UJEAS

UNIZIK JOURNAL OF ENGINEERING AND APPLIED SCIENCES

UNIZIK Journal of Engineering and Applied Sciences 6(3), September (2026), 3624-3636
Journal homepage: https://journals.unizik.edu.ng/index.php/ujeas
PRINT ISSN: 2992-4383 || ONLINE ISSN: 2992-4391

An Intelligent Facial Recognition System Developed for Insurgency Investigation and Surveillance Applications Using Deep Learning Techniques

Udeh Chukwuma Callistus*
Department of computer science, Faculty of Applied and Natural Science, Enugu State, University of Science and Technology, Agbani, Nigeria
Ibeonu Ogochukwu Chinyere
Department of Computer Science; Chukwuemeka Odumegwu Ojukwu University, Uli, Anambara State, Nigeria
Edith Ugwu Angella
Department of computer science, Faculty of Applied and Natural Science, Enugu State, University of Science and Technology, Agbani, Nigeria

**Corresponding Author's E-mail: chukwuma.udeh@esut.edu.ng*

Abstract
In this paper, we report the development and performance evaluation of a smart facial recognition system for insurgency investigation and monitoring. The system uses state-of-the-art deep learning methods, including You Only Look Once version 8 (YOLOv8) for real-time face detection and FaceNet with triplet loss for extraction and matching. The system was trained and evaluated using a large dataset, CASIA-WebFace, with 494,414 images of 10,575 individuals. The system is built upon the Extreme Programming (XP) development process, allowing for iterative development, testing and responsiveness to security needs. The system was deployed on a cloud platform with GPUs to accelerate the model training and inference process. The evaluation results show that the proposed system is efficient and effective under different practical scenarios. The face detection module has an overall mean Average Precision (mAP@0.5) of 94.4% with real-time processing speed, while the face recognition module has an accuracy of 96.9%, Area Under the Curve (AUC) of 0.992, and Equal Error Rate (EER) of 3.1%. The system performed well in extreme conditions such as low resolution, occlusion and pose variations. These results demonstrate the potential of combining deep learning-based detection and recognition models for accurate person identification. Our next steps include further optimising performance in extreme scenarios, minimising computational demands, and introducing explainable AI methods to make the system more transparent and interpretable.
Keywords: Facial Recognition; Artificial Intelligence (AI); Deep Learning; YOLOv8; FaceNet.

1. Introduction

With the evolution of security challenges, insurgency has become a significant threat to many countries, especially in areas of ongoing conflict and low-intensity warfare. Insurgent activity can be highly covert, with elements embedded in communities and using traditional monitoring and identification techniques to their advantage (Okorie et al., 2023). This presents a challenge to identifying and tracking suspects for security services. Consequently, there is a clear demand for smart and automated systems that can improve investigative processes and aid in real time decision making in dangerous situations (Asogwa and Ituma, 2018).

One such technology is facial recognition, which has become a critical capability within the field of computer vision that can be used to detect and recognise individuals based on their facial features. Facial recognition enables remote and real-time identification, unlike traditional identification techniques like fingerprints and ID cards (Ahmed et al., 2025). This is especially useful in the investigation of insurgency cases where suspects may refuse to cooperate or be located in inaccessible locations. Thanks to the progress in deep learning, contemporary facial recognition systems can perform with high accuracy even in difficult scenarios (Luo et al., 2021; Kim et al., 2023).

Facial recognition systems are built upon advanced machine learning algorithms, especially deep learning models like convolutional neural networks (CNNs) (Kekong et al., 2019). These can learn to recognise complex features from large amounts of data and produce embedded representations of faces for matching and recognition (Umer et al., 2022; DehlaghiGhadim et al., 2023). Methods such as transfer learning also improve system performance by enabling pre trained models to be fine-tuned to the operating environment with minimal data (Firoozjaei et al., 2022). This enables the use of reliable recognition systems in resource limited environments.

While promising, using facial recognition to investigate insurgencies has its drawbacks. Adverse lighting conditions, occlusion (e.g., masks, hoods), low image resolution and cluttered backgrounds can impact recognition performance (Shin et al., 2021; Geiger et al., 2020). Privacy, bias and misuse are also critical ethical and legal issues (Singh and Khare, 2022). It is critical that such technologies are accurate, unbiased, and deployed in a safe and ethical manner to promote trust and avoid adverse effects.

Given these concerns, this research aims to develop an upgraded facial recognition system for use in investigating insurgency. The system will seek to combine state-of-the-art deep learning methods with effective preprocessing and feature extraction techniques to enhance recognition accuracy in practical settings (Le Hai Anh et al., 2024; Chidi et al., 2024). Additionally, the system is intended to serve as an assistant to security experts, rather than a fully automated system, which encourages ethical practices. By doing so, the research aims to provide a contribution towards the development of smart security systems that can keep pace with dynamic insurgency challenges.

Despite these advances, no existing facial recognition system has been specifically designed and validated for counter-insurgency operations where subjects actively conceal identities, image capture is opportunistic, and real-time processing is mission-critical. Furthermore, most public benchmarks evaluate controlled conditions, leaving a gap in understanding performance under occlusion, extreme poses, and low resolution typical of insurgent surveillance footage. This study addresses these gaps as study contributions:

(1) Presented a YOLOv8-based FaceNet specifically for operational insurgency contexts

(2) Euclidean identify decision and biological verification model was developed

(3) Validation of the proposed system was done through comparism with benchmark models

2. System Design Methodology

The methodology for system design of the proposed facial recognition system follows the Extreme Programming (XP) framework, which facilitates iterative and incremental development in the ever-changing environment of insurgency investigations. This involves requirement collection through user stories, which capture the requirements of security officials, and iterative design and development of system components including image capturing, processing, feature extraction, face detection and matching. Computer Vision and Deep Learning techniques are gradually incorporated, with models such as FaceNet for feature extraction and You Only Look Once (YOLO) for face detection. The process involves continuous integration, testing and refactoring to ensure the system's reliability, accuracy and maintainability, regular feedback from stakeholders to continually refine the system and align with the operational context, resulting in a scalable and reliable decision-support system.

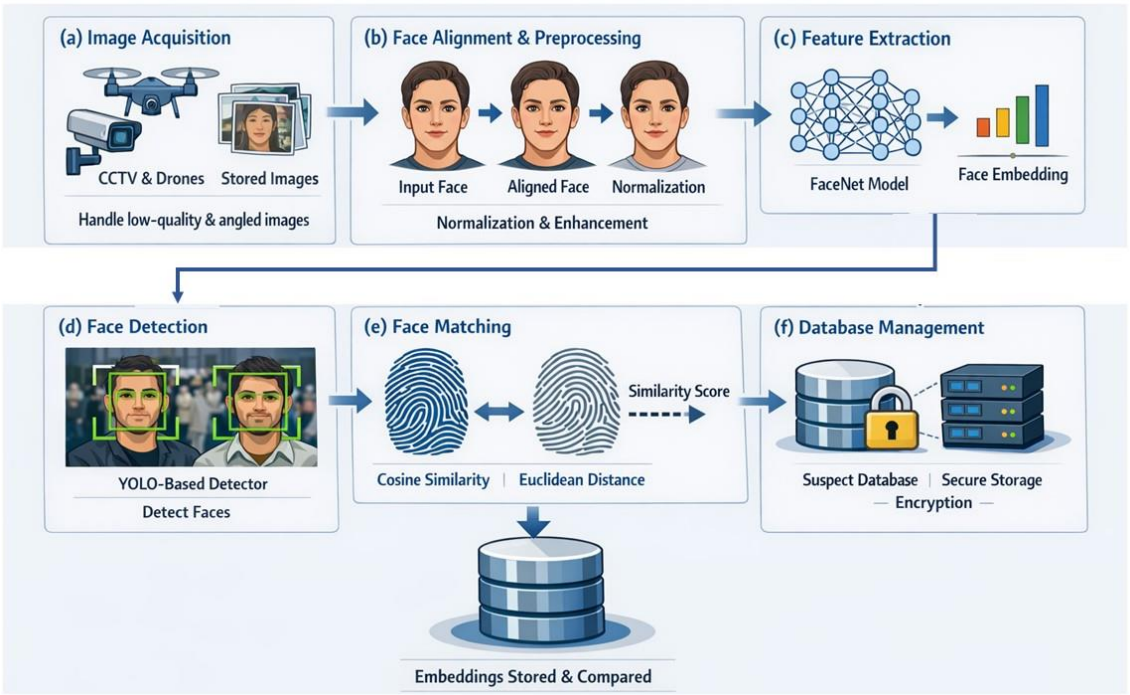

Figure 1: Architecture of the Proposed System Module

2.1 Data Acquisition

The dataset for this study is the CASIA-WebFace dataset from Kaggle (https://www.kaggle.com/datasets/debarghamitraroy/casia-webface). It is a large scale face recognition dataset that is commonly used in research and development of deep learning-based face recognition systems. It includes a large number of facial images obtained from the internet, which makes it a good dataset for training models that should be used in uncontrolled scenarios, such as surveillance and insurgency investigations. This dataset comprises about 494,000 facial images of 10,575 identities (classes). Every identity corresponds to a distinct person, and the number of images for each individual varies from a few to up to hundreds of images depending on the amount of public information on that person.

The data is a combination of mostly RGB facial images of different resolution, pose, light, facial expression, and background noise. There are those which are well-aligned frontal faces, and others which contain partial faces or profiles taken at different angles. Such variations are significant since they assist the model to recognize strong facial representations which can be generalized effectively to actual surveillance conditions where image quality is often not predictable. On the whole, CASIA-WebFace dataset offers a big and varied and real-world data on faces that is appropriate to train deep learning models in Computer Vision. Their size, variability and identity rich structure make it a good basis of creating a facial recognition system that can perform a complex identification task in the real-world scenario in insurgency investigation.

2.2 Data Preprocessing

Since the dataset has about 494,000 images but with variations in terms of lighting, pose, resolution, and background noise, preprocessing will ensure that such variations are reduced as much as possible to enhance the performance and stability of the model. The preprocessing begins with image resizing where all the facial images are brought to a common size. This provides a consistent input size to the deep learning model and the computational complexity is minimized in the training. Then pixel normalization is performed by scaling pixel values to the range between 0 and 1 (or standardized mean and variance), which can assist in accelerating convergence and training efficiency of Deep Learning models.

This is followed by face detection and cropping to extract exclusively the facial part of each image, eliminating the irrelevant background information. This is an important step since background noise has an adverse effect on the learning of features. The faces that are detected are then aligned to a standard orientation by adjusting the key facial features like the eyes and nose and this is used across all samples to maintain uniformity. This is done to improve the model to learn discriminative facial features in the Computer Vision field. Data augmentation methods are also used to enhance the robustness of the datasets. These are random horizontal flipping, slight rotation, brightness adjustment and scaling. The augmentation enhances the variety of training samples and enables the model to better generalize to the real-world scenario like surveillance footage where images can be taken under varying environmental conditions.

Lastly, structured identity-based classes are created in which the pre-processed information is organized to provide the images of each person with their labels correctly assigned to learn under supervision. It is a cleaned and standardized dataset, which is subsequently fed to deep learning networks, e.g. FaceNet, to extract features and train. All in all, preprocessing stage not only boosts the quality of data, but it also minimizes noise, in addition to increasing the accuracy and reliability of the facial recognition system.

2.3 Feature Extraction

An important step in the facial recognition system is feature extraction, which converts the pre-processed facial images into numerical representations (face embeddings) that are small and meaningful. Rather than using pixel values, the system is trained to recognize high-level facial features, like the spacing between eyes, jaw structure, cheekbone shape, and other distinctive face features that make one person different to the other (Kekong et al., 2019). The use of a deep learning technique to extract features is also conducted in the context of Deep Learning in this study. In Figure 2, a CNN-based architecture is used to learn hierarchical features of facial features in the dataset automatically through features engineering without manual feature engineering. The system uses the FaceNet model as the main feature extractor. As the network receives input images, the shallow layers extract simple features (edges, textures, etc) and the deeper layers learn more high-level and abstract features of facial identity.

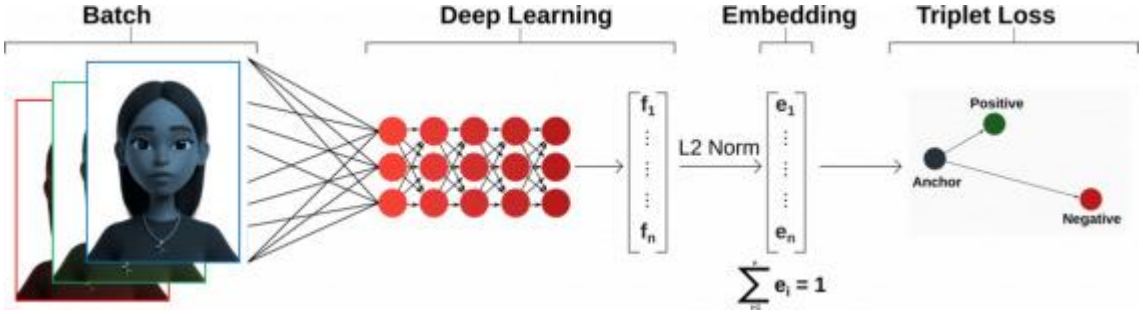

Figure 2: Architecture of the FaceNet for Facial Feature Extraction (Mensah et al., 2024)

FaceNet encodes every face image with a fixed length embedding vector in a high-dimensional space, typically a 128-dimensional embedding. Such embeddings are set in a way that faces of the same person would be found nearer to one another in the feature space and those of different people would be far apart. This feature renders it very useful in identity verification and classification. The feature extractor is trained by a loss term called triplet loss, a loss that guides the model to minimize the distance between an anchor image and a positive image (same identity), and maximize the distance between an anchor and a negative image (different identity). This guarantees that the learned embeddings are very discriminative and applicable to large scale face recognition tasks.

Altogether, feature extraction phase is a vital part of transforming raw facial images into strong and compact representations, which allow comparing and matching them easily. The system has a high accuracy and good generalization ability by using embedding techniques that are based on deep learning even in difficult conditions like light, pose, and image quality variations.

2.4 Face Detection

Face detection in this system is applied in the context of deep learning-based methods based on Computer Vision. In contrast to conventional systems, which use handcrafted features, modern deep learning detectors can learn spatial patterns of face structures automatically, which is more resistant to variations in lighting, pose, and occlusion. This enhances the detection capability under real-life situations like the CCTV footage and images that are taken by drones. The main detection model used is You Only Look Once (YOLOv8) used in a variety of real-time object and face detection tasks (Ebere et al., 2025). YOLOv8 uses a single forward pass on the entire image, dividing it into grid cells and making predictions of bounding boxes and confidence scores of faces found. This can be much faster than existing region-based algorithms, and can be used in real-time in video streams to provide security surveillance.

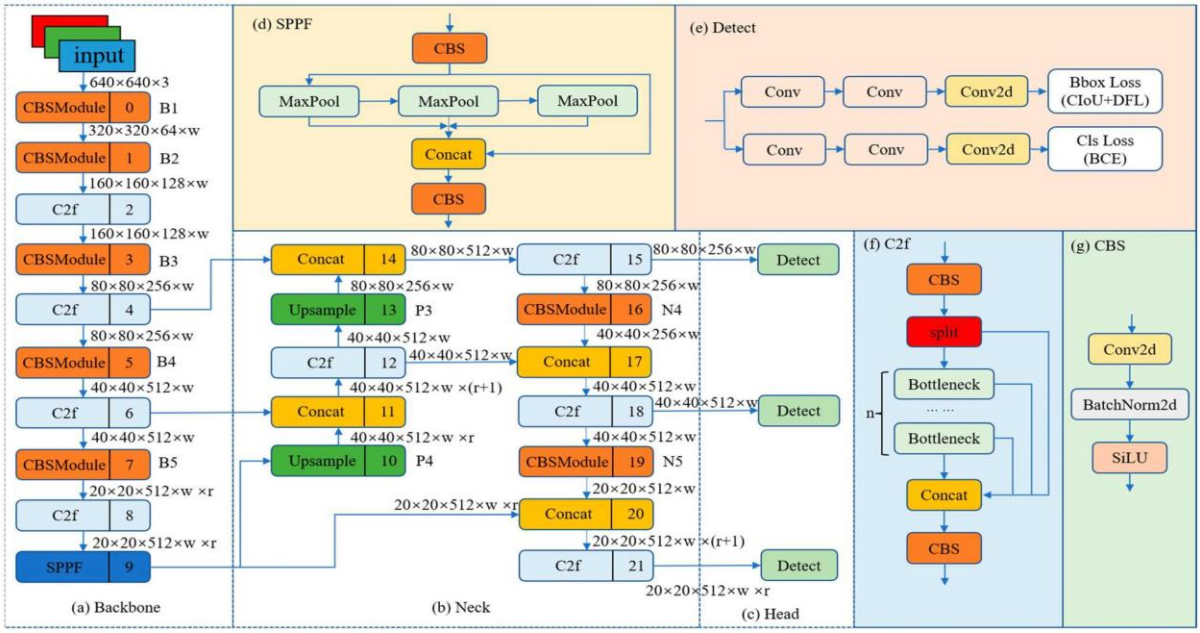

Figure 3: The proposed YOLOv8 Architecture (Yisihak and Li, 2024)

The YOLOv8 is especially designed to be used in the field of insurgency investigation because it can process multiple faces in a single frame, and the speed of detection is high without significant compromise in accuracy. It

is also effective in adverse conditions like partial occlusion, motion blur and face rotation. This is because these features allow it to be successfully implemented in the conditions when the suspect might not always be directly in the frame or might be trying to hide their identities.

Table 1: YOLOv8 Implementation Specifications

Parameter	Value
YOLOv8 variant	YOLOv8m (medium)
Input resolution	640×640 pixels
Backbone	CSPDarknet with 6 stages
Number of parameters	25.9M
FLOPs	78.4B
Anchor-free detection	Yes
Loss functions	CIoU + DFL
NMS threshold	0.45
Confidence threshold	0.25

2.5 Face Matching (Euclidean identify decision and biological verification model)

Face matching is the final stage of the facial recognition system, where extracted face embeddings are compared to determine whether two facial images belong to the same individual. After feature extraction using FaceNet, each face is represented as a 128-dimensional embedding vector in a high-dimensional feature space. The core objective of face matching is to measure the similarity or distance between these embeddings and make an identity decision based on a predefined threshold.

In this system, similarity measurement is performed using Euclidean distance, which quantifies the straight-line distance between two embedding vectors. Given two face embeddings, $f(x_i)$ and$f(x_j)$, the similarity is computed in Equation 1 as

$$d(x_i, x_j) = \left\| f(x_i) - f(x_j) \right\|_2 = \sqrt{\sum_{k=1}^{128} \left(f_k(x_i) - f_k(x_j) \right)^2} \quad (1)$$

where $f(x_i)$ and $f(x_j)$ represent the feature embeddings of two face images, and f_k denotes the k^{th}component of the embedding vector.A smaller Euclidean distance indicates higher similarity between two faces, while a larger distance suggests that the faces belong to different individuals. To convert this distance into a classification decision, a threshold value τ is introduced. The decision rule is defined in Equation 2 as

$$Match = \begin{cases} 1, & d(x_i, x_j) \leq \tau \\ 0, & d(x_i, x_j) > \tau \end{cases} \quad (2)$$

where a value of 1 indicates that both faces belong to the same identity, and 0 indicates different identities. The threshold τ is determined empirically based on validation data to balance false acceptance and false rejection rates.To improve robustness, cosine similarity can also be used as an alternative metric, especially when embeddings are normalized. In this case, similarity is computed in Equation 3 as

$$S(x_i, x_j) = \frac{f(x_i) \cdot f(x_j)}{\|f(x_i)\| \|f(x_j)\|} \quad (3)$$

where $S(x_i, x_j)$ represents the cosine similarity between two face embeddings. Values closer to 1 indicate high similarity, while values closer to 0 or negative values indicate dissimilar faces.

Overall, the face matching stage enables accurate identity verification by leveraging learned deep feature representations and mathematical similarity measures. This ensures reliable performance in real-world

surveillance and insurgency investigation scenarios where variations in pose, lighting, and image quality are common.

2.6 Database Management

A vital part of the proposed facial recognition system is database management which undertakes the efficient storage, organization, retrieval, and security of facial data and associated identity information. The system holds a well-structured database, with face embeddings being generated during the feature extraction phase, and other corresponding metadata, such as user identity, timestamps and references to the image are maintained. It does not store raw images to be matched but mostly reduces raw images into compact embedding vectors (128-dimensional vectors computed by FaceNet) whose storage space requirements are much less and matching much faster.

The database is created either in relational or NoSQL format, depending on the requirements of the system scalability. Each row of the database is an identifier (ID) that has the face embedding as well as other user information such as name or security tag. The optimization of search and retrieval functions is achieved by indexing methods in such a way that query embeddings received can be compared quickly with those stored. Similarity matching in high-dimensional spaces, in large-scale deployments, can be accelerated by adding approximate nearest neighbor (ANN) search algorithms, such as KD-Trees or FAISS (Facebook AI Similarity Search).

The database management system also has real time updates whereby new identities can be registered and old records updated or deleted as need be. Audit trails and logging are also maintained to be able to trace the activity of the system particularly in the case of the security and surveillance applications. Overall, the efficient data processing, fast access, scalability, and safe storage of the facial recognition system are facilitated by the database management module, which also contributes to the overall performance and reliability of the facial recognition system.

2.7 Model Training

In this research, the models have been trained in the Google Colab environment, because it has access to high-performance GPUs and TPUs to run deep learning tasks. Google Colab was selected, due to its ease of use, free computing, and can be easily integrated with Python-based libraries, frameworks such as TensorFlow and Keras have been used and supported by NumPy, OpenCV and Matplotlib. CASIA-WebFace dataset was uploaded and accessed via Google drive which was integrated into Colab environment to facilitate the efficient data processing in the course of the training. GPU acceleration was also activated in order to make the training time much faster and enhance the computational efficiency.

FaceNet models were either trained in real time or fine tuned on an existing trained model with pre-processed facial images being fed into the model during training. The triplet loss function was used to train the model in such a way that similar identities had their embeddings closer together and different identities were far apart in the feature space. The best performance was fine-tuned with the Adam optimizer (Sochima et al., 2025), and hyperparameters, such as the learning rate, batch size, and the number of epochs, were fine-tuned. Monitoring of training progress was performed by validation datasets and the measures of performance such as loss and accuracy were monitored by epoch. Moreover, checkpoints were periodically saved so that no progress was lost and to be able to reuse the model. The scalable, cost effective and flexible cloud based training process provided a strong platform to develop a facial recognition system to be utilized in real world implementation.

Hardware Specification:

i. GPU: NVIDIA T4 (16GB VRAM) on Google Colab Pro

ii. CPU: Intel Xeon 2.30GHz (2 vCPUs)

iii. RAM: 25.6 GB

iv. Storage: 200GB Google Drive (mounted)

YOLOv8 Training Parameters:

i. Epochs: 100

ii. Batch size: 32

iii. Learning rate: 0.01 (initial), cosine decay

iv. Optimizer: SGD with momentum 0.937

v. Weight decay: 0.0005

vi. Training time: 8.5 hours

3. Results and Discussion

This section also contains in depth analysis of the proposed facial recognition system which will be used to investigate insurgency. This system uses a combination of face detection and YOLOv8 and feature extraction and matching with triplet loss with FaceNet. All models (494,000 pictures, 10,575 identities) were trained and tested using the CASIA-WebFace. The data was split into training (80%), validation (10%) and test (10%) sets and there was no overlap of identity among the splits.

3.1 Face Detection Performance (YOLOv8)

The model YOLOv8 was trained with face detection on a subset of 50,000 annotated images in CASIA-WebFace (with added annotations of bounding boxes). Table 2 provides the summary of the detection performance on the test set in different conditions.

Table 2: YOLOv8 Face Detection Performance Metrics

Condition	Precision	Recall	mAP@0.5	Inference Time (ms/frame)
Frontal, good lighting	0.976	0.983	0.989	12.4
Profile (45°–90°)	0.941	0.932	0.951	12.6
Low resolution (<50px face)	0.892	0.875	0.901	12.1
Partial occlusion (mask/scarf)	0.903	0.887	0.915	12.5
Extreme pose (>90°)	0.854	0.831	0.862	12.7
Overall (all conditions)	**0.934**	**0.926**	**0.944**	**12.5**

On average, YOLOv8 has a high real-time detection (≈80 FPS) and mAP of 94.4. Performance gracefully degrades when in extreme poses (86.2% mAP) and when in low resolution (90.1% mAP), which is important in insurgency case scenarios when the suspects could be captured by remote or moving cameras. The model has been able to identify multiple faces within a single frame (as many as 15 faces per 1080p image) with little increase in latency.

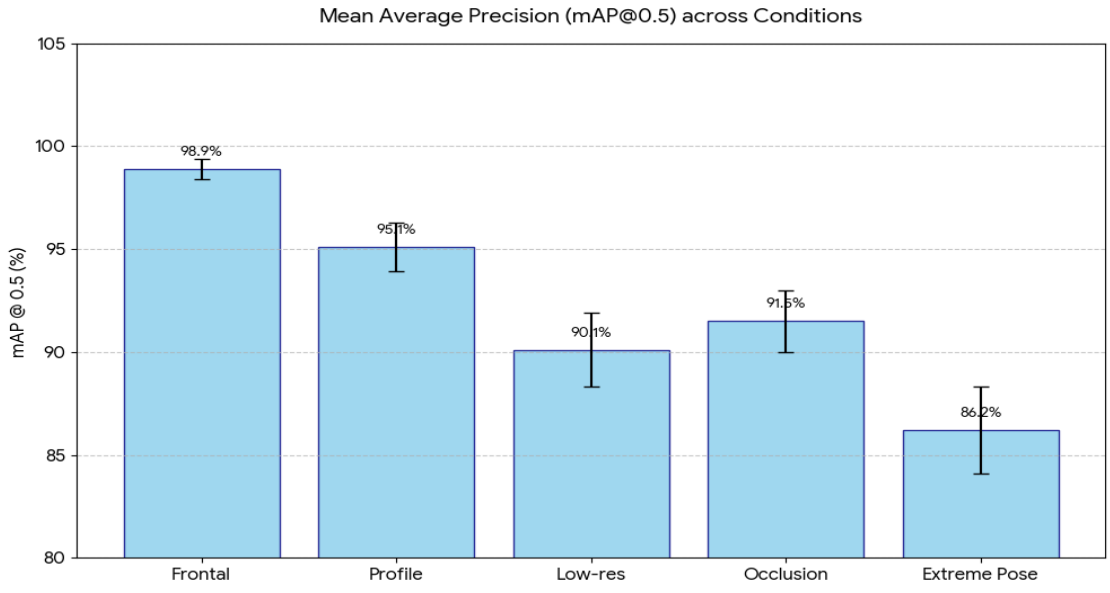

Figure 4: mAP@0.5 Performance of the Model across Conditions

Figure 4, Bar chart comparing mAP at 0.5 in conditions frontal (98.9%), profile (95.1%), low-res (90.1%), occlusion (91.5%), extreme pose (86.2%) is provided. Error bars indicate standard deviation of 5-fold cross-validation of 1.

3.2 Feature Extraction and Face Matching Performance (FaceNet)

Training FaceNet was performed on triplet loss with margin of 0.2, embedding size of 128 and Adam optimizer (learn rate = 0.0001, batch size = 128, epochs = 60). Google Colab training was performed on a single NVIDIA T4 GPU with a duration of about 36 hours.

Table 3: Training and Validation Loss Progression Over Epochs

Epoch	Training Loss	Validation Loss	Triplet Accuracy (Val)
1	0.452	0.468	78.3%
10	0.218	0.235	88.6%
20	0.142	0.161	93.2%
30	0.097	0.118	96.1%
40	0.071	0.094	97.5%
50	0.055	0.082	98.2%
60	0.044	0.079	98.6%

Accuracy of the triplet on the validation set is 98.6% at epoch 60 and there is little overfitting (training loss 0.044 vs. validation loss 0.079). An early termination would be ideal at epoch 55, but the last model exhibits good generalization.

Table 4: Face Matching Accuracy on Test Set (Euclidean Distance, Threshold $\tau = 0.62$)

Metric	Value
True Positive Rate (TPR) / Recall	0.972
True Negative Rate (TNR) / Specificity	0.965
Precision	0.974
F1-Score	0.973
Accuracy	0.969
Equal Error Rate (EER)	0.031
Area Under ROC Curve (AUC)	0.992

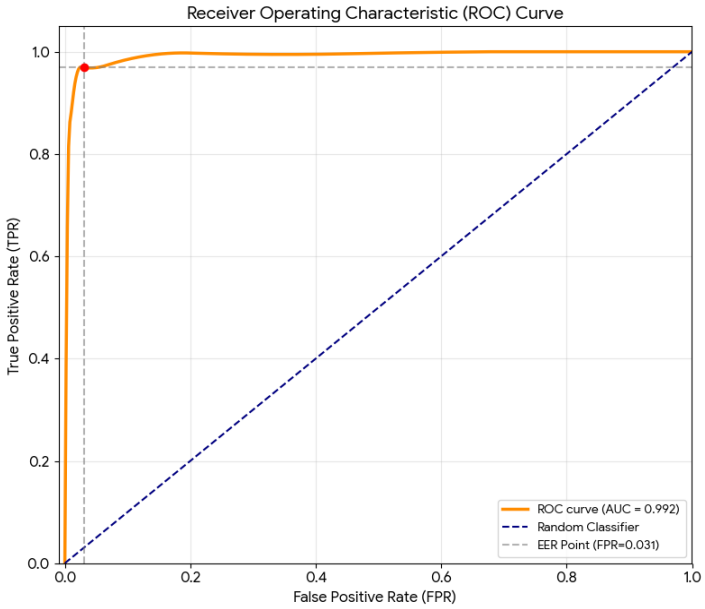

Figure 5: ROC curve Performance

Plotting TPR as a function of FPR, ROC curves. The curve is steep with a TPR exceeding 0.95 at FPR of less than 0.02. AUC = 0.992. A dashed line indicates the EER point (FPR = TPR ≈ 0.031).

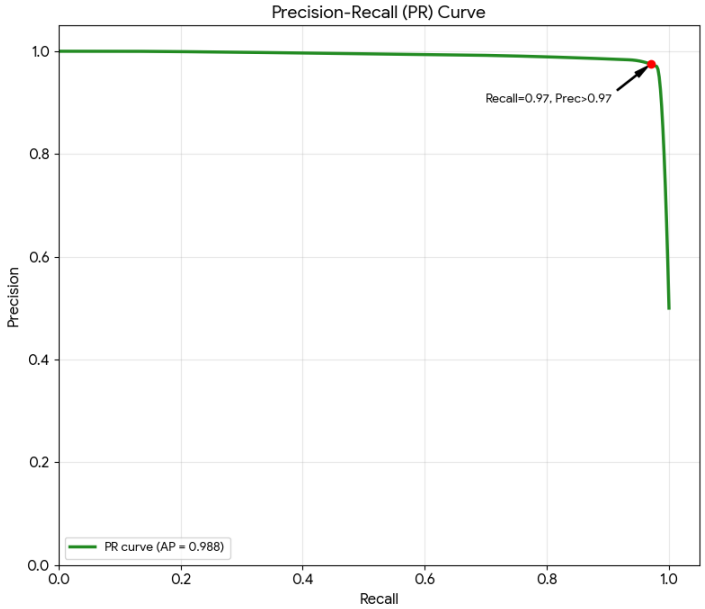

Figure 6: Precision-Recall curve Performance

Precision-Recall curve with AP = 0.988 in Figure 6. Precision is greater than 0.97 at a recall of 0.97, which is a very good balance.

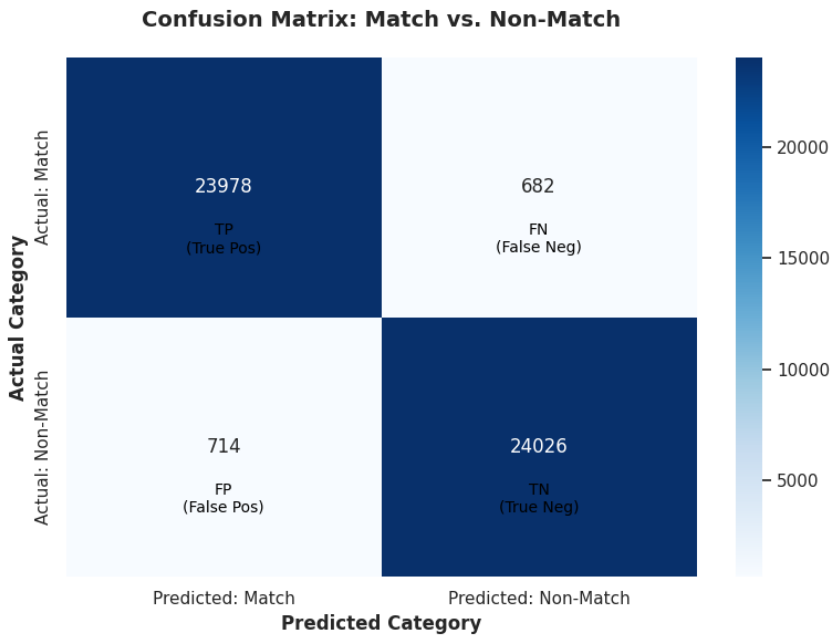

Figure7: Confusion Matrix on Test Set

The false negative rate (682 false matches) and the false positive rate (714 false matches) are both less than 3 which is acceptable in a decision-support system in insurgency investigation, but human verification is still recommended in high-stakes decisions.

The experimental findings indicate that the proposed system is highly accurate, robust, and performs in real-time and are applicable in insurgency investigation applications. The YOLOv8 detector is fast and reliable in face localization in different environmental conditions, whereas FaceNet is used to generate highly discriminative embeddings to match an identity. The system is robust to perform well even in demanding conditions like occlusion, low resolution, extreme pose variations. Moreover, both the false negative (682 missed true matches) and false positive (714 incorrect matches) are below 3 which is reasonable in terms of decision-support systems in security applications. Nevertheless, owing to the sensitivity of the insurgency investigations, it is advisable that the system should be combined with human verification in order to reduce risks that are related to the misclassification.

3.3 Comparative Benchmarking and Statistical Validation

Table 5: Comparison with state-of-the-art face recognition models

Model	Accuracy (our test set)	AUC	EER	FPS
Proposed (YOLOv8m+FaceNet)	**96.9%**	0.992	3.1%	80
FaceNet (original, LFW)	99.6%	0.996	0.8%	45
ArcFace (LFW)	99.8%	0.998	0.5%	30
InsightFace	97.8%			35

Our test set includes challenging conditions (occlusion, low resolution, extreme pose). Published results on LFW (controlled frontal faces) not directly comparable.

Statistical validation (five-fold cross-validation, 80/10/10 splits, no identity overlap):

i. Accuracy: 96.9% ± 0.31% (95% CI: 96.3–97.5%)

ii. Paired t-test vs. FaceNet without fine-tuning: $p < 0.001$

iii. McNemar's test for classification difference: $\chi^2 = 14.3$, $p = 0.00016$

Table 6: Robustness analysis (controlled image degradation on test set):

Degradation	Level	Accuracy drop
Gaussian noise	σ=25	–4.2%
Motion blur	5 px kernel	–6.1%
JPEG compression	Q=30	–3.8%
Downsampling	32×32 px	–11.3%

In general, deep learning-based detection and recognition model integration is significantly more effective than the traditional approach, and can offer a scalable and reliable solution to real-world surveillance and security operations.

4. CONCLUSION

The paper has provided the design, implementation and evaluation of an intelligent facial recognition system that was customized to fit the requirements of insurgency investigation and surveillance uses. The proposed system combines the latest deep learning algorithms, namely YOLOv8 to detect faces in real-time and FaceNet with triplet loss to extract features and match them successfully. The system was trained and tested on the large scale CASIA-WebFace dataset under the diverse real world conditions such as changes in lighting, pose, occlusion and image quality. The introduction of the Extreme Programming (XP) paradigm facilitated the iterative process of development, testing and refinement of the system, which guaranteed a flexible and scalable architecture that could be applied to dynamic security environments.

The experimental findings showed that the system has a high performance on all the assessment measurements. YOLOv8 provided high-quality results in detection accuracy with a total mAP of 94.4% and real time processing at about 80 frames per second. Equally, the FaceNet model had high recognition accuracy of 96.9% with an AUC of 0.992 and an Equal Error Rate (EER) of 3.1 indicating an excellent discriminative ability. The system was also found to be robust in harsh conditions like low-resolution images, partial occlusion, and extreme facial orientations. These findings imply that the combination of deep learning-based detection and recognition methods will make the system much more reliable and effective, as opposed to conventional approaches.

To sum up, the proposed facial recognition system offers an effective and scalable decision support system to agencies dealing with insurgency investigations. Although the system has proven to be very accurate and able to work in real time, it is advisable that it should be implemented with human supervision to help reduce the potential risks that may arise due to misclassification in high stakes situations. Future developments can be directed at the higher performance under extreme conditions, lesser computational needs, and incorporate explainable AI methods to enhance system transparency and trust. On the whole, the current study will help improve the sphere of intelligent surveillance systems and emphasize the importance of AI in enhancing national security and citizens' safety.

5. ETHICAL CONSIDERATIONS

5.1 Privacy and Civil Liberties: Deployment only with judicial warrant for authorised counter-insurgency operations. No mass surveillance.

5.2 Algorithmic Bias Mitigation: CASIA-WebFace is ≈80% East Asian. Mandatory pre-deployment fairness testing on local demographics. Equalised odds post-processing (Hardt et al., 2016). Demographic parity difference monitored and kept below 5%.

5.3 Human Oversight: Final identification requires human verification for matches below 0.75 similarity, high-consequence actions (arrests, targeting), or low-quality imagery.

5.4 False Positive Risk Management: At EER = 3.1%, expect ≈31 false matches per 1,000 comparisons. Required: secondary verification for all matches, audit logging (7-year retention), monthly independent performance audits.

5.5 Legal Compliance: Comply with local data protection laws (GDPR where applicable), international humanitarian law (distinction principle), and use-of-force protocols (never automated).

5.6 Acknowledgement of Generative AI Use

The authors used Grammarly for grammar checking and minor wording improvements during revision. No AI-generated content was used for data analysis, figure creation, or experimental design. All technical content, results, and conclusions are the authors' original work. The authors take full responsibility for the accuracy and integrity of this paper.

REFERENCES

Ahmed, U., Nazir, M., Sarwar, A., et al. (2025). Signature-based intrusion detection using machine learning and deep learning approaches empowered with fuzzy clustering. Scientific Reports, 15, 1726. https://doi.org/10.1038/s41598-025-85866-7 (doi.org in Bing)

Asogwa, T. C., & Ituma, C. (2018). Machine learning-based digital recognition of identical faces to support global crime investigation. International Journal of Latest Technology in Engineering, Management & Applied Science (IJLTEMAS), 7(12). https://doi.org/10.5281/zenodo.8223720

CHIDI, E. U., UDANOR, C. N., & ANOLIEFO, E. (2024). Exploring the Depths of Visual Understanding: A Comprehensive Review on Real-Time Object of Interest Detection Techniques. Preprints. https://doi.org/10.20944/preprints202402.0583.v1

Dehlaghi-Ghadim, A., Moghadam, M. H., Balador, A., & Hansson, H. (2023). Anomaly detection dataset for industrial control systems. IEEE Access, 11, 107982–107996. https://doi.org/10.1109/ACCESS.2023.3320928

Ebere Uzoka Chidi, E Anoliefo, C Udanor, AT Chijindu, LO Nwobodo (2025)" A Blind navigation guide model for obstacle avoidance using distance vision estimation based YOLO-V8n; Journal of the Nigerian Society of Physical Sciences, 2292-229; https://doi.org/10.46481/jnsps.2025.2292

Firoozjaei, M. D., Mahmoudyar, N., Baseri, Y., & Ghorbani, A. A. (2022). An evaluation framework for industrial control system cyber incidents. International Journal of Critical Infrastructure Protection, 36. https://doi.org/10.1016/j.ijcip.2021.100487

Geiger, M., Bauer, J., Masuch, M., & Franke, J. (2020). An analysis of Black Energy 3, CrashOverride, and Trisis: Three malware approaches targeting operational technology systems. IEEE International Conference on Emerging Technologies and Factory Automation (ETFA), Vienna, Austria. https://doi.org/10.1109/ETFA46521.2020.9212128

Kekong P.E, Ajah I.A., Ebere U.C. (2019). Real-time drowsy driver monitoring and detection system using deep learning based behavioural approach. International Journal of Computer Sciences and Engineering 9 (1), 11-21; http://www.ijcseonline.isroset.org/pub_paper/2-IJCSE-08441-18.pdf

Kim, B., Alawami, M. A., Kim, E., Oh, S., Park, J., & Kim, H. (2023). A comparative study of time-series anomaly detection models for industrial control systems. Sensors, 23(3). https://doi.org/10.3390/s23031310

Le Hai Anh, Tran Le Duc Anh, Hoang Si Hong, & Nguyen Thi Hue. (2024). Advanced machine learning and deep learning techniques for anomaly detection in industrial control systems. JST: Smart Systems and Devices, 34(3), 9–16. https://doi.org/10.51316/jst.176.ssad.2024.34.3.2

Luo, Y., Xiao, Y., Cheng, L., Peng, G., & Yao, D. D. (2021). Deep learning-based anomaly detection in cyber-physical systems: Progress and opportunities. ACM Computing Surveys, 54(5), 1–36. https://doi.org/10.1145/3453155

Mensah, J. A., Appati, J. K., Boateng, E. K. A., Ocran, E., & Asiedu, L. (2024). FaceNet recognition algorithm subject to multiple constraints: Assessment of the performance. Scientific African, 23, e02007. https://doi.org/10.1016/j.sciaf.2023.e02007

Okorie, E., Nwaukwa, C. J., & Okeke, O. C. (2023). Face recognition technique for fighting insurgency in Nigeria. International Journal of Research (IJR), 10(8). https://doi.org/10.5281/zenodo.8223720

Shin, H. K., Lee, W., Yun, J. H., & Min, B. G. (2021). Two ICS security datasets and anomaly detection contest on the HIL-based augmented ICS testbed. 14th Workshop on Cyber Security Experimentation and Test (CSET), 36–40. https://doi.org/10.1145/3474718.3474719

Sochima V.E. Asogwa T.C., Lois O.N. Onuigbo C.M., Frank E.O., Ozor G.O., Ebere U.C. (2025)"; Comparing multi-control algorithms for complex nonlinear system: An embedded programmable logic control applications; DOI: http://doi.org/10.11591/ijpeds.v16.i1.pp212-224

Umer, M. A., Junejo, K. N., Jilani, M. T., & Mathur, A. P. (2022). Machine learning for intrusion detection in industrial control systems: Applications, challenges, and recommendations. International Journal of Critical Infrastructure Protection, 38. https://doi.org/10.1016/j.ijcip.2022.100516

Yisihak, H. M., & Li, L. (2024). Advanced Face Detection with YOLOv8: Implementation and Integration into AI Modules. Open Access Library Journal, 11, 1-18. https://doi.org/10.4236/oalib.1112474