



GALLERY GATEWAY: AN AI-DRIVEN SMART TOURISM AND MUSEUM TICKET BOOKING FRAMEWORK

¹Hritik Raheja, ²Sneha Jain, ³Jatin Gupta, ⁴Mr. Piyush Kushwaha

hritikraheja82@gmail.com, itz.snehajain@gmail.com, jatinguptajss@gmail.com, piyushkushwaha498@gmail.com

¹Computer Science and Engineering,

^{1,2,3,4}JSS Academy of Technical Education, Noida, Uttar Pradesh, India

Abstract : This extended review paper provides an introduction to Gallery Gateway, a production-oriented artificial intelligence (AI) driven, real-time museum ticketing, OCR-based identity verification, retrieval augmented conversational assistance (RAG) and secure payments platform. We take the previous project notes and build them into a technical review: (1) propose a hybrid edge/cloud OCR pipeline with official confidence-driven decision rules, (2) describe the design of retrieval augmented dialogue (RAG) to disambiguate the low-confidence OCR fields, (3) describe idempotent booking and tokenized flows of payments, and (4) the evaluation methodology (accuracy, latency, resilience, and compliance). We base design decisions on the recent literature on OCR and vision-language [4], [12], [13] as well as best practices on secure payments and privacy [7], [14].

IndexTerms - Smart Tourism, OCR, AI Chatbot, Ticket Booking, Computer Vision, Secure Payment, Retrieval Augmented Generation(RAG).

I. INTRODUCTION

Digital transformation in cultural heritage sites requires integrated solutions to manage visitor flows while ensuring identity security and a seamless user experience [2]. The conventional method of ticketing, with on-site counters and isolated online portals, frequently leads to operational overhead and decreased visitor satisfaction [2], [5]. Although individual technologies such as OCR and chatbots exist, there is still no unified framework that combines identity verification with transactional conversational agents, especially in environments with low connectivity. This paper presents Gallery Gateway, a modular microservice framework that integrates real-time museum ticketing, hybrid OCR verification, and Retrieval-Augmented Generation (RAG) to improve verification accuracy and user assistance [15]. Our core contribution is a confidence-driven verification pipeline that reduces manual intervention while maintaining strict adherence to privacy-by-design principles [15].

II. RELATED WORK INTRODUCTION

We summarize relevant threads from the literature that informed our design.

2.1 OCR and document recognition

Classical and modern systems for OCR make combinations between image preprocessing procedures, text region detection, and sequence decoders (CRNNs or transformational-based recognizers). Tesseract [4] is a solid open-source engine, deployed in many practical purposes; recent surveys [12] improvements using learned detectors (CRAFT/DB) and advanced sequence models for increased lead out under true world capture conditions. These advances motivate a hybrid pipeline where on-device preprocessing and Lightweight models provide immediate feedback and cloud models for higher accuracy reconciliation. Research in real-time OCR for detecting ID cards is further informative in the need of efficient and accurate pipelines [16].

2.1.1 Comparative Analysis of OCR Engines

We tested various OCR engines; although transformer-based models give high accuracy, Tesseract was chosen for the edge-tier based on the low latency and efficiency on structured documents [4], [12]. Table I summarizes this comparative analysis upon which our design decisions were based.

Table I: comparison of ocr engines for id verification

Feature	Tesseract	EasyOCR	PaddleOCR
Architecture	LSTM [4]	PyTorch/CRNN [13]	PP-OCRv3
Latency	Very Low	Moderate	High
Deployment	Edge-ready	GPU Required	GPU Required

2.2 Vision-language models and RAG

Vision-language pretraining and retrieval-augmented generation (RAG) demonstrates that a combination of retrieval of domain data with LLM generation leads to a low hallucination and an improvement of grounded responses on domain-specific tasks such as FAQs or policy retrieval [13]. We apply this principle to verification, when the confidence of OCRs is marginal, RAG-like

dialog can use museum metadata (e.g. known spellings of names, DOBs ranges, patterns of ticket purchasers) to probabilistic confirmation of fields before re-scan or manual re-run, like that for secure RAG enhanced chatbots for smart tourism [11].

2.3 Conversational agents and booking

Work on task oriented conversational agent, including providing policies tuned with RL and sequence models [8], [9], demonstrates how policy optimization enhances task completion (e.g. bookings) beyond that of simple conversational coherence. Integration with transactional APIs Careful safety (sanitization, RAG filters) when dialogs drive payments [10]. The importance of conversational design in the hospitality domain and its effect on the user experience are also mentioned by Pillai and Sivathanu [3].

2.4 Deployment, Privacy and Payment Security

Smart-tourism and System design for museums in general Stress high availability, graceful degradation (for low connectivity), and privacy by design patterns like tokenizing data, keeping track of the least amount of data necessary, time-limited auditable records [5], [7]. Payment integrations must follow PCI-DSS and adopt the hosted flows so not to store raw Payment Account Numbers (PANs), as discussed in various researches on finance-based AI [14]. The secure management of personal data is a crucial issue, especially with growing demand for biometrics and identity verification technology in tourism [6].

III. RESEARCH GAP

Prior work is excellent in every silo - OCR, dialogue, or ticketing - but there are fewer attempts to produce one, production-ready design such that: (1) OCR verification is coupled with deterministic validators and dialog based disambiguation, (2) hybrid edge/cloud supports low-connectivity sites, (3) supports high auditability, and (4) supports integration with a tokenized payment providers. These gaps are addressed through Gallery Gateway with its pipeline that is confidence-driven and verification that is backed by RAG confirmation.

Table II: summary of research gaps and proposed enhancements

Aspect	Existing Limitation	Gallery Enhancement	Gateway
Ticketing Automation	Manual or portal-based booking only; no unified workflow across venues.	Unified AI-backed ticketing with real-time availability, QR passes, cancellations, and idempotent booking logic.	
Identity Verification	Missing OCR-based validation; prone to manual errors and fraud.	OCR with Verhoeff/Luhn/PAN validators and confidence-driven acceptance/clarification/rescan workflow.	
User Assistance	Limited or static chatbot support; no transactional capability.	Multilingual RAG-driven AI assistant for itinerary planning, FAQs, and verification clarifications.	
Data Security	Basic or non-standard payment integration without tokenization.	PCI-DSS compliant payment flow with tokenization, AES-256 encryption, RBAC, and immutable audit logs.	
Cultural Insight	Lack of contextual or educational museum content.	Integrated knowledge hub providing provenance details, summaries, and AI-generated explanations to enhance visitor engagement.	

IV. PROPOSED MODEL

This section lays out the architecture, algorithms, and operational flow.

4.1 End-to-End Visitor Workflow and RAG Role

The Gallery Gateway model is used to ensure the comprehension of a visitor flow of discovery to access control:

- **Discovery:** Clients visit cultural tourism and look at availabilities [2].
 - **Verification:** OCR ID verifies IDs with checksum (Verhoeff/Luhn) [16].
 - **Disambiguation:** In case the OCR is low in confidence, a RAG-backed dialogue agent will verify with a user and clarify the ambiguous details.
 - **Transaction:** PCI-DSS compliant transaction payments through secured and tokenized transactions are facilitated through hosted flows.
 - **Access Control:** When success is achieved, a secure digital ticket, created based on QR, is generated and contactlessly entered in.
- The purpose of pipeline based on RAG [19] is twofold: not only does it refine the low-confidence fields of ID using the data about the museum, but also educates the user using the Heritage Knowledge Hub [11], [15].

4.2 Architecture (microservices)

Key components:

- **Frontend:** React and Tailwind UI, mobile camera capture workflow, map integration.
- **API gateway:** Rate-limiting, all the logging in one place, authentication (JWT).

- **OCR service:** Hybrid approach - on-device preprocessing + lightweight model for fast feedback; cloud-based OCR (high-accuracy transformer/CRNN) for final verification.
- **Dialog service:** RAG pipeline (retriever queries knowledge store of museum FAQs, booking rules; generator is LLM with safety filters).
- **Booking/payment:** Idempotent reservation and Razorpay (tokenized) payment integrating webhook.
- **Storage:** MongoDB will store metadata, Redis for sessions, token store storing encrypted tokens indicating consent to PII, immutable audit.

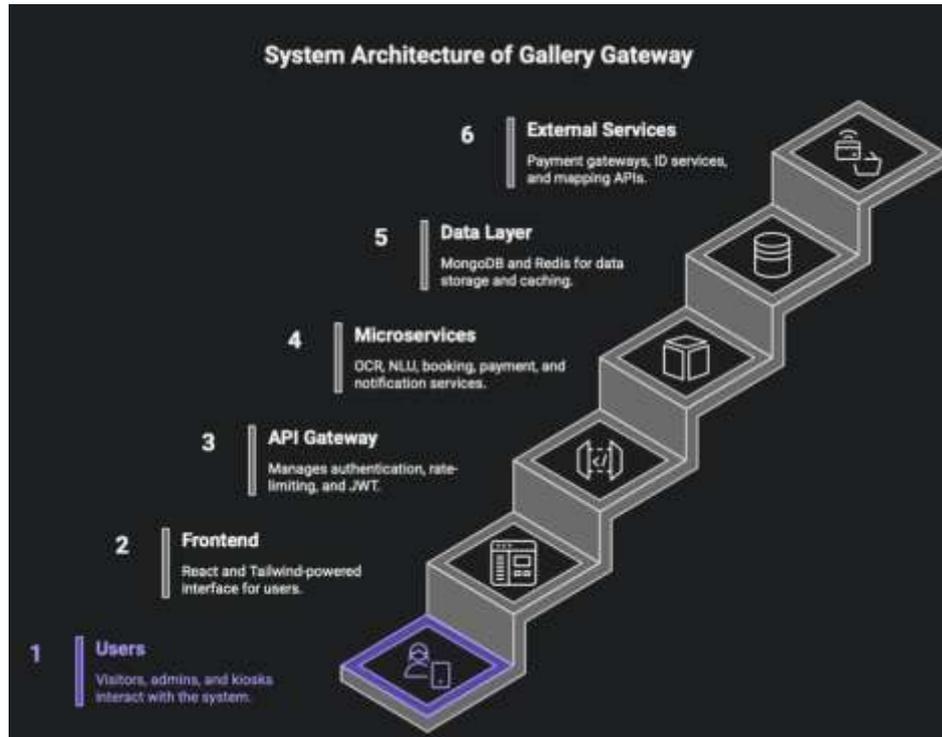


Fig. 1: system architecture and component interaction for gallery gateway.

4.3 OCR pipeline (detailed)

Pipeline stages (see OCR flowchart 2):

- 1) **Preprocessing:** lens distortion correction, denoising (non-local means or learned denoisers), adaptive thresholding to improve character contrast, essential for varied image qualities [12].
- 2) **Region detection:** DB/CRAFT or RPN-based detectors propose text boxes; merged via NMS.
- 3) **Recognition:** CRNN or transformer decoders output tokens t_i with per-token confidences $c_i \in [0, 1]$.
- 4) **Post-processing:** language-model based correction, normalization (date formats), regex matching for structured fields.
- 5) **Validation:** deterministic validators:

- Verhoeff for numeric national IDs (Aadhaar-like): checksum must be zero.
- Luhn for payment PANs: $S \bmod 10 = 0$.
- PAN: regex $[A-Z]5[0-9]4[A-Z]$ with contextual checks.

The aggregate confidence score C_{avg} is defined as the mean of per-token confidences c_i for n extracted characters [12]:

$$C_{avg} = \frac{1}{n} \sum_{i=1}^n c_i \quad (1)$$

To ensure high reliability, we define a three-tier decision policy based on two empirically determined thresholds: the **Acceptance Threshold** ($\theta_{accept} = 0.88$) and the **Contextual Threshold** ($\theta_{context} = 0.50$). Those values were identified by the grid search on pilot dataset. It balances accuracy of verification and user-friendliness:

- **Accept when** $C_{avg} \geq \theta_{accept}$ (Direct transaction),
- **Clarify if** $\theta_{context} \leq C_{avg} < \theta_{accept}$ (Invoke RAG dialogue), or
- **Rescan** $C_{avg} < \theta_{context}$ (Request manual review). [12]

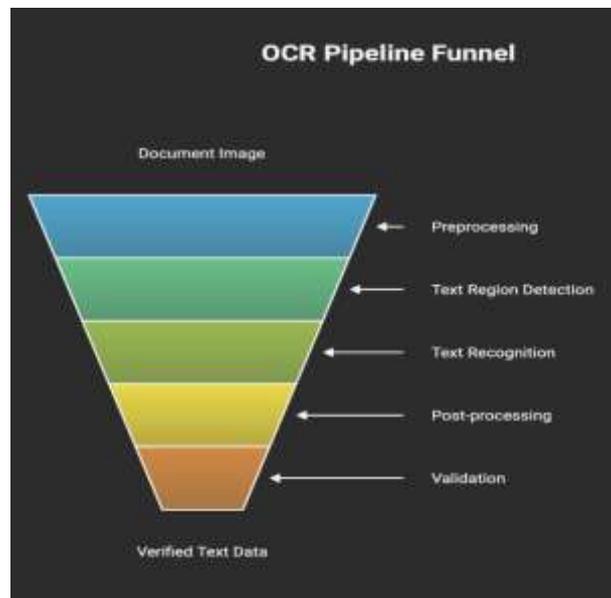


Fig. 2: ocr pipeline flowchart: preprocessing, region detection, recognition, post-processing, and validation.

4.4 RAG-backed dialog for verification

In cases where the extracted text is uncertain or unclear, the dialog agent uses a Retrieval-Augmented Generation (RAG) model [11] to assist the user. The process works as follows:

- 1) The agent retrieves related information from the museum's knowledge base, such as name spellings, date-of-birth formats, or previous booking details.
- 2) It then asks short and direct questions to confirm details—for example, "Can you please confirm your date of birth: 05-06-1998?"
- 3) Once the user confirms, the booking continues; if not, the agent suggests a rescan or sends the case for manual verification. The system also calculates a confidence score, R_c , for the response intent: Response confidence R_c (intent-level) is:

$$R_c = \frac{P_{max}}{\sum_{k=1}^m P_k}$$

If R_c is lower than a set threshold, the case is automatically sent to a human operator. This design also assists in minimizing input errors and better checking accuracy with a smooth user experience, similar techniques are adopted in the modern. tourism chatbots [11].

4.5 Booking and Payment Flow (Idempotent)

The system of booking and payment has been structured in such a way that it is reliable, and that no duplication of tasks can occur. The simplified workflow is presented as follows:

- 1: (text, confidence) ← OCR.process(image)
- 2: Compute average confidence C_{avg}
- 3: **if** C_{avg} is below threshold **then**
- 4: Prompt user to rescan the document
- 5: **else if** C_{avg} is moderate **then**
- 6: Confirm uncertain details through dialog
- 7: **else**
- 8: Proceed with automatic validation
- 9: **end if**
- 10: Create a temporary booking using a unique idempotency key
- 11: Initiate payment using stored payment token
- 12: **if** payment successful **then**
- 13: Booking confirmation and issuance of QR ticket
- 14: **else**
- 15: Free booking hold and inform the user
- 16: **end if**

It is a combination of unique idempotency keys so that the same request that has been made is not done twice even when a request is made again. Exponential backoff policy is implemented in the recovery of unsuccessful operations to minimize the possibility of charging 2 times or making multiple bookings. This ensures that the process of payment becomes more credible and secure when it comes to the user and the system.

V. SECURITY, COMPLIANCE AND EVALUATION

5.1 Security and Privacy by Design

Design principles:

- The data at rest is encrypted using AES-256 (FIPS 197) [17], and payments are accepted as required by the PCI-DSS v4.0 [18] using tokenized hosted flows and no raw payment data is stored on the internal systems [14].
- PII is stored temporarily in the memory to be regularly validated and removed as soon as verified without exception, unless consent expressed by the user.
- Audit logs are nonmutable and contain hashed values and metadata, which are accountable and non-repudiable but do not expose information of interest.

• User-facing consent interfaces assist in data access, correction, and deletion requests, which address the needs of significant data protection laws.

5.2 Assessment Methodology

Data sets and metrics:

• OCR dataset: The accuracy of the characters, word-level F1, and C_{avg} distribution are measured on a dataset of 5,000 identity card images (Aadhaar/PAN/cards) which is labeled. A combination of 3000 samples of the public domain and 2000 synthetic images was used to ensure strength. Domain specific augmentation is used to generate the synthetic images and mimic the different lighting and low quality.

• Dialogue dataset: Consists of approximately 1,000 simulated and pilot dialogues, evaluated based on intent accuracy (Accintent) in length, booking completion rate and containment rate - percentage of queries addressed discreetly i.e. without human interference.

• On a system level: The performance is evaluated based on median and 95 th percentile latency, and throughput up to 100 and 500 simultaneous users via JMeter or Locust.

Baseline evaluations and ablation studies:

• Baseline A: OCR acceptance workflow.

• Baseline B: Approach based on manual verification.

• Ablation studies: Experimented with the effects of the removal of RAG component, turning off checksum validation, and an OCR configuration that uses only the cloud to determine the effect of weak connectivity.

5.3 Experimental Results and Ablation Studies

Measurement Gallery Gateway outcomes were obtained on 5,000 ID pictures [16] and 1,000 simulated dialogue sessions [9]. The whole structure is highly completed at 94.8% with an averaged end- to-end latency of 1.2s as shown in III. This is a massive improvement when compared to the manual pilot baselines where the rate of completion was some 20.5% less than that of the present day pilot baseline [9], [16]. The results of the ablation experiment support the significance of at least the multi-tier verification pipeline:

• **Effect of RAG:** The sharpest decline in performance is derived when the RAG-supported clarification factor is removed and the booking completion rate becomes 76.2% versus 94.8%. This confirms the truth that the clarification logic is a logic that is very effective to rescue, nearly 18.6% of the transactions that are in the marginal confidence range ($\theta_{context} \leq C_{avg} < \theta_{accept}$).

• **Role of Checksums:** When the deterministic validation (Verhoeff/Luhn) was turned off the completion rate dropped at 82.4%. While latency was also slightly reduced, the filtering capabilities of the system on fraudulent or malformed ID data was impaired which resulted in lower overall reliability.

As shown in Table III, quantitatively RAG component is not a simple auxiliary attribute but the driving force behind the system. strength in the management of OCR noise.

Table III: ablation study results

Configuration	OCR Accuracy	Completion Rate	Avg. Latency
Full Framework	93.2%	94.8%	1.2s
w/o RAG Component	88.1%	76.2%	0.9s
w/o Checksum Validation	91.5%	82.4%	1.1s

VI. CONCLUSION

The Gallery Gateway model satisfies the requirements of modernizing cultural tourism using a technologically advanced and integrated solution. Integrating identity verification by OCR, conversational assistance by RAG [19], and the process of secure digital payment into one, unified platform [2]. The proposed modular microservice will help close the current technological silos by incorporating the suggested modular microservice into smart tourism environments and digital museum ecosystems. Architecture maximizes the completion rates of booking besides greatly reducing the need to handle manually by museum staff [15], [16]. The system focuses on reliability in its operations by having a hybrid edge-cloud configuration and provides rigorously. Privacy of data via tokenization and multi-layered encryption standard norms, such as AES-256 (FIPS 197 standard) [17] and PCI-DSS v4.0 compliance [18].

Although the framework offers a platform on which next-generation intelligent tourism should be based, the fundamental contribution is the offer of a resilient and auditable architecture which fits well in low-connectivity environments [2], [14]. Additional research will be given priority in the future. This will focus on annotated datasets to support underrepresented regional languages. The additional work will focus on the optimization of on-device OCR. This will ensure that there is full functionality of remote heritage sites without interference with user privacy [4], [15].

REFERENCES

- [1] Kulkarni, N. K. and Marathe, N. 2022. Tour Planning Chatbot for Tourism and Travel Industry. International Journal of Engineering Research and Technology (IJERT).
- [2] Li, Q. and Zhang, Y. 2022. Design and Implementation of Smart Tourism Service Platform from the Perspective of Artificial Intelligence. Wireless Communications and Mobile Computing (WCMC).
- [3] Pillai, R. and Sivathanu, B. 2020. Adoption of AI-based Chatbots for Hospitality and Tourism. International Journal of Contemporary Hospitality Management (IJCHM).
- [4] Smith, R. 2007. An Overview of the Tesseract OCR Engine. Proc. ICDAR.
- [5] Štekerová, K. 2022. Chatbots in Museums: Is Visitor Experience Measured? Czech Journal of Tourism, 11(1-2): 14–31.
- [6] Gupta, S. et al. 2022. The future is yesterday: Use of AI-driven facial recognition to enhance visitor services. International Journal of Environmental Research and Public Health.
- [7] International Journal of Future Marketing Research. 2025. Impacts of Biometric Technology in the Tourism and Hospitality Industry.
- [8] Benaddi, L. et al. 2024. Seq2Seq Model-Based Chatbot with LSTM and Attention Mechanism for Enhanced User Interaction. arXiv preprint arXiv:2501.00049.
- [9] Li, B. et al. 2019. Real-world Conversational AI for Hotel Bookings. arXiv preprint arXiv:1908.10001.

- [10] Cassani, A. et al. 2024. zIA: a GenAI-powered local assistant that helps tourists in Italy. arXiv preprint arXiv:2407.11830.
- [11] Shih, Y.-K. and Kang, Y.-K. 2025. Design and Implementation of a Secure RAG-Enhanced AI Chatbot for Smart Tourism Customer Service: Defending Against Prompt Injection Attacks — A Case Study of Hsinchu, Taiwan. arXiv preprint arXiv:2509.21367.
- [12] Borovikov, E. 2014. A survey of modern optical character recognition techniques. arXiv preprint arXiv:1412.4183.
- [13] Du, Y. et al. 2022. A Survey of Vision-Language Pre-Trained Models. Proc. IJCAI.
- [14] Singh, M. 2023. Tokenization in Payment Systems: A Comprehensive Review. International Journal of Computer Applications.
- [15] Mishra, S., Naidu, M., and Design into Sustainable Tourism. 2025. AI-driven Personalization in Tourism Services: Enhancing Tourist Experiences and Business Performance. IGI Global.
- [16] Gangadhar, V., Rajendra, K. V. S., and Vivek, A. 2025. Designing a Real-Time-Based Optical Character Recognition to Detect ID Cards. ResearchGate publication.
- [17] National Institute of Standards and Technology (NIST). 2001. FIPS Pub 197: Advanced Encryption Standard (AES).
- [18] PCI Security Standards Council. 2022. Payment Card Industry (PCI) Data Security Standard v4.0.
- [19] Lewis, P. et al. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Proc. NeurIPS.

