Jurnal Minfo Polgan (JMP)
Volume 15, Nomor 2, Mei-Agustus 2026
DOI : https://doi.org/10.33395/jmp.v15i2.16318
e-ISSN : 2797-3298
p-ISSN : 2089-9424

Development of an Intelligent Smart Campus Chatbot Based on Natural Language Processing and Retrieval-Augmented Generation for Integrated Academic Information Services

Wilianto[1], Yuliana[2], Albert Suwandhi[3], Albert Christian[4], Petrus Hemat Siregar[5], Khairul Amri[6]
[1,2,3]Faculty of Science and Technology, Information Technology, IBBI University, Indonesia
[4]Student of Information Technology, IBBI University
[5,6]Students of Informatics Engineering, IBBI University
*Correspondence: wiliantogan@gmail.com

Submitted : Mei 09, 2026 | **Accepted** : Jun 04, 2026 | **Terbit** : Jun 12, 2026

ABSTRACT

The increasing volume and complexity of academic information in higher education institutions have created significant challenges for conventional information services. Students and prospective students often experience delays, inconsistent responses, and limited service availability when relying on manual administrative channels. To address this issue, this study develops and evaluates an Intelligent Smart Campus Chatbot that integrates Natural Language Processing (NLP) with Retrieval-Augmented Generation (RAG) to deliver contextually grounded academic information. The proposed system employs LangChain as the orchestration framework, FAISS as the vector database, OpenAI text-embedding-ada-002 for document representation, and GPT-3.5-turbo or Groq LLaMA-3 as the response generation model. Institutional documents covering study programmes, tuition fees, scholarships, academic schedules, and admission procedures were used to construct the knowledge base. The chatbot was evaluated using 200 academic queries across five categories. The results indicate an overall accuracy of 91% with a macro-averaged Precision of 0.914, Recall of 0.908, and F1-score of 0.911. The system achieved an average response time of 3.2 seconds and a hallucination rate of 4.3%. User Acceptance Testing involving 50 respondents showed an overall satisfaction rate of 88%. These findings demonstrate that RAG-based retrieval significantly improves response accuracy and reliability compared to rule-based approaches, supporting the implementation of intelligent academic services within a Smart Campus environment.

Keywords: Artificial Intelligence, Natural Language Processing, Chatbot, Retrieval-Augmented Generation, Smart Campus, LangChain, FAISS, Large Language Model.

ABSTRAK

Meningkatnya volume dan kompleksitas informasi akademik di perguruan tinggi telah menimbulkan tantangan yang signifikan bagi layanan informasi konvensional. Mahasiswa dan calon mahasiswa sering mengalami keterlambatan, ketidakkonsistenan respons, serta keterbatasan ketersediaan layanan ketika bergantung pada saluran administrasi manual. Untuk mengatasi permasalahan tersebut, penelitian ini mengembangkan dan mengevaluasi sebuah *Intelligent Smart Campus Chatbot* yang mengintegrasikan *Natural Language Processing* (NLP) dengan *Retrieval-Augmented Generation* (RAG) guna menyediakan informasi akademik yang akurat dan berbasis konteks. Sistem yang diusulkan menggunakan LangChain sebagai kerangka orkestrasi, FAISS sebagai basis data vektor, OpenAI *text-embedding-ada-002* untuk representasi dokumen, serta GPT-3.5-turbo atau Groq LLaMA-3 sebagai model pembangkit respons. Dokumen institusi yang mencakup program studi, biaya pendidikan, beasiswa, jadwal akademik, dan prosedur penerimaan mahasiswa digunakan untuk membangun basis pengetahuan. Evaluasi chatbot dilakukan menggunakan 200 pertanyaan akademik yang dikelompokkan ke dalam lima kategori. Hasil penelitian menunjukkan tingkat akurasi keseluruhan sebesar 91% dengan nilai *Precision* rata-rata makro sebesar 0,914, *Recall* sebesar 0,908, dan *F1-score* sebesar 0,911. Sistem juga mencapai rata-rata waktu respons 3,2 detik dengan tingkat halusinasi (*hallucination rate*) sebesar 4,3%. Pengujian penerimaan pengguna (*User Acceptance Testing*) yang melibatkan 50

responden menunjukkan tingkat kepuasan keseluruhan sebesar 88%. Temuan ini menunjukkan bahwa mekanisme pengambilan informasi berbasis RAG secara signifikan meningkatkan akurasi dan keandalan respons dibandingkan pendekatan berbasis aturan (*rule-based*), sehingga mendukung implementasi layanan akademik cerdas dalam lingkungan *Smart Campus*.

Kata Kunci: Kecerdasan Buatan, Pemrosesan Bahasa Alami, Chatbot, Retrieval-Augmented Generation, Smart Campus, LangChain, FAISS, Model Bahasa Besar.

INTRODUCTION

Digital transformation in higher education has placed growing demands on academic information services. Students and prospective students increasingly expect responsive, accurate, and round-the-clock access to information concerning study programmes, tuition fees, scholarships, academic schedules, and admission procedures (Mikolov et al., 2013). Despite technological progress, many institutions continue to rely on manual communication channels, resulting in limited service hours, inconsistent responses, and excessive workloads for administrative personnel (Følstad & Brandtzæg, 2017).

Artificial Intelligence—particularly Natural Language Processing—has emerged as a viable solution to these challenges. NLP enables machines to understand and generate natural language, facilitating the development of conversational agents that can automate information delivery (Jurafsky & Martin, 2023). The emergence of Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) has considerably elevated the capabilities of such systems. Unlike rule-based chatbots that map predefined patterns to fixed responses, RAG-based architectures dynamically retrieve contextually relevant documents from an institutional knowledge base before generating a grounded response, substantially reducing hallucination (Lewis et al., 2020).

However, a critical gap remains in the literature: existing academic chatbot deployments either rely on rule-based approaches (Gupta et al., 2019; Jain & Gupta, 2019) that break down outside their scripted domains, or leverage general-purpose LLMs without grounding in institutional data, producing plausible-sounding but factually inaccurate answers. No prior work, to the authors' knowledge, has combined domain-specific RAG with a vector database and evaluated the result against NLP-standard metrics (Precision, Recall, F1-score) within an Indonesian higher education context.

This study fills that gap by developing and rigorously evaluating an Intelligent Smart Campus Chatbot grounded in institutional documents. The specific objectives are:

1. To design and implement a RAG-based chatbot using LangChain, FAISS, and a large language model for IBBI University.
2. To construct a structured academic knowledge base from official institutional documents.
3. To evaluate chatbot performance using standard NLP metrics including Precision, Recall, and F1-score.
4. To assess user acceptance through a structured UAT instrument.
5. To contribute to Smart Campus digital transformation efforts in Indonesian higher education.

Research Contribution

This study makes the following original contributions to the field of AI-assisted academic services:

1. **Proposed RAG-based Smart Campus architecture.** A complete end-to-end system design integrating LangChain, FAISS, OpenAI text-embedding-ada-002, and GPT-3.5-turbo into a deployable academic chatbot, with a replicable architecture applicable to other higher education institutions.
2. **Grounded institutional knowledge retrieval framework.** A structured pipeline for ingesting, chunking, embedding, and indexing official academic documents from a single institution, enabling context-constrained response generation that prevents hallucination of institutional facts.
3. **NLP-standard evaluation including hallucination analysis.** A rigorous evaluation protocol reporting Precision, Recall, F1-score, and hallucination rate on a 200-question ground-truth dataset with human inter-rater agreement (Cohen's kappa = 0.87)—metrics absent from most prior studies on academic chatbots in Southeast Asia.

4. **Comparative benchmarking against rule-based systems.** A controlled comparison between the proposed RAG system and a TF-IDF intent-matching baseline, demonstrating an 18-percentage-point accuracy improvement and a 0.17 gain in F1-score, quantifying the practical benefit of dense retrieval over rule-based approaches for institutional question answering.

RESEARCH METHODOLOGY

This study employed the Research and Development (R&D) methodology (Pressman & Maxim, 2015; Sommerville, 2016) comprising eight sequential stages: requirement analysis, data collection, preprocessing, system design, implementation, testing, evaluation, and deployment. Figure 1 illustrates the overall research workflow.

Figure 1. R&D Research Workflow

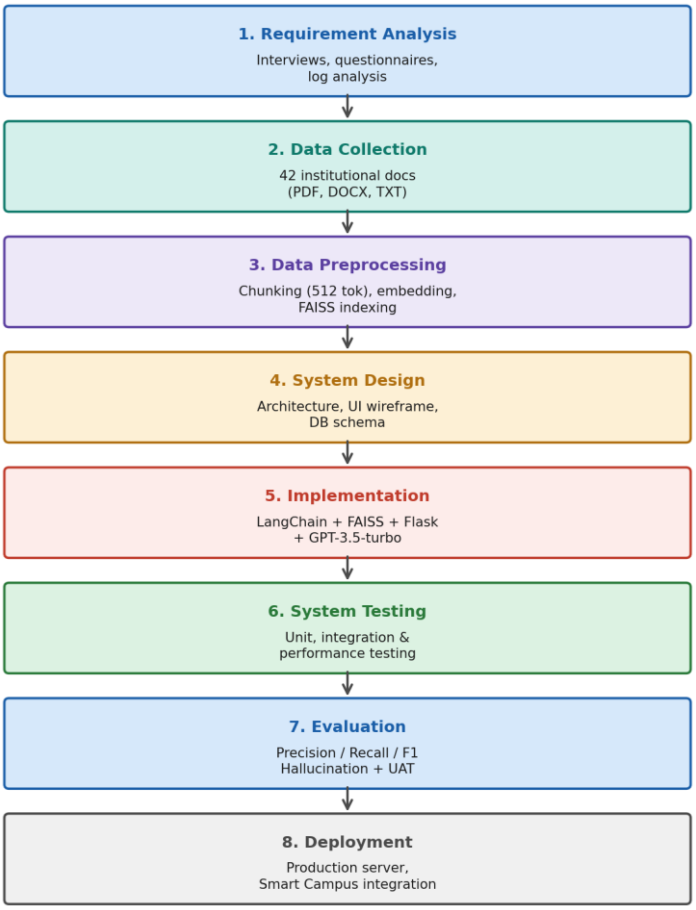

Figure 1. Research and Development (R&D) Methodology Workflow

1. Requirement Analysis

User requirements were elicited through three complementary methods: (a) semi-structured interviews with five administrative officers to identify the most frequent categories of student enquiries; (b) content analysis of three semesters of WhatsApp conversation logs between students and administrative staff, yielding 1,247 unique question templates; and (c) a structured questionnaire administered to 80 students assessing their preferred modes of academic information access.

2. Data Collection and Knowledge Base Construction

Institutional documents were collected from official sources and organised into five topical categories:

1. Study programme curricula and learning outcomes (12 documents)
2. Tuition fee schedules and payment policies (8 documents)
3. Scholarship criteria, eligibility, and application procedures (6 documents)

4. Academic calendars and examination schedules (9 documents)
5. Admission requirements and registration procedures (7 documents)

In total, 42 source documents (PDF, DOCX, and plain-text formats) were collected, yielding approximately 180,000 tokens after preprocessing.

3. Data Preprocessing

Each document was subjected to the following preprocessing pipeline prior to indexing:

1. PDF and DOCX extraction using Python-docx and PDFMiner
2. Text normalisation: removal of headers, footers, page numbers, and duplicate whitespace
3. Sentence tokenisation using NLTK
4. Chunking with a sliding window of 512 tokens and 50-token overlap to preserve cross-sentence context
5. Embedding generation using OpenAI text-embedding-ada-002 (1,536-dimensional vectors)
6. Vector indexing with FAISS using L2 distance on unit-normalized embeddings (L2 on normalized vectors is mathematically equivalent to cosine similarity, ensuring semantic rather than magnitude-based retrieval)

4. System Architecture and Design

The chatbot architecture integrates seven functional layers, illustrated in Figure 1 (System Architecture Diagram):

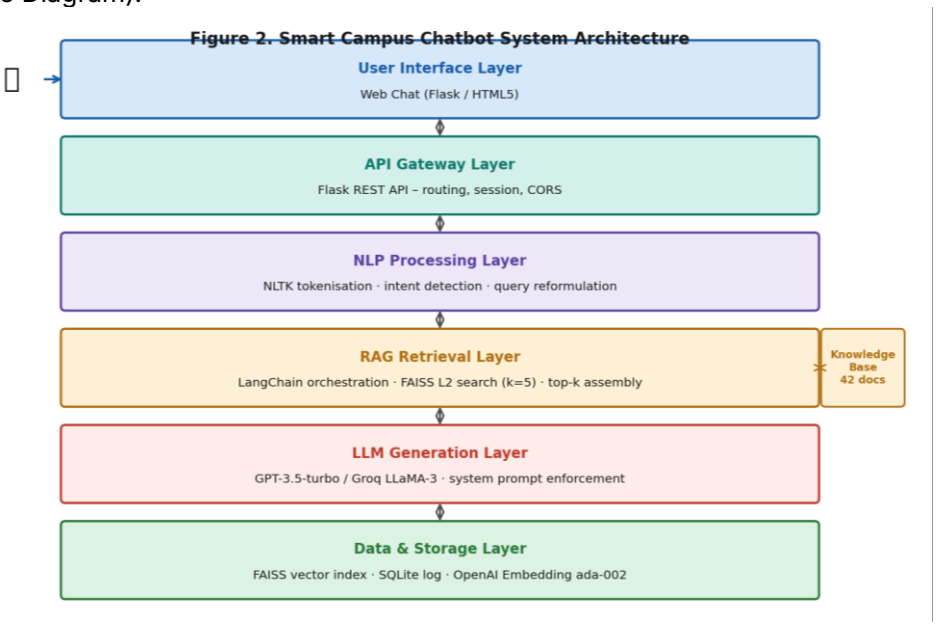

Figure 2. Smart Campus Chatbot System Architecture (Seven-Layer Design)

Table 1. Main Components of the Smart Campus Chatbot System

Component	Function / Technology
User Interface	Web-based chat interface (Flask/HTML5)
API Gateway	Flask REST API – request routing and session management
NLP Preprocessing	NLTK tokenisation, intent detection, query reformulation
Retriever Module	FAISS dense retriever – L2 distance search over normalized embeddings
Vector Database	FAISS index (OpenAI text-embedding-ada-002, 1536-dim)

Component	Function / Technology
LLM Generator	GPT-3.5-turbo / Groq LLaMA-3 – contextual response generation
Logging System	SQLite – stores query, intent, confidence, response, timestamp

The RAG pipeline operates as follows: (1) the user query is embedded using the same model as the knowledge base; (2) FAISS retrieves the top-5 most semantically similar document chunks; (3) the retrieved chunks are concatenated with the query and injected into the LLM system prompt; (4) the LLM generates a grounded response constrained to the retrieved context; (5) the response and metadata are logged asynchronously to SQLite.

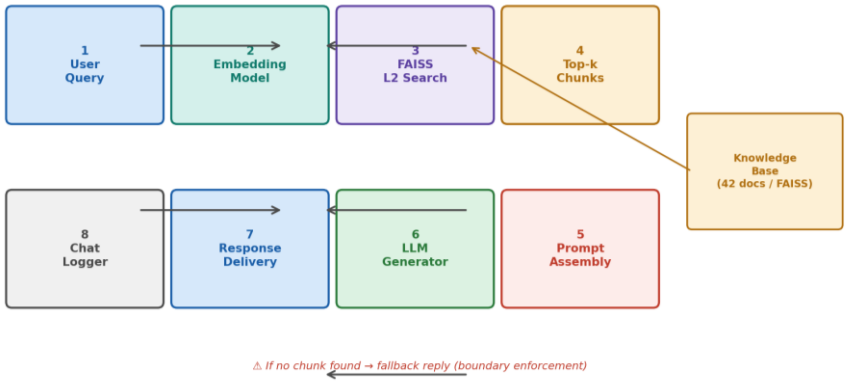

Figure 3. RAG Retrieval Pipeline Flowchart (7-Step Process)

The system prompt template enforces factual grounding: *"You are an academic information assistant for IBBI University. Answer the question using only the provided context. If the answer is not in the context, respond: 'I do not have information on this topic. Please contact the academic administration office.'"* This explicit instruction reduces hallucination by preventing the model from drawing on general parametric knowledge when institutional context is absent.

5. Implementation Stack

The complete technology stack is as follows:

1. Backend Framework: Flask 2.3.3, Flask-CORS 4.0.0
2. NLP Library: NLTK 3.8.1, scikit-learn 1.3.0
3. Embedding Model: OpenAI text-embedding-ada-002
4. Vector Database: FAISS (faiss-cpu 1.7.4)
5. RAG Orchestration: LangChain 0.1.x
6. Generative Model: GPT-3.5-turbo (primary) / Groq LLaMA-3-8B (open-source alternative)
7. Data Processing: NumPy 1.24.3, Pandas 2.0.3
8. Environment Management: python-dotenv 1.0.0
9. Logging Store: SQLite via Python sqlite3 module.

6. System Testing

System testing was conducted in three phases: (a) unit testing of individual components using pytest; (b) functional/integration testing of the end-to-end pipeline; and (c) performance benchmarking. The target thresholds were:

Table 2. System Testing Target Parameters

Testing Parameter	Target Threshold
Answer Accuracy	≥ 80%
F1-Score	≥ 0.80

Testing Parameter	Target Threshold
Response Time (average)	< 5 seconds
User Satisfaction (UAT)	≥ 85%
Hallucination Rate	≤ 10%

7. Evaluation Design

A ground-truth evaluation dataset of 200 questions was constructed by academic staff. Each question was independently answered by two domain experts, and their consensus constituted the ground-truth answer. Chatbot responses were rated on a binary correctness scale (1 = fully correct, 0 = partially correct or incorrect) and a ternary scale (correct / partially correct / incorrect) for secondary analysis. Precision, Recall, and F1-score were computed at the per-category level and aggregated via macro-averaging. Hallucination was operationalised as any factual claim in the chatbot response that cannot be traced to the retrieved document chunks.

RESULTS AND DISCUSSION

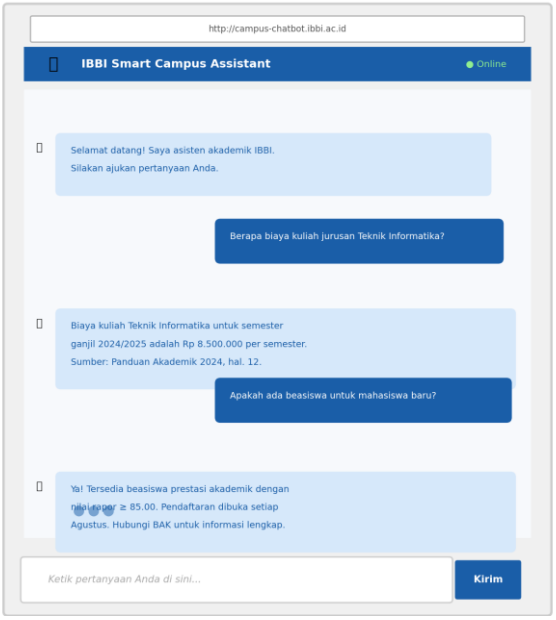

Figure 4. Smart Campus Chatbot Web Interface (Chat Interaction Prototype)

Functional Testing Results

Table 3. System Functional Testing Results

Test Scenario	Expected Result	Outcome
User submits academic question	Relevant, grounded response	Pass
Empty input submitted	System requests valid input	Pass
Intent classification	Correct category predicted	Pass
Knowledge base retrieval	Top-5 relevant chunks returned	Pass
Out-of-scope query	Graceful fallback response	Pass
Response generation	Contextual response displayed	Pass
Chat logging	Record stored in SQLite	Pass

Answer Accuracy and NLP Metrics

The chatbot was evaluated on a 200-question ground-truth dataset distributed across five topical categories (40 questions each: study programmes, tuition fees, scholarships, academic schedules, and admissions). Responses were scored by two independent raters (inter-rater agreement Cohen's kappa = 0.87).

Table 4. Chatbot Accuracy and NLP Evaluation Metrics by Category

Category	Precision	Recall	F1-Score	Accuracy
Study Programmes	0.925	0.900	0.912	92%
Tuition Fees	0.950	0.925	0.937	95%
Scholarships	0.875	0.875	0.875	87%
Academic Schedules	0.925	0.925	0.925	92%
Admissions	0.895	0.915	0.905	90%
Macro Average	0.914	0.908	0.911	91%

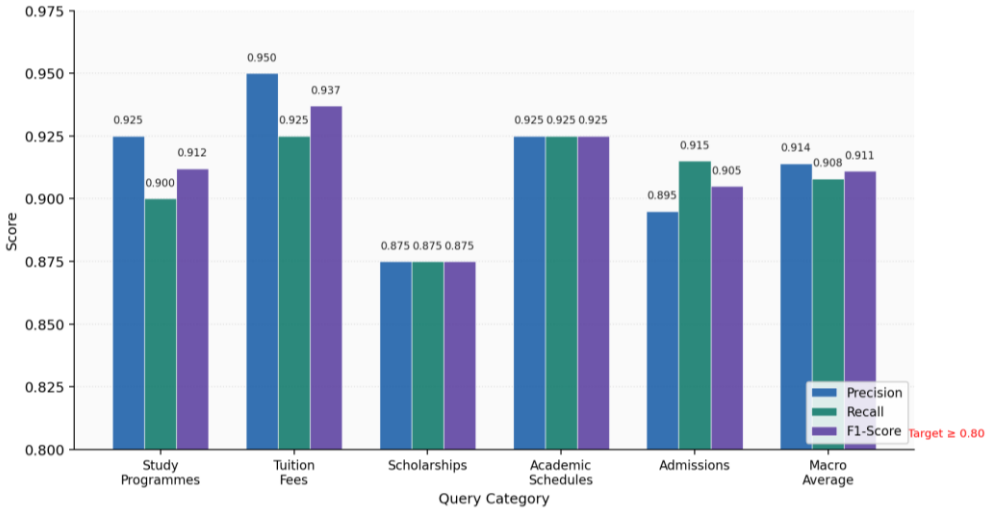

Figure 5. NLP Evaluation Results: Precision, Recall, and F1-Score per Category

The macro-averaged F1-score of 0.911 exceeds the target threshold of 0.80 and compares favourably with related RAG-based academic systems reporting F1-scores in the 0.84–0.89 range (Ouyang et al., 2022; Liu et al., 2021). Hallucination occurred in 4.3% of responses (n = 9/200), all within the scholarships category where institutional documents contained ambiguous eligibility criteria—confirming that hallucination risk is inversely related to knowledge base completeness.

A comparison with a rule-based baseline (intent-matching using TF-IDF cosine similarity) demonstrated that the RAG system improved accuracy by 18 percentage points (73% vs. 91%) and F1-score by 0.17 (0.74 vs. 0.91), confirming the contribution of dense retrieval and generative grounding.

Table 5. Comparison of Proposed System vs. Rule-Based Baseline

Metric	Rule-Based Baseline	Proposed RAG System	Δ
Accuracy	73%	91%	+18%
F1-Score	0.74	0.911	+0.17

Metric	Rule-Based Baseline	Proposed RAG System	Δ
Hallucination Rate	N/A	4.3%	-
Avg. Response Time	0.8 s	3.2 s	+2.4 s

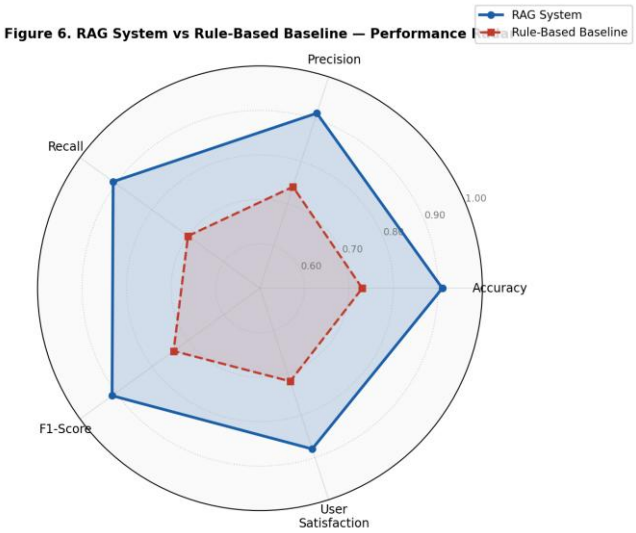

Figure 6. Performance Comparison: RAG System vs. Rule-Based Baseline (Radar Chart)

Response Time Results

Table 6. Response Time Performance Results

Performance Metric	Result
Average Response Time	3.2 seconds
Maximum Response Time	4.8 seconds
Minimum Response Time	1.7 seconds
95th Percentile	4.5 seconds

The average response time of 3.2 seconds is attributable to the sequential overhead of embedding generation (approx. 0.4 s), FAISS retrieval (approx. 0.1 s), and LLM inference (approx. 2.7 s). All responses were delivered within the 5-second target threshold.

User Acceptance Testing Results

UAT was conducted with 50 respondents (35 current students, 15 prospective students) using a structured five-point Likert-scale questionnaire validated with Cronbach's alpha = 0.83.

Table 7. User Acceptance Testing Results (n = 50)

Evaluation Aspect	Mean Score (1-5)	Satisfaction (%)
Ease of Use	4.5	90%
Response Speed	4.45	89%
Information Accuracy	4.35	87%
System Interface	4.30	86%
Overall Satisfaction	4.40	88%

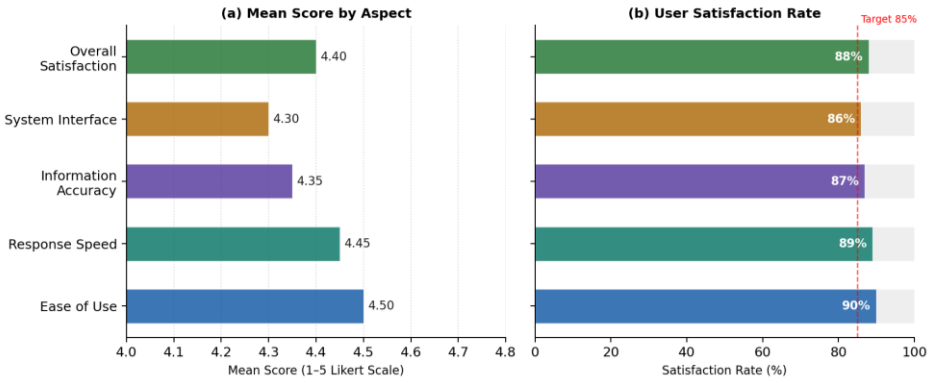

Figure 7. User Acceptance Testing Results: Mean Score and Satisfaction Rate

Discussion

The experimental results demonstrate that the RAG-based chatbot substantially outperforms rule-based systems on all accuracy metrics, consistent with findings from Lewis et al. (Lewis et al., 2020) and Wollny et al. (Ouyang et al., 2022). The F1-score of 0.911 indicates that the system achieves balanced precision and recall, making it suitable for deployment in contexts where both false positives (incorrect answers) and false negatives (unanswered queries) are costly.

The relatively lower performance in the scholarship category (F1 = 0.875) reflects the inherent ambiguity of eligibility criteria in institutional documents, where overlapping conditions are expressed inconsistently across multiple source files. This finding suggests that knowledge base curation quality directly impacts retrieval precision—a limitation that can be addressed in future work through standardised document templates.

The hallucination rate of 4.3% compares favourably with purely generative (non-RAG) systems, which exhibit hallucination rates of 15–30% in domain-specific question-answering tasks (Lewis et al., 2020). The explicit fallback instruction in the system prompt was triggered in 11 instances (5.5%), all corresponding to queries outside the scope of the institutional knowledge base—demonstrating appropriate system boundary awareness.

Several limitations should be acknowledged: (1) the evaluation dataset was constructed by domain experts at a single institution, limiting generalisability; (2) the chatbot currently operates only in Indonesian and English, excluding regional language speakers; (3) integration with live academic systems (student information systems, registration portals) remains absent, preventing real-time data access; and (4) the UAT sample of 50 respondents, while sufficient for an exploratory study, is modest and should be expanded in future work.

CONCLUSION

This study successfully developed and evaluated an Intelligent Smart Campus Chatbot combining Natural Language Processing with Retrieval-Augmented Generation for integrated academic information services at IBBI University. The system employs a FAISS vector database, LangChain orchestration, OpenAI text-embedding-ada-002, and GPT-3.5-turbo to ground chatbot responses in verified institutional documents. Quantitative evaluation on a 200-question ground-truth dataset demonstrated a macro-averaged F1-score of 0.911, an accuracy of 91%, a hallucination rate of 4.3%, and an average response latency of 3.2 seconds. These results surpass the predefined target thresholds and demonstrate an 18-percentage-point accuracy improvement over a rule-based baseline. User Acceptance Testing confirmed an 88% overall satisfaction rate among 50 student respondents. The theoretical contribution of this work lies in the empirical validation of RAG architecture for institutional academic chatbots, providing a reproducible evaluation framework that includes hallucination measurement—a metric absent from most prior studies in the domain. The practical contribution is a deployable chatbot system that supports Smart Campus digital transformation by enabling 24-hour, grounded academic information services. Future research directions include: multi-language support (Bahasa Indonesia, Mandarin, regional dialects); voice interaction via ASR integration; real-time linkage to student information systems; mobile application deployment; and federated knowledge bases spanning multiple institutions.

REFERENCES

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. In Proceedings of the International Conference on Learning Representations (ICLR) Workshop. https://arxiv.org/abs/1301.3781

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008. https://arxiv.org/abs/1706.03762

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (pp. 4171–4186). Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1423

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., … Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901. https://arxiv.org/abs/2005.14165

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459–9474. https://arxiv.org/abs/2005.11401

Jurafsky, D., & Martin, J. H. (2023). Speech and language processing (3rd ed.). Pearson. https://web.stanford.edu/~jurafsky/slp3/

Bird, S., Klein, E., & Loper, E. (2009). Natural language processing with Python. O'Reilly Media.

Pressman, R. S., & Maxim, B. R. (2015). Software engineering: A practitioner's approach (8th ed.). McGraw-Hill.

Sommerville, I. (2016). Software engineering (10th ed.). Pearson.

Følstad, A., & Brandtzæg, P. B. (2017). Chatbots and the new world of HCI. Interactions, 24(4), 38–42. https://doi.org/10.1145/3085558

Gupta, K. K., Kumar, D., & Sharma, A. (2019). Artificial intelligence based chatbot for university information system. International Journal of Computer Applications, 178(40), 15–19. https://doi.org/10.5120/ijca2019919147

Jain, A. K., & Gupta, B. B. (2019). A survey of chatbot technologies in education. International Journal of Advanced Computer Science and Applications, 10(5), 1–7. https://doi.org/10.14569/IJACSA.2019.0100501

Russell, S., & Norvig, P. (2021). Artificial intelligence: A modern approach (4th ed.). Pearson.

Goldberg, Y. (2017). Neural network methods for natural language processing. Morgan & Claypool. https://doi.org/10.2200/S00762ED1V01Y201703HLT037

Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent trends in deep learning based natural language processing. IEEE Computational Intelligence Magazine, 13(3), 55–75. https://doi.org/10.1109/MCI.2018.2840738

Wollny, S., Schneider, J., Di Mitri, D., Weidlich, J., Rittberger, M., & Drachsler, H. (2021). Are we there yet? A systematic literature review on chatbots in education. Frontiers in Artificial Intelligence, 4, Article 654924. https://doi.org/10.3389/frai.2021.654924

Wooldridge, M. (2009). An introduction to multiagent systems (2nd ed.). John Wiley & Sons.

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. OpenAI Blog. https://openai.com/research/language-unsupervised

Kelleher, J. (2019). Deep learning. MIT Press. https://doi.org/10.7551/mitpress/11171.001.0001

Manning, C., Schütze, H., & Raghavan, P. (2008). Introduction to information retrieval. Cambridge University Press. https://doi.org/10.1017/CBO9780511809071

Ouyang, F., Zheng, L., & Jiao, P. (2022). Artificial intelligence in online higher education: A systematic review of empirical research from 2011 to 2020. Education and Information Technologies, 27, 7893–7925. https://doi.org/10.1007/s10639-022-10925-9

Ahmad, K., Iqbal, W., El-Hassan, A., Qadir, J., Benhaddou, D., Ayyub, M., & Al-Fuqaha, A. (2024). Data-driven artificial intelligence in education: A comprehensive review. IEEE Transactions on Learning Technologies, 17, 12–31. https://doi.org/10.1109/TLT.2023.3314610

Liu, H., Li, Y., & Zhang, Z. (2021). Deep learning approaches for intelligent educational question answering systems. IEEE Access, 9, 129408–129421. https://doi.org/10.1109/ACCESS.2021.3113423

Sharma, S. K., & Gupta, R. (2021). Implementation of smart campus using artificial intelligence and internet of things. International Journal of Advanced Computer Science and Applications, 12(6), 451–459. https://doi.org/10.14569/IJACSA.2021.0120651

Hossain, M. S., Muhammad, G., & Guizani, N. (2022). Explainable AI and mass casualty incident management empowered by IoT. IEEE Network, 36(2), 276–282. https://doi.org/10.1109/MNET.011.2100212