Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

Digital KYC System for Verification Using OCR and Machine Learning

Bhushan Patil[1], Laxmi Konakeri[2], Laxmi Kalyani[3], Naveen Pattar[4], Dr. Pritam Dhumale[5]

[1,2,3,4] Student, Dept. of Computer Science, Jain College of Engineering and Research Udyambhag, Belagavi, Karnataka, India
[5] Project Guide; Professor and HOD, Dept. of Computer Science, Jain College of Engineering and Research, Udyambhag, Belagavi, Karnataka, India

DOI: https://doi.org/10.63363/aijfr.2026.v07i06.8713

Abstract

Know Your Customer (KYC) is an essential identity verification process used by banks, financial institutions, insurers, telecommunications operators, and digital service providers. Conventional KYC relies on manual document inspection and data entry, which increases processing time, operational effort, and the risk of human error. This paper presents a Digital KYC system that integrates Optical Character Recognition (OCR) and machine learning to automate the major stages of customer identity verification. Authenticated users upload identity documents such as Aadhaar and PAN cards; OCR extracts the text, and data-processing rules convert the raw output into structured fields such as name, date of birth, and identification number. The extracted fields are validated through document specific patterns, format checking, normalization, and consistency verification, and a face-detection stage prepares the document photograph and the user's selfie for biometric comparison. The application layer is implemented with a FastAPI backend in Python, SQLA lchemy, and MySQL, and is secured with JWT authentication, bcrypt password hashing, role-based access control, rate limiting, secure file handling, notifications, and audit functionality. A working web interface demonstrates the complete workflow, including image-quality assessment and OCR confidence reporting. Dataset-level accuracy has not yet been measured and is identified as the next experimental phase. The proposed approach reduces manual KYC effort and provides a modular, scalable foundation for digital identity verification.

Keywords: Digital KYC, face detection, face verification, FastAPI, identity verification, machine learning, optical character recognition.

1. Introduction

Know Your Customer (KYC) is a fundamental procedure for establishing and verifying the identity of customers before providing regulated or identity-sensitive services. Financial institutions use KYC to collect customer information, verify identity evidence, and assess the risks associated with a customer relationship. As services move from physical branches to online platforms, identity verification must become faster, more accessible, and more reliable. Digital identity guidelines describe identity proofing

as the collection and evaluation of evidence to establish an individual's identity at an appropriate assurance level; the latest NIST guidelines, SP 80063A-4, define requirements for identity proofing and enrollment and stress security, privacy, fraud prevention, and user experience [1].

Traditional KYC requires customers to submit identity documents and personal information, after which an employee manually examines the documents and enters the data into the organization's database. Such procedures become expensive and time-consuming when thousands of applications are processed. Research on digital KYC identifies cost, inefficiency, repeated verification, and customer inconvenience as the main motivations for automation [2].

Optical Character Recognition (OCR) converts text in document images into machine-readable information. Tesseract, for example, is a widely studied OCR engine that includes text-line detection, classification, and adaptive recognition techniques [5]. Deep-learning-based OCR has further improved identity-document processing: Suddul and Seguin separated text detection from text recognition using a UNet segmentation model and a CRNN-LSTM recognizer, and reported 98% pixel accuracy for text detection and approximately 97% mean word recognition accuracy on the main fields of Mauritian identity cards [3].

Extracting information from a document, however, does not prove that the person submitting it is its owner. Face verification supplies an additional identity signal by comparing the photograph on the document with a live or newly captured selfie. DeepFace reported 97.35% accuracy on the Labeled Faces in the Wild (LFW) dataset [6], and FaceNet, which maps faces into a compact Euclidean embedding space, reported 99.63% on LFW under its evaluation setup [7]. Face detection and alignment, an important earlier stage of the pipeline, can be performed jointly by Multi-task Cascaded Convolutional Networks (MTCNN) [8]. Studies focusing specifically on identity-document-to-selfie matching show that this task is harder than conventional face recognition and that state-of-theart matchers can underperform on it unless document-specific augmentation or fine-tuning is applied [9], [10].

A practical Digital KYC system should therefore not treat OCR and face recognition as isolated technologies. Document acquisition, information extraction, validation, biometric verification, authentication, authorization, secure storage, and final decision-making must be integrated into one controlled workflow. This paper addresses that requirement through a modular architecture. The implemented backend provides user authentication, role-based access control, Aadhaar/PAN document upload, selfie upload, notifications, database management, token revocation, file validation, and administrative functions, and the OCR and machine-learning modules are integrated with it through the document-processing workflow.

The main contributions of this work are:

- A modular Digital KYC architecture combining document processing and biometric verification.
- OCR-based extraction of information from identity documents, with structured extraction and validation of key KYC fields.
- A face-detection stage for the document photograph and the submitted selfie.
- A secure backend using FastAPI, JWT, bcrypt, and rolebased access control (RBAC), with secure handling of Aadhaar, PAN, and selfie uploads.
- Database support for users, documents, notifications, authentication logs, and token management.

- A working web interface and a framework for evaluating OCR and biometric performance independently.

2. RELATED WORK

A. Digital KYC and Digital Identity

Schlatt *et al.* studied digital KYC based on blockchain-based self-sovereign identity (SSI), identified the traditional process as costly, inefficient, and inconvenient, and proposed a framework that treats privacy and data protection as design requirements [2]. Their work shows that digital KYC should be designed as a complete process rather than a simple document upload application; however, the blockchain/SSI approach concentrates on identity management and interoperability rather than OCR-based extraction and face matching. NIST SP 800-63A-4 provides the broader technical foundation by defining identity-assurance requirements and emphasizing the reliability of identity evidence, security, privacy, and fraud resistance [1]. The proposed system applies these concepts at the application level by collecting identity documents, extracting information, performing biometric comparison, and controlling access to sensitive data.

B. OCR-Based Identity Document Extraction

Smith described the Tesseract engine and its text-line finding, feature/classification, and adaptive-classifier techniques, which provide a practical foundation for converting document images into text [5]. Suddul and Seguin used U-Net for text-field localization and a CRNN with LSTM cells for recognition on identity cards [3]. Li *et al.* proposed TrOCR, an end-to-end OCR system built from a pre-trained image transformer and text transformer, and showed that it can outperform earlier approaches on printed, handwritten, and scene-text tasks [4]. These studies trace the evolution of OCR from traditional engines to deep-learning and transformer architectures. For KYC, however, OCR accuracy alone is insufficient: extracted values must also be validated against document-specific rules.

C. Face Detection and Recognition

Taigman *et al.* introduced DeepFace, which combines face alignment and representation learning and reached 97.35% accuracy on LFW [6]. Schroff *et al.* proposed FaceNet, in which the distance between embeddings measures face similarity; it reported 99.63% on LFW and 95.12% on the YouTube Faces Database [7]. Zhang *et al.* proposed MTCNN, a three-stage cascade that progressively identifies facial regions and landmarks [8]. These methods form the theoretical basis of the proposed pipeline: MTCNN can perform detection and alignment, while an embedding model such as FaceNet can compare the document face with the selfie.

D. Document-to-Selfie Verification

Albiero *et al.* found that matching identity-document faces to live images is challenging, particularly when the document photograph and the current face differ significantly in age, and showed that specialized fine-tuning outperforms generic matchers [9]. Moussa *et al.* reported that document-specific, face-swapping-based data augmentation improved verification accuracy by 7.96% in their setup and also examined demographic performance [10]. Fatima *et al.* evaluated commercial remote identity-verification systems with 3,991 test subjects and found that performance and equity vary across systems and

demographic groups, so biometric KYC must be assessed for false rejection and demographic consistency as well as average accuracy [11].

TABLE I
SUMMARY OF RELATED WORK

Work	Task	Reported result
Smith [5]	OCR engine (Tesseract)	Text-line finding and adaptive classification
Suddul & Seguin [3]	ID-card text extraction	98% pixel accuracy; ~97% mean word accuracy
Li *et al.* [4]	Transformer OCR (TrOCR)	Outperforms earlier OCR on printed, handwritten, scene text
Taigman *et al.* [6]	Face verification (DeepFace)	97.35% on LFW
Schroff *et al.* [7]	Face embedding (FaceNet)	99.63% on LFW; 95.12% on YouTube Faces
Zhang *et al.* [8]	Detection/alignment (MTCNN)	Joint detection and landmark localization
Albiero *et al.* [9]	ID-to-selfie matching	Fine-tuning improves over generic matchers
Moussa *et al.* [10]	ID-to-selfie verification	+7.96% accuracy with document-specific augmentation
Fatima *et al.* [11]	Commercial remote ID verification	Performance varies across systems and demographics (3,991 subjects)

E. Research Gap

The reviewed literature shows strong progress in individual technologies: digital-identity research addresses identity proofing and privacy, OCR research addresses automated text extraction, face-recognition research addresses biometric comparison, and remote identity-verification research addresses document-to-selfie matching. A practical Digital KYC system, however, must combine these technologies with secure application services. This work therefore focuses on an integrated pipeline: Document Upload → OCR → Data Extraction → Data Validation → Face Detection → KYC Decision, supported by a backend that provides the required security and application infrastructure.

3. PROPOSED METHODOLOGY

The proposed methodology consists of six stages, described below.

A. Document Upload

The user first registers and authenticates with the system. After successful authentication, the user uploads a valid identity document. The current backend supports Aadhaar and PAN documents in JPG, JPEG, PNG, and PDF formats, and applies file-size and file-type validation before storing the file (Fig. 1). A newly uploaded document is assigned an initial processing status such as *pending*.

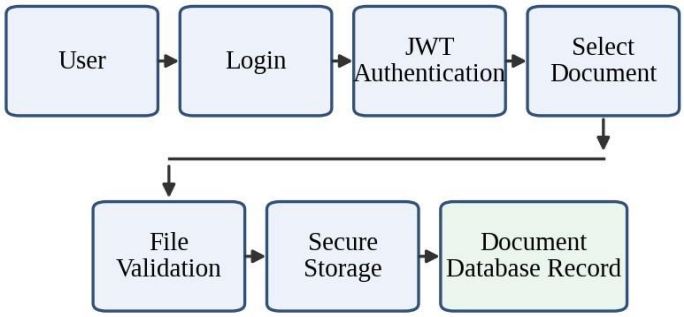


Fig. 1. Document upload workflow. *B. OCR Processing*

The uploaded document is passed to the OCR component, which performs image acquisition, image pre-processing, text region identification, character recognition, and raw text generation. Pre-processing can include resizing, grayscale conversion, noise reduction, thresholding, and image enhancement. The OCR output is then sent to the data extraction stage.

C. Data Extraction

The raw OCR output is converted into structured information. For Aadhaar and PAN documents, the system identifies fields such as name, date of birth, identification number, and other relevant fields present on the document, using regular expressions, keyword-based rules, and document specific patterns:

$$\textit{Structured Data} = \textit{Extract}\ (\textit{Raw OCR Text}) \qquad (1)$$

The result is a set of structured attributes rather than an unorganized block of text.

D. Data Validation

Extracted data is validated before it is used for KYC decision-making (Fig. 2). Validation includes required-field verification, format checking, identification-number pattern validation, date validation, character normalization, comparison with the information submitted by the user, and detection of missing fields. For example, a PAN value must conform to its expected alphanumeric structure, and an Aadhaar number must conform to the expected numerical format.

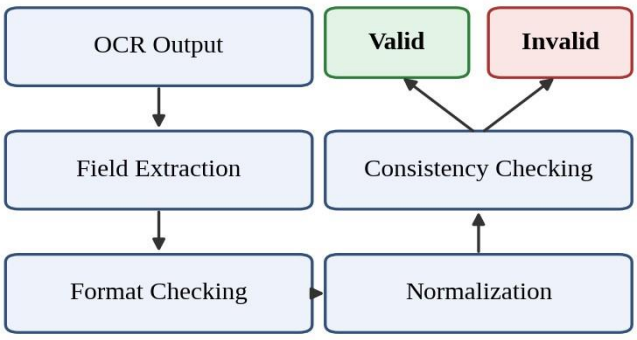


Fig. 2. Data validation pipeline.

E. Face Detection

The biometric stage begins after the document has been processed. The system identifies the face in the photograph on the identity document and in the user's submitted selfie. A face detection model such as MTCNN can be used to detect facial regions and landmarks, and the detected faces are aligned before feature extraction.

F. KYC Decision

The final decision combines document validation and face verification. Let $V \in \{0, 1\}$ indicate whether the document passed validation, C the OCR confidence, S the face-similarity score, τ the similarity threshold, δ a review margin, and C_{min} the minimum acceptable OCR confidence. A simplified decision rule is:

$$KYC = \text{Approved, if } V = 1 \ \wedge C \geq C_{min} \ \wedge S \geq \tau + \delta \quad (2)$$

$$KYC = \text{Rejected, if } S < \tau - \delta \quad (3)$$

$$KYC = \text{Manual Review, otherwise} \quad (4)$$

Manual review is therefore triggered when OCR confidence is low, the document contains inconsistent information, or the face-similarity score is close to the selected threshold. The values of τ, δ, and C_{min} are configurable and must be tuned on a representative evaluation dataset.

4. SYSTEM ARCHITECTURE

The proposed system follows a layered, modular architecture (Fig. 3). The AI components are deliberately separated from the backend so that the OCR engine or face verification model can be replaced without redesigning the application.

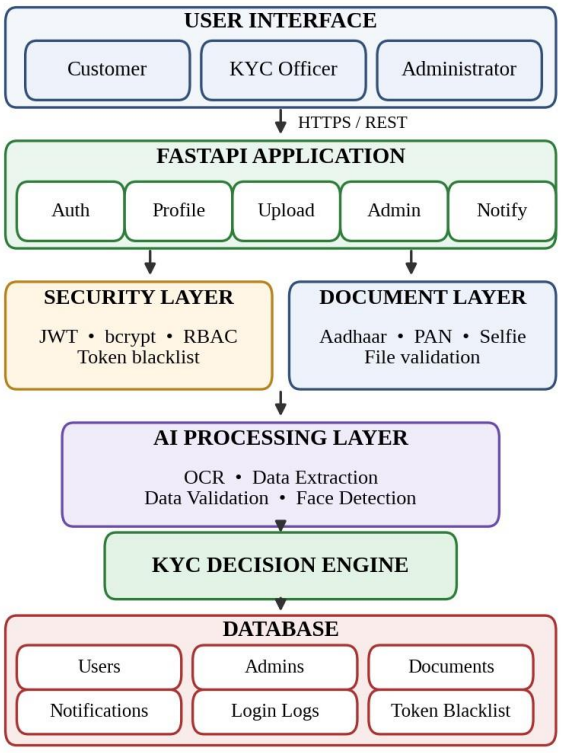


Fig. 3. Overall system architecture of the Digital KYC system.

A. Presentation Layer

The presentation layer is the interface through which customers, KYC officers, and administrators interact with the application. Customers can register, log in, upload documents and selfies, and receive KYC-related notifications.

B. Backend/API Layer

FastAPI provides RESTful endpoints for registration, login, token refresh, logout, profile management, document upload, selfie upload, notifications, administrative operations, and health monitoring.

C. Security Layer

The security layer comprises JWT authentication, bcrypt password hashing, role-based access control, refresh-token blacklisting, file validation, and exception handling.

D. AI Processing Layer

The AI layer contains OCR, data extraction, data validation, and face detection, followed by the KYC decision engine.

E. Database Layer

The backend defines database models for administrators, users, documents, login logs, notifications, password-reset tokens, and the token blacklist. The system is designed for MySQL using SQLAl chemy and aiomysql, while the test environment can use SQLite.

5. IMPLEMENTATION

A. Development Environment

The backend is implemented in Python using the FastAPI framework. Table II lists the technology stack.

TABLE II
TECHNOLOGY STACK

Component	Technology
Programming language	Python 3.11+
Web framework	FastAPI
ORM	SQLAlchemy 2.x
Database	MySQL (SQLite for testing)
Authentication	JWT
Password hashing	bcrypt
Data validation	Pydantic v2
OCR	OCR processing component
Face verification	ML face-verification component

B. Authentication

Authentication uses access and refresh JWT tokens with the lifetimes shown in Table III. Passwords are hashed using bcrypt and are never stored as plaintext. On logout, refresh tokens are added to a blacklist so that they cannot be reused.

TABLE III
BACKEND CONFIGURATION

Parameter	Value
Access-token lifetime	30 minutes
Refresh-token lifetime	7 days
Maximum upload size (default)	10 MB
Accepted file types	JPG, JPEG, PNG, PDF
Automated API test cases	27

C. Role-Based Access Control

The system defines three roles: Customer, KYC Officer, and Administrator. Customers manage their own profiles and upload KYC documents, while administrative operations are protected by role-based authorization.

D. Document Management

The backend provides one endpoint for Aadhaar/PAN document upload and another for selfie upload. The file service performs extension validation, file-size validation, filename sanitization, UUID-based filename generation, directory creation, file storage, and database record creation, following recognized file-upload security practice [12].

E. Notification System

An in-application notification mechanism lets users retrieve notifications, check their status, and mark them as read. It informs users when a document is received, processed, approved, rejected, or requires further action.

F. Testing

Client contains 27 test cases. They cover user registration, duplicate registration, login, invalid login, token validation, token refresh, logout, profile operations, Aadhaar upload, PAN

6. RESULTS AND DISCUSSION

A. Functional and Security Results

The implemented backend provides the infrastructure required for the Digital KYC workflow. Its security mechanisms—JWT access and refresh tokens, bcrypt hashing, RBAC, token revocation, rate limiting, and validated, sanitized file storage—together form a stronger foundation than a simple document-upload application.

B. Web Application Walkthrough

The system was implemented as a web-based application for processing and verifying KYC documents. The user accesses the application, selects and uploads a document such as Aadhaar or PAN, and proceeds through document processing and verification. The document is checked for image quality and processed by OCR to extract name, date of birth, address, and identification details; the extracted information supports document classification and identity verification, and machine-learning-based analysis can support fraud detection. The system finally displays the KYC result and stores the relevant information

for record keeping. Figs. 4–10 present the screens of the implemented application in the order in which a user encounters them, and each is explained below.

1) Home page. The home page (Fig. 4) introduces the system and provides entry points to start verification or sign in. A top bar carries the VerifyAI brand with *Sign in* and *Get Started* buttons, and the central panel presents the headline, a one-line description of the service (document verification, face matching, and fraud detection in one flow), and two call-to-action buttons: *Start Free Verification* for new users and *Sign in to Dashboard* for returning users. The four tiles at the bottom (verification count, success rate, average time, and uptime) are static indicative values shown on the landing page; they are not measurements from this study and must not be read as evaluated performance.



Fig. 4. Home page of the Digital KYC application.

2) Registration. The registration page (Fig. 5) collects the full name, email address, phone number, password, and password confirmation, and offers a *Sign in* option for existing users. Registration is the first step of the workflow because every later operation, such as document upload, is performed by an authenticated user: after the account is created and the user logs in, the backend issues the JWT access and refresh tokens described in Section V-B. The password field includes a show/hide toggle, and the confirmation field prevents mistyped passwords from being stored.

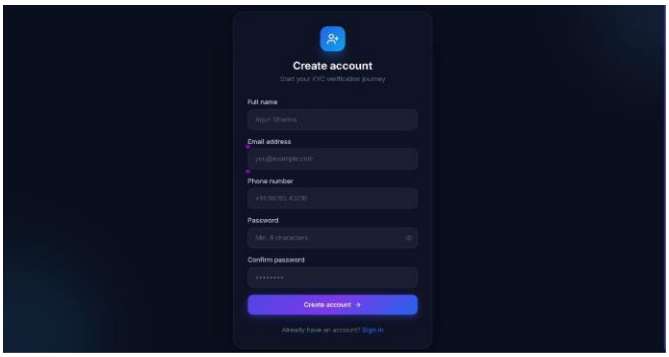


Fig. 5. User registration page.

3) User dashboard. After login, the user dashboard (Fig. 6) summarizes the total number of verifications and their status (successful, suspicious, or failed), offers quick actions such as *Start KYC* and *Update Profile*, shows notifications about submitted documents, and lists recent verifications with verification ID, date, document type, score, and status. The side menu gives access to Dashboard, New KYC,

Notifications, Profile, and Help & FAQ, and the account card at the bottom shows the signed-in user with a logout control. In the sample shown, the user has five verifications (three successful, one suspicious, one failed). The notification panel reports a successful Aadhaar verification with a score of 96.3% and a PAN card that has been received and is being processed, which corresponds to the notification events of Section V-E. In the recent-verifications table, the Aadhaar record of 21 Jul 2026 (96.3%) and the PAN record of 15 Jun 2026 (93.8%) are marked *Verified*, whereas the Aadhaar record of 03 May 2026 (72.1%) is marked *Suspicious*. The *Suspicious* status plays the same role as the Manual Review outcome of (4): a score that is neither high enough for approval nor low enough for rejection. These records illustrate the interface and are sample data.

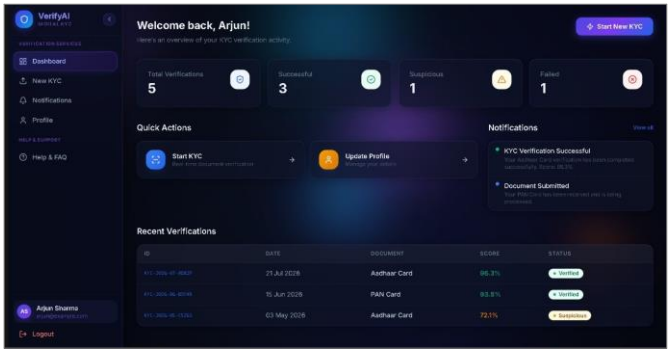


Fig. 6. User dashboard showing verification summary and recent verifications.

4) Document selection and progress indicator. The document selection page (Fig. 7) lets the user choose Aadhaar Card (12-digit unique ID issued by UIDAI) or PAN Card
(Permanent Account Number issued by the Income Tax Department). A progress indicator at the top shows the nine stages of the process: Upload, Quality Check, Preview, OCR Data, Verification, Fraud Check, Face Match, Risk Score, and Result. Completed stages are shown as ticks and the current stage is highlighted, so the user always knows how far the verification has progressed. The *Continue* button appears dimmed until a document type has been chosen, which prevents the user from entering the upload stage without a selection.

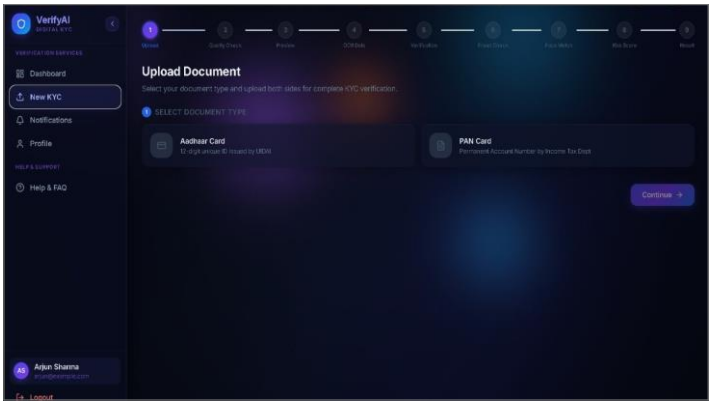


Fig. 7. KYC document selection and progress indicator.

5) *Aadhaar document upload.* For Aadhaar (Fig. 8), separate sections accept the front side (name, date of birth, gender, photograph, Aadhaar number) and back side (address, pincode, UIDAI QR code), either by file upload or by scan, with instructions to keep the images clear, flat, and fully visible. Both sides are marked *Required*, so that identity details and address information are collected for complete verification. Each section provides an *Upload File* tab and a *Scan* tab, and a drag-and-drop area with an upload limit of 10 MB, matching the default in Table III. The selected document type is confirmed with a tick in step 1 of the page, and the guidance banner reminds the user that image clarity at this stage directly affects the quality check and OCR that follow.

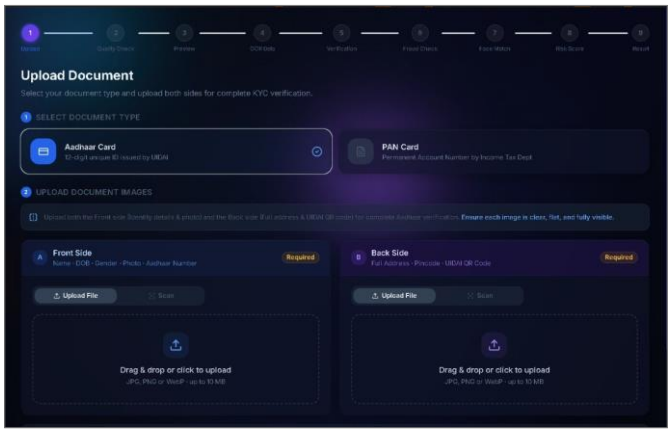


Fig. 8. Aadhaar document upload page (front and back).

6) *Image-quality assessment.* The image-quality assessment page (Fig. 9) reports an overall quality score and individual checks for blur, brightness, resolution, contrast, glare, and color balance. For the sample shown, the overall quality was 76% and the image passed. The individual scores were blur detection 42%, brightness 100%, resolution 90%, contrast 58%, glare detection 99%, and color balance 78%, for an image of 720×1280 pixels. Each check is drawn as a coloured bar: green for a satisfactory value, amber for a borderline value, and red with a warning icon for a weak value. In this sample, blur and contrast were flagged even though the overall score passed, which shows that the pass decision is based on the aggregate score and that individual weaknesses remain visible to the user. A blurred or low-contrast capture is a common cause of OCR errors, so this page gives the user the opportunity to go *Back* and re-capture the document before *Continue* leads to the OCR stage.

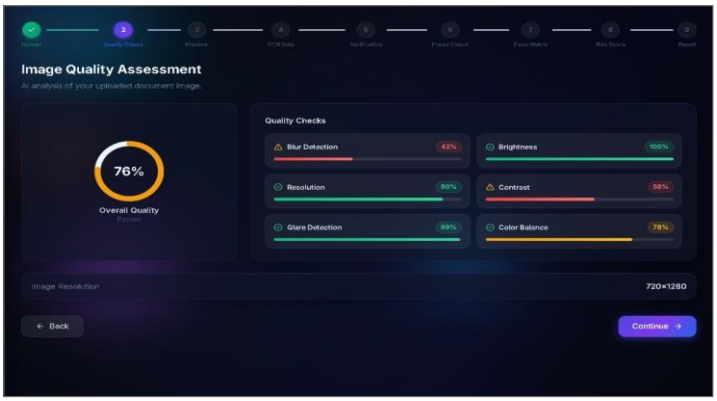


Fig. 9. Image-quality assessment page.

7) OCR extracted data. The OCR extracted-data page (Fig. 10) reports an OCR confidence of 98% (attributed on the page to the Tesseract engine [5]) and lists the extracted full name, date of birth, gender, Aadhaar number, and document type, together with a summary of the quality parameters, so that the user can review the data before the next stage. Each field is shown in its own labelled card, which reflects the structured output of the data-extraction step in (1) rather than a block of raw text, and the user can return to the previous stage or choose *Verify Document* to send the data to the validation and verification stages. The quality summary at the bottom of the page lists four of the parameters (blur, brightness, resolution, and contrast) with an overall value of 70% (pass), whereas Fig. 9 aggregates all six checks to 76%; the two aggregation rules should be harmonised in a future version. The remaining stages (Preview, Verification, Fraud Check, Face Match, Risk Score, and Result) are not shown in this paper.

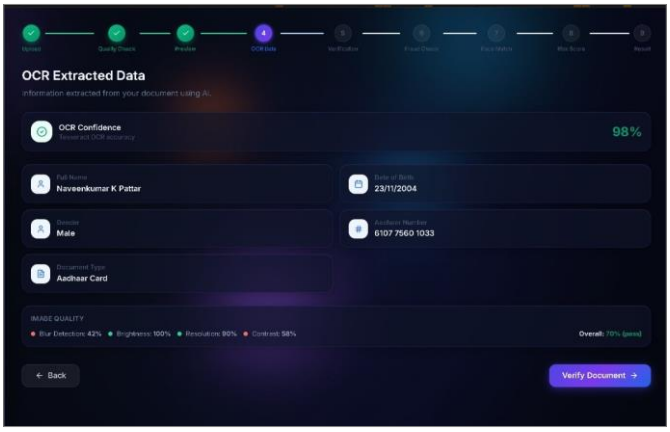


Fig. 10. OCR extracted-data page with confidence value.

The quality score and OCR confidence in Figs. 9 and 10 are outputs for a single uploaded sample. They illustrate the workflow and must not be read as dataset-level accuracy. Likewise, the summary values and records on the home page and dashboard (Figs. 4 and 6) are interface examples. Table IV summarizes the screens and the workflow stage that each one demonstrates.

TABLE IV
SUMMARY OF INTERFACE SCREENS

Fig.	Screen	Purpose / workflow stage
4	Home page	Entry point: start verification or sign in
5	Registration	Account creation before authenticated use
6	Dashboard	Status summary, notifications, history
7	Document selection	Choose Aadhaar or PAN; progress indicator
8	Aadhaar upload	Stage 1: front and back image upload
9	Quality assessment	Stage 2: blur, brightness, resolution, etc.
10	OCR extracted data	Stage 4: structured fields and confidence

C. OCR Evaluation Methodology

OCR performance should be evaluated at field level using

OCR Accuracy = $(N_{correct} / N_{total}) \times 100\%$ (5) where $N_{correct}$ is the number of correctly extracted fields and N_{total} is the total number of expected fields. A more detailed evaluation should also report the Character Error Rate, CER = $(S + D + I)/N$, where S, D, and I are the numbers of substituted, deleted, and inserted characters and N is the number of reference characters; the Word Error Rate (WER), computed analogously at word level; field-level accuracy; precision; recall; and processing time. No dataset-level OCR accuracy is claimed in this paper.

D. Discussion

The main strength of the proposed architecture is its modularity. Authentication, document management, AI processing, and database services are separated into independent layers, so an OCR engine or face-recognition model can be replaced without redesigning the backend. The web interface makes the process easy to follow by showing each verification stage clearly, and the quality-assessment and OCR-review pages let the user detect a poor capture or an extraction error before the data is used for verification.

The literature also indicates that document-to-selfie verification is harder than conventional face recognition [9], [10]. Accuracy values reported by FaceNet or DeepFace on general benchmarks [6], [7] should therefore not be interpreted as the expected accuracy of a KYC system: document photographs differ from selfies in image quality, pose, lighting, age, compression, and background. Likewise, OCR performance on general datasets cannot be transferred automatically to Aadhaar and PAN documents; identity-document OCR must be evaluated on the actual document types and image conditions the system encounters. A separate experimental phase using a controlled dataset is required to measure OCR and biometric performance.

7. CONCLUSION

This paper presented a Digital KYC system for automated identity verification using OCR and machine learning. The proposed system integrates document upload, OCR-based text extraction, structured data extraction, validation, face detection, face matching, and KYC decision-making into a single workflow, supported by a secure FastAPI backend providing JWT authentication, bcrypt password hashing, role-based access control, secure document handling, rate limiting, token revocation, notifications, database management, and administrative functions. The methodology reduces dependence on manual document processing by extracting relevant information automatically, and face matching adds a biometric layer by comparing the document photograph with a submitted selfie.

The literature shows substantial progress in OCR, deep-learning face recognition, and digital identity, but also that real-world KYC presents challenges not captured by general benchmarks, in particular document-to-selfie matching and demographic fairness. The modular architecture allows different OCR and face-verification algorithms to be tested and compared. The current backend should be regarded as the secure application and integration foundation of the complete system; final claims about OCR accuracy, face-matching accuracy, and end-to-end performance should be made only after running the full AI pipeline on a representative evaluation dataset.

8. FUTURE SCOPE

Several improvements can be incorporated in future versions:

1. *Advanced OCR:* transformer-based models such as TrOCR [4] for difficult document images.
2. *Document layout understanding:* identify fields from spatial positions and relationships as well as text.

3. *Face liveness detection:* distinguish a real person from a photograph, video, or presentation attack.
4. *Document forgery detection:* detect forged or digitally manipulated identity documents.
5. *Improved document-to-selfie models:* fine-tune generic face-recognition models using document-to-selfie data and document-specific augmentation [9], [10].
6. *Fairness evaluation:* evaluate verification across age, gender, skin tone, and other demographic groups [11].
7. *Cloud deployment:* object storage, managed MySQL, containerization, load balancing, automated backups, and monitoring.
8. *Encryption:* protect identity documents and biometric data in transit and at rest.
9. *Fraud detection:* machine-learning models for duplicate documents, suspicious identities, manipulated documents, repeated failed attempts, and unusual application patterns.
10. *Privacy-preserving identity:* decentralized identity, verifiable credentials, or selective disclosure, which can reduce repeated verification and improve user control, provided privacy and scalability are carefully considered [2].

References

1. D. Temoshok, Y.-Y. Choong, R. Galluzzo, M. LaSalle, A. Regenscheid,C. Abruzzi, J. L. Fenton, and N. Lefkovitz, *Digital Identity Guidelines: Identity Proofing and Enrollment*, NIST Special Publication 800-63A-4, National Institute of Standards and Technology, 2025.
2. V. Schlatt, J. Sedlmeir, S. Feulner, and N. Urbach, "Designing aframework for digital KYC processes built on blockchain-based self-sovereign identity," *Information & Management*, 2022.
3. G. Suddul and J. F. L. Seguin, "A custom-built deep learning approach fortext extraction from identity card images," *Int. J. Informatics and Communication Technology*, vol. 13, no. 1, pp. 34–41, 2024, doi: 10.11591/ijict.v13i1.pp34-41.
4. M. Li, T. Lv, J. Chen, L. Cui, Y. Lu, D. Florencio, C. Zhang, Z. Li, and F.Wei, "TrOCR: Transformer-based optical character recognition with pre-trained models," in *Proc. AAAI Conf. Artificial Intelligence*, vol. 37, no. 11, 2023, pp. 13094–13102.
5. R. Smith, "An overview of the Tesseract OCR engine," in *Proc. 9th Int. Conf. Document Analysis and Recognition (ICDAR)*, 2007, pp. 629–633.
6. Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing thegap to human-level performance in face verification," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 1701–1708, doi: 10.1109/CVPR.2014.220.
7. F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unifiedembedding for face recognition and clustering," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 815–823, doi: 10.1109/CVPR.2015.7298682.
8. K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection andalignment using multitask cascaded convolutional networks," *IEEE Signal Process. Lett.*, vol. 23, no. 10, pp. 1499–1503, 2016, doi: 10.1109/LSP.2016.2603342.
9. V. Albiero, N. Srinivas, E. Villalobos, J. Perez-Facuse, R. Rosenthal, D.Mery, K. Ricanek, and K. W. Bowyer, "Identity document to selfie face matching across adolescence," arXiv:1912.10021, 2019.
10. E. M. Moussa, I. Sarridis, E. Krasanakis, N. Ramoly, S. Papadopoulos,A.-M. Awal, and L. Younes, "Face-swapping based data augmentation for ID document and selfie face verification," in *Proc.*

IEEE/CVF Winter Conf. Applications of Computer Vision Workshops (WACVW), 2025, pp. 1421–1428.

11. K. Fatima, M. Schuckers, G. Cruz-Ortiz, D. Hou, S. Purnapatra, T.Andrews, A. Neupane, B. Marshall, and S. Schuckers, "A large-scale study of performance and equity of commercial remote identity verification technologies across demographics," *IEEE Conference Publication*, 2024.
12. OWASP Foundation, "File Upload Cheat Sheet," *OWASP Cheat Sheet Series*, 2024.