IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510129

A Survey on Privacy-Preserving Face Recognition: Techniques, Challenges and Future Directions

Gopika S Nair[1], Sebin Sibichan[2], Sri Sai G Shyam[3], Stanly Antony V W[4], Sumi Rehim[5], Dr. Shani Raj[6]

Student, Department of Computer Science and Engineering, College of Engineering Kottarakkara, Kottarakkara, Kerala, India[1-4]

Assistant Professor, Department of Computer Science and Engineering, College of Engineering Kottarakkara, Kottarakkara, Kerala, India[5]

Associate Professor, Department of Computer Science and Engineering, College of Engineering Kottarakkara, Kottarakkara, Kerala, India[6]

Abstract: Face recognition (FR) is now widely deployed across authentication, surveillance, and identity management systems. However, because facial data is inherently sensitive, storing, transmitting, and processing it introduces severe privacy and security risks. Cloud-hosted FR architectures, in particular, leave raw facial images and deep feature embeddings vulnerable to data breaches, adversarial manipulation, and unauthorized inference. To address these vulnerabilities, this survey delivers a systematic review of privacy-preserving face recognition (PPFR) techniques published between 2019 and 2026. Analyzing primary studies alongside more than related works, we organize current research into a five-part taxonomy: cryptographic approaches (fully homomorphic, selective, and searchable encryption), distributed and perturbation methods (federated learning and GAN-based anonymization), hardware-level protection (lensless imaging), adversarial robustness (attack and defense methods), and domain-specific applications (blockchain, IoT, and e-voting). For each category, we unpack the core methodology, benchmark datasets, reported accuracy, and the operational tradeoffs among privacy, accuracy, and computational overhead. Our analysis highlights key trade-offs: homomorphic encryption retains near-perfect recognition accuracy but incurs high computational overhead; federated learning and lensless imaging deliver stronger privacy at a moderate cost to accuracy; and adversarial robustness remains a widespread open gap across most PPFR designs. We conclude by outlining key open challenges and future research directions.

Keywords: Privacy-Preserving Face Recognition, Homomorphic Encryption, Federated Learning, Adversarial Robustness, Biometric Template Protection, Lensless Imaging.

I. INTRODUCTION

Face recognition (FR) has emerged as one of the most widely deployed biometric modalities, powering everyday applications from smartphone unlocking and digital payments to border security and large-scale identity verification. Its rapid adoption stems largely from its non-invasive nature and seamless user experience. Modern deep learning architectures, such as FaceNet [1] and ArcFace [2], have elevated verification accuracy beyond 99% on controlled benchmarks, making large-scale deployment practical and effective. However, the high-dimensional embeddings that drive this accuracy introduce a severe privacy risk: these feature vectors can be inverted to reconstruct identifiable face images or probed to extract sensitive demographic attributes like age, gender, and ethnicity [3].

Deploying FR on cloud infrastructure exacerbates these security concerns, as raw images and embeddings are stored and processed on third-party servers outside the user's direct control. Biometric data breaches have already compromised millions of records—a critical threat because, unlike passwords, biometric credentials cannot be reset or reissued once stolen. These vulnerabilities highlight the urgent need for privacy-preserving face recognition (PPFR), a domain focused on safeguarding facial data across storage, transmission, and computation while preserving high recognition utility [27].

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510129

Most existing surveys on face recognition focus heavily on model architectures and accuracy optimization, treating privacy considerations as a secondary priority. Although the PPFR field has expanded rapidly, research remains siloed across isolated domains, including cryptographic protocols, distributed learning, image anonymization, specialized camera hardware, and adversarial robustness. Because each domain relies on different evaluation metrics and datasets, existing literature lacks a unified benchmark for comparison. This survey addresses that disconnect by systematically analyzing PPFR literature published between 2019 and 2026. We consolidate these diverse research streams into a clear taxonomy and evaluate representative methods side-by-side based on accuracy, privacy guarantees, computational overhead, and experimental datasets.

We establish a comprehensive, five-branch taxonomy of privacy-preserving face recognition techniques, categorizing methods into cryptographic, distributed, hardware-based, adversarial, and application-specific approaches. We systematically analyze key primary studies across this taxonomy, evaluating each framework's underlying methodology, experimental datasets, and quantitative performance metrics. We highlight persistent open challenges and future research directions, focusing on the core privacy-accuracy-complexity trade-off and the widespread lack of adversarial robustness in current PPFR systems.

The remainder of this paper is structured as follows: Section 2 outlines our survey methodology, and Section 3 details the proposed taxonomy. Section 4 reviews the literature across each taxonomy branch. Section 5 provides a comparative performance analysis, Section 6 explores open challenges and future research directions, and Section 7 presents our concluding remarks.

II. TAXONOMY OF PRIVACY-PRESERVING FACE RECOGNITION

Fig. 1 illustrates our proposed taxonomy, which categorizes current PPFR research into five distinct operational branches. The cryptographic branch focuses on performing secure computations over encrypted representations; the distributed and perturbation branch retains raw facial data locally or applies irreversible transformations; the hardware-level branch prevents identifiable optical images from ever being captured at the sensor level; the adversarial branch fortifies recognition models against malicious inputs; and the application branch integrates PPFR mechanisms into functional, end-to-end deployment architectures.

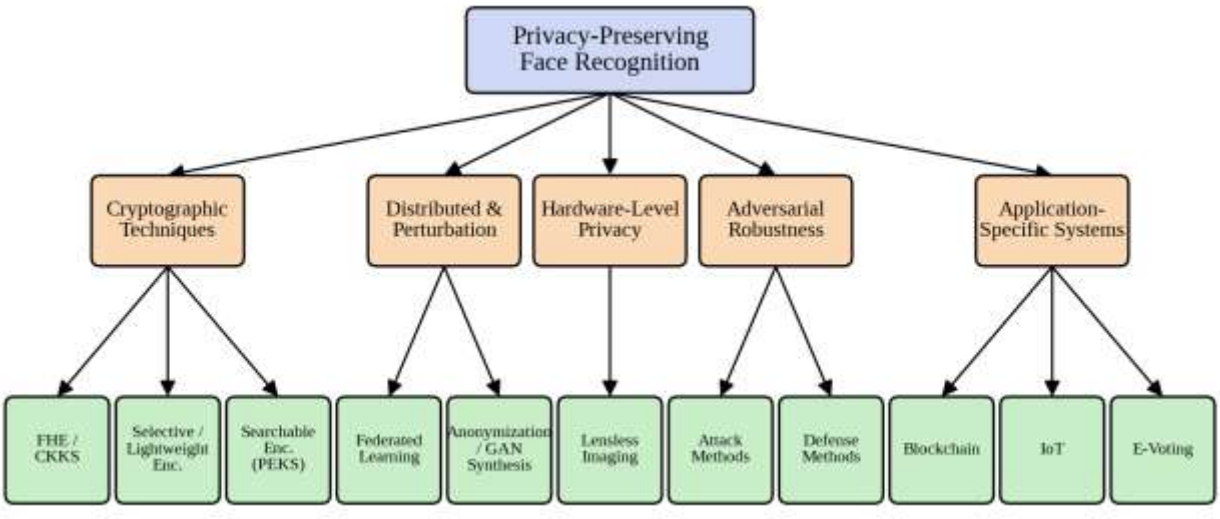


Fig. 1. Taxonomy of privacy-preserving face recognition techniques.

The cryptographic methods encrypt facial features or templates so that matching can be performed without decryption. They include fully homomorphic encryption (FHE) [24], which supports arbitrary arithmetic on ciphertexts; selective and lightweight encryption, which encrypt only the sensitive portion of the data or use efficient symmetric schemes; and searchable encryption, which enables fast preselection in the encrypted domain.

The distributed and perturbation techniques are used to avoid centralizing raw facial data. Federated learning trains models collaboratively while keeping images on user devices, while anonymization and GAN-based synthesis replace or perturb identity-bearing features before any processing.

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026
DOI: 10.17148/IJARCCE.2026.1510129

Lensless cameras encode light directly onto a sensor without forming a recognizable image, providing privacy at the physical acquisition layer rather than through software post-processing.
Attack methods craft imperceptible perturbations that fool FR models, while defense methods detect or neutralize such attacks. This branch addresses the security of FR systems against malicious inputs rather than data confidentiality.
Application-specific works integrate PPFR into complete systems such as blockchain-based e-voting, IoT security, and embedded platforms, evaluating end-to-end performance.

III. LITERATURE REVIEW

A. Fully Homomorphic Encryption (FHE)

Chitrapu et al. propose a PPFR framework combining FaceNet and ArcFace embeddings with Locality Preserving Projection (LPP) for dimensionality reduction from 512-D to 256-D, followed by CKKS homomorphic encryption [23] for encrypted feature matching [4]. Evaluated on CASIA-3D, 105PinsFace, LFW, and Faces94, ArcFace achieves 99.30%, 98.06%, 97.07%, and 99.82% accuracy respectively, and an ablation study confirms the equal error rate (EER) remains effectively unchanged after encryption, showing that CKKS-FHE preserves the discriminative structure of deep embeddings. The main limitation is the computational cost of encrypted arithmetic, which the LPP compression mitigates.
A doorbell-oriented system, Hybrid-HE LLE, masks raw embeddings with a non-invertible sparse transform, encrypts only the sensitive coefficients with CKKS, and offloads heavy Locally Linear Embedding operations to an on-premise hub [5]. On ORL and LFW it preserves more than 99% of baseline accuracy while cutting client CPU time by 92%, reducing uplink traffic fivefold, and keeping door-unlock latency below 120 ms, with formal IND-CPA, EUF-CMA, and optional IND-CCA guarantees. Its selective-encryption design is a notable step toward practical latency.

A third HE-based work proposes dual-layer similarity evaluation that incorporates both facial landmarks and color information during encrypted matching, addressing two limitations of prior HE methods: the risk of original image leakage and the neglect of color cues that affect accuracy [6]. A softmax-based formula fuses the two similarity layers. The work is methodologically strong but does not report results on a standard public benchmark.

HERS (Homomorphically Encrypted Representation Search) compresses a representation toward its intrinsic dimensionality, encrypts it with a custom FHE scheme, and searches an encrypted gallery directly [7]. On face, fingerprint, and ImageNet galleries it demonstrates, for the first time, accurate encrypted search at scale, achieving a 275× speedup over prior encrypted search against a 100-million gallery in about 500 seconds. This establishes the scalability ceiling for FHE-based search.

B. Searchable Encryption and Template Protection

A preselection scheme combines public-key encryption with keyword search (PEKS) for fast candidate preselection with FHE for long-term template protection, both lattice-based and hence post-quantum [8]. The encrypted identification workload is reduced to 8.4% of exhaustive search, identifying 1062 subjects in 210 ms while preserving the accuracy of the unprotected system.

A multimodal system fuses iris and face feature vectors and performs matching on encrypted templates using FHE on the QFIRE-I database [9]. It achieves 96.41% TAR for iris alone, 81.19% TAR for face alone, and 100% TAR at 0.1% FAR for face and iris fusion, demonstrating that multimodal fusion compensates for the weaker face-only performance.
A retina-based framework, “For Your Eyes Only,” combines retina feature extraction with template protection and homomorphic encryption, aligned with the ISO/IEC 24745:2022 standard [10], [26]. On the RIDB database it achieves a genuine acceptance rate above 99% at FAR = 10^{-6}, indicating strong robustness in security-sensitive applications.

C. Lightweight and Novel Encryption

A novel approach encrypts Dlib facial encodings using Internet Protocol addresses as a dynamic key generation source, producing a unique key per communication [11]. On a pins_dataset of 10,000 images of 100 identities it achieves 90.1% accuracy, identical to the unencrypted baseline, showing that encryption can be added without accuracy loss, though the accuracy itself is moderate.

For resource-constrained edge devices, a study compares AES-Light, SPECK, and SIMON on ARM Cortex-M microcontrollers for IoT security [12]. SPECK is 15% faster than AES-Light and SIMON reduces power consumption by 20%, establishing preferred lightweight schemes for different IoT scenarios. This work is not FR-specific but informs the embedded deployment of FR encryption.

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026
DOI: 10.17148/IJARCCE.2026.1510129

D. Federated Learning

Woubie et al. apply federated learning to supervised and unsupervised face recognition on the CelebA dataset, generating imposter data with GANs [25] at the edge so raw images never leave user devices [13]. The federated model without a secure aggregator achieves a superior average EER compared to individual local models, whereas the model with a secure aggregator yields slightly less favorable EER, quantifying the privacy-accuracy trade-off. The study is limited to 1000 simulated devices and simple averaging aggregation.

E. Anonymization and Synthetic Data

A real-time face and hair swapping method uses a GAN-generated synthetic face rather than a reference image of a real person, concealing identity while preserving essential facial attributes for video interviews and teleconferencing [14]. The three-phase pipeline runs in real time from a standard camera and is evaluated on CelebA-HQ and LFW. Because it produces a synthetic identity, it is suited to anonymization rather than to verification tasks that require the true identity.

A masked face recognition method augments only the lower part of masked faces, detecting the occlusion boundary and matching the visible upper part with a pre-trained CNN, then applying SURF-based geometric refinement to generate the lower face [15]. On LFW, CASIA-WebFace, AR, and FACES it improves recognition accuracy by 4% to 6% and significantly raises SSIM while reducing runtime and computational cost relative to full-face inpainting methods.

F. Hardware-Level Privacy (Lensless Imaging)

Cai et al. present an end-to-end optimized lensless system that verifies faces directly on encoded sensor captures, avoiding the reconstruction step that leaks recognizable faces in conventional lensless pipelines [16]. The optimized point spread function is designed to be effective for verification but poor for general image reconstruction, enhancing privacy; under an adversarial U-Net reconstruction attack the system shows lower reconstruction quality than FlatCam and PhlatCam designs. It achieves 96.78% aligned verification accuracy and 95.42% on LFW under a random protocol, narrowing the gap to the RGB upper bound. The approach provides hardware-level privacy but requires specialized imaging hardware.

G. Attack Methods

Adv-Eye constructs natural eye makeup attacks using a makeup generation module, a multi-view visual similarity blending module, and an ensemble attack module based on fine-grained meta-learning and input decay [17]. On the LADN and MT datasets it improves attack success rates by 1.63% and 1.05% over baselines on local black-box models and raises confidence scores by 5.33 and 5.22 points on an online commercial FR platform, while improving visual quality.

An invisible adversarial attack method uses facial landmark detection and superpixel segmentation to confine perturbations to the eyebrows, eyes, nose, and mouth regions, producing samples with high SSIM that are indistinguishable to the naked eye [18]. Printed adversarial images presented through a camera still fool the FR model, demonstrating real-world physical attack capability.

A study of adversarial robustness against preprocessing shows that the choice of face detection model can degrade attack success rate by up to 78% in black-box settings, while the interpolation method during downsampling has minimal impact [19]. A preprocessing-invariant method using input transformations improves attack transferability by up to 27%, highlighting preprocessing as a critical and often overlooked component of FR adversarial robustness.

H. Defense Methods

3FG-Defense detects adversarial faces by modeling cross-domain consistency across geometric (3DMM coefficients), frequency (FFT statistics), and structural (landmark graphs) representations, projecting them to a common space and computing pairwise cosine similarity [20]. On CelebA-HQ, CASIA-WebFace, LFW, and IJB-C it reaches 91.8% AUC with 72.4% TPR at 1% FPR, outperforming single-domain and generic multi-domain baselines, and is about two orders of magnitude faster than diffusion-based purification. Performance degrades under fully adaptive detector-aware attacks, an acknowledged limitation.

I. Privacy Validation

Zhang et al. validate the privacy of state-of-the-art decompose-and-remove (DnR) PPFR systems under a minimum 1k1c assumption, where the adversary has only one image of one identity, using a Map2V reconstruction attack [21]. Reconstructed faces from DuetFace, DCTDP, and PartialFace fool both target systems and human observers, with a subjective privacy score of 93.97% and attribute matching rates above 70% on LFW, showing that current PPFR systems remain vulnerable to reconstruction.

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510129

J. Application-Specific Systems

A blockchain-integrated multimodal biometric e-voting system combines facial and fingerprint authentication with deep-learning image processing on an ARM Cortex-A76 processor, using Ethereum blockchain with homomorphic encryption for immutable, tamper-resistant records [22]. It achieves 1424.501 transactions per second for 10,000 transactions with a 7.02-second mining time, and biometric authentication with 0.00% FAR and FRR and 100% TAR. This demonstrates an end-to-end deployment where privacy-preserving biometrics meets distributed ledger integrity.

IV. COMPARATIVE ANALYSIS

Table I compares the reviewed works on approach, technique, dataset, key result, strengths, and limitations. Fig. 2 shows a generic PPFR system block diagram synthesized across the reviewed works: a client-side capture and feature-extraction stage, a protection stage (encryption, perturbation, or anonymization), an untrusted transmission/storage stage, and a server-side matching stage that never observes raw facial data.

Across the corpus, three consistent patterns emerge. First, cryptographic methods (rows 1–9) achieve the highest fidelity to the unprotected baseline, frequently reporting accuracy or EER that is statistically indistinguishable from the plaintext system, at the cost of encrypted-arithmetic latency that ranges from subsecond (selective encryption, searchable preselection) to hundreds of seconds at web scale (full FHE search). Second, distributed and hardware approaches (rows 10–13) trade a modest accuracy reduction for structural privacy guarantees that do not depend on any cryptographic assumption, which is attractive against future decryption threats such as quantum computing. Third, the adversarial-robustness branch (rows 14–18) is comparatively immature: attack methods (Adv-Eye, invisible perturbation) are more numerous and more thoroughly validated than defenses (3FG-Defense), and the one privacy-validation study in the corpus [21] shows that even purpose-built PPFR systems can be reconstructed under realistic single-image adversary assumptions.

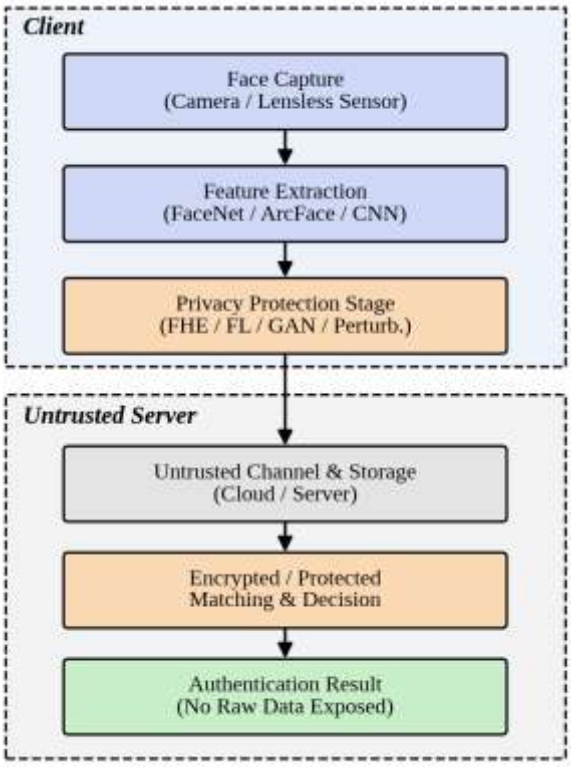


Fig. 2. Generic system block diagram synthesized across the reviewed PPFR architectures.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510129

TABLE I COMPARATIVE ANALYSIS OF PRIVACY-PRESERVING FACE RECOGNITION TECHNIQUES

Ref	Approach	Technique	Dataset	Key Result	Strengths	Limitations
[4]	FHE + LPP	FaceNet/ArcFace + LPP + CKKS	CASIA-3D, 105PinsFace, LFW, Faces94	ArcFace 97.07–99.82%; EER unchanged	High accuracy, privacy preserved	FHE compute cost
[5]	Selective HE	Sparse mask + CKKS + LLE offload	ORL, LFW	>99% accuracy; <120 ms latency	Low latency, verifiable	Split-edge complexity
[6]	HE + fusion	Landmark + color dual-layer similarity	Not specified	Improved similarity accuracy	Handles color, no leakage	No public benchmark
[7]	FHE search	Intrinsic-dim compression + FHE	Face, fingerprint, ImageNet	275× speedup, 500 s / 100M gallery	Scalable, exact matching	High setup cost
[8]	PEKS + FHE	Keyword-search preselection + FHE	Custom (1062 subjects)	8.4% of exhaustive search; 210 ms	Post-quantum, fast preselection	Two-stage complexity
[9]	HE fusion	Multimodal iris+face FHE matching	QFIRE-I	100% TAR @ 0.1% FAR (fused)	Multimodal robustness	Iris capture dependency
[10]	HE + template protection	Retina features, ISO/IEC 24745-aligned	RIDB	GAR >99% @ FAR = 10^{-6}	Standard-compliant, high security	Requires retina imaging
[11]	Lightweight encryption	Dynamic IP-based key + Dlib encoding	pins_dataset (10k img, 100 id)	90.1% accuracy (= baseline)	No accuracy loss from encryption	Moderate baseline accuracy
[12]	Lightweight crypto	AES-Light vs. SPECK vs. SIMON	N/A (ARM Cortex-M)	SPECK +15% speed; SIMON −20% power	Embedded suitability	Not FR-specific
[13]	Federated learning	FL + GAN-generated imposters	CelebA	Better avg. EER w/o secure aggregator	No raw data leaves device	Aggregator trade-off; 1000 sim. devices
[14]	Anonymization	GAN face/hair swap (3-phase, real-time)	CelebA-HQ, LFW	Real-time synthetic identity swap	Real-time, preserves attributes	Not suited to verification
[15]	Perturbation-tolerant matching	CNN + SURF lower-face generation	LFW, CASIA-WebFace, AR, FACES	+4–6% accuracy; higher SSIM	Faster than inpainting	Limited to masked-face scenario
[16]	Hardware privacy	Optimized-PSF lensless verification	Custom lensless, LFW	96.78% aligned; 95.42% LFW (random)	Hardware-level privacy	Requires specialized hardware
[17]	Adversarial attack	Makeup-based ensemble attack	LADN, MT	+1.63%/1.05% ASR; +5.33/5.22 conf.	Natural-looking, transferable	Attack, not defense
[18]	Adversarial attack	Landmark + superpixel-confined perturb.	Not specified	High SSIM; physical-world attack success	Imperceptible, physically realizable	Demonstrates risk only
[19]	Adversarial analysis	Detector / interpolation sensitivity study	Not specified	Detector choice cuts ASR by up to 78%	Reveals preprocessing vulnerability	No new defense proposed

 This work is licensed under a Creative Commons Attribution 4.0 International License

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510129

Ref	Approach	Technique	Dataset	Key Result	Strengths	Limitations
[20]	Adversarial defense	Geometry+frequency+structure fusion	CelebA-HQ, CASIA-WebFace, LFW, IJB-C	91.8% AUC; 72.4% TPR @ 1% FPR	Fast, multi-domain detection	Degrades under adaptive attacks
[21]	Privacy validation attack	Map2V reconstruction under 1k1c	LFW (DuetFace, DCTDP, PartialFace)	93.97% privacy score; >70% attr. match	Exposes DnR-PPFR vulnerability	Attack-only, no mitigation proposed
[22]	Application system	Face+fingerprint + HE + Ethereum	Custom (10,000 tx)	1424.5 TPS; 0% FAR/FRR; 100% TAR	End-to-end deployment	Blockchain overhead

V. CHALLENGES AND FUTURE DIRECTIONS

A. Privacy-Accuracy-Complexity Trade-off

Every technique in our taxonomy navigates an inevitable trade-off among recognition accuracy, computational overhead, and communication costs. For instance, FHE-based protocols maintain near-perfect accuracy but demand massive computational resources, whereas federated learning and anonymization techniques reduce this burden at the expense of a measurable drop in accuracy. Because no single approach excels across all three dimensions—and few studies report all three metrics on a unified benchmark—cross-method comparisons remain difficult.

B. Vulnerability of PPFR Systems to Reconstruction

The recent Map2V study [21] exposes a critical vulnerability: decompose-and-remove (DnR) PPFR architectures, despite their explicitly privacy-centric designs, can still be inverted from just a single protected sample. This reveals that theoretical privacy claims, when validated solely against a system's internal threat model, often fail against motivated adversaries operating under realistic assumptions. Moving forward, rigorous reconstruction-attack evaluations must become a mandatory component of PPFR benchmarking.

C. Underexplored Adversarial Robustness

Currently, most PPFR architectures prioritize data confidentiality while ignoring adversarial vulnerabilities. Conversely, literature focused on adversarial attacks and defenses rarely considers underlying data privacy. Because very few studies address both confidentiality and adversarial resilience simultaneously, the intersection of these dual threats remains a largely unexplored gap in the literature.

D. Standardization and Benchmarking

The field suffers from highly fragmented evaluation practices. Across the reviewed literature, datasets, performance metrics (such as accuracy, EER, TAR/FAR, and AUC), and threat models vary wildly, making direct comparisons nearly impossible. Only a single system in our corpus explicitly conforms to an international standard (ISO/IEC 24745:2022) [10], [26]. Broadly adopting standardized biometric template protection protocols is essential to ensure reproducibility and fair evaluation.

E. Deployability on Constrained Devices

While hardware-level innovations (like lensless imaging) and lightweight cryptographic ciphers (such as SPECK and SIMON) prove that PPFR can operate on resource-constrained edge devices, holistic solutions remain rare. Fusing hardware-level optical protections with robust cryptography into a single, cohesive pipeline—without violating the low-latency requirements of real-time authentication—stands as a major open engineering challenge.

F. Future Research Directions

Developing holistic evaluation frameworks that concurrently measure recognition accuracy, data leakage under reconstruction attacks, and adversarial resilience. Engineering hybrid architectures that merge FHE-based matching with federated learning or hardware-level safeguards to better navigate the accuracy-privacy-latency trade-off. Advancing post-quantum-secure PPFR ecosystems by expanding current lattice-based searchable encryption from isolated components into comprehensive end-to-end recognition pipelines. Establishing standardized, publicly accessible PPFR benchmarks that feature rigorous reconstruction and adversarial attack test suites, moving beyond simple clean-accuracy metrics.



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026
DOI: 10.17148/IJARCCE.2026.1510129

VI. CONCLUSION

This survey systematically reviewed primary studies on privacy-preserving face recognition, mapping the current landscape across a five-branch taxonomy: cryptographic protocols, distributed and perturbation methods, hardware-level privacy, adversarial robustness, and application-specific architectures. Our comparative analysis highlights a persistent divide: homomorphic encryption delivers the highest fidelity to unprotected baseline accuracy but imposes severe computational penalties, whereas federated learning and lensless imaging provide robust structural privacy at the cost of moderate accuracy degradation. Meanwhile, adversarial robustness remains the least mature and least integrated aspect of PPFR system design. Furthermore, recent breakthroughs in reconstruction attacks underscore that theoretical privacy guarantees often collapse under realistic, single-image adversarial conditions.

To advance the field, future PPFR research must prioritize joint evaluations of accuracy, reconstruction resilience, and adversarial robustness using standardized, post-quantum-aware benchmarks. Ultimately, the next generation of secure biometric systems will demand hybrid architectures that strategically synthesize the complementary strengths of the five branches identified in this taxonomy.

REFERENCES

[1] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 815–823.
[2] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 4690–4699.
[3] G. Mai, K. Cao, P. C. Yuen, and A. K. Jain, "On the reconstruction of face images from deep face templates," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 41, no. 5, pp. 1188–1202, 2018.
[4] Chitrapu et al., "A privacy-preserving face recognition framework using homomorphically encrypted deep features with locality preserving projection," [Uploaded corpus paper; venue/year to be confirmed by authors].
[5] [Authors], "A hybrid homomorphic-encryption and locally-linear-embedding scheme for privacy-preserving smart doorbell face recognition," [Uploaded corpus paper; venue/year to be confirmed by authors].
[6] [Authors], "Dual-layer homomorphically encrypted face similarity evaluation using landmark and color features," [Uploaded corpus paper; venue/year to be confirmed by authors].
[7] J. Engelsma, A. K. Jain, and V. N. Boddeti, "HERS: Homomorphically encrypted representation search," *IEEE Trans. Biometrics, Behavior, and Identity Science*, vol. 4, no. 3, pp. 349–360, 2022.
[8] [Authors], "Post-quantum searchable and fully homomorphic encryption for fast face-template preselection and long-term protection," [Uploaded corpus paper; venue/year to be confirmed by authors].
[9] [Authors], "Homomorphically encrypted multimodal iris and face template fusion on the QFIRE-I database," [Uploaded corpus paper; venue/year to be confirmed by authors].
[10] [Authors], "For Your Eyes Only: Retina-based biometric template protection with homomorphic encryption," [Uploaded corpus paper; venue/year to be confirmed by authors].
[11] [Authors], "IP-address-driven dynamic key generation for encrypting Dlib facial encodings," [Uploaded corpus paper; venue/year to be confirmed by authors].
[12] [Authors], "A comparative evaluation of AES-Light, SPECK, and SIMON on ARM Cortex-M microcontrollers for IoT security," [Uploaded corpus paper; venue/year to be confirmed by authors].
[13] A. Woubie and et al., "Federated learning for privacy-preserving face recognition using edge-generated adversarial imposter data," [Uploaded corpus paper; venue/year to be confirmed by authors].
[14] [Authors], "Real-time GAN-based face and hair swapping for identity-concealing video conferencing," [Uploaded corpus paper; venue/year to be confirmed by authors].
[15] [Authors], "Masked face recognition via lower-face augmentation with CNN matching and SURF-based refinement," [Uploaded corpus paper; venue/year to be confirmed by authors].
[16] X. Cai et al., "End-to-end optimized lensless face verification without image reconstruction," [Uploaded corpus paper; venue/year to be confirmed by authors].
[17] [Authors], "Adv-Eye: Natural eye-makeup adversarial attacks against face recognition," [Uploaded corpus paper; venue/year to be confirmed by authors].
[18] [Authors], "Invisible adversarial face attacks via landmark-guided superpixel perturbation," [Uploaded corpus paper; venue/year to be confirmed by authors].
[19] [Authors], "On the sensitivity of adversarial face-recognition attacks to preprocessing choices," [Uploaded corpus paper; venue/year to be confirmed by authors].
[20] [Authors], "3FG-Defense: Cross-domain geometric, frequency, and structural consistency for adversarial face detection," [Uploaded corpus paper; venue/year to be confirmed by authors].

 This work is licensed under a Creative Commons Attribution 4.0 International License



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.1510129

[21] Zhang et al., "Map2V: Validating the privacy of decompose-and-remove face recognition systems under a minimum-knowledge adversary," [Uploaded corpus paper; venue/year to be confirmed by authors].

[22] [Authors], "A blockchain-integrated multimodal biometric e-voting system with homomorphic encryption," [Uploaded corpus paper; venue/year to be confirmed by authors].

[23] J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *Advances in Cryptology – ASIACRYPT*, 2017, pp. 409–437.

[24] C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proc. 41st Annu. ACM Symp. Theory of Computing (STOC)*, 2009, pp. 169–178.

[25] I. Goodfellow et al., "Generative adversarial networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2014, pp. 2672–2680.

[26] ISO/IEC 24745:2022, *Information security, cybersecurity and privacy protection – Biometric information protection*, Int. Org. for Standardization, 2022.

[27] N. K. Ratha, J. H. Connell, and R. M. Bolle, "Enhancing security and privacy in biometrics-based authentication systems," *IBM Systems Journal*, vol. 40, no. 3, pp. 614–634, 2001.

 This work is licensed under a Creative Commons Attribution 4.0 International License