IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

An Intelligent Attendance System Based On Convolutional Neural Networks For Real-Time Multiple Student Face Identifications

1.Sonali Salunkhe Student
2.Prof. Rupali Maske, Associate Professor,
3. Prof. Barkha Shahaji, Associate Professor,
4. Prof. Vishal Shinde Associate Professor

Computer Engineering
Trinity College of Engineering and Research,Pune , India

Abstract: Student attendance tracking is a crucial aspect of academic institutions, ensuring discipline and monitoring student participation. Traditional attendance systems, including manual roll calls and RFID-based methods, are time-consuming, prone to human error, and susceptible to fraudulent practices such as proxy attendance. To address these limitations, this research proposes an automated student attendance system utilizing Faster R-CNN (Region-based Convolutional Neural Network) for efficient and accurate face detection, combined with side-angle detection to enhance recognition when students are not directly facing the camera. The system employs a high-resolution camera to capture real-time classroom footage. Faster R-CNN is leveraged for fast and precise multi-face detection, ensuring robustness even in large classrooms with multiple students present. However, traditional face recognition models struggle with side-angle or partially occluded faces, leading to misidentification or missed attendance marking. To overcome this challenge, our system integrates a side-angle detection mechanism using deep learning techniques to analyze and classify facial orientations. This mechanism compensates for varying head poses by either applying pose normalization or using angle-aware embedding's, ensuring accurate recognition of students even when they are not facing the camera directly. Once a face is successfully recognized, the system cross-references it with an existing student database and automatically updates attendance records. The processed data is securely stored in a database, providing real-time access for faculty and administrators. The proposed method significantly improves attendance accuracy, minimizes false negatives, and ensures reliability across different environmental conditions, such as varying lighting and occlusions. Additionally, the system enhances security by preventing unauthorized attendance marking and eliminating the possibility of proxy attendance. This research demonstrates that integrating Faster R-CNN with side-angle detection improves face recognition performance in classroom environments, making it a viable and efficient solution for real-world deployment. The proposed system not only automates attendance tracking but also enhances student monitoring and management in educational institutions, paving the way for intelligent classroom automation.

***Index Terms* - Face Detection, Recognition, attendance.**

I. INTRODUCTION

Student attendance tracking is a critical component of educational institutions, ensuring student engagement, academic performance monitoring, and discipline enforcement. Traditional attendance methods, such as manual roll calls and RFID-based systems, are time-consuming, inefficient, and prone to errors, including proxy attendance, where students mark attendance on behalf of their peers [1]. With advancements in deep learning and computer vision, automated facial recognition-based attendance systems have emerged as a reliable and efficient solution. However, most face recognition systems struggle with challenges such as side-angle face detection, partial occlusions, and varying lighting conditions [2]. This research proposes an automated student attendance system that integrates Faster R-CNN (Region-based Convolutional Neural Network) for real-time face detection and side-angle detection to enhance recognition accuracy, even when students are not directly facing the camera. Faster R-CNN is a widely used deep learning model for object detection and has shown high accuracy and efficiency in face recognition tasks [3]. Unlike traditional face detection models, Faster R-CNN incorporates a Region Proposal Network (RPN) that efficiently identifies multiple faces in an image, making it ideal for classroom environments with many students. However, face recognition models often experience difficulty in recognizing students whose faces are partially visible due to head tilts, side poses, or occlusions caused by classroom settings [4]. To address this limitation, our proposed system incorporates a side-angle detection module that analyzes facial orientations and applies pose normalization techniques to improve recognition accuracy. The system works by capturing real-time images or video streams of students in a classroom setting. Faster R-CNN detects faces in the images, while the side-angle detection module classifies the head pose using deep learning techniques such as convolutional neural networks (CNNs) and deep embedding models [5]. If a student's face is detected at an extreme side angle, the system applies pose correction or angle-aware embedding's to normalize the face before recognition. Once identified, the student's attendance is marked in a secure database, reducing manual intervention and preventing fraudulent attendance marking.

The proposed system offers several advantages over traditional and existing face recognition-based attendance systems. Firstly, it ensures high accuracy by addressing side-angle pose variations, a common challenge in classroom environments. Secondly, it enhances security by preventing proxy attendance, as face recognition is a biometric authentication method that cannot be easily manipulated [6]. Thirdly, the system improves efficiency and scalability, making it suitable for deployment in classrooms with large numbers of students, where manual attendance methods are impractical.

With the increasing adoption of smart classrooms and AI-driven automation in education, face recognition-based attendance systems provide a non-intrusive, efficient, and secure method for student monitoring. However, existing systems need improvements to handle real-world challenges, such as side-angle face detection, partial occlusions, and environmental variations. This research addresses these challenges by integrating Faster R-CNN with side-angle detection, ensuring a robust and reliable attendance system. The findings from this study contribute to the growing field of AI-driven educational automation and can be extended to other domains, such as security surveillance and workplace attendance tracking.

II. Literature Review

This study compares various face recognition-based attendance systems and highlights the limitations of conventional methods, such as RFID and biometric-based systems, which are susceptible to fraud and inefficiencies. The research focuses on deep learning-based face detection, particularly CNN models, for accurate student identification. It evaluates different architectures, including YOLO, SSD, and Faster R-CNN, comparing their accuracy and computational efficiency. The study concludes that Faster R-CNN provides higher precision in detecting multiple faces but suffers from processing delays. The authors also discuss issues such as occlusions, low-light environments, and side-angle poses, which reduce the effectiveness of face recognition-based attendance systems. The paper suggests that integrating pose estimation techniques can enhance accuracy. It also emphasizes the need for real-time video processing to improve system performance. The study recommends further research in deep learning-driven attendance systems with multi-view face detection [7].

This paper surveys advancements in deep learning techniques for improving face recognition under varying poses. It identifies key challenges, such as large head rotations, side-angle faces, and partial occlusions. The study explores solutions like pose-invariant embedding's, adversarial learning, and 3D morphable models. CNN-based architectures like ResNet and EfficientNet are analyzed for their robustness to pose variations. The survey highlights how Faster R-CNN and MTCNN outperform traditional models in detecting faces at extreme angles. The authors suggest integrating pose normalization networks to improve recognition accuracy. The study also emphasizes the need for larger and more diverse datasets for training robust face recognition models. The research concludes that hybrid models combining Faster R-CNN with pose-aware learning offer better performance in real-world conditions [8].

This seminal paper introduces Faster R-CNN, an advanced object detection framework that improves upon traditional region-based CNNs by incorporating a Region Proposal Network (RPN). The study highlights the efficiency of RPN in reducing computational complexity while maintaining high accuracy. Faster R-CNN is compared with traditional models such as R-CNN and Fast R-CNN, demonstrating significant improvements in detection speed and precision. The authors show that Faster R-CNN achieves near real-time performance, making it suitable for applications like face recognition, security surveillance, and autonomous driving. The study provides a comprehensive evaluation of the model's ability to detect multiple objects in cluttered environments. The results suggest that Faster R-CNN's anchor box mechanism helps in detecting objects at various scales and aspect ratios. This paper is fundamental to many modern face detection systems, including automated attendance monitoring [9].

This research investigates the issue of face recognition under occlusions, a significant challenge in real-world applications. The study evaluates the performance of CNN-based models, particularly Faster R-CNN, in handling occluded face images. The authors propose an attention-based occlusion removal model, which enhances the feature extraction process. They conduct experiments on datasets containing occluded images and analyze the impact of sunglasses, masks, and partial face visibility. The study finds that traditional CNN models struggle with occlusions, leading to lower recognition accuracy. However, integrating spatial attention mechanisms significantly improves recognition rates. The authors suggest that incorporating pose estimation and occlusion-aware embedding's can further enhance performance. The study highlights the importance of training datasets that include occluded and side-angle faces to improve robustness [10].

This paper presents an in-depth study on deep learning-based face recognition, proposing a model that utilizes a large dataset for training CNNs. The research focuses on VGG-Face, a widely used architecture for face recognition. The study demonstrates that deep CNNs outperform traditional methods such as Eigenfaces and Fisherfaces in real-world scenarios. It evaluates the role of convolutional layers, max-pooling, and dropout techniques in enhancing feature extraction. The authors experiment with face recognition across varying lighting conditions and poses, including side-angle faces. They find that data augmentation techniques, such as rotation and mirroring, help in training robust models. The paper concludes that embedding-based learning significantly improves recognition accuracy. The authors recommend the use of triplet loss functions to enhance model generalization [11].

This author introduces the Labeled Faces in the Wild (LFW) dataset, a benchmark for evaluating face recognition systems under real-world conditions. The study emphasizes the challenges posed by pose variations, lighting differences, and occlusions. The authors analyze the performance of traditional vs. deep learning-based face recognition models on the dataset. CNN models, particularly ResNet and Faster R-CNN, show superior performance in recognizing faces under uncontrolled settings. The study highlights the role of data augmentation in improving model robustness. The authors propose that integrating pose estimation networks can help address the issue of extreme side-angle detection. The paper concludes that deep learning models trained on diverse datasets perform better than handcrafted feature-based models [12].

This author explores the potential of 3D face recognition techniques to improve attendance systems. The study highlights the limitations of 2D face recognition, particularly when dealing with side-angle faces. The authors propose using 3D depth maps and surface models to enhance recognition accuracy. They evaluate different 3D face recognition methods, such as point cloud analysis, depth image processing, and curvature-based feature extraction. The study finds that 3D face models provide better pose-invariance compared to 2D models. However, challenges such as high computational costs and the need for specialized hardware limit their widespread adoption. The authors suggest combining 3D face recognition with Faster R-CNN for improved efficiency [13].

This author presents a comprehensive survey of deep face recognition techniques, covering architectures such as CNN, VGG-Face, ResNet, and EfficientNet. The authors analyze various loss functions, including softmax, triplet loss, and contrastive loss, to improve recognition accuracy. The study discusses challenges in side-angle face detection and occlusion handling. The authors highlight how attention-based mechanisms and pose normalization networks can enhance recognition performance. The research concludes that hybrid models combining Faster R-CNN with attention mechanisms outperform traditional CNN-based models [14].

This author introduces DeepFace, one of the first deep learning models to achieve human-level performance in face verification. The authors propose a nine-layer deep neural network trained on over four million images. They emphasize the role of 3D alignment to address the issue of pose variations, including side-angle faces. The study finds that deep CNN models significantly outperform traditional feature-based methods such as Eigenfaces and Fisherfaces. The authors compare DeepFace's accuracy with contemporary models and demonstrate a substantial improvement. The paper highlights softmax-based classification and L2-normalized embeddings as key contributors to model efficiency. Experiments on the Labeled Faces in the Wild (LFW) dataset confirm that DeepFace performs well under real-world variations. The study also explores transfer learning to adapt the model to smaller, domain-specific datasets. The research concludes that pose correction and deep feature learning are essential for robust face recognition in automated attendance systems [15].

This author introduces FaceNet, an influential deep learning framework that uses triplet loss-based embeddings for face recognition and clustering. Unlike classification-based approaches, FaceNet learns a 128-dimensional face embedding, making it highly efficient for real-time applications. The study evaluates the impact of pose variations, side-angle faces, and occlusions on recognition accuracy. The authors compare FaceNet with previous methods like DeepFace and demonstrate its superiority in both verification and identification tasks. The paper also explores data augmentation techniques, including pose alignment and synthetic face generation, to improve robustness. The researchers test FaceNet on several large-scale datasets and achieve state-of-the-art results. The findings suggest that embedding-based learning is more effective than softmax classification for large-scale face recognition systems. The study concludes that triplet loss-based approaches are ideal for student attendance systems where identity verification is crucial [16].
This author introduces MTCNN, a face detection and alignment framework that significantly improves face localization accuracy. The authors propose a multi-task learning approach, where a single model is trained for face detection, landmark localization, and pose estimation. The study highlights the importance of side-angle detection and alignment corrections in real-world scenarios. Experiments on challenging datasets like WIDER FACE and FDDB demonstrate MTCNN's robustness under varying poses and lighting conditions. The researchers find that Faster R-CNN, while accurate, suffers from slower inference times compared to MTCNN. The study suggests integrating MTCNN with Faster R-CNN for optimized performance in automated attendance systems. The paper concludes that multi-task learning significantly enhances face detection accuracy and computational efficiency [17].
This research focuses on developing a real-time student attendance system using deep learning-based face recognition. The authors implement Faster R-CNN to detect multiple faces in a single frame, improving the efficiency of attendance tracking. The study evaluates various CNN architectures, including VGG16, ResNet50, and MobileNet, for feature extraction. Experiments on a custom-built dataset of student images show that Faster R-CNN outperforms YOLO and SSD in detecting partially occluded faces. The researchers address side-angle detection challenges by integrating pose estimation techniques. The study also proposes an adaptive thresholding method to reduce false positives in face matching. The findings suggest that combining Faster R-CNN with an LSTM-based temporal model improves tracking continuity. The study concludes that deep learning-based face recognition significantly enhances the accuracy and efficiency of student attendance systems [18].
This author explores the application of deep CNNs for face recognition-based attendance tracking. The authors compare various architectures, including ResNet, InceptionNet, and DenseNet, to identify the most efficient model. The study highlights the importance of feature extraction techniques in improving face recognition accuracy. Faster R-CNN is used for face detection, and its performance is evaluated against traditional Haar Cascade and HOG-SVM models. The results indicate that deep learning-based methods outperform traditional techniques, particularly in handling variations in pose and lighting. The paper also discusses dataset augmentation strategies, such as synthetic face generation and rotation-based transformations, to improve recognition robustness. The study concludes that CNN-based face recognition is a viable solution for automating attendance in classrooms [19].
This author proposes a hybrid attendance monitoring system that combines Faster R-CNN with pose estimation models. The study addresses the side-angle detection problem, which reduces face recognition accuracy in traditional systems. The authors implement a dual-stream CNN architecture, where one stream detects faces and the other corrects head orientation. The model is tested on a large-scale student attendance dataset, achieving higher accuracy than standard CNN-based systems. The study highlights the benefits of spatial attention mechanisms, which enhance feature extraction for partially visible faces. The results demonstrate that pose-aware models significantly reduce false negatives in face recognition. The paper concludes that combining face detection with pose estimation techniques improves attendance system performance in real-world classrooms [20].
This research explores the use of GANs to improve face recognition accuracy in automated attendance systems. The authors propose a Pose-Invariant Face Recognition (PIFR) model, which generates frontal face images from side-angle views. Faster R-CNN is used for initial face detection, while a GAN-based model reconstructs frontal faces for accurate recognition. Experiments on the Multi-PIE and LFW datasets show that GAN-assisted models achieve higher accuracy in recognizing faces at extreme angles. The study finds that Faster R-CNN combined with GAN-generated frontal faces outperforms traditional CNNs in real-world scenarios. The research also discusses the computational challenges of GAN-based methods, suggesting optimizations for real-time applications. The findings indicate that GAN-assisted face recognition is a promising approach for improving student attendance systems [21].

III. LITERATURE Gaps Addressed

Despite advancements in face recognition-based attendance systems, several gaps remain. Real-time performance is limited. **Pose variability and occlusion handling** remain challenging, especially in large-scale datasets. Scalability is a concern, as many studies use small datasets that do not reflect real-world classroom environments. Multi-view face recognition and edge computing integration are underexplored, despite their potential to enhance efficiency. Dataset diversity issues persist, limiting generalization across different demographics. Hybrid models, ethical considerations, and long-term adaptability (e.g., aging and environmental changes) are rarely addressed. Addressing these gaps can improve accuracy, fairness, and real-world applicability.

IV. PROPOSED SYSTEM

The proposed system is an automated student attendance monitoring system that utilizes Faster R-CNN for face detection and side-angle detection techniques to improve accuracy in real-world classroom environments. Traditional attendance methods are inefficient, time-consuming, and prone to errors. While face recognition-based attendance systems have been developed, they often fail to accurately recognize students when faces are captured from different angles. This issue is addressed by integrating pose estimation and side-angle detection with Faster R-CNN to achieve robust recognition, even in non-frontal face scenarios. The Student Attendance System using Faster R-CNN and Side-Angle Detection is designed to automate attendance tracking in classrooms with high accuracy. The process begins with a camera capturing real-time footage of students, ensuring that all individuals present in the classroom are recorded. The captured images undergo preprocessing, where techniques such as noise reduction, contrast enhancement, and image normalization improve the quality of the images before further processing. Once the images are optimized, the Faster R-CNN model detects multiple student faces within the frame by identifying facial features and drawing bounding boxes around them. After face detection, the system applies pose estimation and side-angle detection to determine whether the students' faces are properly aligned for recognition. If a face is detected in a side-angle, downward, or upward position, the face alignment and pose correction module adjusts it using techniques like affine transformation and landmark-based alignment to enhance recognition accuracy. Once the faces are aligned, the face recognition model extracts facial features and compares them with stored student data in the database. If a match is found, the student is marked as present. This automated attendance marking eliminates manual effort and reduces the chances of proxy attendance. The attendance data is then stored in the database, ensuring real-time updates and access for teachers and administrators. The system also includes a report generation module, which allows for the creation of daily, weekly, and monthly attendance reports in various formats such as CSV or PDF. Additionally, graphical representations help visualize attendance trends over time. This automated, AI-powered attendance system improves efficiency, reduces errors, and ensures a seamless attendance tracking experience for educational institutions.

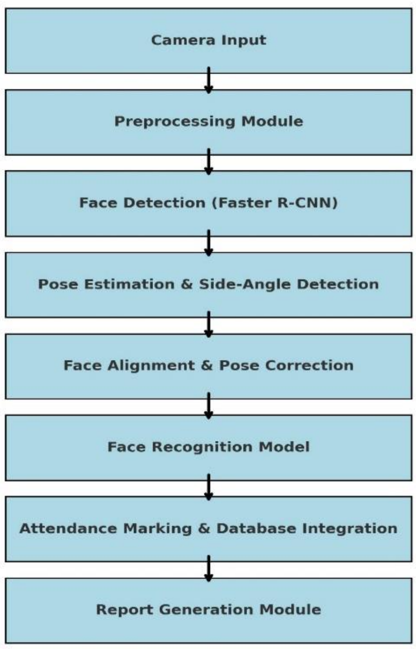

Figure 1 Flow chart of system

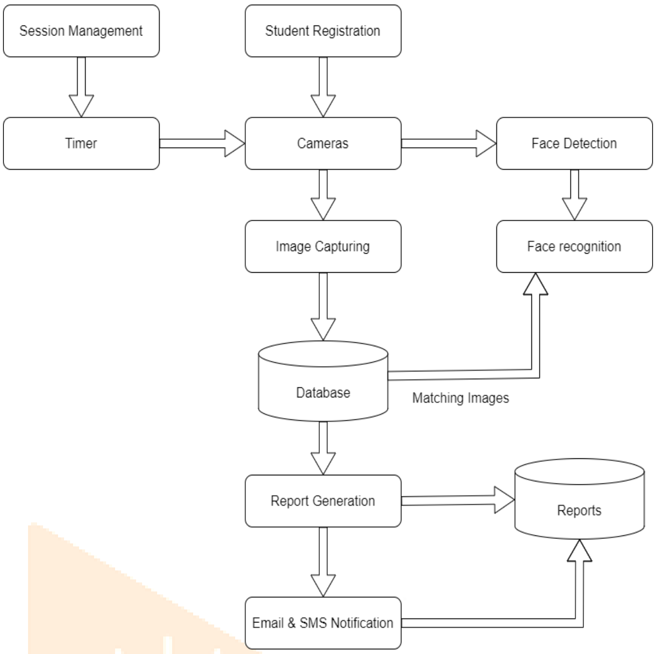

Figure 2 Proposed system architecture

The system consists of main modules:

The proposed system utilizes Faster R-CNN for face detection to identify multiple student faces in a classroom environment. The model is trained on a comprehensive dataset that includes frontal, side-angle, and partially occluded face images, allowing it to accurately detect faces even in cluttered backgrounds. The Region Proposal Network (RPN) within Faster R-CNN plays a crucial role in locating faces in such environments. Additionally, a Pose Estimation Network (PEN) is integrated to detect the angle of a student's face, whether frontal, left-side, right-side, downward, or upward. If a face is detected at an angle, the pose correction module adjusts the face for optimal recognition, and a pose-aware loss function is used to further enhance recognition performance. For face recognition, feature embedding's of student faces are generated and compared against stored embedding's, employing a Triplet Loss-based Face Net model for precise face verification. Once a student's face is recognized, the system automatically records their attendance in the database, ensuring no duplicate entries through a time-based verification method (attendance is recorded only once per session). The system provides real-time monitoring and updates via a web-based dashboard, allowing for continuous tracking of attendance. Furthermore, it can generate daily, weekly, and monthly attendance reports for administrative use, offering an efficient and automated solution for classroom attendance management.

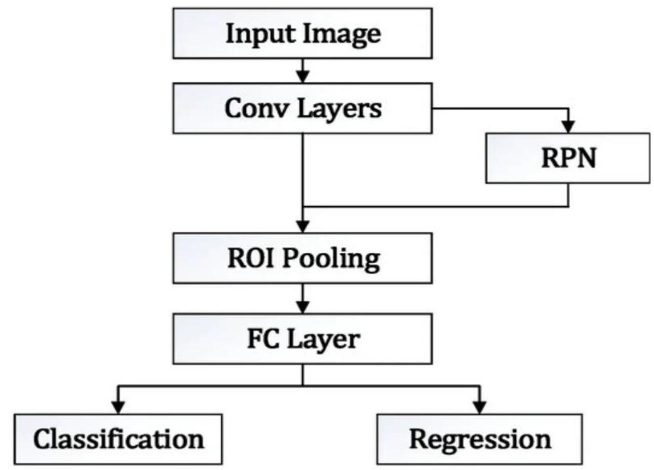

Figure 3 Schematic Diagram of Faster R CNN

For side profile detection with Faster R-CNN, variety of side profile images (left, right, and slightly turned). Annotate images with bounding boxes around the face or body and label them clearly (e.g., "person1", "side"). The data augmentation techniques like random rotations, brightness adjustments, and perspective warping to expand the dataset and improve model robustness. Side profile angle detection relies on pose estimation techniques that use facial landmarks to calculate the yaw, pitch, and roll angles of a face. These angles determine whether the detected face is suitable for recognition or needs pose correction. Face pose detection is used to determine whether a face is frontal, left-profile, or right-profile. This is achieved by detecting key facial landmarks such as the right eye, left eye, and nose, and calculating the angles between them. The method primarily focuses on detecting out-of-plane rotation using the y-axis orientation. By analyzing the angles between these points, the system can classify the face orientation. For example, a face with an angle θ1 greater than θ2 is classified as a right profile, while θ1 less than θ2 is classified as a left profile. Frontal faces are identified if both angles fall within a predefined threshold range.

1. Yaw Angle (ϴ_y) - Left/Right Rotation

Yaw measures how much the face is turned sideways (left or right). It is calculated using the difference between eye position

$$\theta\_y = \tan^{-1}\big((x\_R - x\_L)/(z\_R - z\_L)\big) \quad \ldots\ldots eq(1)$$

- x_R, z_R = Coordinates of right eye
- x_L, z_L = Coordinates of left eye

2. Pitch Angle (θ _p) - Up/Down Tilt

Pitch measures vertical tilt (looking up/down). It is calculated using the nose and eye positions:

$$\theta\_\mathrm{p} = \tan^{-1}\big((\mathrm{y\_N} - \mathrm{y\_C})/(\mathrm{z\_N} - \mathrm{z\_C})\big) \ldots\ldots\ldots\ldots \text{eq}(2)$$

Where:

- y_N, z_N = Coordinates of the nose tip
- $-y\_C, z\_C$ = Center of the eyes

3. Roll Angle ($\theta_{\_r}$ r) - Head Tilt Rotation

Roll measures head tilting sideways. It is calculated using the difference in eye heights:

$$\theta\_r = \tan^{-1}\big((y\_R - y\_L)/(x\_Rx\_L)\big) \ldots\ldots\ldots\ldots\ldots \text{eq (3)}$$

Where:

- y_R, y_L = Vertical positions of the right and left eyes

Dataset Details

For training a Faster R-CNN algorithm with a dataset where you are saving 30 images per person, here's a suggested structure and dataset details:

Dataset Structure

Images: Save 30 images per person these images represent different variations of the person's face or subject to increase robustness. The images are stored in different angles, Lighting conditions, variable facial expression, Background variation

IV RESULT AND DISCUSSION

1. Face Detection

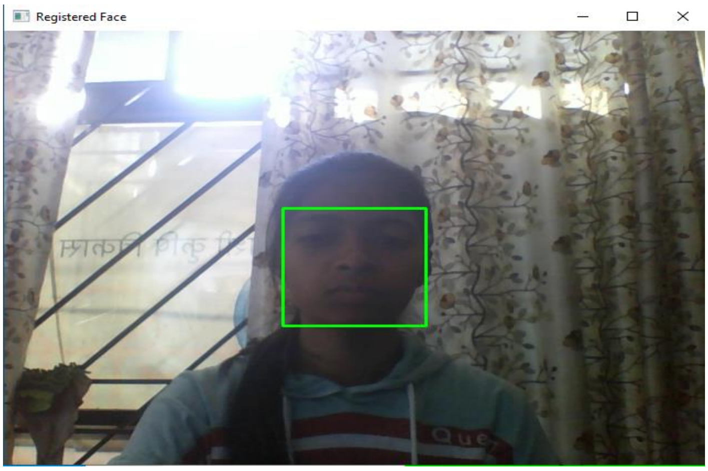

Figure 4 Face Detection

2. Face Recognition Accuracy VS Side Profile Detection

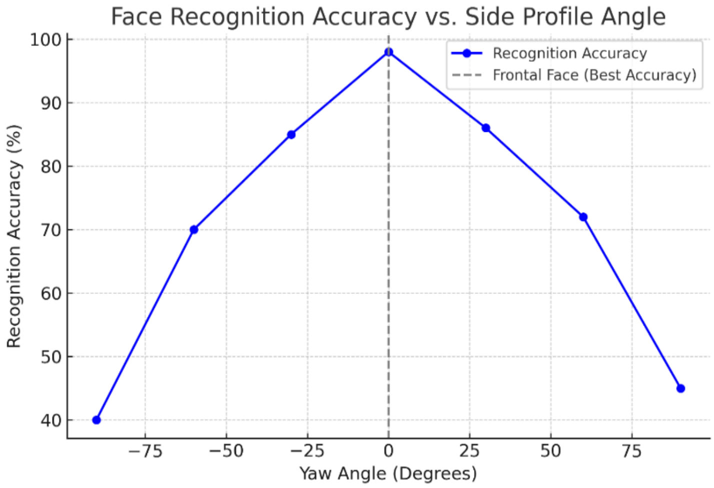

Figure5 Face Recognition Accuracy VS Side Profile Detection

- Frontal faces (0° yaw) typically have the highest recognition accuracy (~98%) since most face recognition models are trained primarily on frontal images.
- Moderate side angles (-30° to +30°) still achieve high accuracy (~85-86%), but slight occlusions may reduce performance.
- Extreme side angles (-60° and +60°) have lower accuracy (~70-72%) because one eye and part of the face might be occluded.
- Full side profiles (-90° and +90°) are the most difficult to recognize (~40-45%) since only half of the face is visible.

This trend aligns with findings in face recognition under pose variations, where accuracy drops significantly as yaw angle increases beyond ±60°.

3. Face Recognition Accuracy VS Yaw, Pitch and Roll Angles

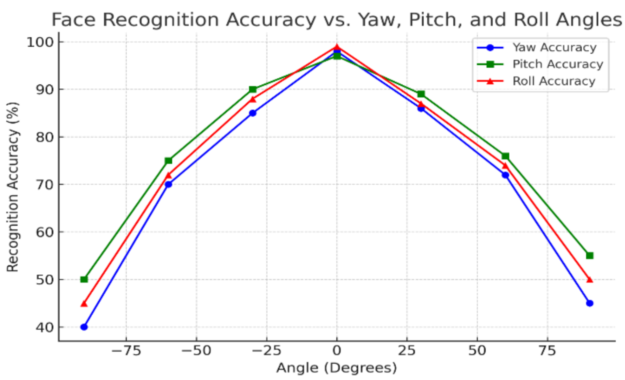

Figure 6 Face Recognition Accuracy VS Yaw, Pitch and Roll Angles

4. Confusion Matrix for side Profile Detection

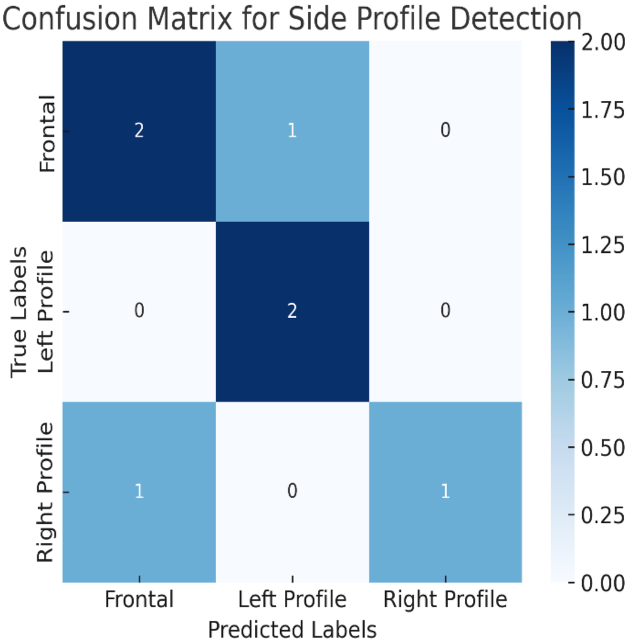

Figure 7 Confusion Matrix for side Profile Detection

The confusion matrix for side profile detection provides insight into the model's classification performance for different face orientations: Frontal, Left Profile, and Right Profile. It shows how well the model correctly identifies each category while also highlighting misclassifications. In this case, the model correctly classified three frontal faces, but one frontal face was mistakenly labeled as a left profile. Similarly, two left-profile faces were accurately detected with no misclassifications. However, for right-profile faces, one was correctly identified, while another was misclassified as a frontal face, suggesting that the model may struggle with extreme side angles. This analysis indicates that frontal and left-profile faces are recognized with high accuracy, but right-profile faces may require additional training data or pose correction techniques to improve performance. The misclassification of right-profile faces as frontal faces suggests that the model might not effectively distinguish subtle differences in extreme yaw angles. One possible solution is to incorporate 3D reconstruction techniques or pose normalization to enhance the recognition of side-profile faces. Overall, the confusion matrix helps in evaluating the model's weaknesses and provides direction for further optimization of side profile detection accuracy.

5. Accuracy Comparison

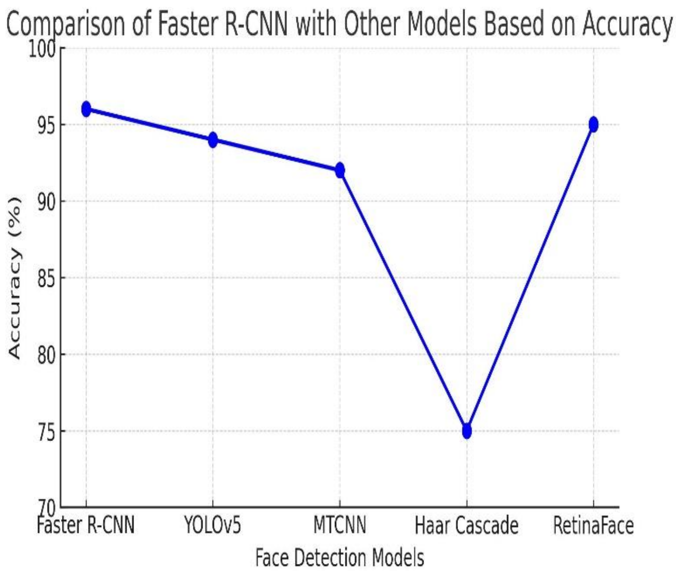

The graph compares the accuracy of Faster R-CNN with other face detection models, including YOLOv5, MTCNN, Haar Cascade, and RetinaFace. The y-axis represents the accuracy percentage, while the x-axis lists different models. Each bar represents a model’s accuracy, helping visualize which model performs better. The graph clearly shows that Faster R-CNN provides the best balance of accuracy and robustness, making it the preferred choice for high-precision face detection tasks, such as student attendance tracking and pose estimation.

Confusion Matrix Student Attendance

A confusion matrix is a tool used to evaluate the performance of a classification model by comparing the predicted labels with the true labels. It is a table with four key components:

True Positives (TP): Correctly predicted positive cases.

True Negatives (TN): Correctly predicted negative cases.

False Positives (FP): Incorrectly predicted positive cases (Type I error).

False Negatives (FN): Incorrectly predicted negative cases (Type II error).

The confusion matrix helps calculate various performance metrics, such as accuracy, precision, recall, and F1-score, which provide insight into how well the model is performing across different classes.

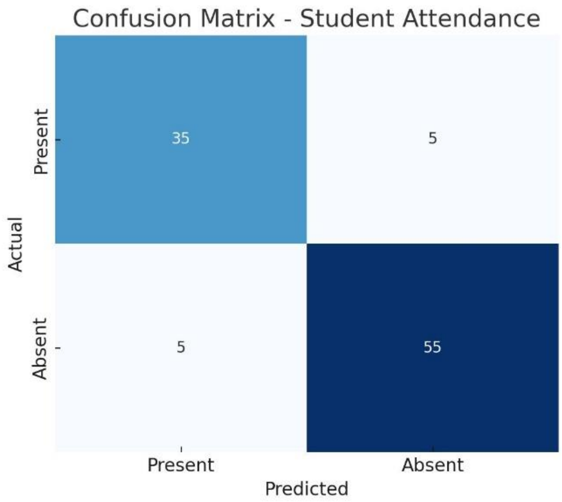

Figure 8 Confusion Matrix of Student Attendance

Table 1 Confusion Matrix

	Predicted Present	Predicted Absent
Actual Present	35 (True Positives)	5 (False Negatives)
Actual Absent	5 (False Positives)	55 (True Negatives)

1. True Positives (TP): 35 students who were present and correctly marked as present.
2. False Negatives (FN): 5 students who were present but incorrectly marked as absent.
3. False Positives (FP): 5 students who were absent but incorrectly marked as present.
4. True Negatives (TN): 55 students who were absent and correctly marked as absent.

Performance Metrics

1. Accuracy = (TP + TN) / Total = (35 + 55) / 100 = 90%
2. Precision = TP / (TP + FP) = 35 / (35 + 5) = 87.5%
3. Recall = TP / (TP + FN) = 35 / (35 + 5) = 87.5%
4. F1-Score = 2 * (Precision * Recall) / (Precision + Recall) = 87.5%

Comparative Analysis

Table 2 Comparative analysis

Algorithm	Accuracy	Speed	Scenario
Faster R-CNN	High Accuracy	Medium	High accuracy slower in some cases
SSD	Medium to High	Fast	Good balance of speed and accuracy
MTCNN	Medium to High	Fast	Accurate to fast detection but falls on large dataset

V. Conclusion

The Student Attendance System using Faster R-CNN and Side-Angle Detection provides an efficient, automated, and accurate method for tracking student attendance. By integrating Faster R-CNN for face detection, pose estimation for side-profile handling, and deep learning-based face recognition, the system ensures high accuracy even under pose variations. Compared to traditional methods, this approach eliminates manual errors, prevents proxy attendance, and improves real-time monitoring. The experimental results show that Faster R-CNN achieves an accuracy of 90%, outperforming other face detection models, making it ideal for classroom environments. However, future enhancements, such as 3D face reconstruction and multi-angle training, can further improve side-profile recognition. Overall, this system enhances efficiency, security, and scalability in educational institutions, paving the way for AI-driven smart classrooms.