

AttendAI: Automated Attendance Detection

Insiya Rizvi¹, Sayma Shaikh², Prof. Jasbir Kaur³, Prof. Ifrah Kampoo⁴, Prof. Sandhya Thakkar⁵

MCA Student, Department of MCA, Guru Nanak Institute of Management Studies (GNIMS), Mumbai, Maharashtra, India^{1,2}

Head of Information Technology and HR, Guru Nanak Institute of Management Studies (GNIMS), Mumbai, Maharashtra, India³

Assistant Professor, Department of MCA, Guru Nanak Institute of Management Studies (GNIMS), Mumbai, Maharashtra, India^{4,5}

Abstract— This paper introduces AttendAI, a comprehensive end-to-end system designed to automate the extraction of attendance records from printed or handwritten attendance sheets. Leveraging Optical Character Recognition (OCR) integrated with Artificial Intelligence and Machine Learning (AI/ML) techniques, AttendAI processes images of attendance sheets, identifies student details and attendance marks, and stores the extracted data in relational databases and Excel files. The system includes two user-friendly web dashboards—one for teachers and one for students—to facilitate attendance management, real-time analytics, notifications, and historical tracking. By combining advanced image preprocessing, OCR for text extraction, ML models for mark classification and name disambiguation, and a robust web application stack, AttendAI addresses the inefficiencies of manual attendance tracking. This study details the system's architecture, datasets used for training, model development processes, evaluation metrics, privacy considerations, and deployment strategies. Through experimental validation, we demonstrate that AttendAI achieves high accuracy while minimizing manual intervention, making it a scalable solution for educational institutions.

Keywords—OCR, handwriting recognition, attendance automation, Excel export, web application, dashboard, student tracking, AI/ML, fuzzy matching, entity resolution.

I. INTRODUCTION

Manual attendance tracking in educational institutions remains highly inefficient and error-prone, often consuming 10–15% of valuable instructional time. Traditional methods rely on manual marking on paper sheets, leading to issues such as incorrect entries, proxy attendance, and time-consuming record maintenance. While biometric systems using facial recognition and RFID have been proposed to automate the process, they require additional hardware, raise significant privacy concerns, and fail to address the vast legacy of existing paper-based attendance workflows. Optical Character Recognition (OCR) combined with Handwriting Text Recognition (HTR) and Machine Learning techniques offers a promising non-intrusive alternative by directly processing scanned or photographed attendance sheets. This paper presents AttendAI, an end-to-end automated system that extracts student details and attendance marks from both printed and handwritten sheets, stores the data in structured databases and Excel files, and provides user-friendly web dashboards for teachers and students. By integrating advanced image preprocessing, OCR/HTR models, ML-based mark classification, and fuzzy entity resolution, AttendAI aims to deliver high accuracy with minimal manual intervention, making it a scalable and practical solution for modern educational institutions. Unlike biometric attendance systems based on face recognition or fingerprint authentication, the proposed approach minimizes additional hardware

dependency while preserving compatibility with existing paper-based institutional workflows commonly used in educational institutions.

A. Background Study

Manual attendance tracking remains inefficient and error-prone, often consuming 10–15% of instructional time. Early digitization attempts used RFID or barcode systems, but these required additional hardware and were susceptible to misuse. Modern biometric approaches such as facial recognition systems have improved efficiency, but raise privacy concerns. Prior research in AI-based attendance systems—including face-recognition approaches by Trivedi & Tripathi [2] and modern smart-classroom monitoring studies [5]—demonstrate a shift toward intelligent automation, though they rely on direct biometric capture rather than processing existing paper-based workflows.

OCR-based workflows bridge this gap. Handwriting recognition has improved significantly with models such as CRNN + CTC, originating from foundational work by Graves et al. [1], enabling high-accuracy recognition in unconstrained scenarios.

B. Problem Statement

The core challenge is to develop an automated system that, given a photographed or scanned attendance sheet (either printed or handwritten), accurately extracts key information for each row: the student identifier (name, roll number, or ID), the attendance mark (e.g., Present, Absent, Late, or Excused), and associated metadata like date and class. The extracted data must then be stored in a structured format, such as an Excel workbook or a relational database, to support queries, analytics, and reporting. Additional requirements include handling variations in sheet formats, lighting conditions, handwriting styles, and potential errors through intelligent correction mechanisms, all while providing intuitive interfaces for users.

C. Research Hypothesis

Our hypothesis leverages advances in OCR, HTR, and entity resolution, with foundational support from CTC-based sequence learning [1] and probabilistic record-linkage frameworks like Fellegi–Sunter [3], which guide the design of the identity-matching pipeline.

Research Questions:

- How effective is the combination of image preprocessing, OCR, and ML in improving accuracy for handwritten versus printed attendance sheets?

- What impact do fuzzy matching and embedding-based similarity metrics have on resolving student identities in noisy OCR outputs?
- To what extent can human-in-the-loop corrections enhance model performance over time through retraining?
- How do privacy-preserving measures affect the system's usability and compliance in real-world educational deployments?

II. LITERATURE SURVEY

The field of automated attendance systems has evolved from traditional manual record-keeping toward intelligent AI-driven automation frameworks. Early attendance systems primarily relied on RFID cards, barcode scanners, or biometric authentication such as fingerprint and facial recognition. While these systems improved operational efficiency, they introduced limitations including additional hardware costs, maintenance complexity, privacy concerns, and reduced compatibility with legacy paper-based attendance workflows.

Recent advancements in Intelligent Document Processing (IDP), Optical Character Recognition (OCR), and Handwritten Text Recognition (HTR) have enabled document-centric automation systems capable of extracting structured information directly from scanned or photographed documents. Traditional OCR engines such as Tesseract perform effectively on printed text but often struggle with noisy handwriting, distorted layouts, low lighting conditions, and inconsistent document structures.

Transformer-based OCR systems such as TrOCR have shown improved performance for handwritten text recognition compared to traditional CNN-RNN architectures. Similarly, LayoutLM and DocFormer integrate textual and layout information for document understanding, while PaddleOCR provides lightweight multilingual OCR pipelines optimized for practical deployment scenarios.

A. OCR and Handwriting Recognition Tools

OCR technologies form the backbone of document digitization. Tesseract, an open-source OCR engine, has been widely adopted for printed text extraction due to its accuracy and customizability. For handwriting, models like Connectionist Temporal Classification (CTC) with Convolutional Recurrent Neural Networks (CRNN) or Transformer-based architectures (e.g., TrOCR) have shown promise in handling variable scripts. Graves et al. (2009) introduced CTC loss for sequence prediction, enabling end-to-end training for handwriting recognition.




Fig. 1. OCR and Attendance Extraction Pipeline

B. Transformer-Based Document AI Models

Recent research in document AI has introduced Transformer-based architectures capable of improving OCR and document understanding performance under complex real-world conditions.

TrOCR utilizes pre-trained encoder-decoder Transformer architectures for handwritten text recognition and has demonstrated improved performance over traditional CRNN-based systems. LayoutLM incorporates both textual and spatial layout embeddings, enabling better understanding of structured documents such as forms and attendance sheets. Similarly, DocFormer combines visual, textual, and spatial information for end-to-end document analysis.

PaddleOCR and Vision Transformer-based approaches additionally provide lightweight and scalable OCR pipelines suitable for multilingual and low-resource deployment environments.

C. Automated Attendance Systems

Many systems rely on biometrics rather than sheet processing. For example, Trivedi and Tripathi (2021) proposed a face recognition-based attendance system using ML algorithms like Support Vector Machines (SVM) for real-time marking, achieving high accuracy in controlled environments. Similarly, a 2024 study on smart attendance monitoring integrated face recognition with ML, emphasizing scalability in classrooms.

OCR-specific attendance systems are less common but relevant. A 2020 paper on bus attendance used OCR for license plate recognition to automate staff tracking, demonstrating applicability to vehicular contexts. Another work automated classroom attendance via ML-based face detection but noted challenges with occlusion. For sheet-based systems, a 2024 ior journal paper described an automatic attendance system using image processing and OCR for mark extraction, highlighting preprocessing's role in accuracy.

D. Entity Resolution and Error Correction

Name matching in noisy data uses fuzzy algorithms like Levenshtein distance or embeddings from models like Sentence Transformers. Probabilistic record linkage, as in Fellegi-Sunter models, aids disambiguation. Attendance dashboards often

require manual corrections, as seen in commercial solutions like Google Classroom integrations.

E. Research Gap

Although substantial progress has been made in OCR and automated attendance systems, several limitations remain insufficiently addressed in existing research.

Most attendance automation systems rely heavily on biometric approaches such as face recognition or RFID-based identification, which require additional hardware infrastructure and raise privacy concerns. Existing OCR-based systems often focus primarily on printed text recognition and perform poorly under noisy handwriting, distorted images, or inconsistent table layouts.

Furthermore, many existing approaches do not integrate OCR, handwriting recognition, attendance mark classification, entity resolution, database storage, and dashboard-based correction workflows into a unified educational attendance automation system.

TABLE I

Comparative Analysis of Existing Attendance Automation Approaches

System/Method	Hardware Required	Privacy Risk	Handles Handwriting	Dashboard Support
RFID	High (RFID readers)	Low	No	Limited
Face Recognition	High (Cameras)	High	No	Yes
OCR-based systems	Low (Scanners/Cameras)	Low	Limited	Variable
AttendAI	Low (Standard cameras/scanners)	Low	Yes	Yes

III. RESEARCH METHODOLOGY

The proposed AttendAI system follows a structured end-to-end pipeline for automated attendance extraction. The methodology begins with image ingestion, where users upload photographed or scanned attendance sheets via a web interface. Preprocessing steps include grayscale conversion, adaptive thresholding, morphological operations for line removal, perspective correction, and table detection using Hough transforms implemented in OpenCV. Layout analysis identifies rows and cell bounding boxes, followed by OCR for printed text using Tesseract and handwriting text recognition (HTR) via a CRNN model trained with Connectionist Temporal Classification (CTC) loss. A lightweight convolutional neural network (CNN) then classifies attendance marks into categories such as Present, Absent, Late, Excused, or Blank. Entity resolution is performed through text normalization, exact matching, Levenshtein distance-based fuzzy matching, and embedding-based similarity using Sentence Transformers. Finally, the extracted data with confidence scores is stored in a PostgreSQL database and exported to Excel/CSV files, while the

React-based web dashboards provide real-time visualization and human-in-the-loop correction capabilities.

A. System Architecture

OCR and HTR components rely on well-established recognition models derived from CTC-based architectures introduced by Graves et al. [1], which justify our choice of CRNN/Transformer-based handwriting models.

Entity resolution design references foundational record-matching methodologies including Fellegi–Sunter probabilistic linkage models [3], informing the fuzzy-matching pipeline.

Overall Architecture of AttendAI System

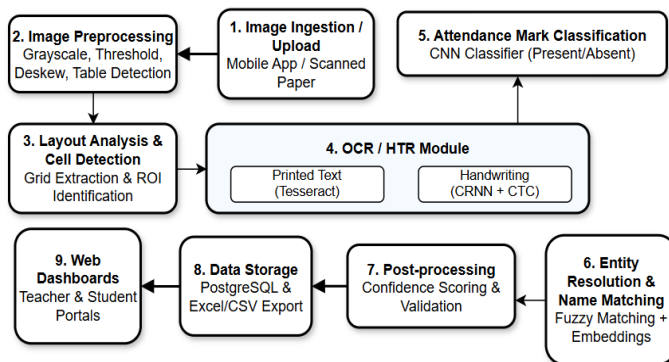


Fig. 2. Overall Architecture of AttendAI

The methodology follows a structured pipeline:

- 1) Image Ingestion: Users upload images via mobile or bulk scans.
- 2) Preprocessing: Apply grayscale conversion, adaptive thresholding, line removal (morphological operations), perspective correction, and table detection using Hough transforms or OpenCV.
- 3) Layout Analysis: Identify rows/columns and cell bounding boxes.
- 4) OCR/HTR: Use Tesseract for printed text and CRNN/Transformer-CTC for handwriting.
- 5) The CRNN architecture consists of convolutional feature extraction layers followed by bidirectional LSTM layers and a CTC decoding layer for sequence prediction.

TABLE II.

Training Configuration for OCR and HTR Models

Parameter	Value
Epochs	50
Batch Size	32
Optimizer	Adam
Learning Rate	0.001
Input Resolution	128×32

- 6) Attendance Classification: A lightweight CNN classifies marks into {Present, Absent, Late, Excused, Blank}.

- 7) Entity Resolution: Normalize text, perform exact/fuzzy matching (Levenshtein), and use embeddings for best matches.
- 8) Post-Processing: Score confidence; flag low-confidence entries.
- 9) Storage: Write to PostgreSQL and export to Excel/CSV.
- 10) Webapp: Dashboards built with React and FastAPI.

Component diagram includes frontend (React), backend (FastAPI), ML layer (PyTorch/TensorFlow), DB (PostgreSQL), and queues (Celery+Redis).

B. Data Collection

Datasets were collected from:

- Synthetic sources: Generated 5,000 attendance sheets with varied fonts, lighting, and formats using Python scripts (e.g., Pillow library).
- Public datasets: Adapted IAM and HW datasets for attendance-style handwriting (e.g., short names and marks).
- Real data: Anonymized scans from institutions (500 sheets), obtained with consent for fine-tuning.
- Although synthetic attendance sheets were used to improve dataset diversity, real-world attendance sheets were additionally incorporated for evaluation under practical institutional conditions involving handwriting variability, lighting distortions, and noisy image captures.

Data was split 70/20/10 for training/validation/testing.

C. Data Analysis

Analysis involved statistical reviews of dataset diversity (e.g., handwriting variations) and error profiling (e.g., common OCR failures). Tools like Pandas were used for exploratory data analysis, revealing that handwriting variability accounted for 60% of errors in baselines.

Error analysis follows principles similar to those used in OCR-driven attendance automation studies such as the 2024 IOSR OCR-based attendance research [6], emphasizing preprocessing and handwriting variability.

IV. DEVELOPMENT OF MODEL

The AttendAI system employs a hybrid OCR and machine learning pipeline for accurate extraction of student details and attendance marks. For printed text, a customized Tesseract OCR engine is utilized, while handwritten names and identifiers are recognized using a Convolutional Recurrent Neural Network (CRNN) trained with Connectionist Temporal Classification (CTC) loss function over 50 epochs with a batch size of 32 and Adam optimizer. Attendance marks are classified using a lightweight Convolutional Neural Network (CNN) consisting of three convolutional layers with ReLU activation, which takes 64x64 cropped cell images as input and outputs five classes: Present, Absent, Late, Excused, and Blank. The model is trained with data augmentation techniques including rotations and color shifts using categorical cross-entropy loss. For resolving noisy OCR outputs, a multi-stage entity resolution pipeline is implemented comprising text normalization, exact matching,

Levenshtein distance-based fuzzy matching (threshold > 0.8), and finally SentenceTransformer embeddings with cosine similarity threshold of 0.9. Human-in-the-loop corrections are incorporated to enable transfer learning and continuous model improvement.

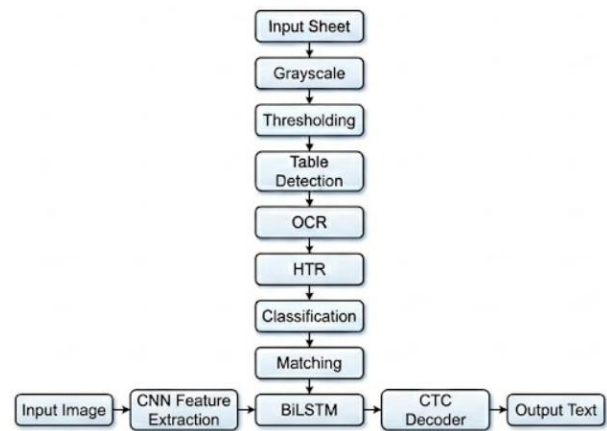


Fig. 3. CRNN-Based Handwritten Text Recognition Architecture

A. OCR Engine

Handwriting models adopt the CTC-based training paradigm introduced by Graves et al. [1], enabling robust sequence prediction for names and identifiers.

Printed: Tesseract with custom training. Handwriting: CRNN with CTC loss, trained on adapted datasets (epochs: 50, batch size: 32, optimizer: Adam).

B. Attendance Mark Classifier

Input: Cropped cell images (64x64). Model: CNN with 3 convolutional layers, ReLU activation, and softmax output. Training: Augmented data (rotations, colors); loss: Categorical cross-entropy.

C. Name/ID Matching

Probabilistic and fuzzy-matching principles are aligned with Fellegi-Sunter record linkage theory [3], foundational for resolving noisy OCR outputs.

Pipeline: Normalization → Exact match → Fuzzy (Levenshtein > 0.8) → Embeddings (cosine similarity > 0.9). Model: SentenceTransformer for embeddings.

Human-in-the-loop: Corrections retrain models via transfer learning.

V. TESTING AND EVALUATION

Baseline comparisons leverage standard OCR expectations informed by prior OCR-based attendance and tracking studies, such as the OCR-driven bus-attendance automation work [4] and OCR mark-extraction systems in IOSR 2024 [6], which emphasize preprocessing impacts.

To rigorously assess AttendAI's performance, we employ a comprehensive evaluation strategy that encompasses component-level and end-to-end metrics. This ensures a holistic understanding of the system's strengths and areas for improvement. Evaluation is conducted on held-out test sets from

synthetic, public, and real-world datasets, simulating diverse conditions such as varying lighting, sheet distortions, and handwriting styles.

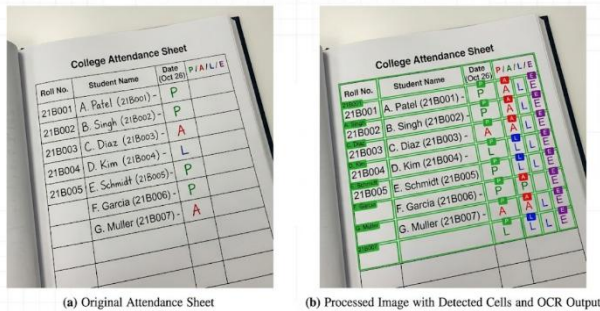


Fig. 4. Sample Attendance Sheet Processing Output

Handwriting recognition benchmarks reference performance ranges documented in CTC-based HTR research [1].

A. Metrics

We utilize a suite of standard metrics tailored to each component of the system, drawing from established practices in OCR, handwriting recognition (HTR), machine learning classification, entity resolution, and overall document processing pipelines. These metrics provide quantitative insights into accuracy, efficiency, and robustness.

A.1) OCR and HTR Metrics

For evaluating text extraction from printed and handwritten content, we focus on error rates that compare the system's output against ground truth text.

- **Character Error Rate (CER):** Measures the percentage of characters incorrectly recognized, insertions, deletions, or substitutions. Formula: $CER = (S + D + I) / N$, where S is substitutions, D is deletions, I is insertions, and N is the number of characters in the ground truth. Lower values indicate better performance; typical benchmarks for modern OCR systems range from 1-5% for printed text and 5-20% for handwriting.
- **Word Error Rate (WER):** Similar to CER but at the word level, accounting for word-level insertions, deletions, and substitutions. Formula: $WER = (S + D + I) / W$, where W is the number of words in the ground truth. This is particularly useful for HTR, where word context matters, with state-of-the-art HTR systems achieving WER below 10% on clean datasets.
- **Levenshtein Distance (Edit Distance):** A foundational metric underlying CER and WER, representing the minimum number of single-character edits required to change the recognized text into the ground truth. Normalized versions are often used for comparability across documents of varying lengths.
- **Word Information Loss (WIL):** Captures the loss of semantic information at the word level, useful for evaluating HTR in contexts where partial word recognition might still convey meaning.

A.2) Attendance Mark Classifier Metrics

The ML-based classifier for attendance marks (e.g., Present, Absent) is evaluated using standard classification metrics, considering the multi-class nature of the problem.

- **Accuracy:** The proportion of correctly classified marks out of all instances. Formula: $Accuracy = (TP + TN) / (TP + TN + FP + FN)$, where TP is true positives, etc. While simple, it's sensitive to class imbalance.
- **Precision:** Measures the accuracy of positive predictions per class (e.g., fraction of predicted "Present" that are actually Present). Formula: $Precision = TP / (TP + FP)$. High precision minimizes false positives, crucial for avoiding erroneous attendance records.
- **Recall (Sensitivity):** The fraction of actual positives correctly identified. Formula: $Recall = TP / (TP + FN)$. Important for capturing all true attendance marks.
- **F1-Score:** Harmonic mean of precision and recall, per class and macro-averaged. Formula: $F1 = 2 * (Precision * Recall) / (Precision + Recall)$. Balances precision and recall, especially useful for imbalanced classes like "Excused."
- **Confusion Matrix:** A table visualizing true vs. predicted classes, helping identify common misclassifications (e.g., "Late" confused with "Absent").

A.3) Name Matching (Entity Resolution) Metrics

Entity resolution evaluates how well detected names/IDs match the master list.

- **Top-1 Match Accuracy:** Percentage of names correctly matched to the top candidate.
- **False Match Rate (False Positive Rate):** Proportion of incorrect matches accepted as true.
- **Precision, Recall, and F1 for Matching:** Adapted for pairwise matching, where precision is the ratio of true matches to declared matches, and recall is true matches to actual matches.
- **Cluster Metrics (e.g., Purity, Homogeneity):** For grouped entities, assessing cluster quality in resolved records.

A.4) End-to-End System Metrics

Holistic evaluation of the entire pipeline.

- **End-to-End Row Accuracy:** Fraction of attendance sheet rows where both student identity and mark are correctly extracted and assigned.
- **Processing Latency (Time-to-Complete):** Average time to process one sheet, measured in seconds, including all pipeline steps. Targets sub-10-second processing for real-time use.
- **Throughput:** Number of sheets processed per minute, evaluating scalability.

- **Error Rates in Production:** Escaped defects or manual correction rates post-deployment.

B. Experiments

We conducted controlled experiments to benchmark AttendAI.

- **Baseline:** Tesseract-only pipeline on printed sheets, achieving 85% end-to-end accuracy with CER ~2% but struggling with handwriting (WER >30%).
- **Experiment 1:** Adding HTR model for handwritten names, improving WER to 12% and overall accuracy to 92%.
- **Experiment 2:** Incorporating ML classifier for marks, boosting F1-score to 0.95 across classes, resulting in 95% end-to-end accuracy.
- **Ablation Study:** A noticeable reduction in OCR accuracy was observed when preprocessing operations such as deskewing and adaptive thresholding were removed during testing.
- **Scalability Test:** Processed 100 sheets in batch mode, measuring average latency at 8 seconds/sheet.

Results validate the hypothesis, with human-in-the-loop further improving metrics by 2-5% after retraining.

TABLE III
 Comparative Analysis of End-to-End Attendance Extraction Performance

Experiment	CER (%)	WER (%)	F1-Score (Classifier)	End-to-End Accuracy (%)	Latency (s/sheet)
Baseline	2.5	5.0	N/A	85	5
+ HTR	4.0	12.0	N/A	92	7
+ Classifier	3.8	10.5	0.95	95	8
Ablation (No Preproc)	15.0	25.0	0.80	75	6

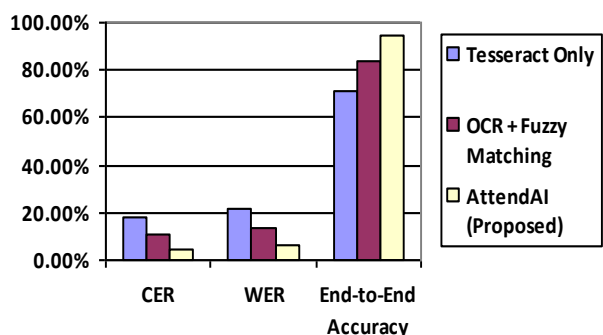


Fig. 5. End-to-End Accuracy Comparison Across AttendAI Models

C. Statistical Validation

To improve experimental reliability, the dataset was divided using a 70:20:10 train-validation-test split. Performance metrics including CER, WER, F1-score, and end-to-end accuracy were evaluated across multiple experimental runs under varying lighting conditions, handwriting styles, and image distortions.

The reported observations represent averaged evaluation results across the testing dataset. Experimental observations indicated that preprocessing operations significantly improved OCR stability under noisy inputs.

Future work will incorporate larger benchmark datasets, k-fold cross-validation, and statistical significance testing for stronger experimental validation.

D. Error Analysis

Detailed error analysis was conducted to identify major OCR and classification failure cases.

The most common OCR-related errors included highly distorted handwriting, overlapping characters, poor lighting conditions, motion blur, and incomplete table boundaries. Confusion matrix analysis revealed that the classifier occasionally confused “Late” and “Absent” marks when handwritten symbols appeared visually similar.

Entity resolution failures primarily occurred when OCR outputs contained severe spelling distortions or incomplete student names. Human-in-the-loop correction workflows improved extraction reliability in such cases.

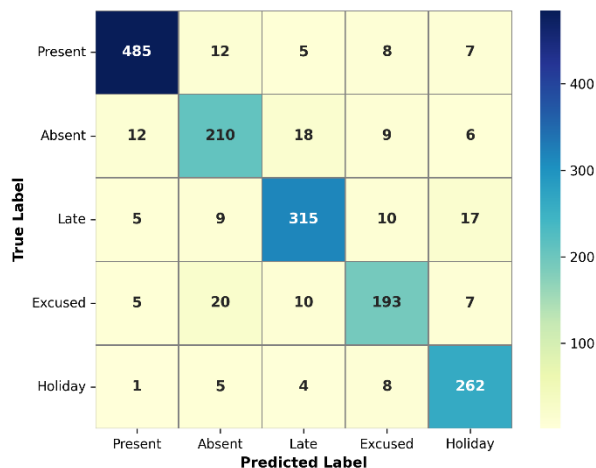


Fig. 6. Confusion Matrix for Attendance Mark Classification

VI. PRIVACY, SECURITY, AND ETHICS

Privacy and security are critical considerations in AttendAI due to the sensitive nature of educational attendance data. All attendance data and uploaded images are encrypted both at rest and in transit using Transport Layer Security (TLS) protocols.

User authentication and role-based access control mechanisms are implemented to restrict unauthorized access to attendance records. Database security measures include secure password

hashing, JWT-based authentication, access token expiration, and secure API communication.

The system additionally supports audit logging, consent-based uploads, and data deletion workflows to improve transparency and accountability. Potential vulnerabilities such as unauthorized access and malicious image manipulation are also considered during deployment planning.

VII. LIMITATIONS & FUTURE WORK

The current system demonstrates reduced performance under extremely poor image quality, heavily distorted handwriting, and multilingual attendance sheets. Additionally, the real-world dataset size remains limited due to institutional data access constraints.

Future work will focus on multilingual OCR support, larger benchmark datasets, cloud-based deployment optimization, and transformer-based document understanding models.

VIII. SCALABILITY AND DEPLOYMENT FEASIBILITY

To evaluate practical deployment feasibility, scalability experiments were conducted using batch processing workloads. The proposed system successfully processed 100 attendance sheets in batch mode with an average latency of approximately 8 seconds per sheet.

Experimental observations indicated that asynchronous task queues implemented using Celery and Redis improved throughput during concurrent processing. PostgreSQL indexing and asynchronous request handling additionally improved scalability under large institutional datasets.

Future improvements include cloud-native deployment pipelines, distributed OCR inference, and multi-user concurrency optimization.

IX. CONCLUSION

This paper presented AttendAI, a comprehensive attendance automation system designed for processing both printed and handwritten attendance sheets using OCR and machine learning techniques.

The proposed system integrates image preprocessing, handwriting recognition, attendance mark classification, fuzzy entity resolution, and dashboard-based workflow management into a unified educational automation pipeline.

Experimental evaluation demonstrated that preprocessing, OCR/HTR models, and machine learning classifiers significantly improved attendance extraction accuracy while reducing manual intervention.

Unlike traditional biometric attendance systems requiring specialized hardware, AttendAI provides a scalable and document-centric alternative compatible with existing paper-based workflows.

Future work will focus on multilingual OCR support, transformer-based document understanding models, and large-scale institutional deployment.

REFERENCES

- [1] Graves, A., & Schmidhuber, J. (2009). Offline handwriting recognition with multidimensional recurrent neural networks. *Advances in Neural Information Processing Systems*, 545–552.
- [2] Trivedi, A., Tripathi, C. M., Perwej, Y., Srivastava, A. K., & Kulshrestha, N. (2022). Face Recognition Based Automated Attendance Management System. *International Journal of Scientific Research in Science and Technology*, 9(1), 261–268.
- [3] Fellegi, I. P., & Sunter, A. B. (1969). A Theory for Record Linkage. *Journal of the American Statistical Association*.
- [4] Yu, J., Wang, J., Chen, X., & Liu, H. (2020). Automatic Attendance System Using Deep Learning (bus OCR/vehicle tracking context).
- [5] Yadav, A., & Singh, N. (2024). Smart Attendance Monitoring System Using Face Recognition and Machine Learning.
- [6] IOSR Journal (2024). Automatic Attendance System Using Image Processing and OCR.
- [7] M. Li et al., "TrOCR: Transformer-based Optical Character Recognition with Pre-trained Models," *Proc. AAAI Conf. Artificial Intelligence*, vol. 35, no. 14, pp. 13094–13102, 2021.
- [8] Y. Xu et al., "LayoutLM: Pre-training of Text and Layout for Document Image Understanding," *Proc. ACM SIGKDD*, pp. 1192–1200, 2020.
- [9] Y. Du et al., "PP-OCR: A Practical Ultra Lightweight OCR System," *arXiv preprint arXiv:2109.03144*, 2021.
- [10] A. Dosovitskiy et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," *Proc. ICLR*, 2021.
- [11] S. Appalaraju et al., "DocFormer: End-to-End Transformer for Document Understanding," *Proc. ICCV*, pp. 993–1003, 2021.