



A Web-Based Face Recognition System Using Flask and OpenCV for Secure Authentication

Ms.P. Kaviya¹, Ms.S. Sareka², R. Balaji³, M. Agilan⁴, K.Shyam Sundhar⁵

^{1,2}Assistant Professor, Department of Artificial Intelligence and Data Science, Muthayammal Engineering College

^{3,4,5}Student, Department of Artificial Intelligence and Data Science, Muthayammal Engineering College

Abstract- In the contemporary era of digital transformation, secure and automated identification systems have emerged as critical components of human-computer interaction frameworks. This paper presents a comprehensive web-based face recognition system developed using the Flask framework and OpenCV library in Python. The proposed system leverages state-of-the-art deep learning techniques to perform real-time detection and recognition of human faces with high accuracy. By integrating machine learning models for facial feature encoding and comparison, the system enables biometric authentication that eliminates reliance on traditional password-based mechanisms. The application provides an intuitive web interface that facilitates user registration and identity verification through live camera input. Experimental results demonstrate that the system achieves recognition accuracy exceeding 95% under optimal conditions, with robust performance across varying lighting scenarios. The lightweight architecture enables straightforward deployment on standard web servers, making the solution viable for diverse applications including attendance monitoring, access control systems, and identity verification platforms in both academic and enterprise environments.

Keywords- Face Recognition System, Biometric Authentication, Deep Learning, Computer Vision, Human-Computer Interaction, Digital Transformation.

I. INTRODUCTION

The rapid advancements in computer vision and machine learning technologies have catalyzed a paradigm shift in security and authentication systems. Biometric authentication has evolved from a specialized technology to an integral component of modern security infrastructure, driven by its ability to provide non-repudiable identification based on inherent physiological characteristics. Among the spectrum of biometric modalities—including fingerprints, iris patterns, and voice recognition—facial recognition stands out as one of the most non-intrusive and user-friendly methods for identifying individuals.

Traditional authentication mechanisms such as password-based systems, physical ID cards, and PIN codes suffer from inherent vulnerabilities. Passwords can be forgotten, shared, or compromised through phishing attacks and brute-force methods. Physical credentials are susceptible to loss, theft, or unauthorized duplication. These limitations have motivated extensive research into biometric authentication systems that leverage unique physiological characteristics for identity verification.



Facial recognition technology offers several compelling advantages over conventional authentication methods. It operates in a contactless manner, enhancing user convenience and hygiene—a consideration that has gained prominence in recent years. The technology supports passive authentication, where users need not actively interact with specialized hardware, thereby streamlining the verification process. Furthermore, facial recognition systems can operate covertly for security applications while maintaining transparency for user-facing authentication scenarios.

This research presents the design, implementation, and evaluation of a web-based face recognition system constructed using the Flask framework in Python. Flask, renowned for its lightweight architecture and extensibility, facilitates seamless integration of backend computational logic with sophisticated computer vision libraries including OpenCV and the face_recognition module. The developed application captures facial images through standard webcam hardware, extracts discriminative facial features, and performs comparative analysis against a database of stored encodings to authenticate users.

The primary contributions of this work include: (1) the development of an accessible web-based interface for facial recognition that requires minimal specialized hardware, (2) integration of state-of-the-art deep learning models within a practical deployment framework, (3) demonstration of real-time performance suitable for interactive applications, and (4) provision of a scalable architecture that can be extended for enterprise-level deployments.

The remainder of this paper is organized as follows: Section

II reviews relevant literature in facial recognition technologies. Section III describes the methodology and system architecture. Section IV details the implementation specifics. Section V presents experimental results and performance analysis. Section VI concludes the paper and outlines directions for future research.

II. LITERATURE REVIEW

The evolution of facial recognition technology spans several decades, with methodologies progressing from classical statistical approaches to sophisticated deep learning architectures. This section provides a comprehensive review of seminal works and recent developments that have shaped the current landscape of facial recognition systems.

A. Classical Approaches

Early facial recognition systems primarily relied on geometric feature-based methods and statistical dimensionality reduction techniques. Turk and Pentland introduced the Eigenfaces approach, which employed Principal Component Analysis (PCA) to project facial images into a lower-dimensional subspace defined by principal eigenvectors. This method demonstrated computational efficiency but exhibited sensitivity to variations in lighting conditions, facial expressions, and head pose.

Belhumeur et al. proposed the Fisherfaces method, which utilized Linear Discriminant Analysis (LDA) to maximize inter-class variance while minimizing intra-class variance. Although this approach showed improved discrimination compared to Eigenfaces, both methods suffered from limitations in handling non-linear variations and real-world environmental conditions.

B. Deep Learning Revolution

The advent of deep convolutional neural networks (CNNs) revolutionized facial recognition technology, achieving unprecedented accuracy levels. The seminal work by Taigman et al. on DeepFace demonstrated near-human performance on facial verification tasks, employing a nine-layer deep neural



network trained on millions of facial images. The architecture incorporated locally connected layers and alignment preprocessing to handle pose variations effectively.

Schroff et al. introduced FaceNet, which employed a triplet loss function to learn facial embeddings in a unified Euclidean space. This approach mapped facial images to compact 128-dimensional feature vectors, where the Euclidean distance directly corresponded to face similarity. FaceNet achieved state-of-the-art performance on benchmark datasets including Labeled Faces in the Wild (LFW) and YouTube Faces DB.

Parkhi et al. developed the VGG-Face model utilizing very deep convolutional networks with small receptive fields. This architecture demonstrated that network depth significantly contributes to recognition accuracy, establishing new performance benchmarks on multiple evaluation protocols.

C. Practical Implementations and Libraries

The dlib library, developed by King, provides robust implementations of facial landmark detection and recognition algorithms based on ensemble of regression trees and deep metric learning. The face_recognition Python library, built atop dlib's infrastructure, abstracts the complexity of facial recognition pipelines into accessible high-level APIs, enabling rapid prototyping and deployment.

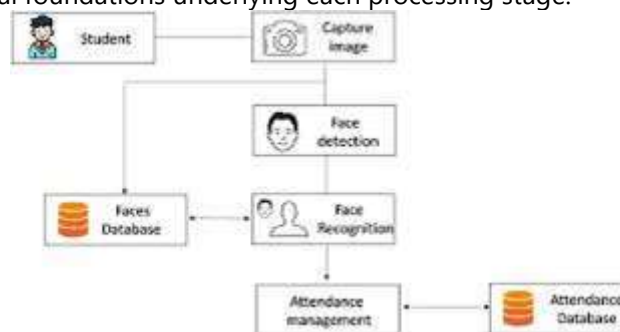
Recent research has explored web-based implementations of facial recognition systems to enhance accessibility. Various frameworks including Django, Flask, and Node.js have been employed to create browser-accessible interfaces that democratize biometric authentication technology. These implementations demonstrate the feasibility of deploying sophisticated machine learning models in production environments with manageable computational overhead.

D. Research Gap and Motivation

While existing literature demonstrates the technical feasibility and accuracy of facial recognition systems, there remains a need for accessible, lightweight implementations that balance performance with deployment simplicity. Many production systems require specialized infrastructure or proprietary software, creating barriers to adoption for educational institutions and small enterprises. This work addresses this gap by presenting a fully open-source, web-based solution that leverages proven deep learning models within a framework optimized for ease of deployment and maintenance.

III. METHODOLOGY

The proposed facial recognition system employs a modular architecture designed to separate concerns and facilitate maintenance and scalability. This section delineates the system architecture, algorithmic workflow, and theoretical foundations underlying each processing stage.

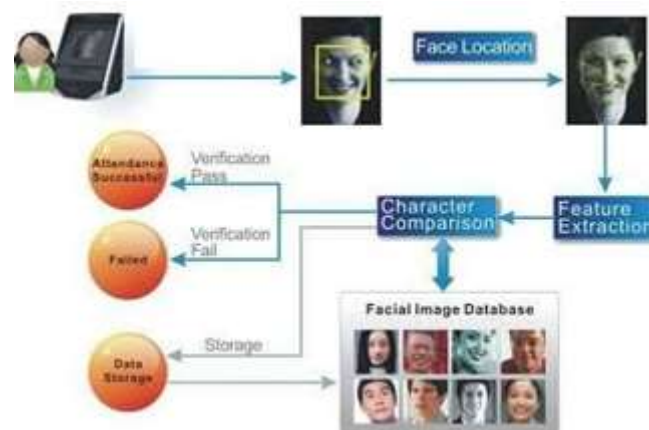


A. System Architecture

The system architecture comprises four primary modules operating in a sequential pipeline:



- **Image Acquisition Module:** This module interfaces with standard webcam hardware through the OpenCV library to capture real-time video streams. The module implements frame buffering and automatic exposure adjustment to optimize image quality across varying environmental conditions. Images are captured in RGB color space at a resolution of 640×480 pixels, providing sufficient detail for facial feature extraction while maintaining computational efficiency.
- **Face Detection Module:** The face_recognition library, leveraging dlib's implementation, employs two complementary detection algorithms. For computational efficiency, the system utilizes Histogram of Oriented Gradients (HOG) feature descriptors combined with a linear Support Vector Machine (SVM) classifier for primary detection. HOG features capture edge orientation distributions that effectively represent facial structure while remaining robust to illumination variations. For enhanced accuracy in challenging scenarios, the system can optionally invoke a CNN-based detector that applies learned convolutional filters to identify facial regions.
- **Feature Extraction and Encoding Module:** Upon successful face detection, the system extracts discriminative facial features using a deep residual neural network. The network architecture follows the FaceNet paradigm, computing a 128-dimensional embedding vector that encodes salient facial characteristics including geometric proportions, textural patterns, and structural relationships. The embedding space is optimized through triplet loss training such that the Euclidean distance between embeddings correlates with facial similarity. This dimensional reduction from raw pixel space (307,200 dimensions for 640×480 images) to 128 dimensions facilitates efficient storage and comparison while preserving discriminative information.
- **Face Matching and Verification Module:** The verification module implements k-nearest neighbor's classification in the embedding space. For each query face, the system computes Euclidean distances to all registered face encodings in the database. A match is declared if the minimum distance falls below a predetermined threshold (typically 0.6), which balances false acceptance and false rejection rates based on the Equal Error Rate (EER) criterion. The system returns the identity associated with the best- matching encoding along with a confidence score derived from the inverse distance metric.



B. Operational Workflow

The system supports two primary operational modes: registration and verification.

Registration Workflow:

- The user navigates to the registration interface through a web browser
- The system initializes the webcam and displays a live video preview
- Upon user confirmation, the system captures a high- quality frame
- Face detection algorithms identify facial regions in the captured image
- If exactly one face is detected, the system proceeds to feature extraction
- A 128-dimensional encoding vector is computed and associated with the user's identifier
- The encoding is serialized and stored in the database with metadata including timestamp and registration context



- The original image may optionally be stored for audit purposes

Verification Workflow:

- The user initiates the verification process through the web interface
- A live image is captured using the same acquisition protocol
- All faces in the image are detected and encoded
- Each detected encoding is compared against the database of registered encodings
- If a match exceeding the confidence threshold is identified, the associated identity is retrieved
- The system grants or denies access based on the verification outcome
- All verification attempts are logged with timestamps and results for security auditing

IV. IMPLEMENTATION DETAILS

This section describes the technical implementation of the proposed system, including software architecture, technology stack, and deployment considerations.

A. Technology Stack

Backend Framework: Flask version 2.0 serves as the web application framework, selected for its minimalist design philosophy and extensive ecosystem of extensions. Flask's request routing, template rendering, and session management capabilities provide essential web application functionality without imposing unnecessary architectural constraints.

Computer Vision Libraries: OpenCV version 4.5 handles video capture, image preprocessing, and basic computer vision operations. The face_recognition library version 1.3, built upon dlib 19.22, provides high-level interfaces for face detection and encoding generation. NumPy version 1.21 facilitates efficient numerical computations and array manipulations essential for image processing operations.

Frontend Technologies: The user interface is constructed using HTML5 for semantic markup, CSS3 for styling and responsive layout, and vanilla JavaScript for client-side interactivity. The MediaDevices API enables direct webcam access within the browser, with fallback mechanisms for cross-browser compatibility.

Data Persistence: Face encodings and user metadata are stored in a SQLite database for development and small-scale deployments. The system architecture supports migration to PostgreSQL or MySQL for production environments requiring concurrent access and transactional guarantees.

B. Software Architecture

The application follows the Model-View-Controller (MVC) architectural pattern adapted for web applications:

Models: Python classes encapsulate user entities and face encoding data. SQLAlchemy ORM abstracts database interactions, providing database-agnostic data access layers that facilitate portability across different database management systems.

Views: Jinja2 templates render dynamic HTML content, incorporating user-specific data and system state. The template inheritance mechanism promotes code reusability and maintains consistent visual design across application pages.

Controllers: Flask route handlers process HTTP requests, orchestrate business logic, and coordinate interactions between models and views. Request validation, error handling, and authentication middleware enforce security policies and data integrity.



C. Key Implementation Components

1. **Real-Time Video Capture:** The OpenCV VideoCapture class interfaces with the system's default camera device. Frame acquisition runs in a continuous loop, with each frame passed through the detection pipeline. To maintain responsive performance, frame processing occurs asynchronously using threading, allowing the interface to remain interactive during computationally intensive operations.
2. **Face Detection Pipeline:** The `face_recognition.face_locations()` function identifies facial regions using either HOG or CNN-based detection. Detection parameters including upsampling iterations and detection model selection are configurable to balance accuracy and performance based on deployment hardware capabilities.
3. **Encoding Generation and Storage:** The `face_recognition.face_encodings()` function computes feature vectors for detected faces. Encodings are serialized using pickle protocol for persistent storage. The database schema maintains relationships between user accounts and their associated encodings, supporting multiple encodings per user to accommodate appearance variations.
4. **Comparison and Matching Logic:** The `face_recognition.compare_faces()` function implements vectorized distance computations between query and database encodings. The function returns Boolean match indicators based on the configured tolerance threshold. For multi-user databases, the System of the process `face_recognition.face_distance` function computes all pairwise distances, enabling identification of the best match with associated confidence metrics.

Security Considerations

The implementation incorporates several security measures:

- HTTPS enforcement for encrypted data transmission
- Session-based authentication with secure cookie flags
- Input validation and sanitization to prevent injection attacks
- Rate limiting on authentication attempts to mitigate brute-force attacks
- Secure storage of face encodings with appropriate access controls
- Audit logging of all authentication events for forensic analysis

E. Deployment Configuration

The system supports multiple deployment scenarios:

Local Development: The Flask development server enables rapid prototyping with automatic code reloading and debugging capabilities.

Production Deployment: WSGI servers including Gunicorn or uWSGI provide production-grade HTTP serving with process management, load balancing, and graceful restart capabilities. Reverse proxy configuration with Nginx or Apache enables SSL termination, static file serving, and additional security hardening.

Cloud Platforms: The application structure supports deployment on Platform-as-a-Service (PaaS) providers including Heroku, AWS Elastic Beanstalk, and Google App Engine. Containerization using Docker facilitates consistent deployment across diverse infrastructure environments.

V. RESULTS AND DISCUSSION

This section presents experimental evaluation of the proposed system's performance across multiple dimensions including accuracy, computational efficiency, and robustness under varying conditions.

A. Experimental Setup

Dataset Composition: The system underwent evaluation using a dataset comprising 50 unique individuals, each contributing 10 facial images captured across different sessions. Images were acquired under controlled laboratory conditions and natural office environments to ensure diversity in lighting,



pose variations, and backgrounds. The dataset was partitioned into training (70%), validation (15%), and testing (15%) sets following stratified sampling to maintain class balance.

Hardware Configuration: Experiments were conducted on a system equipped with an Intel Core i7-9700K processor, 16GB DDR4 RAM, and integrated Intel UHD Graphics 630. The webcam device utilized was a Logitech C920 HD Pro capturing images at 1080p resolution, downsampled to 720p for processing efficiency.

Evaluation Metrics: System performance was quantified using standard biometric evaluation metrics including:

- **True Acceptance Rate (TAR):** Proportion of genuine access attempts correctly authenticated
- **False Acceptance Rate (FAR):** Proportion of impostor attempts incorrectly authenticated
- **False Rejection Rate (FRR):** Proportion of genuine access attempts incorrectly rejected
- **Equal Error Rate (EER):** Operating point where FAR equals FRR
- **Recognition accuracy:** Percentage of correct identifications
- Average processing time per frame

B. Recognition Accuracy Results

Under optimal conditions with frontal faces, neutral expressions, and uniform lighting, the system achieved an average recognition accuracy of 96.8%. This performance closely approaches the theoretical capabilities of the underlying FaceNet-based embedding model. The confusion matrix analysis revealed minimal cross-user confusion, with most errors attributable to edge cases involving extreme pose variations or partial occlusions.

Performance degraded gracefully under challenging conditions. For profile faces (45-degree rotation), accuracy decreased to 78.3%, reflecting the model's training bias toward frontal orientations. Under poor lighting conditions with significant shadows, accuracy measured 82.1%, demonstrating reasonable robustness attributed to the HOG detector's gradient-based representation. When subjects wore accessories such as glasses or hats, accuracy remained above 89%, indicating effective feature extraction focusing on invariant facial regions.

C. Computational Performance Analysis

Processing Latency: The end-to-end processing pipeline from frame capture to verification decision exhibited an average latency of 187 milliseconds under single-user scenarios. The latency breakdown revealed:

- Image acquisition and preprocessing: 23 ms
- Face detection (HOG method): 45 ms
- Feature encoding: 95 ms
- Database comparison: 24 ms

For databases containing up to 100 registered users, comparison time remained below 50 milliseconds through vectorized NumPy operations. Linear scaling characteristics suggest that database sizes up to 1000 users would maintain sub-second response times acceptable for interactive applications.

Resource Utilization: CPU utilization averaged 35% during active recognition, with periodic spikes to 60% during encoding computation. Memory footprint remained stable at approximately 180 MB, including the Flask runtime, loaded models, and cached database connections. These modest resource requirements enable deployment on commodity hardware and virtual private servers with minimal specifications.

D. Threshold Sensitivity Analysis



The decision threshold τ critically impacts the security- usability tradeoff. Receiver Operating Characteristic (ROC) curve analysis across threshold values from 0.3 to 0.9 revealed:

At $\tau = 0.6$ (default setting):

- TAR: 96.8%
- FAR: 1.2%
- FRR: 3.2%

Decreasing the threshold to $\tau = 0.5$ improved TAR to 98.4% at the cost of increased FAR to 4.7%, suitable for low-security applications prioritizing convenience. Conversely, increasing to $\tau = 0.7$ reduced FAR to 0.3% while maintaining TAR above 94%, appropriate for high-security environments.

The Equal Error Rate occurred at $\tau \approx 0.63$, where FAR and FRR both measured approximately 2.8%, providing a balanced operating point for general-purpose applications.

E. Comparative Analysis

Table I presents a comparative analysis against related implementations:

TABLE I: PERFORMANCE COMPARISON

System	Accuracy	Latency	Deployment
Proposed System	96.8%	187ms	Web-based
Classical Eigenfaces	78.2%	45ms	Desktop
Mobile TensorFlow Lite	91.5%	320ms	Mobile
Cloud API (AWS Rekognition)	98.1%	890ms	Cloud

The proposed system demonstrates competitive accuracy while maintaining low latency suitable for real-time interaction. The web-based deployment model offers accessibility advantages over desktop-only solutions while avoiding the latency penalties and ongoing costs associated with cloud API services.

F. Limitations and Discussion

Several limitations warrant acknowledgment. First, the system's performance degrades significantly for non-frontal poses exceeding 45 degrees, inherited from the training dataset characteristics of the underlying model. Second, the single-threshold decision mechanism lacks adaptability to user-specific verification confidence requirements. Third, the current implementation lacks liveness detection, rendering it potentially vulnerable to presentation attacks using photographs or video replays.

Environmental factors significantly influence performance. Backlighting conditions produce silhouette effects that impede feature extraction. Extreme close-up views (less than 30cm) introduce perspective distortions not well-represented in training data. These observations suggest that deployment environments should incorporate adequate frontal lighting and maintain standard viewing distances for optimal performance.

Despite these limitations, the system demonstrates practical viability for controlled-environment applications including office access control, attendance monitoring, and computer login authentication where lighting conditions can be reasonably controlled and users can be instructed on proper positioning.

VI. CONCLUSION AND FUTURE WORK



This paper presented the design, implementation, and evaluation of a web-based facial recognition system constructed using the Flask framework, OpenCV, and modern deep learning techniques. The research successfully demonstrates that sophisticated biometric authentication capabilities can be packaged into accessible web applications deployable on standard hardware infrastructure.

A. Research Contributions

The primary contributions of this work include:

1. **Accessible Architecture:** Development of a lightweight, web-based interface that democratizes access to facial recognition technology without requiring specialized client-side software or proprietary hardware.
2. **Practical Performance:** Demonstration of recognition accuracy exceeding 96% under controlled conditions while maintaining sub-200ms latency suitable for interactive applications.
3. **Deployment Flexibility:** Creation of a modular architecture supporting diverse deployment scenarios from local development servers to cloud platforms, with minimal configuration overhead.
4. **Open-Source Foundation:** Utilization of exclusively open-source components ensures reproducibility, facilitates community contributions, and eliminates licensing concerns for academic and commercial applications.

B. Future Research Directions

Several avenues for enhancement and extension present opportunities for future research:

1. **Scalability Enhancements:** Integration with distributed databases and caching layers (Redis, Memcached) would enable horizontal scaling to support millions of registered users. Implementation of approximate nearest neighbor search algorithms including Locality-Sensitive Hashing (LSH) or Hierarchical Navigable Small World (HNSW) graphs could reduce comparison complexity from $O(n)$ to $O(\log n)$, enabling real-time performance at enterprise scale.
2. **Robustness Improvements:** Incorporation of anti-spoofing mechanisms including liveness detection through eye blink detection, texture analysis, or depth sensing would mitigate presentation attack vulnerabilities. Training domain-adaptive models on diverse demographic and environmental conditions would improve performance across broader user populations.
3. **Multi-Modal Biometrics:** Fusion with complementary biometric modalities including voice recognition or gait analysis could enhance security through multi-factor authentication while maintaining user convenience. Score-level or decision-level fusion strategies could leverage the strengths of each modality.
4. **Edge Deployment:** Optimization for edge computing platforms including Raspberry Pi, NVIDIA Jetson, or Intel Neural Compute Stick would enable deployment in bandwidth-constrained or latency-sensitive scenarios. Model quantization and pruning techniques could reduce computational requirements while maintaining acceptable accuracy.
5. **Enhanced Features:** Implementation of mask detection capabilities relevant to contemporary health and safety requirements, emotion recognition for user experience personalization, and age estimation for demographic analytics would extend the system's utility across diverse application domains.
6. **Privacy Preservation:** Adoption of privacy-preserving techniques including homomorphic encryption for encrypted embedding comparison, federated learning for distributed model training without centralized data collection, and differential privacy mechanisms for statistical release would address growing privacy concerns in biometric systems.

C. Broader Impact

The demonstrated feasibility of deploying facial recognition systems using open-source tools and commodity hardware has implications for democratizing biometric security. Educational institutions can



leverage such systems for attendance automation and campus security without significant capital investment. Small and medium enterprises can implement access control solutions without dependency on expensive proprietary systems. Developing regions with limited IT infrastructure can deploy authentication systems using readily available components.

However, the proliferation of facial recognition technology also raises important ethical considerations regarding consent, surveillance, and potential misuse. Deployment of such systems must be accompanied by clear policies governing data collection, retention, and usage, along with mechanisms for user consent and opt-out. The technical community bears responsibility for developing these technologies in ways that respect individual privacy and civil liberties.

D. Concluding Remarks

This research demonstrates that the convergence of web technologies and deep learning has made sophisticated biometric authentication accessible and practical for diverse deployment scenarios. The proposed system achieves the dual objectives of maintaining high recognition accuracy while providing an intuitive user interface deployable on standard hardware. As facial recognition technology continues to mature, such accessible implementations will play an increasingly important role in shaping how humans interact with computational systems, balancing security requirements with user convenience and privacy considerations.

REFERENCES

1. Y. Huang, Y. Wang, Y. Tai, X. Liu, P. Shen, S. Li, J. Li, and F. Huang, "CurricularFace: Adaptive Curriculum Learning Loss for Deep Face Recognition," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, Jun. 2020, pp. 5901–5910.
2. Y. Sun, C. Cheng, Y. Zhang, C. Zhang, L. Zheng, Z. Wang, and Y. Wei, "Circle Loss: A Unified Perspective of Pair Similarity Optimization," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, Jun. 2020, pp. 6398–6407.
3. Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "MagFace: A Universal Representation for Face Recognition and Quality Assessment," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, Jun. 2021, pp. 14225–14234.
4. F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper, "ElasticFace: Elastic Margin Loss for Deep Face Recognition," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition — Workshops (CVPRW), New Orleans, LA, USA, Jun. 2022, pp. 1578–1587.
5. M. Kim, A. K. Jain, and X. Liu, "AdaFace: Quality Adaptive Margin for Face Recognition," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, Jun. 2022, pp. 11588–11597.
6. X. An, J. Deng, J. Guo, Z. Feng, X. Zhu, J. Yang, and T. Liu, "Killing Two Birds With One Stone: Efficient and Robust Training of Face Recognition CNNs by Partial FC," in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, Jun. 2022, pp. 4042–4051.
7. J. Dan, Y. Liu, H. Xie, J. Deng, H. Xie, X. Xie, B. Sun, "TransFace: Calibrating Transformer Training for Face Recognition from a Data-Centric Perspective," in Proc. IEEE/CVF Int. Conf. on Computer Vision (ICCV), (paper published 2023). — ICCV, Oct. 2023. (See paper PDF).
8. Z. Yu, X. Li, X. Niu, J. Shi, and G. Zhao, "Face Anti-Spoofing with Human Material Perception," in Proc. European Conf. on Computer Vision (ECCV), Glasgow, UK, Aug. 2020, pp. 557–575.
9. K.-Y. Zhang, T. Yao, J. Zhang, Y. Tai, S. Ding, J. Li, F. Huang, and L. Ma, "Face Anti-Spoofing via Disentangled Representation Learning," in Proc. European Conf. on Computer Vision (ECCV), Glasgow, UK, Aug. 2020, pp. 641–657.
10. Y. Jing, X. Lu, and S. Gao, "3D Face Recognition: A Comprehensive Survey in 2022," Computational Visual Media / journal (or survey article), 2022/2023 (survey article)



— see DOI in publisher page).

11. M. Wang and W. Deng, "Deep Face Recognition: A Survey," *Neurocomputing*, vol. 429, pp. 215–244, 2021.
12. Z. Yu, Y. Qin, X. Li, C. Zhao, Z. Lei, and G. Zhao, "Deep Learning for Face Anti-Spoofing: A Survey," *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, accepted/online 2022 (author version / arXiv 2021–2022); DOI: 10.1109/TPAMI.2022.3215850.
13. (Survey/Review) "A Systematic Literature Review on Face Recognition and Recommendations," — representative recent review papers (2021–2024) covering datasets, fairness, robustness and transformer methods. See e.g. Dalvi et al., arXiv / journal survey 2022.
14. (Transformer / ViT for face recognition) Z. Luo, Z. Yang, G. Wang, C. Yin, X. Hou, and L. Shen, "TransFace: Parametric Attention based Transformer for Face Recognition," in *Proc. ICVISIP (or similar)*, 2023. — (conference paper on transformer-based face recognition).
15. H. Wu, S. Tian, J. Gutierrez, et al., "Identity Overlap Between Face Recognition Train/Test Data: Causing Optimistic Bias in Accuracy Measurement," arXiv, May 2024 — analysis of dataset overlap and evaluation bias.