFaceNet, Adversarial Attacks, Masked Face, Biometric Authentication, ArcFace, Loss Functions.

I. INTRODUCTION

Face recognition (FR) is a technology that can identify or authenticate an individual based on their image (or video frame) and their characteristics. The process begins by locating a face in an image and aligning it to a standard position (pose). Next, the image is transformed into a condensed set of numerical values (dataset) that can be compared to the images stored in a database. When two sets of numerical values are considered to be similar, the identity of the person can be confirmed. Advances in FR technology have provided best practices that include: Traditional techniques such as Eigenfaces, Fisherfaces, LBP, and Gabor wavelets could provide moderate FR performance within controlled laboratory environments (limited lighting, specific camera angle: example, frontal camera face to face); however, they cannot be used for generalization in other environmental conditions (i.e., natural environments). In 2014, the introduction of Deep Learning (DL) models such as DeepFace and DeepId allowed FR to advance a new level of performance. With these new DL architectures, the accuracy of FR was improved from ~95% accuracy of traditional labs to better than 99.8% on the LFW benchmark data set, and additionally surpassed human-level recognition on the same Benchmark [1]. In summary, despite excellent performance on benchmark data sets, FR faces three (3) significant challenges for real-world deployment: (1) Face occlusion - Additionally, the performance of FR recognition for systems trained on non-occluded faces and

using images with individuals whose faces are partially covered with a face mask, sunglasses, or other occlusions is generated.

The COVID-19 pandemic created an immediate need for understanding how technology affects us (e.g., telemedicine). This is an ongoing issue as it relates to medicine, security and surveillance (all three areas will need to continuously work together in finding solutions for these types of incidents). Additionally, there are practical problems with high-performance models such as ResNet-100 and ArcFace requiring enormous amounts of memory and thousands of billions of floating-point operations (FLOPs) per image processed, making it impractical to run them on mobile or embedded devices. The third issue is the vulnerability of deep neural networks (DNN) to carefully crafted adversarial perturbations, leading to DNNs' inability to accurately classify images and then subsequently classify images completely opposite to how they should. This raises major security concerns due to the difficulty in identifying whether an image is actually what it seems.

This survey reviews eighteen papers spanning these four themes. Figure 1 illustrates the standard deep FR pipeline, and Figure 2 maps the reviewed papers to the survey taxonomy

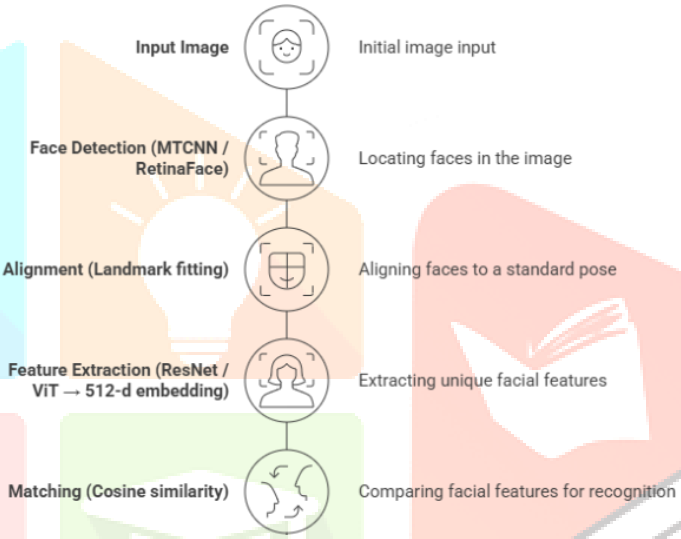

Figure 1: Standard deep face recognition pipeline common to all reviewed systems.

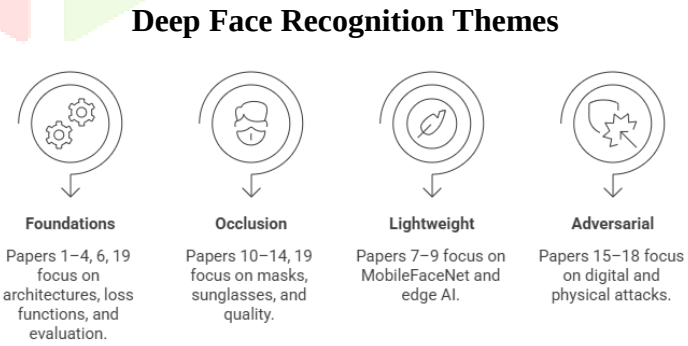

Figure 2: Taxonomy mapping all nineteen reviewed papers to their primary research theme.

II. THEME 1: FOUNDATIONAL ARCHITECTURES AND EVALUATION

A. Literature Review

[1] M. Wang and W. Deng, "Deep face recognition: A survey," Neurocomputing, vol. 429, pp. 215–244, 2021

Wang and Deng provide the most comprehensive baseline, which covers thirty-plus architectures, margin-based loss functions (SphereFace, CosFace, ArcFace), face processing pipelines, benchmarks such as lfw, youtube faces (ytf), megaface, and ijb-c. The authors identify cross-pose, cross-illumination, cross-age, cross-ethnicity, cross-modality, and low-resolution recognition as open challenges. They document that deep learning raised LFW accuracy from 95% 'shallow' methods to 99.8% plus 'deep' models.

[2] S. Kishore, P. H. G., S. C. S., and M. K. B., "Evaluation of deep learning methods in face recognition: Datasets, metrics, and results," Int. J. Sci. Technol. (IJSAT), vol. 16, no. 4, 2025.

An assessment of DeepFace, FaceNet, SphereFace, CosFace, and ArcFace on various datasets by means of the True Acceptance Rate (TAR) at fixed False Acceptance Rate (FAR) and Receiver Operation Characteristic (ROC) area. While ArcFace performs well on controlled benchmarks, nevertheless its performance gap shrinks on unconstrained datasets. The main conclusion is the choice of dataset has a massive impact on seeming model ranking – conclusions from a single benchmark are dubious.

[4] V. Bharadi, R. Sansare, T. Padelkar, and V. Shinde, "Real time face recognition system using convolutional neural network," Int. J. Creative Res. Thoughts (IJCRT), vol. 10, no. 4, Apr. 2022.

A system that refines VGG16 (ImageNet pre-trained) on twenty subjects achieves a real-time system that has 97–99% accuracy using OpenCV. The result shows that transfer learning greatly reduces data need for accurate recognition: A big pre-trained backbone can be tuned to the target domain with far fewer examples than training afresh.

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

An examination of vision transformers for face recognition In order to preserve inter-patch spatial information lost by non-overlapping token partitioning, we utilize overlapping sliding patches. Facial Transformer is trained on MS-Celeb-1M with CosFace loss, which achieves performance comparable to ResNet CNNs with similar #params on LFW, CFP-FP, and IJB-C. The results indicate that convolutional inductive biases are not strictly required as self-attention can learn similar spatial relationships from data.

[19] G. Sujatha et al., "Multi-CNN model to evaluate the performance of face detection and recognition with facial feature detection and recognition," J. Theor. Appl. Inf. Technol., vol. 103, no. 9, May 2025

A 3 stage pipeline consisting of CNN-1 for face detection, CNN-2 for facial feature detection, and CNN-3 for face recognition attains a score of 99.68% for both Face Detection and Recognition (FDR) and Facial Feature Detection and Recognition (FFDR). This intermediate step involving explicit feature detection allows better interpretability for forensic applications

B. Comparative Analysis and Gaps

Table 1 summarizes accuracy across reviewed architectures. The pattern is consistent: ArcFace leads on controlled benchmarks; margins shrink on unconstrained data. The Face Transformer [6] is the most significant architectural development, establishing that attention-based models are viable alternatives to CNNs without sacrificing accuracy.

TABLE I: ARCHITECTURE COMPARISON ON LFW AND RELATED BENCHMARKS

Model	Paper	Benchmark	Accuracy (%)
ArcFace (ResNet-100)	[1][2]	LFW	99.83
FaceNet (Inception)	[2]	LFW	99.63
CosFace (ResNet-50)	[2]	LFW	99.73
Face Transformer (ViT-S)	[6]	LFW	99.80
VGG16 + Transfer Learning	[4]	Custom (20 subjects)	97–99
Multi-CNN (3-stage)	[19]	Multiple datasets	99.68

Key gaps: (1) No paper systematically evaluates cross-demographic performance disparity. (2) Evaluation is almost universally on curated datasets, not surveillance-grade footage combining multiple simultaneous degradations. (3) The ViT vs. CNN question remains open for models larger than ViT-S.

III. THEME 2: OCCLUSION-ROBUST FACE RECOGNITION

A. Literature Review

[10] A. H. Abdul Amir and Z. N. Nemer, "Inclusive review on advances in masked human face recognition technologies," Iraqi J. Intell. Comput. Inform. (IJICI), vol. 4, no. 1, pp. 1–17, Jun. 2025.

A comprehensive review documenting that face masks obstruct the nose, mouth, and chin — features heavily relied upon by networks trained on unmasked faces — causing accuracy drops of 20–50 percentage points. Solutions catalogued include synthetic mask simulation for data augmentation, Siamese networks for similarity learning, and CNN architectures fine-tuned on masked datasets. The authors argue that health crises and security contexts will sustain demand for Masked Face Recognition (MFR) capability well beyond COVID-19.

[11] D. Zeng, R. Veldhuis, L. Spreeuwers, and R. Arendsen, "Occlusion-invariant face recognition using simultaneous segmentation," IET Biometrics, 2021.

The Simultaneous Occlusion-Invariant Deep Network (SOIDN) combines face recognition and occlusion segmentation in a single jointly-trained architecture. A face recognition branch and an occlusion segmentation branch share a CNN backbone and are connected by an occlusion mask adaptor module. Joint training with classification and segmentation losses produces: (1) occlusion-invariant features, (2) explicit occlusion maps, and (3) reliability masks that weight feature contributions by occlusion extent. SOIDN outperforms sequential approaches on both LFW-occ (synthetic) and AR (real) datasets.

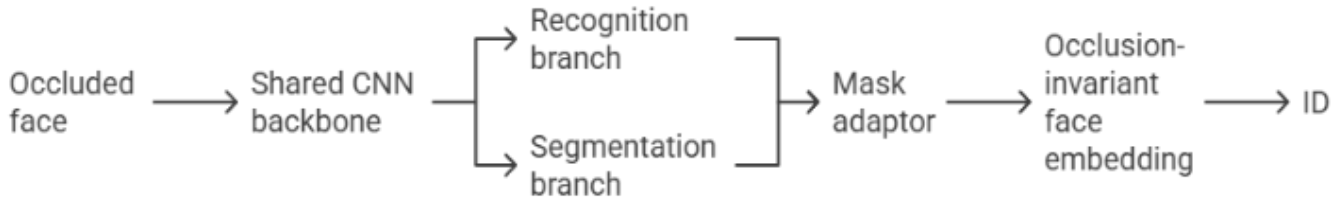

Figure 3: SOIDN architecture (Zeng et al. [11]) showing simultaneous recognition and segmentation branches connected by the occlusion mask adaptor module.

[12] M. E. Erakın, U. Demir, and H. K. Ekenel, "On recognizing occluded faces in the wild," arXiv preprint arXiv:2109.03672, Sep. 2021.

The Real-World Occluded Faces (ROF) dataset contains 6,421 neutral, 4,627 sunglasses, and 678 masked face images of 47–181 subjects collected from real-world photographs. Benchmark experiments demonstrate a critical finding: model performance on real-world occlusions is substantially worse than on synthetic occlusions of equivalent apparent severity. This performance gap means research evaluated exclusively on synthetic occlusion overestimates deployment readiness.

[13] J. Carnap, A. Kurz, O. Henniger, and A. Kuijper, "Occlusion detection for face image quality assessment," Technical University of Darmstadt / Fraunhofer IGD, Tech. Rep.

A preprocessing quality gate that quantifies the percentage of face area occluded using face segmentation and facial landmark estimation. The resulting occlusion percentage score can reject heavily occluded inputs before they enter a recognition pipeline, reducing false non-matches. The approach is architecture-agnostic and works on non-frontal faces. Evaluated on public datasets, it effectively detects both opaque sunglasses and medical masks.

[14] R. H. Khobragade et al., "Occluded face recognition using optimum features based on efficient preprocessing and machine learning," e-Prime – Adv. Electr. Eng. Electron. Energy, vol. 12, 2025.

Motivated by the data scarcity problem for occluded training samples, this paper uses contrast correction and anisotropic filtering in preprocessing, then extracts Gabor features, Linear Binary Patterns based on Haar Wavelet components, Histogram of Gaussian features, and statistical global descriptors. A Support Vector Machine (SVM) classifier performs recognition. On their evaluation datasets, the SVM system outperforms several deep learning baselines, demonstrating that traditional approaches remain competitive when labeled occluded training data is scarce.

B. Comparative Analysis and Gaps

TABLE II: OCCLUSION-ROBUST RECOGNITION METHODS

Paper	Approach	Occlusion Type	Key Contribution
Abdul Amir & Nemer [10]	Review (CNN, Siamese)	Masks	Comprehensive post-pandemic MFR review
Zeng et al. [11]	SOIDN (joint segmentation)	Masks, sunglasses	Simultaneous segmentation + recognition
Erakın et al. [12]	ROF dataset benchmark	Masks + sunglasses (real)	Real-world benchmark; performance gap proof
Carnap et al. [13]	Quality gate (segmentation)	Masks, sunglasses	Occlusion % scoring for pipeline filtering
Khobragade et al. [14]	Gabor + SVM	General	Competitive accuracy with scarce data

The most critical unaddressed gap is that no reviewed paper handles simultaneous occlusion combined with other degradations (low resolution, poor illumination). Real surveillance footage routinely combines all these factors. Additionally, dynamic occlusion is entirely absent from the literature. The ROF dataset [12] is the most significant contribution for enabling valid real-world evaluation.

IV. THEME 3: LIGHTWEIGHT AND EFFICIENT MODELS

A. Literature Review

[7] D. T. Long, "A lightweight face recognition model using convolutional neural network for monitoring students in e-learning," I.J. Mod. Educ. Comput. Sci., vol. 12, no. 6, pp. 16–28, Dec. 2020.

A custom lightweight CNN designed for continuous student monitoring during online examinations on standard laptop hardware. Rather than adapting heavy architectures (VGG, ResNet, FaceNet), the author designs from scratch with reduced filter counts, Batch Normalization, and Dropout. The model achieves competitive accuracy on public datasets and integrates with a Learning Management System (LMS). This paper establishes the design methodology of application-specific lightweight models.

[8] J. Xiao, G. Jiang, and H. Liu, "A lightweight face recognition model based on MobileFaceNet for limited computation environment," EAI Endorsed Trans. Internet Things, Feb. 2022.

Three targeted optimizations reduce MobileFaceNet from 4.9 MB to 3.4 MB while maintaining accuracy: (1) fewer layers, reducing parameter count; (2) h-ReLU6 replacing Parametric Rectified Linear Unit (PReLU) for better hardware efficiency; (3) Efficient Channel Attention (ECA) learning feature-channel importance via lightweight global average pooling instead of fully connected attention. Results: 98.52% on LFW, 97.54% on VGGFace2, 91.33% on a self-built database, at approximately 85 milliseconds per image on target embedded hardware.

[9] A. Hassanpour and Y. Kowsari, "Lightweight face recognition: An improved MobileFaceNet model," arXiv preprint arXiv:2311.15326, Nov. 2023

A detailed comparative study of MobileFaceNet and MMobileFaceNet developed for the Efficient Face Recognition (EFaR-2023) competition. Two methodological contributions: (1) training dataset curation — a carefully selected subset of Webface42M substantially outperforms training on noisier datasets; (2) Sharpness-Aware Minimization (SAM) optimization, which seeks parameters in flat loss landscape regions to improve generalization, provides significant accuracy gains over standard Adam optimization. Models ranked among the top performers in the competition's parameter-restricted categories.

B. Comparative Analysis and Gaps

TABLE III: LIGHTWEIGHT MODEL COMPARISON

Model	Paper	Size / Params	LFW Acc. (%)	Key Optimization
Custom CNN	[7]	Small (custom)	Competitive	Application-specific design
Optimized MobileFaceNet	[8]	3.4 MB	98.52	h-ReLU6 + ECA attention
MMobileFaceNet	[9]	<5M params	Top-ranked (EFaR)	SAM optimizer + data curation
ArcFace ResNet-100 (ref.)	[1][2]	~250 MB	99.83	Large-scale training

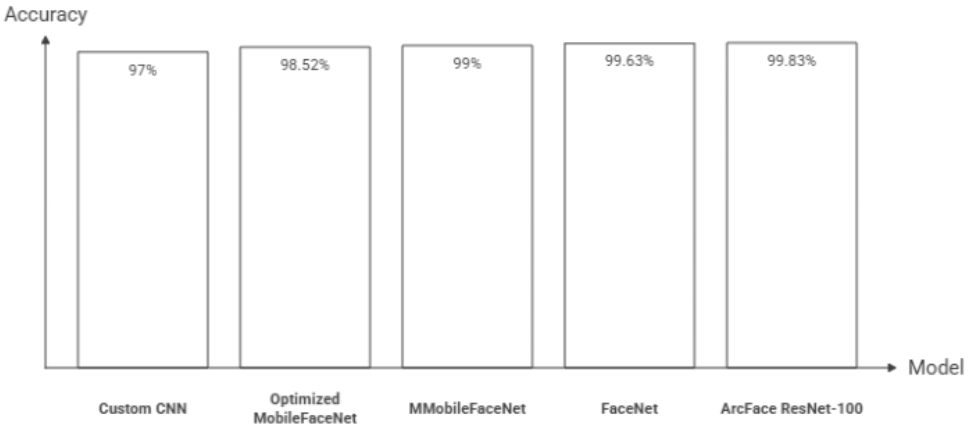

Figure 4: Trade-off between model size and LFW recognition accuracy across reviewed lightweight and large models.

Unaddressed gaps: (1) Lightweight models are evaluated only on clean benchmarks, not on masked or degraded inputs - precisely the conditions where edge-deployed models will operate. (2) Energy consumption, the actual limiting factor for battery-powered devices, is not reported. (3) On-device enrollment of new identities without cloud connectivity is unexplored.

V. THEME 4: ADVERSARIAL ATTACKS ON FACE RECOGNITION

A. Literature Review

[15] S. Hussain et al., "ReFace: Real-time adversarial attacks on face recognition systems," UC San Diego / Peraton Labs / Georgia Institute of Technology, Tech. Rep.

Most adversarial attacks require solving an optimization problem per input, making real-time deployment impractical. ReFace addresses this using an Adversarial Transformation Network (ATN) — a U-Net encoder-decoder that generates adversarial perturbations as a feed-forward function, achieving a 10,000-fold speedup over Projected Gradient Descent (PGD). A novel hybrid architecture closes the white-box accuracy gap between the original ATN and gradient-based methods. Black-box transfer results: face identification accuracy on Amazon Web Services (AWS) SearchFaces reduced from 82% to 16.4%; Azure face verification from 91% to 50.1%.

[16] F. Zhou, Q. Zhou, H. Ling, and X. Lu, "Adversarial attacks on both face recognition and face anti-spoofing models," in Proc. Int. Joint Conf. Artif. Intell. (IJCAI), 2025, pp. 2494–.

Face Anti-Spoofing (FAS) modules detect and filter manipulated images before they reach the recognition model, defeating many adversarial attacks in practice. The Reference-free Multi-level Alignment (RMA) framework crafts adversarial examples that fool both FR and FAS models simultaneously. Three components: (1) Adaptive Gradient Maintenance balances gradient contributions from both models; (2) Reference-free Intermediate Biasing improves FAS transferability; (3) Multi-level Feature Alignment reduces feature discrepancies across representation levels. RMA outperforms all state-of-the-art attacks in the joint FR+FAS setting.

[17] F. Zhou, B. Yin, H. Ling, Q. Zhou, and W. Wang, "Improving the transferability of adversarial attacks on face recognition with diverse parameters augmentation," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2024, pp. 3516–.

Adversarial examples crafted on one surrogate model often fail to transfer to different victim models. DPA addresses this through two stages: Diverse Parameters Optimization (DPO) initializes the surrogate with both pre-trained and random parameters and saves intermediate training checkpoints, creating a diverse ensemble of surrogate models; Hard Model Aggregation (HMA) enhances feature maps of the ensemble with beneficial perturbations. The result: adversarial examples that transfer more reliably across FR architectures.

[18] M. Wang, J. Zhou, T. Li, G. Meng, and K. Chen, "A survey on physical adversarial attacks against face recognition systems," arXiv preprint arXiv:2410.16317, Oct. 2024.

The first comprehensive survey of physical adversarial attacks — attacks carried out by modifying real objects (adversarial glasses, hats, makeup, stickers, infrared illumination) rather than perturbing digital pixels. Physical attacks are more practically relevant since cameras capture the real world. The survey categorizes attacks by medium: wearable patterns, makeup, and light-based methods. Defense strategies reviewed include adversarial training, input preprocessing, and detection methods. The authors document a persistent gap between laboratory demonstrations and realistic deployment conditions.

B. Comparative Analysis and Gaps

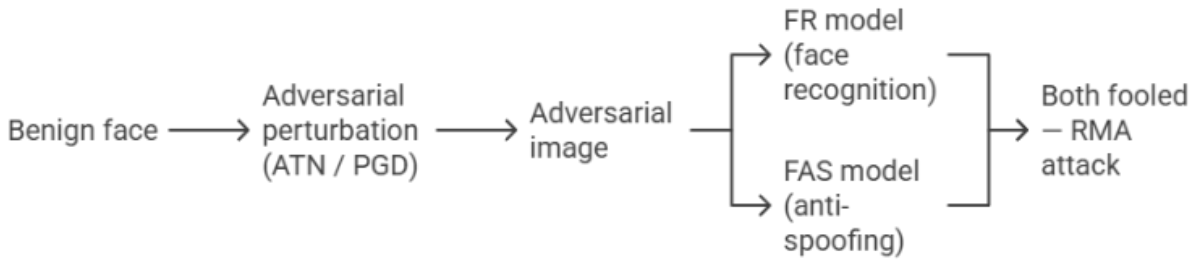

Figure 5: Adversarial attack flow showing simultaneous bypass of FR and FAS models (RMA, Zhou et al. [16]).

TABLE IV: ADVERSARIAL ATTACK METHODS AND KEY RESULTS

Limitations that previously prevented successful attacks in practice are no longer valid: real-time generation [15], bypassing FAS [16], cross-model transfer [17], and physical realization [18]. Considerably less attention is given to: (1) certifiable defenses with proven robustness guarantees; (2) adversarial vulnerability of lightweight models; and (3) attacks against computation-constrained real-world systems.

VI. CROSS-CUTTING RESEARCH GAPS

Analysis of four themes reveals five major cross-domain research gaps that present the most significant challenges for future work:

1) No coherent multi-challenge framework. Every paper presents results in isolation. No system simultaneously optimizes for occlusion robustness, computational efficiency, and adversarial security, and no evaluation benchmark incorporates all three criteria.

2) Synthetic-to-real performance gap. Erakın et al. [12] demonstrated that synthetic occlusion benchmarks are significantly more optimistic than real-world testing. This applies to adversarial attacks (digital perturbations vs. physical modifications) and lightweight models (clean datasets vs. noisy real-world inputs).

3) Demographic fairness is unaddressed. No paper in this review systematically evaluates performance differences across gender, age group, or skin tone, despite this being a critical concern for equitable deployment.

Paper	Attack Type	Setting	Key Result
Hussain et al. [15]	ATN real-time	Black-box (commercial)	AWS: 82→16.4%; Azure: 91→50.1%
Zhou et al. [16]	RMA (FR + FAS bypass)	Black-box transfer	State-of-the-art joint FR+FAS fooling
Zhou et al. [17]	DPA transfer	Black-box transfer	Best transferability across victim models
Wang et al. [18]	Physical (wearables, makeup, light)	Physical world	Survey of physical attack evolution

4) Lightweight model adversarial robustness is unknown. The intersection between Themes 3 and 4 whether depth-wise separable convolutions, ECA attention, and parameter pruning affect adversarial vulnerability - has not been evaluated.

5) Privacy-preserving deployment is absent. None of the reviewed studies address approaches such as federated learning, differential privacy, or on-device processing that would enable deployment under privacy regulations governing biometric facial embeddings

.

VII. CONCLUSION

A. Summary

This survey has reviewed nineteen research papers across four themes of deep learning-based face recognition.

Foundational architectures: Architecture choice matters less than training data scale, loss function design, and evaluation methodology. ArcFace leads on controlled benchmarks; Face Transformer demonstrates that self-attention is a competitive alternative to convolution; transfer learning enables accurate recognition with limited data.

Occlusion robustness: The COVID-19 pandemic stimulated wide-ranging investigation of masked face recognition. The most advanced solution is the SOIDN joint segmentation-recognition method [11]. The ROF dataset [12] demonstrated that real occlusion is harder than synthetic. Quality-based rejection [13] offers an architecture-agnostic defense.

Lightweight models: The accuracy gap between large and small models is substantially reduced. Optimized MobileFaceNet [8] achieves 98.52% at 3.4 MB. The insight from [9] that training methodology can compensate for architectural limitations applies broadly. Adversarial security: Face recognition systems are more vulnerable than benchmark accuracies suggest. ReFace [15] enables real-time commercial attacks; RMA [16] bypasses anti-spoofing defenses; DPA [17] improves cross-model transferability; physical attacks [18] pose real-world threats. Currently, no certified defenses exist for face recognition systems.

B. Future Scope and Directions

Six key areas require immediate research attention:

First, unified multi-task systems optimized jointly for occlusion robustness, edge efficiency, and adversarial security. Multi-task architectures are already established for joint occlusion and recognition.

Second, benchmark datasets with combined simultaneous real-world challenges. The ROF dataset covers only the occlusion dimension; datasets combining multiple simultaneous degradations are needed.

Third, research on the intersection of model compression and adversarial vulnerability, designing lightweight models exhibiting both high accuracy and security.

Fourth, demographic fairness across all evaluations, with disaggregated performance reporting and fairness-aware loss functions that explicitly penalize demographic group performance inequalities.

Fifth, privacy-preserving architectures — federated learning, differential privacy, and on-device processing — are essential for privacy-sensitive deployment contexts.

Sixth, certified adversarial defenses with provable robustness bounds for a given perturbation budget must be extended from image classification to the face recognition setting, accounting for embedding space geometry.

The overall trajectory of deep face recognition is one of significant growth in technology maturity and global deployment. The research community's collective obligation is to guarantee that deployment occurs equitably, with adequate attention to privacy and security as well as accuracy.

ACKNOWLEDGMENT

The authors would like to thank Mr. Laxman S. Naik, Assistant Professor, Department of Computer Engineering, PSPS's Rajendra Mane College of Engineering & Technology, Ambav, for his guidance and support throughout this work.

REFERENCES

[1] M. Wang and W. Deng, "Deep face recognition: A survey," *Neurocomputing*, vol. 429, pp. 215–244, 2021.

[2] S. Kishore, P. H. G., S. C. S., and M. K. B., "Evaluation of deep learning methods in face recognition: Datasets, metrics, and results," *Int. J. Sci. Technol. (IJSAT)*, vol. 16, no. 4, 2025.

[3] A. M. P., S. N., and M. K., "Face recognition using convolutional neural network: A systematic review," *Int. J. Eng. Res. Technol. (IJERT)*, vol. 11, no. 06, Jun. 2022.

[4] V. Bharadi, R. Sansare, T. Padelkar, and V. Shinde, "Real time face recognition system using convolutional neural network," *Int. J. Creative Res. Thoughts (IJCRT)*, vol. 10, no. 4, Apr. 2022.

[5] Z. Sun and G. Tzimiropoulos, "Part-based face recognition with vision transformers," in *Proc. 33rd British Machine Vision Conf. (BMVC)*, London, UK, Nov. 2022, BMVA Press.

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

[7] D. T. Long, "A lightweight face recognition model using convolutional neural network for monitoring students in e-learning," *I.J. Mod. Educ. Comput. Sci.*, vol. 12, no. 6, pp. 16–28, Dec. 2020.

[8] J. Xiao, G. Jiang, and H. Liu, "A lightweight face recognition model based on MobileFaceNet for limited computation environment," *EAI Endorsed Trans. Internet Things*, Feb. 2022.

[9] A. Hassanpour and Y. Kowsari, "Lightweight face recognition: An improved MobileFaceNet model," *arXiv preprint arXiv:2311.15326*, Nov. 2023.

[10] A. H. Abdul Amir and Z. N. Nemer, "Inclusive review on advances in masked human face recognition technologies," *Iraqi J. Intell. Comput. Inform. (IJICI)*, vol. 4, no. 1, pp. 1–17, Jun. 2025.

[11] D. Zeng, R. Veldhuis, L. Spreeuwers, and R. Arendsen, "Occlusion-invariant face recognition using simultaneous segmentation," *IET Biometrics*, 2021.

[12] M. E. Erakın, U. Demir, and H. K. Ekenel, "On recognizing occluded faces in the wild," *arXiv preprint arXiv:2109.03672*, Sep. 2021.

[13] J. Carnap, A. Kurz, O. Henniger, and A. Kuijper, "Occlusion detection for face image quality assessment," Technical University of Darmstadt / Fraunhofer IGD, Tech. Rep.

[14] R. H. Khobragade *et al.*, "Occluded face recognition using optimum features based on efficient preprocessing and machine learning," *e-Prime – Adv. Electr. Eng. Electron. Energy*, vol. 12, 2025.

[15] S. Hussain *et al.*, "ReFace: Real-time adversarial attacks on face recognition systems," UC San Diego / Peraton Labs / Georgia Institute of Technology, Tech. Rep.

[16] F. Zhou, Q. Zhou, H. Ling, and X. Lu, "Adversarial attacks on both face recognition and face anti-spoofing models," in *Proc. Int. Joint Conf. Artif. Intell. (IJCAI)*, 2025, pp. 2494–.

[17] F. Zhou, B. Yin, H. Ling, Q. Zhou, and W. Wang, "Improving the transferability of adversarial attacks on face recognition with diverse parameters augmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2024, pp. 3516–.

[18] M. Wang, J. Zhou, T. Li, G. Meng, and K. Chen, "A survey on physical adversarial attacks against face recognition systems," *arXiv preprint arXiv:2410.16317*, Oct. 2024.

[19] G. Sujatha *et al.*, "Multi-CNN model to evaluate the performance of face detection and recognition with facial feature detection and recognition," *J. Theor. Appl. Inf. Technol.*, vol. 103, no. 9, May 2025.