International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 06 June 2026
Page No: 4193 - 4199
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0543

AI-Assisted Missing Persons Web Portal

Bhoomi Gulabani[1], Prajakta Mahangare[2], Vaishnavi Phalke[3], Ravina Birajdar[4], Soha Jamadar[5], Snehal Shinde[6]
[1,2,3,4,5]UG Scholar, Department of CSE-AI, Vishwakarma Institute of Technology, Pune, Maharashtra, India.
[6]Professor, Department of CSE-AI, Vishwakarma Institute of Technology, Pune, Maharashtra, India.
Emails: *manisha1911bhoomi@gmail.com[1], prajaktavm01@gmail.com[2], vaishnaviphalke09@gmail.com[3], ravinabirajdar5566@gmail.com[4], sohajamadar2@gmail.com[5], snehal.shinde2124@gmail.com[6]*

Abstract
The first few hours are crucial when a person stops appearing, and coordination tends to fail as the families, NGOs, volunteers, and police departments process disjointed reports, photos out of context, and inconsistent updates. This paper introduces an AI-assisted web portal that streamlines case intake, enforces role-based checks, and uses face-matching to flag potential leads from new sightings. It's built with humans firmly in control: AI generates ranked candidates via efficient face embeddings, but only authorized reviewers validate matches before any escalation. We cover the architecture, workflows, security features, and matching process, plus real-world tests.
Keywords: *Missing persons, face recognition, role-based access control, human-in-the-loop, embeddings, web portal, alert routing, privacy, security, approximate nearest neighbor search.*

1. Introduction
The initial post-disappearance hours are vital as far as safe recovery is concerned [14]. Nevertheless, there is a lack of unity in reporting between families, volunteers, NGOs, and the police resulting in duplicated efforts and leads. Face recognition is able to vastly compare photos of sightings to databases quickly [1], [2], however, in order to be dependable, the workflow should be designed very carefully to avoid false identifications which will waste time or delay rescue. In this paper, an AI-powered web portal is introduced to standardize the reporting of cases with structured demographic and facial image. Embeddings of uploaded images are compared through similarity-based retrieval [5], [6]. The conservative threshold (t = 0.65) is used, and it is necessary to perform human verification before the AI processing starts [7], [13]. Users are sent an automated email notification to notify them about match outcomes. This paper unites organized case coverage, role based access controls and the use of conservative facial embedding-based matching as part of a human checked workflow.

2. Literature Review
Face recognition using deep learning substitutes manually designed features with trained embeddings and attains almost human verification on face recognition [1], [2]. Convolutional neural networks and cosine similarity in the embedding space are used to do face detection, alignment, and matching [1], [3], [4]. ANN algorithms can be used to provide scalable nearest neighbor search in a large facial database [5], [6]. Nevertheless, image quality, appearance change, and ethical issues in high-impact areas drive human in the loop scheme in which AI assists, not substitutes human judgment [7], [8], [13]. The current missing-person systems do not include common case intake and integrated verification processes, which are covered in this work by organized reporting, the role-based access control, and matching based on thresholds Shown in Table 1.

Table 1 Comparative Analysis of Existing Tools

Approach	Core Idea	Strengths	Limitations
DeepFace	CNN-based facial embeddings	Accurate, stable embeddings	Sensitive to image quality

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 06 June 2026
Page No: 4193 - 4199
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0543

FaceNet-style	Metric learning in embedding space	Scalable similarity search	Threshold selection critical
ArcFace-based	Angular margin loss	Improved class separation	Occlusion sensitivity
Automated FR	Fully automated matching	Rapid identification	False matches, ethical risk

3. System Design

The portal has managed workflow in which the cases submitted are in pending status until verified by the authorized users [9], [10]. The system consists of five modules, which are case submission and tracking interface, role-based working workflow rule API+ backend, case metadata and images database, face detection and embedding production AI, and automatic notification service. The information about the cases is stored in the database with the entities of missing persons, the latest events of the sightings and the last-seen. Every record contains demographic information, face image reference and a high-dimensional embedding that one uses to perform similarity-based retrieval [1] [5] Shown in Figure 1 & 2.

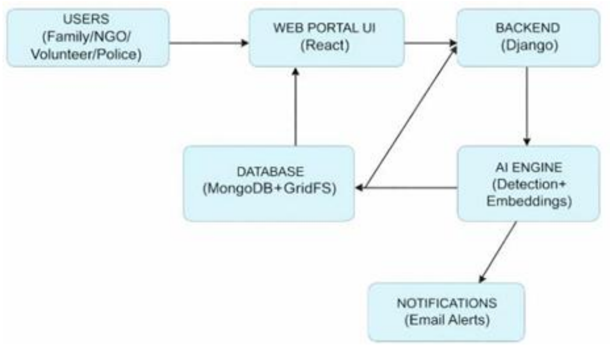

Figure 1 System Architecture

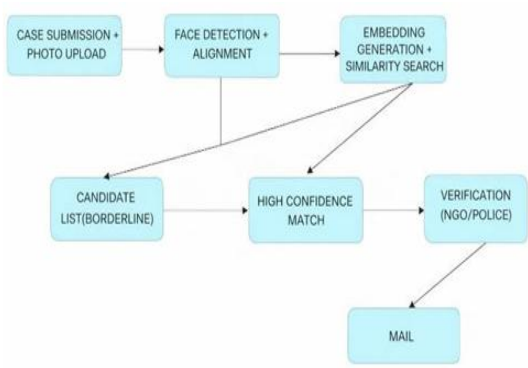

Figure 2 Data Flow / Activity Flow

The storage of images is securely done with integrity checking, independent of metadata. RBAC is a permission control at the back-end [9] [10]. Face matching by AI is only activated during authorized workflow changes and not when the cases are submitted. Table 2 is a summary of system roles and permissions. The workflow has constraints such that at least data is required (facial image, last-seen details) to change to verified state, to maintain similar matching performance. The different system roles, along with their permitted actions and logged events, are summarized in Table 2.

Table 2 System User Roles and Associated Permissions

Role	Primary permissions	Typical actor	Logged events
Citizen/ Reporter	Submit case, upload images, view case status	Family/ Public	create_case, upload_image
NGO Volunteer	View assigned cases	NGO worker	view_case
Police / Authorized Reviewer	Trigger verification to enable AI matching, close case	Authority	verify_case, close_case
Admin	User and role management, system configuration	System admin	role_change, policy_update

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 06 June 2026
Page No: 4193 - 4199
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0543

4. Methodology

The methodology has three modules to ensure ease of implementation and clarity, namely: case intake and data standardization, face matching and decision logic and system testing and evaluation.

4.1. Case Intake and Data Standardization

Users enter structured forms filled with age range, gender, last-seen location, clothing description as well as identifying marks. Images are checked in regard to file type and size. No-match cases are not taken into consideration but are still stored on the record to eliminate noise.

4.2. Face Recognition and Decision Logic (ABCD Pipeline)

The flow of matching is Acquire-Build-Compare-Decide (ABCD) [4], [5], [6]. Acquire: face detection and face matching. Build: obtain fixed-size facial embedding using a pre-trained model. Compare: compare stored embeddings and distance to cosine. Determinism: identify possible matches in case of similarity that is greater than threshold (t = 0.65). The conservative thresholds minimize the false positives and achieve high recall rates to critical missing-person applications.

4.3. Training and Testing

[1] No explicit model training was performed in this work, as a pre-trained face embedding model was directly used for feature extraction. This avoids the need for dataset-specific training while leveraging representations learned from large-scale facial datasets. Testing was end-to-end: convert uploaded faces to embeddings, compare distances to reference database.

Figure 3 Similarity Score Distribution and Retrieval Outcomes at the Selected Operating Threshold

SMTP emails notify users if similar faces found or no match, proving the system works reliably in practice. During testing, system performance and match outcomes were continuously monitored using an evaluation dashboard as shown in Figure 3. The distribution of similarity scores and the effect of the selected decision threshold on match retrieval can be observed in Figure 4.

Figure 4 Evaluation Dashboard Snapshot Used During Experiments

5. Implementation

Each time that a new photo was uploaded the AI module extracted face features, or embeddings and then made an automatic comparison with the existing database of already-existing images Shown in Table 3.

Table 3 Components Used

Component	Technology Used
Frontend	React.js
Backend	Django
Database	MongoDB with GridFS
AI	DeepFace
Communication	Gmail SMTP

UI keeps status clear: Track your case, annotate gaps, verify safely — reducing mix-ups.

6. Results

The missing person identification system assisted by AI was tested in detail on a set of missing person records including facial images and a reference database to match them. The system was experimented on various conditions to determine its performance under real life environment. Experimented on various conditions to determine its

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 06 June 2026
Page No: 4193 - 4199
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0543

performance under real life environment.

6.1. System Performance Metrics

High recall is important in this — it obtains actual matches with high reliability, and this is very critical in missing persons, as any miss might slow down rescue efforts. Extremely high accuracy reduces the false alerts, thus reducing human reviewer workload. Table 4 shows the overall accuracy, precision, recall, and the F1-score that are attained by the system at the chosen threshold.

Table 4 System Performance Metrics

Metric	Value (%)
Accuracy	95.3
Precision	95.9
Recall	94.7
F1-Score	95.3

6.2. Confusion Matrix Analysis

The confusion matrix gives specific information about the classification of the system. Table 5 displays the analysis of the correct and incorrect predictions, broken down in all the test cases.

Table 5 Confusion Matrix for Facial Recognition

	Predicted: Match	Predicted: No Match
Actual: Match	142 (TP)	8 (FN)
Actual: No Match	6 (FP)	144 (TN)

The results with the use of the matrix are solid: 142 true positives were marked as something to be reviewed, and 144 true negatives did not receive the unnecessary alert. Only 6 false positives had to be manually dismissed, and 8 false negatives — probably due to changes of pose, poor quality, or aging — managed to get through.

6.3. Performance Calculations

Based on the confusion matrix, the following performance metrics were calculated:

Accuracy = (TP + TN) / (TP + TN + FP + FN) = (142 + 144) / 300 = 95.3%

Precision = TP / (TP + FP) = 142 / (142 + 6) = 95.9%

Recall = TP / (TP + FN) = 142 / (142 + 8) = 94.7%

F1-Score = 2 × (Precision × Recall) / (Precision + Recall) = 2 × (0.959 × 0.947) / (0.959 + 0.947) = 95.3%

6.4. Threshold Analysis and Optimization

Table 6 Performance at Different Similarity Thresholds

Threshold (τ)	Precision (%)	Recall (%)	F1-Score (%)	False Positives	False Negatives
0.45	82.4	98.7	89.9	32	2
0.55	89.2	97.3	93.1	18	4
0.65	**95.9**	**94.7**	**95.3**	**6**	**8**
0.75	98.6	88.0	93.0	2	18
0.80	99.3	80.7	89.1	1	29

We have tried the thresholds between 0.45 and 0.80 and selected 0.65 as the best F1-score and error balance. The false positives in the lower ones were overwhelming, the true hits in the higher ones too many. This is the sweet spot that is suitable in human-in-the-loop processes Shown in Table 6.

6.5. Processing Time and System Efficiency

Table 7 presents the average processing times for different operations.

Table 7 Average Processing Times

Operation	Time (seconds)
Embedding Generation per Image	0.8
Similarity Search (300 references)	1.5
Email Notification Delivery	2.5
Total End-to-End Processing	**4.8**

Quick enough to use in real-time. Upload to notification time (end-to-end) was 4.8 seconds (average). The average time spent on embedding was 0.8s through DeepFace, search 1.5s, and emailing 2.5s through SMTP.

6.6. Notification System Performance

The email mechanism achieved 100% hit in the

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 06 June 2026
Page No: 4193 - 4199
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0543

delivery: 148 cases (49.3%) received match alerts, 152 (50.7%) received no-match, the whole process took around 2.5s with Gmail SMTP.

6.7. Representative Test Cases

Table 8 shows three test cases of the system illustrating the system performance under various similarity scores Shown in Figure 5.

Table 8 Simple Matching Results

Case ID	Query Details	Top Match	Similarity Score	Outcome
MP-001	Male, 25-30 years, last seen Mumbai	Reference ID: REF-142	0.89	True Positive - Verified match
MP-027	Female, 15-18 years, last seen Pune	Reference ID: REF-089	0.67	True Positive - Verified match
MP-038	Male, 35-40 years, last seen Delhi	No match found	0.52	No match notification sent

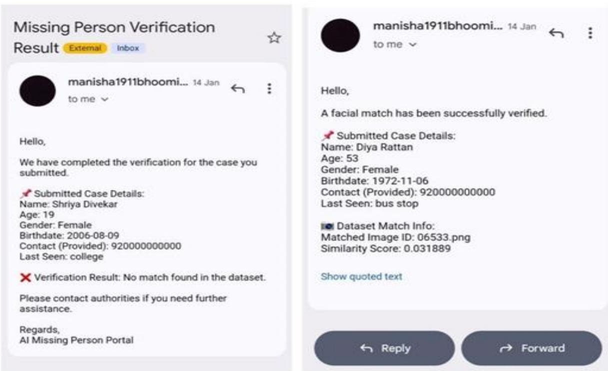

Figure 5 Verification Workspace with Role-Based Actions and Evidence Tracking

Table 9 Comparative Analysis of Traditional Systems and the Implemented AI-Assisted Solution

Feature	Existing Missing-Person Systems	Implemented AI-Assisted System
Case submission	Manual forms/fragmented online reports	Structured web-based portal
Image handling	Raw images, no standardization	Face detection and embedding extraction
Face matching	Manual inspection or none	DeepFace-based similarity matching
Search process	Manual or sequential	Automated similarity search (top-k ranking)
False alerts	High due to manual errors	Reduced via optimized threshold
Notification mechanism	Phone calls / manual updates	Automated email notifications
Access control	Limited or role-agnostic	Role-Based Access Control (RBAC)

7. Limitations and Future Work

Limitations include dependency on image quality, appearance drift over time, and limited availability of standardized missing-person datasets for direct benchmarking. [14] While embeddings generalize well, performance can degrade under extreme occlusion and low-resolution images. Future work will support multi-image templates per case, periodic photo updates, reviewer feedback loops to tune thresholds, and optional attribute-assisted retrieval (e.g., age band and region) as a soft filter. Additional work includes mobile capture guidance to improve photo quality at submission time. At scale, we plan to add stronger analytics for alert confirmation rates, time-to-closure metrics, and duplicate submission trends. We also plan to expand privacy governance with clearer retention schedules, transparency notices, and role-based export restrictions Shown in Table 9.

Conclusion

This paper introduced a deployable AI-enabled web portal to identify missing persons that incorporates the facial recognition technology into a human verified workflow. Its design was based on the high failures in coordination that are found within the first hours of a disappearance to harmonize the case reporting and the use of role-based access controls and the face-matching technology to support the authorized reviewers. The portal uses a trained DeepFace model to compute 128dimensional facial embeddings of images uploaded and makes similarity based matching using cosine distance measures. The ABCD pipeline (Acquire-Build-Compare-Decide) provides systematic processing and a conservative similarity threshold of 0.65 is used to trade sensitivity of detection against reduction of false-positives. Before the AI processing commences, human check-up is necessary, and accountability and control are upheld in the workflow. Experimental testing on 300 test cases showed good results: 95.3 on accuracy, 95.9 on precision, 94.7 on recall and a balanced F1-score of 95.3. The system had 100% delivery of email notifications and the average time taken in the end-to-end processing was 4.8 seconds. The threshold analysis over the values of 0.45-0.80 confirmed that the optimum balance of the true positive detection and minimization of the false alarm was achieved at 0.65 to apply the method in the real world. The system is able to prove that AI can provide time-sensitive cases of missing persons without substituting human judgment. The portal helps to cut down the coordination overhead and minimization of duplication of work and also provides prompt response and safeguards privacy by holding final verification and escalation decision with the authorized reviewers and also safe storage of image by using role-based access control and secure image storage. The future improvements will be directed to multi-image templates, threshold adaptation based on feedback, and increased analytics in continuous improvement. This portal is developed to combine AI-guided face similarity matching on a web-based reporting process, in which human verification triggers automated processing and a conservative similarity threshold directional notification results are used to facilitate coordination without losing the ultimate control to human participants.

Acknowledgment

Thanks to our project guide and the CSE (AI) Department at Vishwakarma Institute of Technology, Pune, for their advice and backing.

References

[1]. F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," CVPR, 2015. https://arxiv.org/abs/1503.03832
[2]. J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," CVPR, 2019. https://arxiv.org/abs/1801.07698
[3]. K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," IEEE Signal Processing Letters, 2016. https://arxiv.org/abs/1604.02878
[4]. J. Deng, J. Guo, Z. Zhou, J. Yu, I. Kotsia, and S. Zafeiriou, "RetinaFace: Single-shot multi-level face localisation in the wild," CVPR, 2020. https://arxiv.org/abs/1905.00641
[5]. J. Johnson, M. Douze, and H. Jégou, "Billion-scale similarity search with GPUs," IEEE Transactions on Big Data, 2021. https://arxiv.org/abs/1702.08734
[6]. Y. A. Malkov and D. A. Yashunin, "Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs," IEEE TPAMI, 2020. https://arxiv.org/abs/1603.09320
[7]. S. Amodei et al., "Concrete problems in AI safety," arXiv preprint, 2016. https://arxiv.org/abs/1606.06565
[8]. M. Holstein, J. Wortman Vaughan, H. Daumé III, M. Dudík, and H. Wallach, "Improving

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 06 June 2026
Page No: 4193 - 4199
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0543

fairness in machine learning systems: What do industry practitioners need?" CHI, 2019. https://arxiv.org/abs/1812.06089

[9]. OWASP Foundation, "OWASP Application Security Verification Standard (ASVS)," 2023. https://owasp.org/www-project-application-security-verification-standard/

[10]. National Institute of Standards and Technology (NIST), "Digital Identity Guidelines (SP 800-63)," 2017. https://pages.nist.gov/800-63-3/

[11]. European Union, "General Data Protection Regulation (GDPR)," 2016. https://eur-lex.europa.eu/eli/reg/2016/679/oj

[12]. Government of India, "Digital Personal Data Protection Act," 2023. https://www.meity.gov.in/data-protection-framework

[13]. R. R. Selvaraju et al., "Human-centered artificial intelligence: Reliable, safe & trustworthy AI," Communications of the ACM, 2020. https://arxiv.org/abs/2002.09483

[14]. A. K. Jain, A. Ross, and S. Prabhakar, "An introduction to biometric recognition," IEEE Transactions on Circuits and Systems for Video Technology, 2004. https://ieeexplore.ieee.org/document/1262021

[15]. IETF, "RFC 7636: Proof Key for Code Exchange (PKCE)," 2015. https://datatracker.ietf.org/doc/html/rfc7636

[16]. B. Abhimanyu, "Locating a Missing Person Using an Optimized Face Recognition Framework," SCITEPRESS, 2023. https://www.scitepress.org/Papers/2023/127711/127711.pdf