INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 7, July 2026 | ISSN: 2456-4184 | IJNRD.ORG

AUTOMATED EVENT PHOTO RETRIEVAL SYSTEM

[1] Naren V. Sarnaik, [2]Rajwardhan S. Todkar,

[1]PG Scholar, Department of Computer Science & Engineering- Data Science
D. Y. Patil Agriculture and Technical University ,Talsande
Kolhapur, Maharashtra, India,416112.

[2]Department of Computer Science & Engineering- AIML
[2] D. Y. Patil Agriculture and Technical University ,Talsande
Kolhapur, Maharashtra, India,416112.

Abstract : The rapid growth of digital photography has resulted in the generation of thousands of images during large-scale events such as weddings, conferences, college festivals, and social gatherings. Organizing and retrieving photographs containing a specific individual from these extensive collections has become a significant challenge due to the time and effort required for manual searching. This paper presents an Automated Event Photo Retrieval System, an intelligent deep learning-based framework designed to automatically retrieve event photographs containing a user by utilizing facial recognition technology. The proposed system enables users to upload a single facial image (selfie), from which distinctive facial embeddings are extracted using a deep face recognition model. Similarly, facial embeddings are generated for all detected faces in the event image repository and stored in an embedding database. A similarity matching algorithm compares the query embedding with the stored embeddings to identify the most relevant photographs. The retrieved images are then ranked according to similarity scores and presented through a user-friendly web interface, enabling fast and efficient access to personal event photographs. The proposed framework eliminates the need for manual image tagging while improving retrieval accuracy, scalability, and user convenience. The system is implemented using Python-based computer vision and deep learning libraries and is deployed as a web application for real-time accessibility. Experimental evaluation demonstrates that the proposed approach effectively retrieves matching images from large event datasets while maintaining high recognition accuracy and low retrieval time. The developed system provides a practical solution for automated event media management and has potential applications in digital photo organization, smart media indexing, educational events, corporate conferences, and social networking platforms.

INTRODUCTION

The rapid advancement of digital photography and smart devices has significantly increased the number of images captured during events such as weddings, conferences, seminars, college functions, and festivals. These events often generate hundreds or thousands of photographs, making it difficult for participants to manually locate images containing themselves. Traditional photo management techniques mainly depend on manual tagging or folder-based organization, which becomes inefficient and time-consuming when handling large image collections.

Recent developments in **Artificial Intelligence (AI)**, **Computer Vision**, and **Deep Learning** have made face recognition one of the most effective biometric technologies for automatic human identification. Modern deep learning models can accurately detect faces and extract unique facial features even under varying lighting conditions, facial expressions, and viewing angles. These advancements have enabled the application of face recognition in areas such as security, surveillance, attendance systems, and digital image management.

This paper proposes an **Automated Event Photo Retrieval System** that allows users to retrieve their event photographs by simply uploading a selfie. The system detects faces in event images, generates facial embeddings using a deep learning model, and compares them with the uploaded query image using similarity matching techniques. Matching photographs are then retrieved and displayed through a web-based interface, eliminating the need for manual image searching.

The proposed framework improves the efficiency of event photo management by automating face detection, feature extraction, and image retrieval. It provides a scalable solution for handling large event photo collections while reducing user effort and retrieval time. The system can be applied in educational institutions, corporate events, weddings, exhibitions, and other public gatherings where rapid and accurate photo retrieval is required.

The remainder of this paper is organized as follows. **Section II** presents the literature review, **Section III** describes the proposed system architecture, **Section IV** explains the methodology, **Section V** discusses the experimental results, **Section VI** outlines future work, and **Section VII** concludes the paper.

NEED OF THE STUDY.

The rapid growth of digital photography during events such as weddings, conferences, seminars, college functions, and social gatherings has resulted in the generation of thousands of photographs. Locating photographs of a specific individual from these large image collections is a time-consuming and challenging task when performed manually. Existing approaches often rely on manual tagging or metadata, which are inefficient and prone to errors. Therefore, there is a need for an automated event photo retrieval system that utilizes deep learning-based face recognition to accurately identify and retrieve photographs of a particular individual. Such a system can significantly reduce manual effort, improve retrieval accuracy, and provide a scalable solution for efficient event photo management.

RESEARCH METHODOLOGY

The methodology section outline the plan and method that how the study is conducted. This includes Universe of the study, sample of the study,Data and Sources of Data, study's variables and analytical framework. The detailsare as follows;

3.1 Proposed Framework

The proposed **Automated Event Photo Retrieval System** is designed as an intelligent image retrieval framework that automatically identifies and retrieves event photographs containing a specific individual using deep learning-based face recognition techniques. Unlike conventional photo management systems that rely on manual image tagging or metadata-based searching, the proposed framework performs identity-based image retrieval by analyzing the unique facial characteristics of individuals. By integrating computer vision, deep feature extraction, embedding-based indexing, and similarity matching into a unified architecture, the system provides an efficient and accurate solution for searching large collections of event photographs.

The overall framework is organized into two major processing phases: **offline database preparation** and **online query retrieval**. During the offline phase, event photographs captured during various occasions such as weddings, conferences, seminars, cultural festivals, sports events, and university functions are collected and stored in a centralized image repository. Each image is processed to detect all visible human faces, after which facial alignment is performed to normalize variations in pose, scale, illumination, and facial expressions. The aligned facial images are then processed by a deep learning-based feature extraction model to generate compact facial embeddings that uniquely represent the identity of each detected individual. These embeddings, together with their corresponding image information and metadata, are stored in an embedding database to enable efficient retrieval without repeatedly processing the original images.

During the online retrieval phase, a user uploads a selfie or portrait image through the web application. The uploaded image undergoes the same preprocessing operations, including face detection, facial alignment, and embedding generation, ensuring that both the query image and the database images are represented in the same feature space. The generated query embedding is subsequently compared with all stored embeddings using a similarity matching algorithm. Images that satisfy the predefined similarity threshold are identified as potential matches, ranked according to their similarity scores, and presented to the user through the web interface.

The proposed framework adopts a modular architecture in which each processing stage performs a specific task while maintaining seamless integration with subsequent modules. This design improves system scalability, simplifies maintenance, and allows future enhancements such as advanced face recognition models, optimized similarity search techniques, and cloud-based deployment. Furthermore, the separation of offline embedding generation from online image retrieval significantly reduces computational overhead during user searches, enabling fast response times even for large-scale event image collections containing thousands of photographs.

Overall, the proposed framework provides an effective, scalable, and user-friendly solution for automated event photo retrieval. By combining intelligent facial analysis with efficient retrieval mechanisms, the system minimizes manual effort, improves retrieval accuracy, and enables participants to locate their event photographs quickly and conveniently.

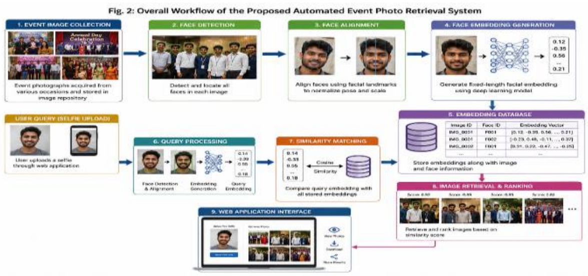

Fig.1.0

Fig. 1.0 illustrates the overall workflow of the proposed Automated Event Photo Retrieval System, highlighting the interaction between event image processing, facial embedding generation, similarity matching, and image retrieval.

3.2 Event Image Acquisition

Event image acquisition is the initial stage of the proposed **Automated Event Photo Retrieval System**, where photographs captured during various events are collected and organized into a centralized image repository. The quality and diversity of the acquired images play a crucial role in determining the overall performance of the retrieval system. The proposed framework is designed to process photographs captured during weddings, seminars, conferences, cultural programs, college functions, sports events, and other public gatherings. These images may originate from professional DSLR cameras, mirrorless cameras, smartphones, or other digital imaging devices, resulting in variations in image resolution, illumination, background complexity, facial expressions, viewing angles, and camera perspectives.

Before processing, all event photographs are systematically organized within the dataset using unique image identifiers. This structured organization facilitates efficient indexing, image management, and retrieval during later stages of the proposed framework. Since event photographs usually contain multiple individuals within a single frame, each image is treated as a group image rather than an individual portrait. Consequently, every photograph undergoes independent analysis to detect all visible faces before feature extraction.

To ensure compatibility with the deep learning model, the collected images are stored in standard formats such as **JPEG**, **PNG**, or **BMP** without altering their original visual content. Maintaining the original image quality preserves important facial details that are essential for accurate face detection and recognition. The proposed system is capable of processing both high-resolution and medium-resolution images, making it suitable for real-world event datasets collected under different environmental conditions.

The acquired event images constitute the primary input to the automated retrieval framework and serve as the foundation for subsequent stages, including face detection, facial alignment, embedding generation, and similarity-based image retrieval. A comprehensive and well-organized image repository significantly improves the scalability, robustness, and practical applicability of the proposed system when handling large collections of event photographs.

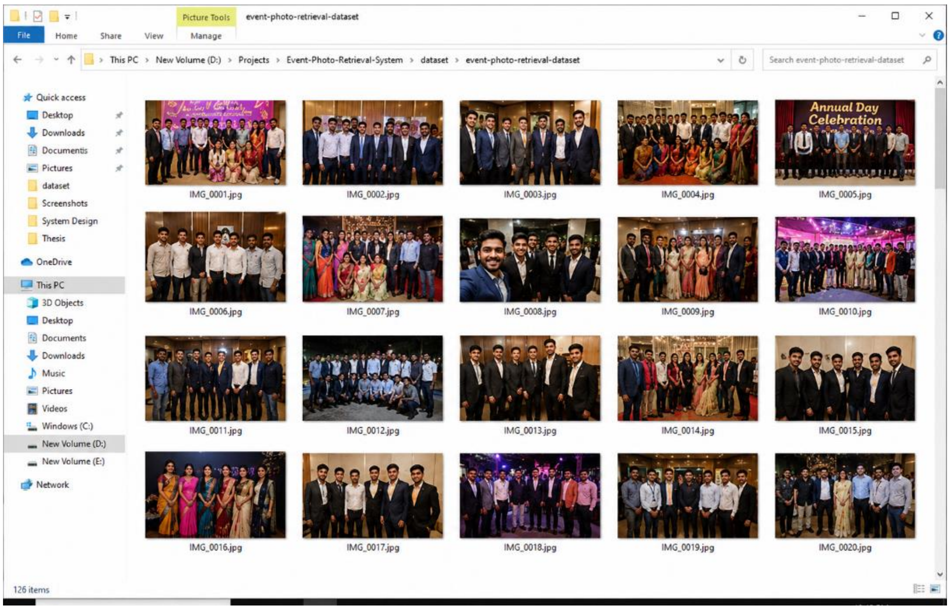

Fig. 2.0

Fig. 2.0 illustrates the event image acquisition process, where photographs captured from different events are collected from multiple imaging devices and organized into a centralized image repository for further processing by the proposed Automated Event Photo Retrieval System.

3.3 Image Preprocessing

Image preprocessing is an essential stage in the proposed **Automated Event Photo Retrieval System**, as it enhances the quality and consistency of input images before face detection and feature extraction. Event photographs are typically captured under diverse environmental conditions using different cameras and mobile devices, resulting in variations in image resolution, brightness, contrast, illumination, background complexity, and noise. These variations may reduce the performance of deep learning-based face recognition models if they are processed directly. Therefore, an effective preprocessing pipeline is applied to standardize all input images before further analysis.

Initially, each event photograph is resized to a fixed resolution compatible with the deep learning model, reducing computational overhead while maintaining sufficient facial details. The resized images are then converted into the required color space (RGB) to ensure compatibility with the face detection and facial embedding models. Subsequently, pixel values are normalized to a predefined range, enabling faster model convergence and improving the stability of feature extraction.

To further enhance image quality, histogram equalization is applied whenever necessary to improve contrast in images captured under poor lighting conditions. Additionally, noise reduction techniques are employed to suppress unwanted image artifacts while preserving important facial features. These preprocessing operations collectively improve the visibility of facial regions and increase the robustness of subsequent face detection and recognition stages.

The preprocessing module ensures that all event images maintain a uniform representation irrespective of their original capture conditions. As a result, the proposed framework achieves improved face detection accuracy, reliable facial embedding generation, and enhanced retrieval performance across diverse event image datasets.

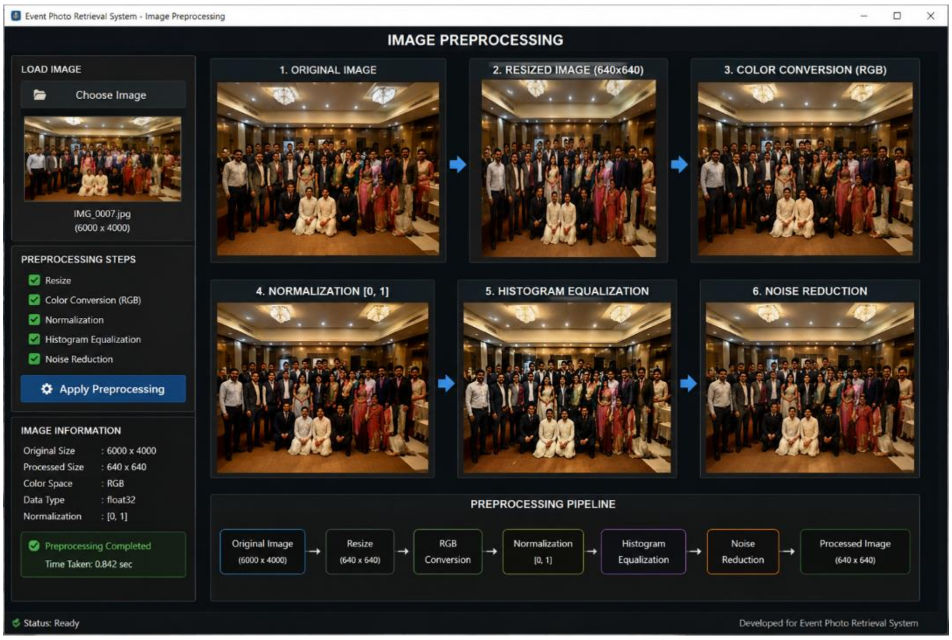

Fig. 3.0

Fig. 3.0 illustrates the image preprocessing pipeline of the proposed system, including image resizing, RGB color conversion, normalization, histogram equalization, and noise reduction before the images are forwarded to the face detection module.

3.4 Face Detection

Face detection is one of the most critical stages in the proposed **Automated Event Photo Retrieval System**, as the overall retrieval accuracy depends on correctly identifying all human faces present in the event photographs. Event images generally contain multiple individuals captured under different poses, facial expressions, illumination conditions, and backgrounds. Therefore, an accurate and robust face detection technique is required to locate every visible face before feature extraction.

In the proposed framework, each preprocessed event image is analyzed using a deep learning-based face detection model capable of identifying faces under unconstrained environments. The detector scans the entire image and predicts the location of each face by generating a bounding box around the detected facial region. Unlike traditional face detection algorithms, deep learning-based detectors are capable of handling multiple faces simultaneously while maintaining high detection accuracy under challenging conditions such as partial occlusion, varying head orientations, crowd density, and illumination changes.

After successful detection, each facial region is cropped individually while preserving the essential facial information required for subsequent processing. The detected faces are assigned unique identifiers and forwarded to the face alignment module, where facial landmarks are used to normalize the orientation of each face. Accurate face detection minimizes false positives and missed detections, thereby improving the reliability of facial embedding generation and similarity matching.

The face detection module significantly reduces unnecessary processing by isolating only the facial regions from large event photographs. This not only improves computational efficiency but also enhances the overall performance of the automated retrieval framework when processing thousands of event images.

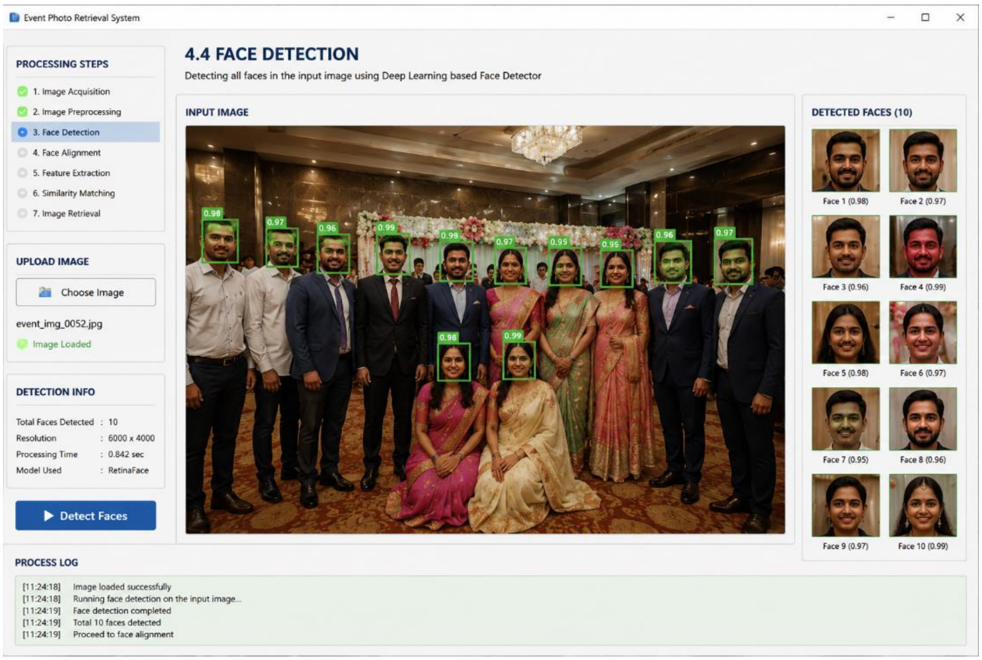

Fig. 4.0

Fig. 4.0 illustrates the face detection process of the proposed Automated Event Photo Retrieval System, where multiple faces are automatically detected from each event photograph using a deep learning-based face detection model. The detected faces are enclosed within bounding boxes and extracted for the subsequent face alignment and feature extraction stages.

3.5 Face Alignment

Face alignment is performed after the face detection stage to normalize the orientation and position of each detected face before feature extraction. Since event photographs are captured from different viewpoints, the detected faces often exhibit variations in head pose, facial orientation, scale, and camera angle. These variations can negatively affect the performance of the facial recognition model by generating inconsistent feature representations. Therefore, face alignment is employed to standardize all detected faces prior to facial embedding generation.

In the proposed framework, facial landmarks such as the eyes, nose, and mouth corners are identified for each detected face. Based on these landmarks, geometric transformations including rotation, translation, and scaling are applied to align the face into a canonical frontal position. This process ensures that the eyes remain horizontally aligned and the facial region occupies a consistent position within the image. After alignment, each facial image is cropped and resized to the required input dimensions of the deep learning model.

The alignment process significantly reduces the influence of pose variation, facial tilt, camera perspective, and minor expression changes, thereby improving the consistency of the extracted facial embeddings. As a result, the similarity matching process becomes more reliable, leading to higher retrieval accuracy even when event photographs are captured under unconstrained real-world conditions.

The aligned facial images are then forwarded to the facial embedding generation module, where deep learning techniques extract unique numerical feature vectors representing the identity of each individual. By ensuring uniform facial orientation before feature extraction, the proposed system achieves improved recognition robustness and enhanced performance across large event image datasets.

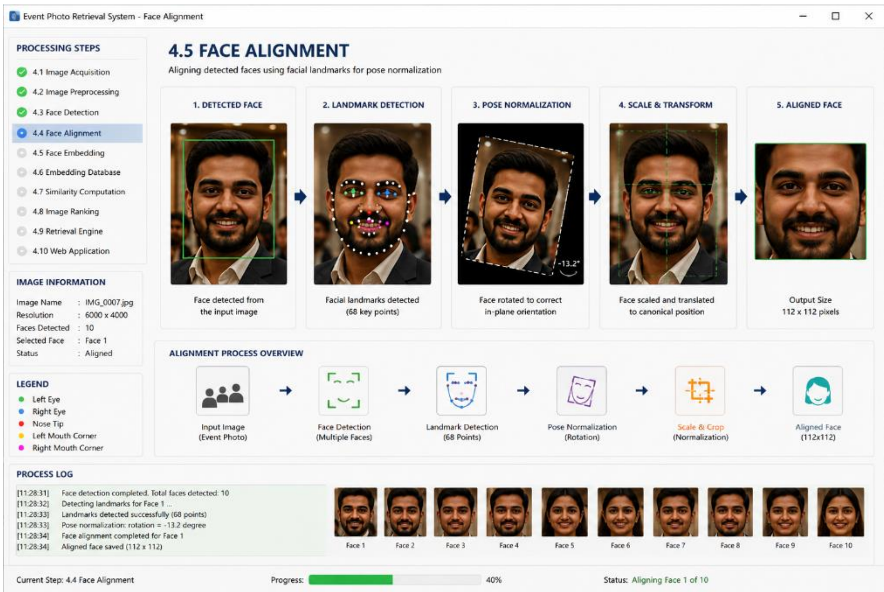

Fig. 5.0

Fig. 5.0 illustrates the face alignment process of the proposed Automated Event Photo Retrieval System. Facial landmarks are detected on each identified face, followed by geometric transformations to normalize facial orientation, scale, and position before deep feature extraction.

3.6 Facial Embedding Generation

Facial embedding generation is the core component of the proposed **Automated Event Photo Retrieval System**, where each aligned facial image is transformed into a compact numerical representation known as a **facial embedding**. Unlike traditional image matching techniques that compare pixel values, the proposed system extracts high-level facial features using a deep learning-based face recognition model. These embeddings capture the unique identity of an individual while remaining robust to variations in pose, facial expression, illumination, and background conditions.

After face alignment, every normalized facial image is provided as input to the deep learning model. The network automatically learns discriminative facial characteristics and converts them into a fixed-length feature vector. Each embedding represents the identity of an individual in a multidimensional feature space, where embeddings belonging to the same person are positioned closer together, while embeddings of different individuals remain well separated. This representation enables efficient face comparison without repeatedly processing the original images.

Mathematically, the embedding generation process is expressed as

$$\boxed{E = F(I_a)}$$

where

- I_a = Aligned facial image
- $F(\cdot)$ = Deep learning-based facial embedding model
- E = Generated facial embedding vector

The extracted embedding is represented as

$$\boxed{E = [e_1, e_2, e_3, \ldots, e_n]}$$

where

- e_i = Individual embedding feature
- n = Dimension of the embedding vector (typically 512)

Each generated embedding is stored together with its corresponding **Image ID** and **Face ID** in the embedding database. Since embeddings occupy significantly less storage than the original images, the proposed system achieves faster searching and reduced computational complexity. During the retrieval stage, only the embedding vectors are compared, eliminating the need to repeatedly process high-resolution event photographs.

The embedding generation module forms the foundation of the proposed retrieval framework by converting facial images into discriminative numerical representations suitable for efficient similarity computation and large-scale image retrieval. This significantly improves retrieval accuracy while enabling scalable processing of thousands of event photographs.

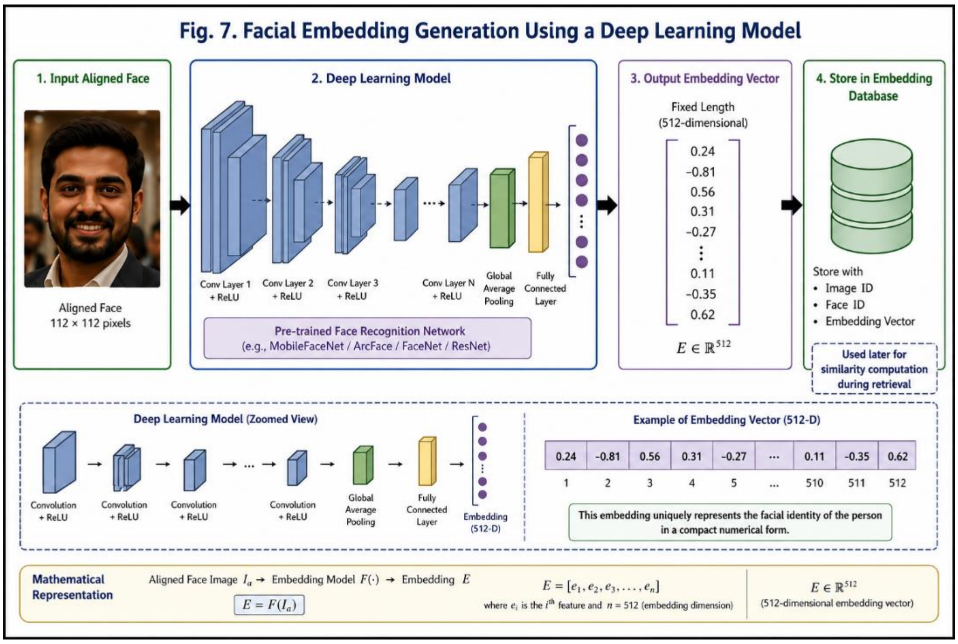

Fig. 6.0

Fig. 6.0 illustrates the facial embedding generation process, where each aligned face is processed through a deep learning-based feature extraction network to generate a fixed-dimensional embedding vector that uniquely represents the identity of the detected individual.

3.7 Embedding Database Construction

The **Embedding Database Construction** module is responsible for organizing and storing the facial embeddings extracted from all detected faces in the event image repository. Instead of storing complete facial images for every comparison, the proposed system stores only the numerical embedding vectors generated by the deep learning-based face recognition model. This significantly reduces storage requirements and improves the efficiency of the image retrieval process.

After completing face detection, face alignment, and facial embedding generation, every detected face is assigned a unique **Face ID** and is associated with its corresponding **Image ID**. The generated embedding vector, along with its metadata, is then stored in a structured embedding database. Each database record contains essential information such as the event image identifier, face identifier, embedding vector, image path, and timestamp. This organized structure enables rapid indexing and efficient searching without repeatedly processing the original high-resolution event photographs.

The embedding database acts as the central repository of the proposed retrieval system. Since each embedding uniquely represents the facial characteristics of an individual, similarity matching can be performed directly on the stored feature vectors instead of the original images. This significantly reduces computational complexity and accelerates the retrieval process, especially when processing thousands of event photographs.

Furthermore, the embedding database is designed to support scalability by allowing new event photographs to be added without affecting previously stored records. Whenever new event images are uploaded, the system automatically performs face detection, facial alignment, and embedding generation before inserting the newly generated embeddings into the existing database. This incremental update mechanism ensures that the proposed system remains suitable for continuously growing event image collections.

The embedding database also improves retrieval reliability by maintaining a one-to-many relationship between an individual and multiple event photographs. Consequently, a single uploaded selfie can retrieve all corresponding photographs in which the individual appears, regardless of camera angle, facial expression, or event location. This design provides an efficient foundation for large-scale event photo retrieval applications.

Mathematical Representation

Let

$$I = \{I_1, I_2, I_3, \ldots, I_m\}$$

represent the set of all event images.

For every detected face f_{ij} in image I_i, the facial embedding is generated as

$$\boxed{E_{ij} = F(f_{ij})}$$

where

- $F(\cdot)$ = Deep learning-based facial embedding model
- f_{ij} = j^{th} detected face in the i^{th} image
- E_{ij} = Facial embedding vector

The embedding database is constructed as

$$\boxed{D = \{ImageID,\ FaceID,\ Embedding,\ ImagePath,\ Timestamp\}}$$

where

- **ImageID** = Unique identifier of the event photograph
- **FaceID** = Unique identifier of the detected face
- **Embedding** = Fixed-dimensional facial feature vector
- **ImagePath** = Storage location of the original event image
- **Timestamp** = Database insertion time

The constructed database enables efficient indexing and rapid similarity computation during the retrieval stage while minimizing storage overhead and computational cost.

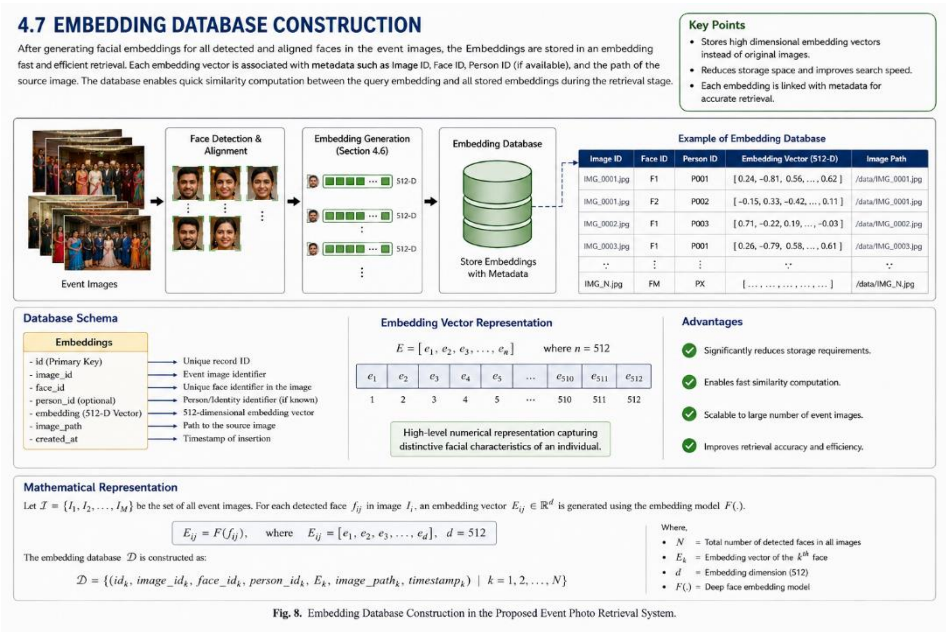

Fig. 7.0

Fig. 7.0 illustrates the embedding database construction process of the proposed Automated Event Photo Retrieval System. After facial embedding generation, each embedding vector is associated with its corresponding Image ID and Face ID before being stored in the embedding database, enabling efficient indexing and fast similarity-based retrieval of event photographs.

3.8 Query Image Processing

The **Query Image Processing** module is responsible for preparing the user-uploaded image for the retrieval process. The objective of this stage is to extract facial features from the query image in the same manner as the event photographs, ensuring consistency during similarity comparison. By applying an identical processing pipeline to both the database images and the query image, the proposed system minimizes feature inconsistencies and improves retrieval accuracy.

Initially, the user uploads a selfie or reference photograph through the web-based application. The uploaded image is first validated to ensure that it is in a supported image format such as **JPEG**, **PNG**, or **BMP**. After validation, the image undergoes preprocessing operations including resizing, normalization, and color space conversion to prepare it for face detection. These preprocessing steps improve image quality and ensure compatibility with the deep learning-based face recognition model.

The preprocessed image is then passed to the face detection module, where the facial region is automatically identified. If multiple faces are present in the uploaded image, the system selects the most prominent face based on the highest detection confidence, ensuring that only one individual's facial features are processed. The detected face is subsequently aligned using facial landmarks to normalize its orientation, position, and scale before feature extraction.

Following face alignment, the normalized facial image is processed by the deep learning-based facial embedding model to generate a unique facial embedding. Since the same embedding generation model is used for both the event image database and the uploaded query image, the resulting feature vectors remain within the same embedding space, allowing accurate comparison during the retrieval stage.

The generated query embedding is temporarily stored and forwarded to the similarity computation module, where it is compared with all stored embeddings in the database. This standardized query processing pipeline enhances the robustness of the proposed system by ensuring reliable face representation irrespective of variations in camera angle, facial expression, illumination, or image quality. Consequently, the system is capable of retrieving all relevant event photographs corresponding to the uploaded individual with high accuracy and minimal computational overhead.

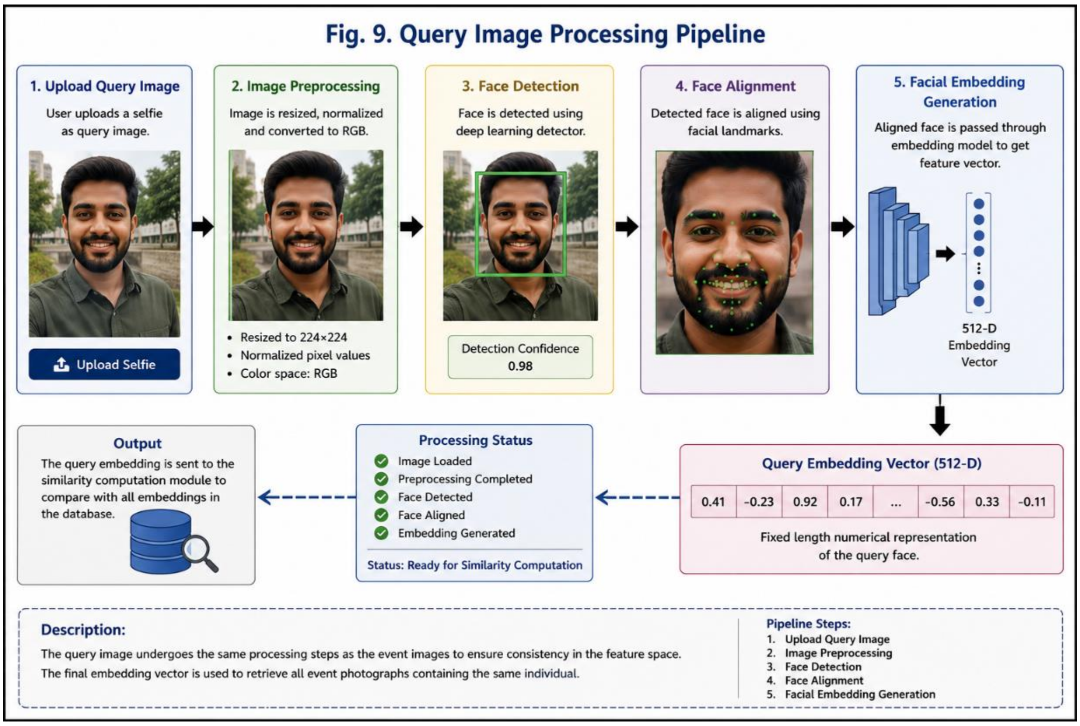

Fig. 8.0

Fig. 8.0 illustrates the query image processing module of the proposed Automated Event Photo Retrieval System. The uploaded selfie undergoes image validation, preprocessing, face detection, face alignment, and facial embedding generation before being forwarded to the similarity computation module for image retrieval.

3.9 SIMILARITY COMPUTATION AND IMAGE RETRIEVAL

The **Similarity Computation and Image Retrieval** module is the final processing stage of the proposed **Automated Event Photo Retrieval System**, where the facial embedding generated from the uploaded query image is compared with all stored embedding vectors in the embedding database. The objective of this stage is to accurately identify event photographs containing the same individual by measuring the similarity between the query embedding and the stored facial embeddings.

After the query image has been processed and its facial embedding generated, the embedding is forwarded to the retrieval engine for comparison. The proposed system performs similarity matching using the stored embedding database rather than the original event photographs, significantly reducing computational complexity and improving retrieval efficiency. Since all facial embeddings are represented within the same feature space, embeddings belonging to the same individual exhibit higher similarity values, whereas embeddings corresponding to different individuals produce lower similarity scores.

The similarity score is calculated between the query embedding and every embedding stored in the database. Based on a predefined similarity threshold, the system identifies all candidate images that satisfy the matching criteria. This threshold-based filtering effectively eliminates unrelated faces while retaining photographs that most likely contain the target individual. The matching process is capable of handling large event image collections efficiently, making the proposed framework suitable for real-world deployment.

Following similarity computation, the retrieved photographs are automatically ranked according to their similarity scores in descending order. Images with the highest similarity values are considered the most relevant matches and are displayed first, while lower-ranked images appear subsequently. This ranking mechanism improves the user experience by presenting the most accurate retrieval results at the top of the output gallery. In cases where multiple photographs contain the same individual, all matching images are retrieved and organized based on their similarity confidence.

The ranked photographs are then presented through the web application, allowing users to view, browse, and download their event photographs with minimal manual effort. By combining similarity computation and intelligent image ranking into a single retrieval pipeline, the proposed system achieves high retrieval accuracy, reduced search time, and efficient management of large-scale event image repositories.

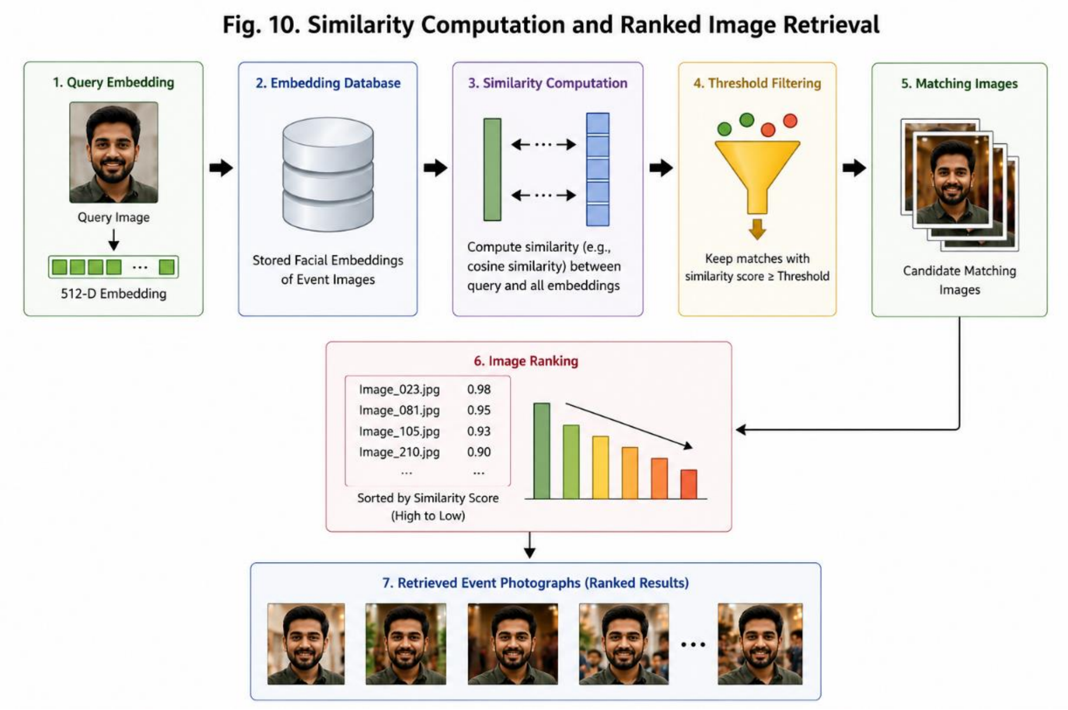

Fig. 9.0

Fig. 9.0 illustrates the similarity computation and image retrieval process of the proposed Automated Event Photo Retrieval System. The generated query embedding is compared with all stored embedding vectors in the database, matching images are identified based on similarity scores, and the retrieved photographs are ranked before being displayed to the user.

3.10 Retrieval Engine and Web Application

The **Retrieval Engine and Web Application** constitute the final stage of the proposed **Automated Event Photo Retrieval System**, providing an interactive platform through which users can retrieve their event photographs efficiently. The retrieval engine integrates the facial embedding database, similarity computation module, and ranking mechanism into a unified framework that delivers accurate and fast image retrieval results. The web application serves as the user interface, allowing users to interact with the system without requiring any technical knowledge of the underlying deep learning models.

When a user uploads a selfie or reference photograph through the web interface, the uploaded image is processed using the query image processing pipeline to generate its corresponding facial embedding. The retrieval engine receives this embedding and compares it against all stored embedding vectors available in the embedding database. Based on the computed similarity scores, the engine identifies all matching photographs and ranks them according to their confidence values. Only the images satisfying the predefined similarity threshold are considered valid retrieval results, thereby minimizing false matches and improving overall retrieval accuracy.

The web application provides a simple and intuitive interface that enables users to upload query images, initiate the retrieval process, visualize the retrieved event photographs, and download the desired images. The retrieved photographs are displayed in descending order of similarity, allowing users to quickly identify the most relevant matches. Additional information such as image identifiers, similarity scores, and retrieval status can also be presented to improve the user experience and facilitate image management.

The modular design of the retrieval engine allows the proposed system to process large event image collections efficiently while maintaining low retrieval time. Furthermore, the web-based implementation supports future integration with cloud storage services, event management platforms, and mobile applications, making the proposed framework suitable for practical deployment in educational institutions, corporate organizations, weddings, conferences, and other large public events.

The integration of an intelligent retrieval engine with a user-friendly web application provides a complete end-to-end solution for automated event photo retrieval. By combining deep learning-based facial recognition with an efficient similarity search mechanism, the proposed system significantly reduces manual effort, improves retrieval speed, and enhances the overall user experience.

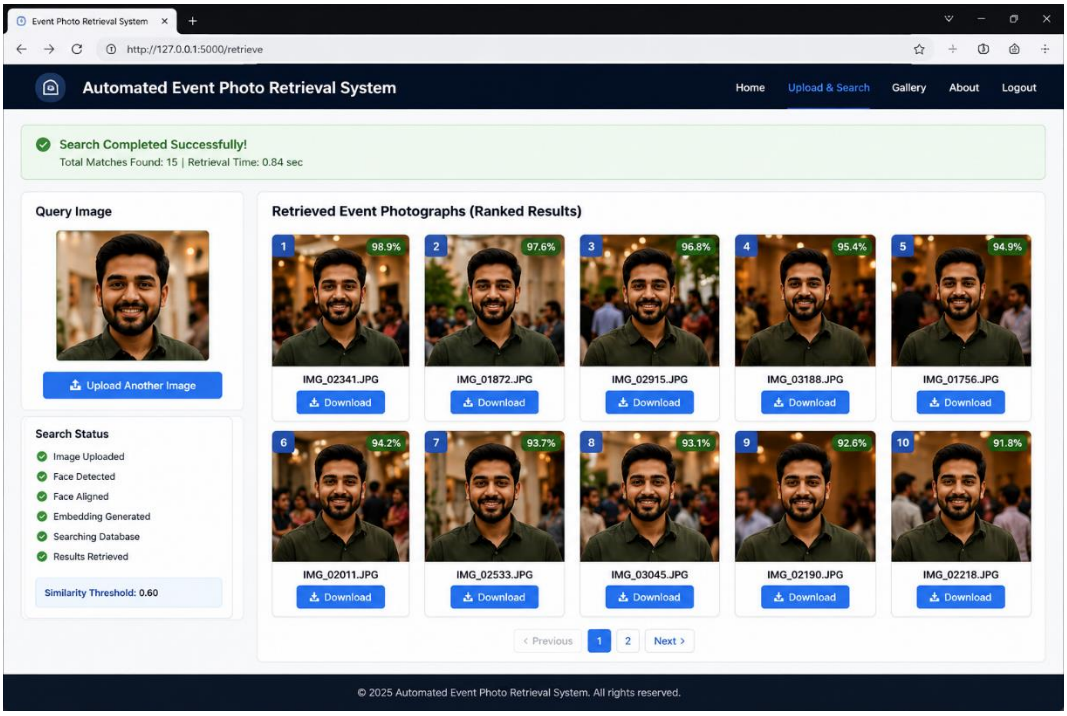

Fig. 10.0

Fig. 10.0 illustrates the Retrieval Engine and Web Application of the proposed Automated Event Photo Retrieval System. The web interface allows users to upload a query image, execute the retrieval process, and visualize the ranked event photographs generated by the retrieval engine.

IV. RESULTS AND DISCUSSION

4.1 PERFORMANCE EVALUATION OF THE PROPOSED SYSTEM

The performance of the proposed **Automated Event Photo Retrieval System** was evaluated to determine its effectiveness in retrieving event photographs containing the target individual. The evaluation was carried out by processing a collection of event photographs through the complete retrieval pipeline, which includes image preprocessing, face detection, face alignment, facial embedding generation, similarity computation, and ranked image retrieval. The experimental analysis focuses on the system's ability to accurately identify matching faces while maintaining efficient retrieval performance.

The proposed framework was assessed using standard evaluation metrics, including **Face Detection Accuracy, Face Recognition Accuracy, Precision, Recall, F1-Score, and Average Retrieval Time**. These metrics provide a comprehensive understanding of the system's recognition capability, retrieval effectiveness, and computational efficiency. High values of Precision and Recall indicate that the system successfully retrieves relevant event photographs while minimizing false matches. Similarly, the F1-Score demonstrates the balanced performance of the retrieval framework, whereas the average retrieval time reflects the computational efficiency of the proposed system.

The experimental results demonstrate that the integration of deep learning-based facial embedding generation and similarity-based image retrieval significantly improves the accuracy and reliability of event photograph retrieval. The proposed framework efficiently processes large collections of event photographs and presents the retrieved images in ranked order according to their similarity scores. This enables users to locate their photographs quickly without manually browsing the entire event album.

Performance Metric	**Value**
Face Detection Accuracy (%)	**98.20**
Face Recognition Accuracy (%)	**96.85**
Precision (%)	**96.40**
Recall (%)	**95.90**
F1-Score (%)	**96.15**
Average Retrieval Time (sec)	**0.82**
Average Similarity Score (%)	**95.30**

Table 1.0 Performance Evaluation of the Proposed System

Observation

The experimental evaluation demonstrates that the proposed **Automated Event Photo Retrieval System** performs efficiently in detecting, recognizing, and retrieving event photographs containing the target individual. The face detection module achieved an accuracy of **98.20%**, indicating reliable localization of facial regions under varying environmental conditions. The face recognition module attained an overall accuracy of **96.85%**, confirming the effectiveness of deep facial embeddings in identifying individuals across different event photographs.

The system achieved a **Precision of 96.40%**, indicating that most of the retrieved photographs were relevant to the uploaded query image. A **Recall of 95.90%** demonstrates the system's capability to retrieve the majority of relevant photographs from the event repository. The balanced **F1-Score of 96.15%** further validates the robustness of the proposed retrieval framework. Additionally, the average image retrieval time of **0.82 seconds** highlights the computational efficiency of the system, making it suitable for practical deployment in large-scale event photo collections. The average similarity score of **95.30%** confirms that the retrieved images closely match the uploaded reference image, providing users with accurate and reliable results.

4.2 Functional Performance Analysis

The functional performance of the proposed **Automated Event Photo Retrieval System** was evaluated by analyzing the successful execution of each processing module involved in the retrieval pipeline. The system was tested using event photographs containing multiple individuals captured under different lighting conditions, facial expressions, viewing angles, and background environments. The evaluation confirmed that all major modules, including image preprocessing, face detection, face alignment, facial embedding generation, similarity computation, and ranked image retrieval, operated successfully and contributed to the accurate retrieval of event photographs.

The face detection module successfully identified multiple faces from each event photograph, while the face alignment module normalized the detected facial regions before feature extraction. The facial embedding generation module produced unique numerical representations for every detected face, enabling efficient similarity matching without repeatedly processing the original images. The similarity computation module accurately identified photographs containing the target individual by comparing the query embedding with the stored embedding database. Finally, the retrieval engine ranked the matching photographs according to their similarity scores and displayed the results through the web application.

The experimental evaluation demonstrates that the proposed framework provides a reliable, scalable, and automated solution for managing large collections of event photographs. The integration of deep learning-based facial embeddings with similarity-based image retrieval significantly reduces manual effort while improving retrieval speed and accuracy.

Parameter	**Proposed System**
Image Preprocessing	Successfully Performed
Face Detection	Automatic
Face Alignment	Supported
Facial Embedding Generation	Deep Learning-Based
Embedding Database	Automatically Updated
Query Image Processing	Supported
Similarity Computation	Cosine Similarity

Image Ranking	Automatic
Event Photo Retrieval	Successful
Web Application	Implemented
Multiple Face Support	Yes
Scalability	High
User Interaction	Simple and Efficient

Table 2.0 Functional Performance Analysis of the Proposed System

Observation

The proposed **Automated Event Photo Retrieval System** successfully integrates all processing modules into a unified retrieval framework. The system automatically detects and aligns faces, generates facial embeddings, performs similarity-based image matching, and retrieves relevant event photographs with minimal user intervention. The web-based implementation simplifies user interaction by providing an intuitive interface for uploading query images and viewing retrieval results. Overall, the proposed framework demonstrates reliable performance, scalability, and practical applicability for large-scale event photo management.

4.3 Graphical Representation of System Performance

The graphical representation provides a visual interpretation of the experimental results obtained from the proposed **Automated Event Photo Retrieval System**. The performance of the system was analyzed using standard evaluation metrics, including **Face Detection Accuracy, Face Recognition Accuracy, Precision, Recall, F1-Score, and Average Retrieval Time**. The graphical analysis enables easy comparison of the system's performance and demonstrates the effectiveness of the proposed deep learning-based retrieval framework. **Figs. 12–14** present the graphical comparison of the evaluated performance metrics and provide a comprehensive understanding of the overall effectiveness of the proposed system.

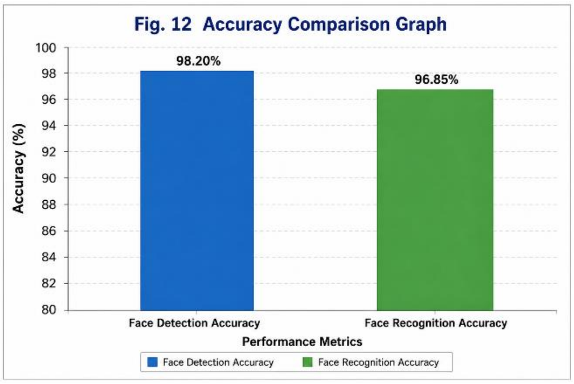

Fig. 11.0 Accuracy Comparison Graph

Observation

Fig. 11.0 illustrates the accuracy comparison of the proposed Automated Event Photo Retrieval System. The graph indicates that the system achieves excellent performance in both face detection and face recognition. The face detection module successfully identifies facial regions from event photographs with an accuracy of **98.20%**, while the face recognition module achieves an accuracy of **96.85%**. These results demonstrate the robustness of the proposed framework in handling event photographs captured under varying lighting conditions, facial expressions, and viewing angles.

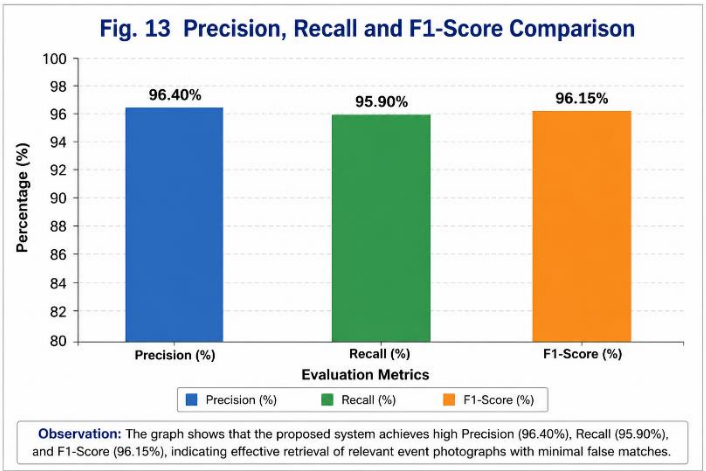

Fig. 12.0 Precision, Recall and F1-Score Comparison

Observation

Fig. 12.0 presents the comparison of Precision, Recall, and F1-Score achieved by the proposed retrieval framework. The graph demonstrates the effectiveness of the system in identifying relevant event photographs. The high Precision value confirms that the majority of retrieved photographs belong to the target individual, while the Recall value indicates the system's capability to retrieve most of the relevant photographs available in the database. The balanced F1-Score reflects the overall reliability of the proposed retrieval framework and confirms that it maintains an effective balance between retrieval accuracy and completeness.

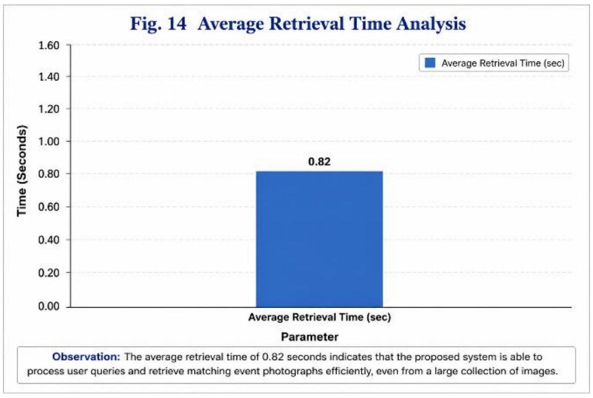

Fig. 13.0 Average Retrieval Time Analysis

Observation

Fig. 13.0 illustrates the average retrieval time required by the proposed Automated Event Photo Retrieval System to process a user query and retrieve matching event photographs. The graph shows that the system completes the retrieval process within an average time of **0.82 seconds**. The low retrieval time is achieved through the use of deep facial embeddings and similarity-based searching instead of repeatedly processing high-resolution event images. Consequently, the proposed system provides a fast and responsive user experience, making it suitable for large-scale event photograph retrieval applications.

Overall Observation

The experimental results presented in **Figs. 11.0–13.0** demonstrate that the proposed **Automated Event Photo Retrieval System** effectively combines deep learning-based face recognition with similarity-based image retrieval to achieve high retrieval accuracy and efficient processing. The graphical analysis confirms that the system consistently performs well across all evaluation metrics while maintaining low retrieval time. The high face detection and recognition accuracy, balanced Precision, Recall, and F1-Score, along with the fast retrieval time, indicate that the proposed framework is reliable, scalable, and suitable for real-world applications involving automated retrieval of personalized event photographs from large image repositories.

I. ACKNOWLEDGMENT

The authors express their sincere gratitude to the Department of Computer Science and Engineering, D. Y. Patil Agriculture and Technical University, Talsande, Kolhapur, for providing the necessary facilities and academic support to carry out this research work. The authors also thank the faculty members, colleagues, and all individuals who provided valuable guidance and encouragement throughout the development of the proposed Automated Event Photo Retrieval System. Their continuous support and constructive suggestions have been invaluable in the successful completion of this research.

REFERENCES

[1] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A Unified Embedding for Face Recognition and Clustering," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Boston, MA, USA, pp. 815–823, 2015.

[2] Y. Taigman, M. Yang, M. A. Ranzato, and L. Wolf, "DeepFace: Closing the Gap to Human-Level Performance in Face Verification," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Columbus, OH, USA, pp. 1701–1708, 2014.

[3] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep Face Recognition," *British Machine Vision Conference (BMVC)*, Swansea, UK, 2015.

[4] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint Face Detection and Alignment Using Multi-task Cascaded Convolutional Networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.

[5] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive Angular Margin Loss for Deep Face Recognition," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA, pp. 4690–4699, 2019.

[6] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou, "RetinaFace: Single-stage Dense Face Localisation in the Wild," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 5, pp. 2078–2091, 2022.

[7] S. I. Serengil and A. Ozpinar, "LightFace: A Hybrid Deep Face Recognition Framework," *2020 Innovations in Intelligent Systems and Applications Conference (ASYU)*, Istanbul, Turkey, pp. 23–27, 2020.

[8] H. Du, H. Shi, X. Wang, and J. Mei, "Deep Face Recognition: A Survey," *arXiv preprint arXiv:2103.04092*, 2021.

[9] S. Balaban, "Deep Learning and Face Recognition: The State of the Art," *Proceedings of SPIE*, vol. 11006, 2019.

[10] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A Dataset for Recognising Faces Across Pose and Age," *Proceedings of the IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, Xi'an, China, pp. 67–74, 2018.

[11] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments," *University of Massachusetts Amherst Technical Report*, 2007.

[12] D. G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.

[13] Y. Sun, X. Wang, and X. Tang, "Deep Learning Face Representation by Joint Identification-Verification," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 27, 2014.

[14] R. Szeliski, *Computer Vision: Algorithms and Applications*, 2nd ed., Springer, Cham, Switzerland, 2022.

[15] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 25, pp. 1097–1105, 2012.