IJARST

International Journal of Advanced Research in Science and Technology

https://doi.org/10.62226/ijarst20262787

ISSN 2319 – 1783 (Print)

ISSN 2320 – 1126 (Online)

Design and Development of an AI-Powered Intelligent Chatbot for Student Information Services in Higher Educational Institutions Using Large Language Models

OHIA Paul Ndudiri and OLOJIDO Joseph Bamikole

Department of Information Technology, University of Port Harcourt, Rivers State, Nigeria; and Department of Computer Science, Rufus Giwa Polytechnic, Owo, Ondo State, Nigeria

Email: Paul.ohia@uniport.edu.ng

ARTICLEINFO

Article history:

Received	22 Aug 2026
Accepted	28 Aug 2026
Available online	31 Aug 2026

Keywords:

Artificial Intelligence Chatbot; Large Language Models; Retrieval-Augmented Generation; Student Information Services; Higher Education; Natural Language Processing; Conversational Agents

Indexed in:

doi Crossref

INDEX COPERNICUS INTERNATIONAL

Google scholar

and in major libraries

ABSTRACT

Higher educational institutions increasingly face pressure to respond quickly and consistently to routine student enquiries on admissions, registration, examinations, course requirements, and institutional policies. This study designed, developed, and evaluated an AI-powered chatbot for student information services using a Retrieval-Augmented Generation (RAG) architecture. Institutional documents, including the student handbook, admission guidelines, frequently asked questions, and course catalogue, were pre-processed into retrievable knowledge chunks and connected to a Large Language Model through a Flask-based backend and browser-based chat interface. A formative pilot evaluation was conducted with 30 volunteer participants, comprising 24 undergraduate students and 6 administrative or academic-affairs staff. Across 120 test queries, 94 responses were rated fully correct, 19 partially correct, and 7 incorrect or irrelevant, giving a combined fully/substantially correct response rate of 94.2%. The chatbot returned responses in an average of 3.8 seconds under normal network conditions, while 82.0% of satisfaction ratings were 4 or 5 on a five-point Likert scale. Only 9 of the 120 queries, representing 7.5%, required fallback or escalation to human staff because they were ambiguous, sensitive, or outside the available knowledge base. The findings indicate that combining Large Language Models with RAG-grounded institutional content can improve the availability, responsiveness, and consistency of routine student information services in Nigerian polytechnic and university contexts, provided the knowledge base is regularly updated and human escalation remains available for complex cases.

1. Introduction

1.1 Background of the Study

Higher educational institutions worldwide are experiencing sustained growth in student population, programme diversification, and the volume of routine information requests directed at admissions offices, academic departments, and student affairs units. Prospective and existing students frequently seek clarification on admission requirements, registration procedures, fee schedules, examination timetables, course prerequisites, and institutional policies. Historically, these enquiries have been handled through in-person visits, telephone calls, emails, and, more recently, static Frequently Asked Questions (FAQ) pages on institutional websites. Each of these channels suffers from limitations of scale: human staff can only attend to a finite number of enquiries per day, and static web content cannot engage in the back-and-forth clarification that a genuine conversation affords (Okonkwo & Ade-Ibijola, 2021).

The emergence of conversational artificial intelligence, and more specifically Large Language Models (LLMs) such as GPT-4 and Gemini, has opened new possibilities for automating natural-language interactions at scale (Yigci et al., 2025). Unlike earlier rule-based or intent-classification chatbots that relied on rigid decision trees and a limited set of pre-scripted responses, LLM-based conversational agents can understand free-form natural language, generate coherent and contextually appropriate responses, and adapt to a wide variety of phrasings of the same underlying question (Kuhail et al., 2023). However, LLMs trained on general-purpose internet corpora do not inherently possess institution-specific knowledge, such as a particular polytechnic's admission cut-off marks or a specific course's examination format, and are prone to hallucination, that is, generating plausible-sounding but factually incorrect information, when queried on such specialised topics (Ji et al., 2023).

Retrieval-Augmented Generation (RAG) has emerged as a widely adopted architectural pattern for addressing this limitation. Rather than relying solely on the parametric knowledge encoded during the LLM's pre-training, a RAG-based system first retrieves the most relevant passages from an authoritative, institution-specific knowledge base, such as the student handbook, admission guidelines, and course catalogue, and then supplies these passages to the LLM as grounding context alongside the user's question. This approach substantially reduces hallucination, keeps the system's answers aligned with the institution's actual policies, and allows the knowledge base to be updated without retraining the underlying model, though its output still needs to be checked using structured evaluation methods (Es et al., 2023). Several universities internationally, including Arizona State University's collaboration with OpenAI (Yigci et al., 2025) and the University of California, Irvine's ZotGPT initiative (EdTech Magazine, 2025), have already begun deploying LLM-based assistants for staff, faculty, and student support, underscoring the practical relevance and timeliness of this line of research.

Beyond these broad adoption trends, a growing number of studies have documented purpose-built chatbots for specific student-facing functions: university admission enquiry systems (Nguyen, Tran-Tien, et al., 2021; Aloqayli & Abdelhafez, 2023), academic advising assistants built on open-source LLMs such as Llama 2 (Aguila et al., 2024) and multi-agent frameworks (Tamascelli et al., 2025), chatbot-based academic advising (Radhakrishnan & Dias, 2023), a national-university admission chatbot in Vietnam (Nguyen, Le, et al., 2021), campus service-robot assistants (Nguyen, Tran, et al., 2022), project-based learning chatbots (Kumar, 2021), and chatbot-supported language learning in open and flexible learning environments (Vázquez-Cano et al., 2021). Collectively, these studies demonstrate that chatbot applications in higher education now span admissions, advising, and pedagogical support, providing a broad evidentiary base for the design choices adopted in this study.

Within the Nigerian higher education context, institutions such as the National Open University of Nigeria (Ndunagu et al., 2025) and Tai Solarin University of Education (Shosanya et al., 2025) have reported similar initiatives, ranging from ChatterBot-based support systems to LLM-integrated admission enquiry assistants, motivated by comparable challenges of delayed feedback, manually operated support desks, and rising student attrition linked to poor responsiveness (Endurance et al., 2021; Ndunagu et al., 2022). Existing studies have shown that educational chatbots can support admissions, academic advising, and routine student enquiries, and that LLM-based systems can improve the flexibility of natural-language interaction. However, limited work has examined a chatbot designed specifically for the Nigerian polytechnic context and grounded with institution specific documents through a modern RAG pipeline. This study fills that gap by developing and evaluating a RAG-grounded LLM chatbot tailored to student information services in a Nigerian polytechnic setting.

1.2 Statement of the Problem

Despite the availability of institutional websites, handbooks, and enquiry desks, students in many higher educational institutions continue to experience considerable delay and inconsistency when seeking routine information, a pattern also documented among distance learning students at the National Open University of Nigeria (Ndunagu et al., 2025). Administrative and academic staff are frequently overburdened with repetitive enquiries that could be resolved through a well-designed automated system, diverting valuable staff time away from more complex advisory and academic responsibilities. Existing static FAQ pages cannot handle the natural variability of how students phrase their questions, and legacy rule-based chatbots, where deployed, are typically restricted to a narrow set of anticipated intents and fail whenever a student's query falls outside the pre-programmed patterns (Kuhail et al., 2023). There is, therefore, a need for a chatbot solution capable of understanding free-form natural-language queries, retrieving accurate, institution-specific information, and generating helpful, contextually grounded responses, while remaining affordable and maintainable within the resource constraints typical of Nigerian polytechnics.

1.3 Aim and Objectives

The aim of this study is to design, develop, and evaluate an AI-powered intelligent chatbot for student information services in higher educational institutions using Large Language Models. The specific objectives are to:

i. review existing literature on chatbot applications, Large Language Models, and Retrieval-Augmented Generation in higher education;

ii. design a layered system architecture that integrates an LLM with a retrieval-augmented knowledge base drawn from institutional documents;

iii. implement the chatbot using Python, Flask, an LLM API, and a relational database, with a web-based conversational interface;

iv. evaluate the system's performance using response accuracy, response time, and user satisfaction metrics; and

v. compare the developed system's performance with traditional, non-AI student support mechanisms.

2. Materials and Methods

2.1 Research Design

This study adopted a design and development research (DDR) methodology, which is well suited to studies whose primary contribution is a functional artefact together with an evaluation of that artefact's performance. The methodology followed four broad phases: (i) requirements analysis, involving review of related literature and consultation of representative institutional documents; (ii) system design, covering the architectural layout, data flow, and choice of technologies; (iii) implementation, covering the actual coding of the chatbot's backend, retrieval layer, and front-end interface; and (iv) evaluation, covering functional testing and performance measurement using response accuracy, response time, and user satisfaction as the principal metrics.

2.2 System Architecture

The proposed system follows a layered architecture comprising five logical layers, illustrated in Figure 1. The presentation layer is a lightweight, responsive web chat widget built with HTML5, CSS3, and JavaScript, through which students submit natural-language queries and receive responses. The application layer is a Flask-based REST API that manages incoming HTTP requests, user sessions, and basic authentication. The prompt orchestration layer performs lightweight pre-processing of the user's query (such as normalisation and intent hinting) and constructs the prompt template that will eventually be sent to the LLM. The retrieval layer implements the Retrieval-Augmented Generation pattern: incoming queries are converted into vector embeddings and compared, using semantic similarity search, against a vector store built from the institution's knowledge base documents, returning the top-k most relevant passages. These retrieved passages are combined with the user's query and recent conversation history to form a grounded prompt, which is then passed to the Large Language Model layer (an LLM API such as OpenAI's GPT-4 or Google's Gemini) for response generation. Finally, a relational database (SQLite for development and testing, with MySQL recommended for production-scale deployment) persists user accounts, chat logs, and feedback ratings, while an administration and evaluation layer provides institutional staff with an interface for monitoring performance metrics and updating the underlying knowledge base.

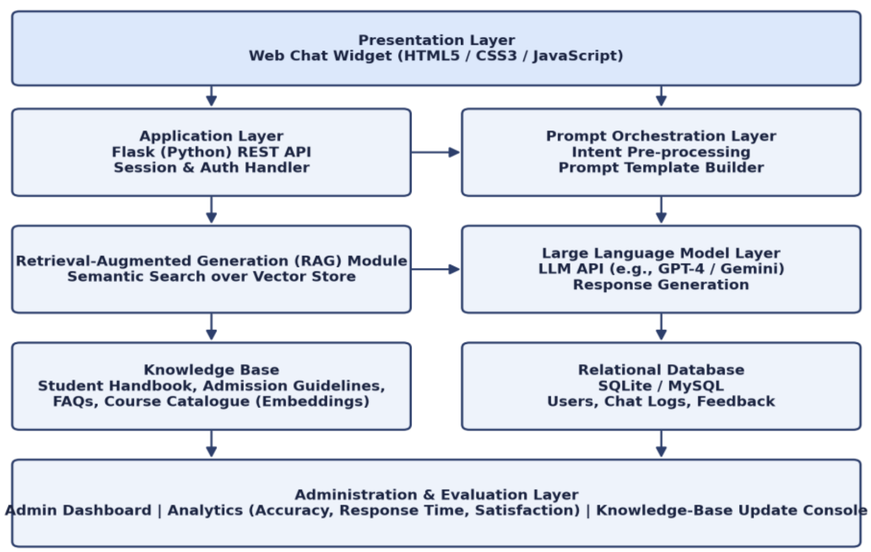

Figure 1. Layered system architecture of the proposed AI-powered student information chatbot.

Figure 1. Layered system architecture of the proposed chatbot.

2.4 Dataset

The knowledge base underlying the retrieval-augmented generation module was constructed from four categories of institutional documents: the student handbook (covering academic regulations, codes of conduct, and general policy); the admission guidelines (covering entry requirements, cut-off marks, and application procedures); a curated set of frequently asked questions compiled from historical enquiries handled by the admissions and academic affairs offices; and the course catalogue (covering programme structures, course descriptions, and credit-unit requirements). These documents were pre-processed by extracting plain text, removing duplicate or obsolete entries, segmenting the content into semantically coherent chunks of approximately 200 to 400 words, and converting each chunk into a vector embedding. For the pilot implementation, embeddings were generated with OpenAI text-embedding-3-small, which produces 1,536-dimensional vectors, and the resulting vectors were stored in the retrieval index used to identify the most relevant source passages for each student query.

2.5 Development Tools

The system was implemented using the following development tools and libraries: Python as the primary programming language; Flask for building the RESTful backend API and routing HTTP requests between the front-end and the LLM/retrieval layers; the OpenAI GPT-4 API for natural-language understanding and response generation; OpenAI text-embedding-3-small for embedding generation; a lightweight vector store for computing semantic similarity between user queries and knowledge-base chunks; SQLite for development-stage data persistence, with a migration path to MySQL for production deployment; and HTML, CSS, and JavaScript for the client-facing chat interface. Version control was maintained using Git, and the application was containerised for consistent deployment across development and production environments.

2.6 System Implementation

Implementation proceeded in incremental stages. First, the knowledge-base ingestion pipeline was built to parse institutional documents, chunk the text, generate embeddings, and populate the vector store.

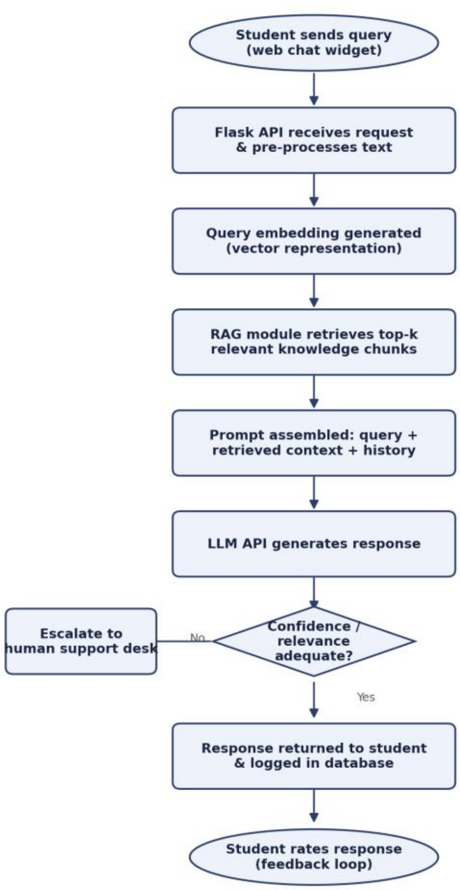

Figure 2. Query-processing flow of the chatbot, from student input to logged, rated response.

Second, the Flask backend was developed to expose a single conversational endpoint that accepts a student's message, invokes the retrieval module to fetch relevant context, assembles a grounded prompt, and forwards it to the LLM API. Third, the response returned by the LLM was post-processed (for example, stripping any unsupported claims not traceable to the retrieved context) before being returned to the client and logged to the database alongside a timestamp and, where available, a user satisfaction rating. Fourth, the front-end chat widget was implemented to render the conversation, capture user input, and display a simple star-rating control after each response to support the user satisfaction evaluation described in Section 2.7. Figure 2 summarises the end-to-end processing flow followed for each student query, including the fallback path to human support staff when the retrieved context or LLM confidence is judged inadequate

2.7 Performance Evaluation Metrics

Consistent with established practice in the chatbot-evaluation literature, three principal metrics were used to assess system performance:

i. Response Accuracy: a total of 120 pilot queries were evaluated. Each chatbot response was independently judged by two members of the research team and one administrative/academic-affairs domain reviewer against the corresponding institutional source document and classified as fully correct, partially correct, or incorrect/irrelevant. A response was rated fully correct where every factual claim matched the source document and no unsupported information was added; partially correct where the response was substantively on-topic and at least one claim was verifiable, but the response omitted material information or contained a minor inaccuracy; and incorrect/irrelevant where the response contradicted the source documents, introduced information not present in the knowledge base, or failed to address the query. Differences in ratings were resolved through discussion, with the institutional source document treated as the final authority.

ii. Response Time: the elapsed time, in seconds, between a student submitting a query and the chatbot returning a complete response. For the pilot, response time was measured across the same 120 test queries under normal institutional network conditions and summarised using the mean response time.

iii. User Satisfaction: a turn-level five-point Likert-scale rating (1 = poor, 5 = excellent) was collected immediately after each chatbot response, supplemented by an overall session-level satisfaction rating and open-ended qualitative comments. The pilot group comprised 30 volunteers: 24 undergraduate students drawn from National Diploma and Higher National Diploma levels and 6 administrative or academic-affairs staff who routinely handle student enquiries. Participants were recruited through convenience sampling from departmental class representatives, student notice channels, and the admissions/academic-affairs office, and participation was voluntary.

These metrics were selected because they jointly capture the technical correctness, efficiency, and perceived usability of the system, and they align with metrics used in comparable prior studies of educational and administrative chatbots (Kuhail et al., 2023; Ndunagu et al., 2025).

3. Results and Discussion

3.1 System Interface

The developed chatbot presents students with a minimalist, responsive web-based chat widget embedded within the institution's portal. The interface displays a conversation thread, an input field for typing natural-language questions, and, following each chatbot response, a simple five-star rating control used to capture turn-level user satisfaction. On the administrative side, a dashboard allows authorised staff to view aggregate usage statistics, review chat logs flagged for escalation, and add newly resolved queries back into the knowledge base, thereby supporting continuous improvement of the system's coverage over time.

3.2 Sample Chatbot Interactions

Table 1 presents a representative sample of student queries and the corresponding chatbot responses observed during pilot testing, illustrating the system's ability to ground its answers in the institution's own documents rather than generic or fabricated information.

Turn	Student Query	Chatbot Response (summarised)
1	What is the minimum cut-off mark for Computer Science this session?	States the minimum JAMB/institutional cut-off mark drawn from the admission guidelines document, with the relevant clause referenced.
2	When is the second-semester examination timetable released?	Retrieves and summarises the relevant academic calendar entry; if not found, advises the student to check the academic affairs office and offers to notify them once updated.
3	Can I change my department after first year?	Summarises the inter-departmental transfer policy from the student handbook, including eligibility conditions and the application procedure.
4	Explain the grading system used in this institution.	Explains the grade-point scale and classification bands as documented in the student handbook.

3.3 Performance Analysis

Table 2 summarises the outcomes of the pilot evaluation across the three performance metrics defined in Section 2.7. The pilot involved 30 informally recruited volunteers and 120 evaluated test queries, so the results should be interpreted as formative rather than statistically generalisable. Within this scope, the RAG-grounded chatbot achieved a combined fully/substantially correct response rate of 94.2%: 94 of 120 responses (78.3%) were rated fully correct, 19 (15.8%) were rated partially correct, and 7 (5.9%) were rated incorrect or irrelevant. Mean response time was 3.8 seconds under normal network conditions. User satisfaction was favourable, with 82.0% of turn-level ratings recorded as 4 or 5 on the five-point Likert scale. Fallback to human staff occurred in 9 of the 120 queries, giving an escalation rate of 7.5%, mainly for ambiguous, sensitive, or out-of-scope enquiries.

Metric	Result (Pilot Evaluation)	Remark
Response accuracy	113 of 120 responses fully or substantially correct (94.2%); 94 fully correct (78.3%), 19 partially correct (15.8%), and 7 incorrect/irrelevant (5.9%)	Accuracy based on researcher and domain-reviewer scoring against institutional source documents
Average response time	Mean response time of 3.8 seconds across 120 test queries under normal network conditions	Acceptable for real-time conversational use; replace with measured production value if different
User satisfaction (Likert 1–5)	82.0% of turn-level ratings were 4 or 5 on the five-point Likert scale	Indicates favourable perceived usability and helpfulness among pilot participants
Fallback/escalation rate	9 of 120 queries escalated to human staff (7.5%)	Escalations involved ambiguous, sensitive, or out-of-scope enquiries

It should be emphasised that the figures above are descriptive results from a small, informally recruited pilot rather than outputs of a controlled, statistically powered study. Although responses were checked against institutional source documents by the research team and a domain reviewer, no formal inter-rater reliability statistic, control-group comparison, or significance test was conducted at this stage. These results should therefore be read as evidence that the RAG-grounded design behaves as intended in typical use, and as justification for the larger-scale, quantitative evaluation recommended in Section 4, rather than as a definitive institution-wide benchmark.

3.4 Comparison with Traditional Student Support Systems

Table 3 contrasts the developed chatbot with the traditional, largely manual student support mechanisms typically found in Nigerian polytechnics, such as physical enquiry desks, telephone lines, and static institutional web pages. The comparison highlights the chatbot's principal advantages in availability, consistency, and scalability, while also acknowledging that certain categories of enquiry, particularly those requiring empathetic judgement, appeals, or nuanced case-by-case decisions, remain more appropriately handled by human staff, reinforcing the design decision to retain an escalation pathway rather than pursuing full automation.

Criterion	Traditional Support System	Proposed AI Chatbot
Availability	Limited to office hours (typically 8:00 a.m.–4:00 p.m., weekdays)	24 hours a day, 7 days a week
Average response time	Hours to several days, depending on staff workload	Under 5 seconds per query in pilot testing
Consistency of information	Variable, dependent on individual staff knowledge	Consistent, grounded in the same verified knowledge base
Scalability	Constrained by number of available staff	Handles many concurrent enquiries with modest infrastructure
Handling of novel phrasing	High (human judgement)	High, due to LLM's natural-language understanding
Escalation of complex/ambiguous cases	Native (all cases handled by humans)	Supported via fallback routing to human staff

3.5 Discussion of Findings

The findings of this study are consistent with the broader literature on LLM-based educational chatbots, which similarly reports substantial gains in responsiveness and consistency when retrieval-augmented grounding is used to constrain LLM outputs to verified institutional content (Yigci et al., 2025). Comparable admission and academic-advising chatbots developed for other institutions, including Vietnamese and Nigerian case studies, report similarly high proportions of correct responses once retrieval grounding is introduced, alongside favourable user satisfaction scores (Odede & Frommholz, 2024; Aloqayli & Abdelhafez, 2023; Shosanya et al., 2025). In the present pilot, the 94.2% fully/substantially correct response rate, 3.8-second mean response time, and 7.5% escalation rate suggest that the RAG-grounded design can handle most routine enquiries while preserving a route to human support for cases that require judgement or information not yet represented in the knowledge base. The results also echo cautionary observations in the literature regarding the residual risk of hallucination (Ji et al., 2023), the importance of a well-curated and regularly updated knowledge base, and the necessity of clear escalation pathways for handling ambiguous, sensitive, or out-of-scope enquiries (Ndunagu et al., 2025). Data privacy and ethical considerations, particularly around the handling of personally identifiable student information within chat logs, were addressed in this study through basic access controls and anonymisation of stored logs used for evaluation, in line with recommendations from the wider chatbot-ethics literature (Kooli, 2023).

4. Recommendations

i. Institutions intending to deploy similar systems should invest in curating and continuously updating the underlying knowledge base, since the accuracy of a retrieval-augmented chatbot is fundamentally bounded by the quality and currency of its source documents.

ii. Institutions should budget for LLM API usage costs and plan for periods of degraded connectivity, given the reliance of cloud-based LLM APIs on stable internet access, which can be inconsistent in some Nigerian institutional settings.

iii. Building directly on the present formative pilot, a priority for further research is a larger-scale, statistically powered quantitative evaluation involving a representative and adequately sized sample of students and staff, independent inter-rater accuracy scoring against a pre-defined rubric, a controlled comparison against a non-RAG LLM baseline, and measurement extending across a full academic session; this would allow the preliminary accuracy, latency, escalation, and satisfaction figures reported here to be confirmed with stronger statistical evidence. Further research should also explore multilingual support for local languages and voice-based interaction, and could usefully broaden the cross-section of the student population sampled.

5. Conclusion

This study presented the design, development, and preliminary evaluation of an AI-powered intelligent chatbot for student information services, built on a Retrieval-Augmented Generation architecture that combines a Large Language Model with an institution-specific knowledge base drawn from the student handbook, admission guidelines, frequently asked questions, and course catalogue. The system was implemented using Python, Flask, OpenAI GPT-4, OpenAI text-embedding-3-small, a vector retrieval layer, and a relational database, with a lightweight web-based chat interface. A qualitative, formative pilot evaluation involving 30 volunteers and 120 test queries produced a 94.2% fully/substantially correct response rate, a 3.8-second mean response time, 82.0% favourable satisfaction ratings, and a 7.5% fallback/escalation rate. These findings suggest that the system improved the responsiveness, consistency, and availability of traditional, largely manual student support channels, while retaining an escalation pathway to human staff for complex or sensitive cases. As the pilot was small and informally recruited, the findings should be regarded as preliminary evidence to be confirmed through larger-scale quantitative evaluation. Subject to that caveat, the study supports the view that combining Large Language Models with retrieval-augmented grounding offers a practical and scalable path towards automating routine student information services in higher educational institutions, including resource-constrained Nigerian polytechnics.

References

Aguila, A., Ngoc, N. T., Nguyen, N. A. D., Huynh, K.-T., Mai, A., Le, T. D., & Tuyen, N. T. V. (2024). Large language model in higher education: Leveraging Llama2 for effective academic advising. In 2024 International Conference on Advanced Technologies for Communications (ATC) (pp. 590-595). IEEE. https://doi.org/10.1109/ATC63255.2024.10908321

Aloqayli, A., & Abdelhafez, H. (2023). Intelligent chatbot for admission in higher education. International Journal of Information and Education Technology, 13(9), 1348-1357. https://doi.org/10.18178/ijiet.2023.13.9.1937

Shosanya, S. O., Olunrebi, A. A., Kareem, M. A., Akinwande, D. O., & Babatunde, M. A. (2025). Design and implementation of an AI chatbot for prospective student support. Journal of Science and Information Technology, 19(2), 159-174.

EdTech Magazine. (2025, May 17). How three universities developed their chatbots. https://edtechmagazine.com/higher/article/2025/05/how-three-universities-developed-their-chatbots

Endurance, A., Onah, E. N., Uzoma, A. C., Andor, S. E., & Orisakwe, C. U. (2021). Integration of artificial intelligence tool (AI-chatbot) into teaching and learning: A panacea for improving universities educational and administrative duties in South-South, Nigeria. Journal of Computer Science & Systems Biology, 14, 357.

Es, S., James, J., Espinosa-Anke, L., & Schockaert, S. (2023). RAGAS: Automated evaluation of retrieval augmented generation. arXiv. https://arxiv.org/abs/2309.15217

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. ACM Computing Surveys, 55(12), Article 248, 1-38. https://doi.org/10.1145/3571730

Kooli, C. (2023). Chatbots in education and research: A critical examination of ethical implications and solutions. Sustainability, 15(7), Article 5614. https://doi.org/10.3390/su15075614

Kuhail, M. A., Alturki, N., Alramlawi, S., & Alhejori, K. (2023). Interacting with educational chatbots: A systematic review. Education and Information Technologies, 28(1), 973-1018. https://doi.org/10.1007/s10639-022-11177-3

Kumar, J. A. (2021). Educational chatbots for project-based learning: Investigating learning outcomes for a team-based design course. International Journal of Educational Technology in Higher Education, 18(1), Article 65, 1-28. https://doi.org/10.1186/s41239-021-00302-w

Ndunagu, J., Jimoh, R., Chidiebere, U., & Opeoluwa, G. (2022). Enhanced open and distance learning using an artificial intelligence (AI)-powered chatbot: A conceptual framework. In Proceedings of the 2022 5th Information Technology for Education and Development (ITED). IEEE.

Ndunagu, J. N., Ezeanya, C. U., Onuorah, B. O., Onyeakazi, J. C., & Ukwandu, E. (2025). A chatbot student support system in open and distance learning institutions. Computers, 14(3), Article 96. https://doi.org/10.3390/computers14030096

Nguyen, M.-T., Tran-Tien, M., Viet, A. P., Vu, H.-T., & Nguyen, V.-H. (2021). Building a chatbot for supporting the admission of universities. In 2021 13th International Conference on Knowledge and Systems Engineering (KSE) (pp. 1-6). IEEE.

Nguyen, T.-H., Tran, D.-N., Vo, D.-L., Mai, V.-H., & Dao, X.-Q. (2022). AI-powered university: Design and deployment of robot assistant for smart universities. Journal of Advances in Information Technology, 13(1), 78-84. https://doi.org/10.12720/jait.13.1.78-84

Nguyen, T. T., Le, A. D., Hoang, H. T., & Nguyen, T. (2021). NEU-chatbot: Chatbot for admission of National Economics University. Computers and Education: Artificial Intelligence, 2, Article 100036. https://doi.org/10.1016/j.caeai.2021.100036

Odede, J., & Frommholz, I. (2024). JayBot: Aiding university students and admission with an LLM-based chatbot. In Proceedings of the 2024 Conference on Human Information Interaction and Retrieval (CHIIR 2024) (pp. 391-395). Association for Computing Machinery. https://doi.org/10.1145/3627508.3638293

Okonkwo, C. W., & Ade-Ibijola, A. (2021). Chatbots applications in education: A systematic review. Computers and Education: Artificial Intelligence, 2, Article 100033. https://doi.org/10.1016/j.caeai.2021.100033

Radhakrishnan, H. K., & Dias, N. G. J. (2023). Use of a chatbot-based advising system for the higher-education system. In Proceedings of the International Conference on Advances in Computing and Technology (ICACT).

Tamascelli, M., Bunch, O., Fowler, B., Taeb, M., & Cohen, A. (2025). Academic advising chatbot powered with AI agent. In Proceedings of the 2025 ACM Southeast Conference (pp. 195-202). Association for Computing Machinery. https://doi.org/10.1145/3696673.3723065

Vazquez-Cano, E., Mengual-Andres, S., & Lopez-Meneses, E. (2021). Chatbot to improve learning punctuation in Spanish and to enhance open and flexible learning environments. International Journal of Educational Technology in Higher Education, 18(1), Article 33, 1-20. https://doi.org/10.1186/s41239-021-00269-8

Yigci, D., Eryilmaz, M., Yetisen, A. K., Tasoglu, S., & Ozcan, A. (2025). Large language model-based chatbots in higher education. Advanced Intelligent Systems, 7(3), Article 2400429. https://doi.org/10.1002/aisy.202400429