PowerGuard

Privacy-Preserving Biometric Attendance System

Prof. Aishwariya Kamat, Aqib Choudhary, Khushnuma Shaikh, Hamza Shaikh, Nasaroddin Kazi

Guide, Student, Student, Student, Student,Student
Artificial Intelligence and Data Science,
Rizvi College of Engineering, India

Abstract : PowerGuard is a privacy-preserving biometric attendance system that eliminates proxy attendance while ensuring biometric data remains cryptographically secure. Unlike conventional facial recognition systems that store raw images or unencrypted templates, PowerGuard performs all face matching operations within an encrypted domain using a dual-encryption architecture: Fully Homomorphic Encryption (FHE) for enrollment and Partial Homomorphic Encryption (PHE) for real-time scanning. The system employs OpenCV for face detection, DeepFace with VGG-Face for embedding extraction, and SQLite3 for local storage. Testing with 15 enrolled subjects achieved 94.7% recognition accuracy and 0.5-second confirmation time with zero false positives. PowerGuard demonstrates that high-accuracy biometric attendance can coexist with uncompromising privacy protection on standard hardware.

IndexTerms **- Biometric Attendance, Homomorphic Encryption, Face Recognition, DeepFace, Privacy-Preserving, Python, OpenCV, FHE, PHE.**

I. INTRODUCTION

Attendance management is a critical administrative function in educational institutions and organizations globally. Despite rapid technological advances, a significant number of institutions still rely on legacy mechanisms such as manual roll-calls, paper registers, or token-based systems like RFID cards. These conventional approaches share a fundamental and pervasive vulnerability: proxy attendance, colloquially known as 'buddy punching,' where one individual fraudulently marks the attendance of an absent peer. This practice undermines academic integrity and operational efficiency.

Biometric systems, particularly those based on facial recognition, have emerged as a highly compelling solution to eliminate proxy attendance entirely. By binding identity to a unique biological trait, they offer a robust deterrent against impersonation. However, the widespread adoption of facial recognition for attendance has been significantly impeded by a critical and unresolved privacy paradox. Conventional facial recognition systems operate by storing raw biometric images or unencrypted mathematical templates (feature vectors) of an individual's face in a central database. Should such a database be compromised, the consequences are catastrophic and irreversible. Unlike a password or an RFID card, a person's facial biometrics cannot be changed, reissued, or revoked. A data breach, therefore, exposes individuals to a permanent, lifelong risk of identity theft and unwarranted surveillance [1, 2].

PowerGuard is proposed as a next-generation biometric attendance platform engineered to directly address this privacy paradox. It is a system that performs all critical face-matching operations entirely within an encrypted mathematical space, ensuring that no raw facial images or unprotected biometric vectors are ever written to persistent storage. This is achieved through a novel dual-encryption architecture that leverages Fully Homomorphic Encryption (FHE) for high-security enrollment and Partial Homomorphic Encryption (PHE) for real-time attendance scanning. By doing so, PowerGuard provides a practical demonstration that robust biometric security and uncompromising individual privacy are not mutually exclusive but can be achieved simultaneously on standard, off-the-shelf computing hardware.

NEED OF THE STUDY.

The establishmentThe primary motivation for this study stems from the escalating global exposure of biometric databases to sophisticated cyberattacks and the inadequacy of current attendance solutions in mitigating these privacy risks. A critical examination of existing systems reveals a clear and urgent need for a privacy-first approach.

First-generation attendance systems, which rely on physical tokens like RFID cards or PIN codes, are inherently vulnerable to loss, theft, and, most critically, proxy attendance. Second-generation systems, which include most commercially available facial recognition terminals (e.g., ZKTeco, ESSL FacePass), successfully mitigate buddy punching but introduce an even greater liability: they typically store biometric templates in weakly encrypted or, in some cases, plaintext formats on local or networked servers. Research in biometric security has demonstrated that facial images can be reconstructed with alarming fidelity from unencrypted feature vectors, turning a compromised attendance database into a goldmine for malicious actors [2, 3]. Educational and corporate institutions that collect and store facial data without adequate cryptographic protections are therefore placing their entire population at a permanent, irreversible risk.

This study is necessitated by the clear gap between the operational need for robust, proxy-proof attendance and the ethical and legal imperative to protect sensitive biometric data. There is a demonstrable need for a third-generation system—one that retains the anti-spoofing benefits of biometrics while rendering the stored data mathematically useless to an attacker. PowerGuard addresses this need by providing a reference implementation that proves the technical feasibility and operational practicality of a fully encrypted biometric workflow. Its design, which is hardware-agnostic and deployable on standard laptops, ensures that this advanced level of privacy protection is accessible without requiring specialized or expensive infrastructure.

RESEARCH METHODOLOGY

The This research adopts a design and development methodology, focusing on the creation of a modular, privacy-preserving software application. The PowerGuard system was engineered using a five-stage sequential pipeline and a dual-encryption strategy to ensure both real-time performance and cryptographic security. The methodology is structured around five core Python modules, each with a distinct responsibility.

3.1 System Architecture and Technology Stack

The system is architected as a modular desktop application using Python 3.10+.
The core components are:
(1) encryption.py, which manages FHE/PHE key generation and cryptographic operations via the CipherFace library;
(2) database.py, which handles all SQLite3 interactions for student records, attendance logs, and classroom data;
(3) camera.py, a hardware-agnostic module for frame acquisition and preprocessing;
(4) gui.py, which implements the Tkinter-based user interface with a multi-threaded task manager; and
(5) matcher.py, responsible for calculating encrypted cosine similarity between biometric vectors. The primary technologies employed include OpenCV for face detection, DeepFace with the VGG-Face backend for generating 128-dimensional facial embeddings, and the CipherFace suite for homomorphic encryption operations.

3.2 Biometric Recognition Pipeline

The core attendance process follows a five-stage pipeline executed on each frame of a live video feed:
1. Adaptive Face Detection: A Haar Cascade Classifier detects facial regions in real-time. To maintain high performance (60+ FPS), frames are downscaled for detection, and bounding box coordinates are mapped back to the display resolution.
2. Biometric Embedding Extraction: The detected face region is passed to the DeepFace framework, which utilizes the pre-trained VGG-Face model to extract a 128-dimensional feature vector that uniquely encodes the geometric characteristics of the face.
3. Immediate Encryption: This 128-dimensional vector is immediately encrypted using the system's public key. In this stage, no plaintext biometric data is ever stored or transmitted.
4. Encrypted Matching: The newly encrypted live vector is compared against all previously stored encrypted vectors in the SQLite3 database. The Matcher module computes the Cosine Similarity between vectors directly within the encrypted domain. A match is confirmed if the similarity score exceeds a pre-defined confidence threshold (set at 0.7). Critically, this comparison requires no decryption of the stored biometric data.
5. Attendance Logging: Upon a successful match, the system retrieves the corresponding student's plaintext metadata (name, roll number, class) and logs a timestamped attendance record into the database.

3.3 Dual-Encryption and Enrollment Strategy

To balance maximum security with operational efficiency, a dual-mode encryption strategy is employed. The enrollment workflow uses Full Homomorphic Encryption (CipherFace FHE mode) to encrypt the averaged vector from three facial samples, prioritizing long-term security over processing speed. Conversely, the real-time scanning mode utilizes Partial Homomorphic Encryption (CipherFaceLite PHE mode), which introduces a minimal latency overhead of approximately 45ms, making

sub-second attendance confirmation practical. This methodology ensures that even in the event of a complete database breach, an attacker would only obtain encrypted mathematical vectors, from which the original facial biometric data cannot be mathematically recovered.

IV. RESULTS AND DISCUSSION

This section presents the empirical outcomes of the PowerGuard system's performance and provides a critical analysis of the results obtained during controlled testing.

4.1 Performance Metrics and Quantitative Results

The system was evaluated in a simulated classroom environment with a cohort of 15 enrolled student profiles. Testing was conducted on a standard laptop configuration (Intel Core i5 processor, 8GB RAM) using both an integrated 720p webcam and a 1080p mobile camera stream via DroidCam. The following quantitative metrics were recorded:

- Detection Frame Rate: The Haar Cascade-based Adaptive Face Detection module consistently operated between 55 and 65 Frames Per Second (FPS) after frame downscaling. This ensured a smooth, real-time visual feed with no perceptible lag in the GUI overlay.
- Embedding Extraction Latency: The DeepFace VGG-Face model required an average of 280 milliseconds to generate a 128-dimensional biometric vector from a detected face region under standard indoor fluorescent lighting.
- Encryption Overhead: In Partial Homomorphic Encryption (PHE) mode (CipherFaceLite), the encrypted matching process added an average overhead of only 45 milliseconds per scan.
- End-to-End Confirmation Time: From the moment a face is detected until the attendance log is written to the SQLite3 database, the average confirmation time was 0.5 seconds.
- Recognition Accuracy: The system achieved a 94.7% true positive recognition rate across all test subjects under standard indoor lighting conditions.
- False Positive Rate: During the entire testing period, the system recorded zero (0) false positives for non-enrolled individuals attempting to spoof the system.

4.2 Feature Validation and Discussion

Beyond raw speed, the system successfully validated its core privacy claims. Post-enrollment inspection of the SQLite3 database confirmed that no raw images were stored; the `Students` table contained only AES-encrypted vector blobs alongside plaintext metadata fields. The "Biometric Vault" admin panel functioned as designed, allowing for the secure deletion of records without leaving recoverable biometric artifacts.

The results confirm that PowerGuard successfully reconciles the tension between security and privacy. The 94.7% accuracy is competitive with commercial second-generation systems; however, PowerGuard achieves this benchmark while providing a cryptographic guarantee of privacy that conventional systems lack. The 0.5-second confirmation time demonstrates that Homomorphic Encryption (PHE) is a viable, practical solution for real-time attendance scenarios and is no longer confined to theoretical or batch-processing applications.

4.3 Limitations and Observations

A notable limitation identified during testing pertains to environmental robustness. While accuracy was high under standard indoor lighting, performance degraded in extreme backlighting or when subjects wore facial masks covering significant portions of the lower face. This is an inherent limitation of the underlying VGG-Face model and the Haar Cascade detector rather than the encryption pipeline. Furthermore, while the FHE enrollment process was secure, it took approximately **8 seconds** per user due to the computational intensity of full homomorphic operations. This is acceptable for a one-time enrollment workflow but highlights why PHE is necessary for the real-time scanning loop.

The discussion of these results establishes that PowerGuard is not merely a proof-of-concept but a functional, deployable system that significantly raises the bar for privacy in biometric attendance infrastructure.

I. ACKNOWLEDGMENT

The authors wish to express their sincere gratitude to Prof. Aishwariya Kamat, Department of Artificial Intelligence and Data Science Engineering, Rizvi College of Engineering, for their invaluable guidance, persistent encouragement, and constructive feedback

throughout the development of this project. Their expertise was instrumental in shaping both the technical architecture and the privacy-centric ethos of PowerGuard.

We extend our deepest appreciation to Dr. Varsha Shah, Principal, Rizvi College of Engineering, Mumbai, and Prof. Junaid Mandviwala, Head of the Department of Artificial Intelligence and Data Science Engineering, for providing the necessary infrastructure, resources, and institutional support that made this work possible.
Finally, we thank all faculty members and technical staff of the department whose direct and indirect support contributed to the successful completion of this research endeavor.

REFERENCES

[1] A. K. Jain, A. A. Ross, and K. Nandakumar, *Introduction to Biometrics*. New York, NY, USA: Springer, 2011, pp. 1–49.
[2] S. Marcel, M. S. Nixon, and S. Z. Li, Eds., *Handbook of Biometric Anti-Spoofing: Trusted Biometrics under Spoofing Attacks*. London, UK: Springer, 2014, pp. 15–30.
[3] C. Gentry, "A Fully Homomorphic Encryption Scheme," Ph.D. dissertation, Dept. Comput. Sci., Stanford Univ., Stanford, CA, USA, 2009.
[4] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep Face Recognition," in Proc. British Machine Vision Conference (BMVC), Swansea, UK, 2015, pp. 41.1–41.12.
[5] G. Bradski and A. Kaehler, Learning OpenCV: Computer Vision with the OpenCV Library. Sebastopol, CA, USA: O'Reilly Media, 2008.
[6] P. Viola and M. Jones, "Rapid Object Detection using a Boosted Cascade of Simple Features," in Proc. IEEE Computer SocietyConference on Computer Vision and Pattern Recognition (CVPR), Kauai, HI, USA, 2001, vol. 1, pp. 511–518.
[7] M. van Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan, "Fully Homomorphic Encryption over the Integers," in Advances in Cryptology – EUROCRYPT 2010, Lecture Notes in Computer Science, vol. 6110. Berlin, Heidelberg: Springer, 2010, pp. 24–43.