

Face Recognition-Based Attendance Management Systems: A Review of Haar Cascade, LBPH, and Deep-Learning Approaches with Evaluation and Governance Considerations

Mr. Suraj Janvekar¹, Dr. Sangram Patil²

Department of Computer Science and Engineering

D. Y. Patil Agriculture and Technical University, Talsande, India

Abstract: Recording attendance by verbal roll call, paper register, or manual spreadsheet entry consumes teaching time and remains open to transcription error and proxy marking. Face recognition offers a contactless alternative that links identity verification directly to digital record creation. This paper reviews the techniques that underpin face recognition-based attendance management systems: face detection (Haar Cascade, MTCNN, RetinaFace), recognition (Local Binary Pattern Histogram, deep embeddings, margin-based losses, and transformer models), and published attendance systems built from them. A Python-based reference implementation using OpenCV, Tkinter, NumPy, Pandas, and CSV storage is examined as a case study, covering its registration, training, attendance, and reporting modules and its threshold-based decision rule. The review finds that attendance systems consistently report faster marking and lower administrative effort, but many say little about independent test sets, threshold selection, false-positive rates, or demographic performance. The paper therefore consolidates an evaluation protocol, a privacy and governance framework, and a set of open research gaps. No new experimental accuracy results are claimed; reported benefits are treated as design objectives until validated in the deployment environment.

Keywords: Face recognition, attendance management, Haar Cascade, LBPH, deep learning, OpenCV, biometric identification, privacy, review.

I. INTRODUCTION

Attendance recording is a routine but consequential administrative task in schools, colleges, and training institutes. A teacher typically calls names aloud, circulates a paper register, or opens a spreadsheet at the start of every session. Each format costs a few minutes of contact time and invites errors such as a misheard name, a skipped row, or a slip while the record is later consolidated elsewhere. These methods share a deeper weakness: they confirm only that someone answered on a student's behalf, not that the student was actually present.

Face recognition closes part of this identity gap without physical contact and without hardware beyond a camera. A webcam captures a frame, a detector isolates the facial region, and a recognition model compares that region with previously enrolled samples. If the match clears a confidence threshold, the system logs the enrollment number, name, subject, date, and time automatically. Because of this simplicity, the approach has become a frequent subject of student projects and small institutional prototypes [24]–[27]. Many of these prototypes use a lightweight classical pipeline in which Haar Cascade detection, descended from the Viola-Jones framework [1], is paired with Local Binary Pattern Histogram (LBPH) recognition [2]; both stages ship with OpenCV, which makes the design practical for a small team working in Python [3], [4]. In parallel, learned-embedding methods such as DeepFace, FaceNet, and ArcFace have raised the performance ceiling for face recognition considerably [5]–[7].

This paper reviews that landscape from the viewpoint of attendance management. It makes four contributions. First, it surveys the detection and recognition techniques most relevant to classroom attendance. Second, it synthesizes published attendance systems and examines, as a case study, a Python and OpenCV reference implementation described in a project synopsis and its companion design paper. Third, it consolidates an evaluation protocol together with privacy, security, and demographic-fairness considerations. Fourth, it identifies open gaps and future directions. The remainder of the paper is organized as follows: Section II frames the problem, Sections III and IV review detection and recognition methods, Section V reviews attendance systems, Section VI presents the case study, Sections VII to IX discuss robustness, bias, and governance, Section X sets out evaluation methodology, and Sections XI to XIV close with comparison, gaps, future work, and conclusions.

II. PROBLEM CONTEXT

The per-session problems described above compound across a full timetable, a term, and multiple instructors, producing institutional-level difficulties that an individual teacher cannot solve. Completed paper registers are hard to search, consolidate, or analyze without re-entering them by hand. Records become scattered across registers, spreadsheets, and departments with no single authoritative copy. Subject-wise and semester-wise reports turn into a repetitive manual task for administrative staff in every reporting cycle. Verifying a disputed entry afterwards usually means locating a specific paper page or spreadsheet row, with no independent evidence beyond the entry itself.

A face recognition system can tie identity verification directly to digital record creation, reducing repetitive data entry and improving the consistency of records. It cannot, however, guarantee error-free identification. The outcome depends on the quality of the capture environment, the variety of the enrollment images, the detector's performance, the recognition model's configuration, and the position of the decision threshold. This is the same trade-off between recognition performance and the practical and ethical cost of collecting, storing, and matching physical characteristics that characterizes biometric systems generally [8].

III. FACE DETECTION TECHNIQUES

Detection locates one or more facial regions in an image before any recognition step can run. Two families are relevant to attendance systems.

A. Viola-Jones and Haar Cascade

The Viola-Jones framework remains the reference point for fast detection. It combines an integral-image representation, Haar-like features, AdaBoost-based feature selection, and a cascade of classifiers that rejects most non-face regions in the first few stages, which allows real-time processing on ordinary hardware [1]. Haar Cascade detection still performs well in controlled settings where faces point toward the camera and lighting is reasonably even. It tends to struggle once faces are rotated, distant, partially covered, or captured in poor light.

B. Deep-Learning Detectors

More recent detectors narrow that gap. Multi-task Cascaded Convolutional Networks (MTCNN) combine detection with facial-landmark localization in a cascaded deep-learning pipeline [9], while RetinaFace extends single-stage dense detection by predicting landmarks jointly with bounding boxes across scales [10]. Both carry a higher computational cost, which matters when the target hardware is a classroom laptop and not a GPU workstation.

IV. FACE RECOGNITION TECHNIQUES

A. Local Texture-Based Recognition (LBPH)

LBPH builds on Local Binary Patterns, which encode how a pixel's intensity compares with its immediate neighbors as a short binary code. Aggregating these codes into regional histograms yields a texture description that can be compared between a stored template and a new facial image with a simple distance measure [2]. OpenCV provides Eigenfaces, Fisherfaces, and LBPH recognizers out of the box [3], and its LBPH class exposes the training and prediction calls that attendance applications rely on [4]. The appeal for a prototype is practical: LBPH trains locally on the institution's own hardware, with no pretrained deep network or external service.

Its main weakness is sensitivity to the variation a classroom introduces between registration and attendance: changes in illumination, pose, expression, and partial occlusion. Illumination variation in particular has long been recognized as one of the most disruptive factors for appearance-based recognition, motivating a substantial body of work on illumination-invariant preprocessing [12]. An older but related line of work, joint Bayesian face verification, models intra-personal and inter-personal similarity as latent Gaussian variables and estimates a likelihood ratio between two faces; it predates deep embeddings but anticipates some of the same identity-versus-variation reasoning [11].

B. Deep Embedding Methods

Deep methods learn facial features directly from large labeled image collections instead of relying on hand-crafted descriptors. DeepFace combined face alignment with a deep neural representation to substantially narrow the gap to human-level verification [5]. FaceNet learned an embedding space in which Euclidean distance directly reflects facial similarity, folding verification, recognition, and clustering into a single representation [6]. VGGFace followed a similar route at larger scale for general-purpose recognition [13], DeepID2 combined identification and verification objectives in one network [14], and DeepID3 pushed the idea through deeper network stacks [15].

C. Margin-Based Loss Functions

A second thread of work concentrated on the loss used to separate identities in embedding space. Large-margin softmax introduced an angular margin into the softmax loss [16]; SphereFace extended the idea into a hypersphere embedding [17];

additive margin softmax [18] and CosFace [19] proposed simpler and more stable cosine-based margins; and ArcFace added an additive angular margin with a direct geometric interpretation that has become one of the most widely adopted choices [7].

D. Recent Developments

AdaFace adjusts the margin per sample according to estimated image quality, so that low-quality images are down-weighted during training and not penalized as if they were clear [20]. Transformer architectures have also been adapted to faces: part-based vision-transformer models split the face into regions before attention is applied [21], and TransFace argues that data-centric choices in training matter as much as architecture [22]. A survey by Wang and Deng traces how alignment, representation learning, loss design, dataset scale, and evaluation protocol have jointly shaped progress through the convolutional era, and cautions that reported gains are tied closely to the conditions under which a method was trained and tested [23]. That caution applies equally to newer transformer-based results.

Deep methods generally outperform LBPH under unconstrained conditions. Even so, the comparison should be verified against the intended attendance environment and not assumed from benchmark leaderboards, which are typically built from very different image distributions than a classroom webcam feed.

V. FACE RECOGNITION IN ATTENDANCE SYSTEMS

Published attendance systems follow a common pipeline: register students, capture images, detect and recognize faces, and write a record. Sabbir et al. describe a smart attendance system built around a deep-learning recognizer and report reduced manual effort compared with manual methods [26]. Yadav et al. describe a related smart attendance system with an OpenCV-centered pipeline combined with a backend database [27]. At the classical end, Akash et al. combine LBPH recognition with a Haar Cascade detection stage in an automated classroom attendance tool [24], and Beri et al. report a comparable Haar Cascade plus LBPH pipeline together with a measured accuracy evaluation [25]. Table I summarizes these works.

TABLE I
ATTENDANCE SYSTEMS REVIEWED

Ref.	Approach	Reported emphasis
[24]	Haar Cascade detection + LBPH recognition	Automated classroom attendance tool
[25]	Haar Cascade + LBPH	Comparable pipeline with a measured accuracy evaluation
[26]	Deep-learning recognizer	Smart attendance with reduced manual effort
[27]	OpenCV-centered pipeline + backend database	Smart attendance with persistent storage

Two observations follow. First, the Haar Cascade plus LBPH design inside a Python and OpenCV desktop application is a very well represented one in recent literature, with several functionally similar systems published independently in 2024 and 2025 [24], [25]; the recognition pipeline itself is therefore not a novelty. Second, reports of this kind consistently claim faster attendance and less administrative work, but many say comparatively little about independent test sets, how the decision threshold was chosen, measured false-positive rates, or whether performance was checked across different student demographics. A successful demonstration in front of an evaluator is not the same as measured accuracy, which requires images the model has never seen, collected under the lighting and camera conditions the system will actually face.

VI. CASE STUDY: A HAAR CASCADE AND LBPH REFERENCE SYSTEM

The reference system studied here is a Python desktop application for recording student attendance in a single classroom or small laboratory. It uses OpenCV and OpenCV-Contrib for detection and recognition, NumPy for array operations, Pandas for record handling, Pillow for rendering images inside the interface, Tkinter for the graphical interface, and optionally pyttsx3 for spoken feedback and Requests if external services are added. The assumed platform is Windows 10 or 11 with an Intel Core i3 processor or better, at least 4 GB of RAM, a webcam, and about 500 MB of free disk space. It is organized into four modules.

A. Student Registration

An administrator enters the student's enrollment number and name, and the webcam captures roughly 50 facial images that are stored in the training dataset. This count is a provisional default, not a validated requirement; the right value should be set by the trade-off between accuracy and registration time measured in a proper evaluation. Before an image is accepted, the module should verify that the enrollment number is unique, the name is not empty, a face is actually detected, and the image is not empty or corrupted. Captured images should include deliberate variation in expression, distance, and head angle while keeping the face clearly visible.

As a storage estimate, about 50 grayscale face crops of roughly 200 by 200 pixels at around 15 to 20 KB each occupy on the order of 1 MB per student, so a cohort of about 500 students needs on the order of 500 MB before the model file, which is typically much smaller than the image set.

B. Face-Model Training

Training reads the captured images, detects the facial region in each, converts it to grayscale, and attaches the student's numeric label. The LBPH recognizer is trained on the labeled regions and the model is written to a local file. Good practice includes validating the dataset for missing files, duplicate identifiers, and empty samples, recording the number of valid images per student, and tagging the saved model with a version or creation date so that later behavior can be traced to a specific training run. Adding a new student requires retraining.

C. Attendance Marking and Decision Rule

During marking, each webcam frame is converted to grayscale, faces are detected with the Haar Cascade classifier, each facial region is passed to the LBPH recognizer, and the predicted label and its distance value are compared with a configured threshold. The decision is expressed as

$$Accept(d) = 1 \text{ if } d \leq \tau; 0 \text{ if } d > \tau \quad (1)$$

where d is the recognition distance and τ is the threshold. In OpenCV's LBPH implementation a lower distance indicates a closer match, so τ acts as an upper bound; because the direction of `predict()` output is not guaranteed across major releases, this convention should be re-verified for the version in use. Typical accepted-match distances are often reported in roughly the 40 to 60 range for face-sized crops at moderate resolution, but this depends on dataset and preprocessing and is only a starting point for tuning τ on a held-out validation set. If the prediction clears the threshold the interface shows the identity; otherwise it reports an unknown face instead of guessing. Attendance is written only for accepted predictions, after a check for an existing record for the same student, subject, and date. A low-confidence rejection and a duplicate rejection are different outcomes and should be shown differently to the operator, because the first may require a retry or manual entry while the second needs no action.

D. Attendance Reporting and Storage

Attendance is stored in CSV files per subject, with the subject name, date, and time. A richer record would also hold the enrollment number, student name, the recognition distance, a status of accepted, rejected, or manually corrected, and an operator identifier. The recognition distance should always be read together with the threshold active at that time, not as a percentage confidence. CSV suits a small prototype but is weak on concurrent access, complex queries, transactional consistency, access control, backup, and audit trails; an institutional deployment is better served by a relational database. Within CSV, a manual correction should be appended as a new row that references the original, so the correction history is preserved.

E. Claimed Benefits and Their Limits

The synopsis lists fast marking, fewer manual errors, digital records, ease of operation, and low cost as advantages, and lists dependence on lighting, reduced performance with partially covered faces, the need for a webcam, and the need for initial registration and training as limitations. One claim deserves restraint. The synopsis states that the system eliminates proxy attendance, but a more defensible statement is that tying the record to a biometric check may reduce some forms of proxy attendance without removing every way a determined student could circumvent it. Likewise, a single accepted recognition does not prove continuous presence for an entire session. Finally, a working prototype interface has been demonstrated, yet no formal accuracy trial on a registered cohort has been reported, so accuracy and time-saving figures remain design objectives.

VII. ROBUSTNESS TO OCCLUSION AND DIFFICULT CONDITIONS

Masks, spectacles, shadows, expression changes, and inconsistent camera angles all degrade recognition to some degree, and students arrive with masks, glasses, scarves, or turned three-quarters away from the camera far more often than curated benchmarks suggest. Work on masked face recognition has examined synthetic mask-augmented training data, feature learning focused on the visible upper face, and methods that preserve accuracy when much of the face is hidden [28]. Unsupervised representation learning is relevant because it points toward robustness gains without collecting large amounts of newly labeled data at the deployment site [29]. The FRCSyn challenge goes further by using entirely synthetic training data to improve generalization and fairness without collecting additional images of real people, which sidesteps some consent and retention questions [30].

Better augmentation or a stronger model narrows the gap but does not remove the need for an operational safeguard: the system must be able to signal low confidence and defer to a person instead of forcing a guess.

VIII. ACCURACY, BIAS, AND RESPONSIBLE USE

Recognition accuracy is not uniform across demographic groups or operating conditions. An early and influential audit by Buolamwini and Gebru found substantial accuracy disparities by skin type and gender across commercial gender-classification systems, which helped establish demographic auditing as an expectation for face-analysis systems [32]. At much larger scale,

the NIST Face Recognition Vendor Test found demographic differentials in both false-positive and false-negative rates across a large share of the evaluated algorithms [31], continuing a tradition of third-party vendor testing that dates back at least to 2002 [34]. NIST's later reports on summarizing demographic differentials find broadly similar patterns in newer algorithms [35]. Albiero et al. further show that race-related error differences vary with the algorithm, task, threshold, and image difficulty, so a differential measured for one system cannot simply be assumed for another [33].

For attendance, the consequences are concrete. A false positive can credit attendance to the wrong student, and a false negative can fail to credit a student who was present. Both matter if attendance feeds into grading or eligibility. Sensible designs therefore retain the recognition distance for every decision, keep a manual-correction path open, and report performance across the groups relevant to the deployment site wherever that is ethically and legally appropriate.

IX. PRIVACY, SECURITY, AND ETHICAL CONSIDERATIONS

Facial images and the templates derived from them are biometric data. Before deployment an institution should settle the purpose of collection, the legal or institutional basis, the retention period, who may access the data, and how a student can request correction. Safeguards should be mapped onto the regulation that actually applies. A deployment in the European Union would need to satisfy the General Data Protection Regulation's treatment of biometric data as a special category requiring explicit consent and a documented legal basis. In India, the Digital Personal Data Protection Act governs processing of personal data, with extra care warranted if any enrolled student is a minor. At a United States K-12 institution, FERPA and, for students under 13, COPPA would need to be checked. These three are illustrative, not exhaustive.

Recommended safeguards include telling students plainly what the system does, obtaining consent or institutional authorization before enrollment, storing images, models, and records securely with role-based access and encryption where feasible, collecting no more images than the task needs, defining a deletion procedure for students who leave, logging model updates and manual corrections, requiring human review for disputed or low-confidence results, and never repurposing the system for unrelated monitoring. A non-biometric alternative should exist for students who cannot or do not want to be enrolled, because a failed match must never be read automatically as an absence.

X. EVALUATION METHODOLOGY

Because published attendance systems often omit rigorous testing, a consistent protocol is needed. Table II lists metrics that together describe detection, recognition, error behavior, and operational value.

TABLE II
EVALUATION METRICS

Metric	What it captures
Face-detection rate	Share of frames in which a present face is located
Recognition accuracy	Correct identifications over total predictions
Precision / recall / F1	Per-identity correctness and coverage
False-acceptance rate	Rate of an incorrect identity being accepted
False-rejection rate	Rate of a correct identity being rejected
Attendance success rate	Share of accepted predictions correctly logged
Duplicate-entry rate	Repeated records for the same student, subject, and date
Processing time per frame	Latency of detection plus recognition
Class-level marking time	Total time to mark attendance for a full class
Cross-condition performance	Accuracy across lighting, pose, distance, occlusion, and demographic group

For a multiclass task, accuracy and the per-identity precision and recall are

$$\text{Accuracy} = \text{Correct predictions} / \text{Total predictions} \quad (2)$$

$$\text{Precision} = TP / (TP + FP) \quad (3)$$

$$\text{Recall} = TP / (TP + FN) \quad (4)$$

where TP, FP, and FN denote true positives, false positives, and false negatives. Training and test samples must not overlap, and the report should record the number of students enrolled, images per student, camera used, lighting conditions, threshold applied, hardware configuration, and number of sessions covered.

A controlled experiment can proceed by registering a defined group of students with a fixed number of images each; capturing independent test images under normal classroom lighting; repeating tests at different distances and head poses; introducing occlusions such as spectacles or a partial mask; tabulating true and false positives and negatives; comparing automated and manual attendance time for the same class; classifying the likely cause of each failure; checking whether performance differs across relevant user groups in line with the demographic-testing rationale of [31]–[33]; examining whether recognition errors translate into incorrect attendance records; and assessing how useful manual verification is for cases the system flags as uncertain.

XI. COMPARATIVE ANALYSIS

Table III compares, at a qualitative level, the classical pipeline with deep-learning alternatives. The entries describe reported tendencies in the cited literature and are not measurements from the case study.

TABLE III
QUALITATIVE COMPARISON OF RECOGNITION APPROACHES

Approach	Data and compute need	Reported sensitivity
Haar Cascade + LBPH [1], [2]	Small local dataset (tens of images per identity); low compute, real-time on CPU	Sensitive to illumination, pose, and occlusion [12]
MTCNN / RetinaFace + deep embedding [9], [10]	Large labeled corpora; moderate to high compute, GPU beneficial	More robust to pose and scale; heavy occlusion still hard [28]
FaceNet / ArcFace [6], [7]	Large labeled corpora; moderate to high compute	Stronger in unconstrained conditions; needs bias evaluation [23], [31]
Transformer-based, AdaFace [20]–[22]	Large labeled or synthetic corpora; high compute	Competitive or superior accuracy; occlusion and demographic robustness still open [30]

The practical conclusion is that LBPH remains defensible for a controlled single-classroom prototype where cost, local training, and simplicity dominate, whereas embedding-based recognition is the natural upgrade once the classical pipeline has been validated and the institution can support the additional data, compute, and governance load. Added model complexity should be justified by measured improvement in the target environment, not by results from a different setting.

XII. RESEARCH GAPS AND CHALLENGES

Several gaps emerge from this review. Most attendance studies lack independent test sets and report accuracy without stating how the threshold was set or what false-acceptance and false-rejection rates it implies [24]–[27]. Demographic performance is rarely examined at the deployment site even though large-scale evaluations show that differentials are common [31], [35]. Robustness to masks, glasses, and off-angle faces is usually discussed qualitatively and not tested [28]. Storage and auditability are weak in CSV-based prototypes, which limits dispute resolution. Spoofing resistance is seldom addressed, so a photograph or video could in principle be presented to the camera. Finally, direct head-to-head comparisons of LBPH and embedding-based recognition on the same classroom data are scarce, which leaves the choice of method to benchmark leaderboards that do not reflect classroom conditions [23].

XIII. FUTURE DIRECTIONS

Future systems could replace CSV files with a relational or cloud-hosted database, add web or mobile dashboards, email or SMS notifications, and multi-camera support, and incorporate liveness detection, role-based access control, automated backup, an audit trail for corrections, and attendance analytics. On the recognition side, promising directions include better handling of masks and partial occlusion [28]–[30], embedding-based recognition using FaceNet, ArcFace, or related methods [6], [7], quality-adaptive training [20], and synthetic training data that reduces the need to collect real student images [30]. More on-device processing would also keep biometric data on the local machine. Each addition should be accompanied by the measurement and governance steps described above.

XIV. CONCLUSION

This review examined face recognition-based attendance management from detection through recognition, deployed systems, evaluation, and governance. The Haar Cascade and LBPH pipeline is accessible, runs on modest hardware, and is well suited to a controlled academic prototype, as the reference implementation shows, but it is not state of the art and is already well represented in the literature. Its real-world value depends on registration-image quality, camera placement, lighting, pose, threshold selection, and how low-confidence predictions are handled. Because the underlying project reports no experimental measurements, its accuracy and time-saving claims remain design objectives that require controlled validation. Responsible deployment should include privacy protections, access controls, retention rules, demographic testing, and a manual-correction mechanism from the start, and should never rest on the assumption, unjustified for any current system, that automated recognition is infallible.

ACKNOWLEDGMENT

The authors thank the Department of Computer Science and Engineering, D. Y. Patil Agriculture and Technical University, for support during this work.

REFERENCES

- [1] P. Viola and M. J. Jones, "Rapid object detection using a boosted cascade of simple features," in Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR), vol. 1, 2001, pp. I-511-I-518.
- [2] T. Ahonen, A. Hadid, and M. Pietikainen, "Face description with local binary patterns: Application to face recognition," IEEE Trans. Pattern Anal. Mach. Intell., vol. 28, no. 12, pp. 2037-2041, Dec. 2006.
- [3] OpenCV, "Face recognition with OpenCV," OpenCV Documentation. [Online]. Available: https://docs.opencv.org/4.x/da/d60/tutorial_face_main.html. Accessed: Aug. 17, 2026.
- [4] OpenCV, "LBPHFaceRecognizer class reference," OpenCV Documentation. [Online]. Available: <https://docs.opencv.org/3.4/javadoc/org/opencv/face/LBPHFaceRecognizer.html>. Accessed: Aug. 17, 2026.
- [5] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2014, pp. 1701-1708.
- [6] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2015, pp. 815-823.
- [7] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2019, pp. 4690-4699.
- [8] A. K. Jain, A. Ross, and S. Prabhakar, "An introduction to biometric recognition," IEEE Trans. Circuits Syst. Video Technol., vol. 14, no. 1, pp. 4-20, Jan. 2004.
- [9] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," IEEE Signal Process. Lett., vol. 23, no. 10, pp. 1499-1503, Oct. 2016.
- [10] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, "RetinaFace: Single-stage dense face localisation in the wild," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2020, pp. 5203-5212.
- [11] D. Chen, X. Cao, L. Wang, F. Wen, and J. Sun, "Bayesian face revisited: A joint formulation," in Computer Vision - ECCV 2012 Workshops and Demonstrations. Berlin, Germany: Springer, 2013, pp. 566-579.
- [12] X. Zou, J. Kittler, and K. Messer, "Illumination invariant face recognition: A survey," in Proc. 7th IEEE Int. Conf. Autom. Face Gesture Recognit. (AFGR), 2007, pp. 296-301.
- [13] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in Proc. Brit. Mach. Vis. Conf. (BMVC), 2015, pp. 41.1-41.12, doi: 10.5244/C.29.41.
- [14] Y. Sun, X. Wang, and X. Tang, "Deep learning face representation by joint identification-verification," in Adv. Neural Inf. Process. Syst. (NeurIPS), vol. 27, 2014.
- [15] Y. Sun, Y. Liang, X. Wang, and X. Tang, "DeepID3: Face recognition with very deep neural networks," arXiv:1502.00873, 2015.
- [16] W. Liu, Y. Wen, Z. Yu, and M. Yang, "Large-margin softmax loss for convolutional neural networks," in Proc. Int. Conf. Mach. Learn. (ICML), 2016, pp. 507-516.
- [17] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, "SphereFace: Deep hypersphere embedding for face recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2017, pp. 212-220.
- [18] F. Wang, J. Cheng, W. Liu, and H. Liu, "Additive margin softmax for face recognition," arXiv:1704.08063, 2017.
- [19] H. Wang et al., "CosFace: Large margin cosine loss for deep face recognition," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2018, pp. 5265-5274.
- [20] M. Kim, A. K. Jain, and X. Liu, "AdaFace: Quality adaptive margin for face recognition," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2022, pp. 18750-18759.
- [21] Z. Sun and G. Tzimiropoulos, "Part-based face recognition with vision transformers," arXiv:2212.00057, 2022.
- [22] J. Dan et al., "TransFace: Calibrating transformer training for face recognition from a data-centric perspective," in Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV), 2023, pp. 20642-20653.
- [23] M. Wang and W. Deng, "Deep face recognition: A survey," Neurocomputing, vol. 429, pp. 215-244, Mar. 2021.
- [24] J. Akash, V. Kalaiselvan, and A. Harish, "Automated attendance system using face recognition with LBPH," TIJER - Int. Res. J., vol. 12, no. 3, pp. a885-a889, 2025.
- [25] N. Beri, V. Srivastava, and N. Malik, "Face recognition attendance management system using LBPH and Haar cascade," J. Trends Comput. Sci. Smart Technol., vol. 6, no. 3, pp. 257-273, 2024.
- [26] M. M. Sabbir, M. S. Hossain, et al., "Face recognition smart attendance system using deep learning," Procedia Comput. Sci., 2022.
- [27] A. Yadav, S. Singh Malhi, P. Pushp, R. Raushan, S. Singh Yadav, S. Verma, and S. Pandey, "Smart attendance system using face recognition," in Proc. KILBY 100 7th Int. Conf. Comput. Sci. (ICCS), 2023. [Online]. Available: <https://ssrn.com/abstract=4487964>
- [28] Z. Wang et al., "Masked face recognition dataset and application," IEEE Trans. Biometrics, Behav., Identity Sci., vol. 3, no. 2, pp. 1-13, Apr. 2021.
- [29] F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper, "Unsupervised face recognition using deep learning," Electronics, vol. 11, no. 23, p. 3909, 2022.
- [30] I. DeAndres-Tame et al., "FRCSyn challenge at WACV 2024: Face recognition challenge in the era of synthetic data," in Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV) Workshops, 2024, pp. 892-901.
- [31] P. Grother, M. Ngan, and K. Hanaoka, "Face recognition vendor test (FRVT) part 3: Demographic effects," Nat. Inst. Standards Technol. (NIST), Gaithersburg, MD, USA, NISTIR 8280, 2019.
- [32] J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," Proc. Mach. Learn. Res., vol. 81, pp. 1-15, 2018.
- [33] V. Albiero, K. Zhang, and K. W. Bowyer, "How does race bias in face recognition arise?," in Proc. IEEE Int. Conf. Comput. Vis. Workshops (ICCVW), 2020, pp. 1-8.

- [34] P. J. Phillips, P. Grother, R. J. Micheals, D. M. Blackburn, E. Tabassi, and J. M. Bone, "Face recognition vendor test 2002: Evaluation report," IEEE Trans. Pattern Anal. Mach. Intell., vol. 27, no. 1, pp. 1-15, Jan. 2005.
- [35] P. Grother, M. Ngan, and K. Hanaoka, "Face recognition vendor test (FRVT) part 8: Summarizing demographic differentials," Nat. Inst. Standards Technol. (NIST), Gaithersburg, MD, USA, NISTIR 8429, 2022.

Copyright & License:



© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.