Design and Implementation of a Retrieval-Augmented Chatbot for the International Burch University Website

Student paper

Lazar Matic
Department of Information Technologies
Faculty of Engineering, Natural and Medical sciences
International Burch University
Sarajevo, Bosnia and Herzegovina
lazarmatic91@gmail.com

Sanin Zajmovic
Department of Information Technologies
Faculty of Engineering, Natural and Medical sciences
International Burch University
Sarajevo, Bosnia and Herzegovina
sanin.zajmovic@gmail.com

***Abstract*— The rapid growth of digital information combined with traditional navigation methods has made it increasingly difficult for students to efficiently locate accurate and relevant content on large university websites. This study presents the design, development, and evaluation of an intelligent chatbot for the International Burch University (IBU) website, using Natural Language Processing (NLP) and Retrieval-Augmented Generation (RAG).**

The system was trained exclusively on official IBU web content collected through automated web scraping of more than 470 webpages. The data was cleaned, segmented into semantic chunks, embedded using the Sentence-BERT all-MiniLM-L6-v2 model, and indexed with FAISS for similarity-based retrieval. User queries are processed through a RAG pipeline, where relevant documents are retrieved and provided as context for grounded answer generation.

The chatbot was evaluated using retrieval metrics, token-level F1 scores, and human usefulness ratings. The results show strong retrieval performance, achieving a Hit Rate@5 of 100%, while answer quality reached an average F1 score of 0.35 and a usefulness rating of 60%.

Keywords*— *Natural Language Processing, Retrieval-Augmented Generation, University Chatbot, Semantic Search, Sentence-BERT, FAISS

I. Introduction

The rapid growth of digital information has significantly increased the complexity of accessing accurate and relevant content on large institutional websites. University web platforms typically contain hundreds of interconnected pages covering academic programs, administrative procedures, regulations, student services, and campus resources. While these information is publicly available, students often experience difficulties locating specific answers efficiently due to complex navigation structures, fragmented content, and time-consuming manual searching. As a result, students frequently rely on administrative offices for routine inquiries, which increases response time and workload for university staff.

In recent years, NLP and conversational Artificial Intelligence (AI) systems have been proved as effective solutions for improving information accessibility across various domains. Chatbots, in particular, provide an intuitive interface that allows users to ask questions in natural language and receive immediate responses. However, many existing university chatbots rely on rule-based systems or limited FAQ datasets, which restrict their coverage and reliability. More advanced large language model (LLM)-based systems, while flexible, may generate incorrect or hallucinated responses if not properly grounded in verified institutional data. This highlights the need for chatbot solutions that combine conversational flexibility with factual accuracy based on official university sources.

The main objective of this paper is to design, develop, and evaluate an intelligent chatbot for the IBU website that improves students' access to accurate and relevant institutional information. The proposed system leverages a RAG approach, integrating semantic retrieval with LLM-based response generation. The chatbot is trained exclusively on official IBU website content collected through automated web scraping, ensuring that all generated answers are grounded in verified university documentation. By providing instant, context-aware responses to student queries, the system aims to enhance user experience, reduce information retrieval time, and lower the workload on administrative departments.

This study is guided by two central hypotheses. First, an NLP-based chatbot utilizing semantic search and retrieval-augmented generation can significantly improve students' ability to access accurate and relevant information from the university website compared to traditional navigation methods. Second, a RAG-based chatbot trained solely on scraped data from official university sources can generate reliable and trustworthy responses while minimizing misinformation and reducing the need for direct administrative support. These hypotheses are evaluated through both quantitative retrieval and answer quality metrics, as well as qualitative human usefulness assessments.

The significance of this research lies in its practical and academic contributions. From a practical perspective, the proposed chatbot offers a scalable and automated solution for enhancing digital student support and standardizing information delivery within the university. From an academic perspective, the study contributes empirical evidence on the effectiveness of RAG-based chatbot systems in educational environments and demonstrates a fully automated pipeline for large-scale institutional data collection, semantic indexing, and grounded response generation.

The remainder of the paper is organized as follows. First section provides the introduction, while section II reviews related work on university chatbots and RAG systems. Section III describes the dataset, system architecture, and methodology used for chatbot development. Section IV presents the experimental results and system evaluation. Finally, Section V concludes the paper and outlines directions for future work.

II. Literature Review

The purpose of this literature review is to establish a solid understanding of the existing research papers relevant to this study and to take a look at the current status within the university and other related chatbots. This section summarizes key findings from five selected studies, examining how each contributes to the development of core concepts and methodological approaches in the field. By critically analyzing these works, this review identifies areas of pros and cons, evaluates the strengths and limitations of various research designs, and highlights unresolved issues. Through this comparative assessment, the review provides the analytical foundation upon which the present study is built.

A. Critical analysis of related literature

In related literature on university chatbots, several systems have been reported, including LiSA, a rule based chatbot prototype delivered by Facebook Messenger, used in a survey of 100 students to express their needs like enrollment info, schedules and observe usage behaviors including anthropomorphism and occasional inappropriate language, but never implemented as a fully functional system. [1] A second study by UTEHY Admission Chatbot which is built for a Vietnamese university admissions context compared multiple NLP pipelines (BERT, PhoBERT, GPT-2..), achieving good intent/entity detection metrics and deploying on Facebook, but suffered from only around 51 % conversation accuracy and lacked any usability or user-satisfaction evaluation. [2] On the other side, BarkPlug v.2, a state-of-the-art LLM based system using RAG over documents from 42 departments at a USA university, delivered strong quantitative performance (high RAGAS scores) and was tested for usability, but still had moderate user satisfaction, occasional hallucinations, limited accessibility and no long-term usage analysis. [3] On the simpler end, a rule-based FAQ chatbot by Manipal University (India) using AIML plus a proposed but never implemented semantic component LSA for campus related questions remains undocumented beyond a few example responses, offering no performance or user evaluation. [4] Finally, a low resource language effort by the University of Computer Studies, Yangon (Myanmar) implemented an AIML based chatbot handling around 970 QA pairs, demonstrating feasibility in Myanmar but evaluation was limited to only four sample dialogues with no user feedback or comprehensive accuracy assessment. [5]

B. Synthesis and Research Justification

The reviewed literature highlights several critical gaps that justify the approach and contributions of the present study. First, many existing university chatbot systems lack sufficient evaluation. AIML based systems such as this developed for Manipal University and Myanmar institutions provides little to no empirical validation, while more advanced implementations like UTEHY report limited dialogue performance, achieving only 51.2% conversation accuracy, largely due to insufficient training data. Second, the practical utility of prior systems is restricted by narrow domain coverage. For example, UTEHY's chatbot focuses solely on admissions, and manually curated datasets such as Myanmar's 970 question answer pairs or Manipal's unspecified collections cannot scale to cover the full range of information students typically seek. Third, no existing study demonstrates a chatbot trained on a complete set of university resources drawn directly from the institution's website, leaving a significant gap in coverage comprehensiveness. Fourth, the disconnect between technical performance and user satisfaction which is evident in BarkPlug v.2, which achieved a strong RAGAS score of 0.96 but only moderate usability (SUS 67.75) shows that retrieval accuracy alone is not enough to ensure a positive user experience. Finally, the absence of cross institutional generalizability and the lack of standardized evaluation frameworks make it difficult to compare systems meaningfully or support informed adoption decisions by university administrators.

The present study addresses these limitations through several key innovations. Unlike LiSA, which concluded at the requirements gathering stage, this study develops a fully operational RAG based chatbot trained on a comprehensive dataset scraped from more than 470 official IBU URLs, covering admissions, academic programs, facilities, regulations, services, and administrative units. This breadth far exceeds the narrow admission focus of UTEHY and the limited manually authored QA datasets used in AIML based systems. Building on the architectural strengths of BarkPlug v.2 while addressing its reported issues with hallucinations and user experience.

Methodologically, the system integrates Sentence Transformers (*all-MiniLM-L6-v2*) for efficient semantic embedding, pickle based state management, and FAISS similarity retrieval of the top three relevant chunks to provide grounded responses that reduce hallucination risks. Unlike systems dependent on costly manual annotation reflected in UTEHY's $\kappa = 0.67$ inter annotator agreement, which illustrates annotation challenges and our fully automated web scraping pipeline offers scalability, maintainability, and complete coverage as university information evolves. Since SBERT handles all necessary preprocessing internally during embedding, the retrieved text remains in its neutral, unnormalized form, which works better with LLMs to generate more readable answers which improves user experience. By implementing this system for IBU with comprehensive data coverage, planned evaluation through accuracy and performance metrics, and a deployment strategy that ensures immediate accessibility for students via website integration, this study contributes new empirical evidence on the effectiveness of RAG based educational chatbots and provides a replicable framework suitable for other institutions seeking to enhance information accessibility and reduce administrative burden.

III. Design and Implementation of the IBU Website Chatbot

Although the International Burch University website provides comprehensive and up-to-date information, its increasing size and structural complexity create practical challenges for effective information retrieval. With more than 470 webpages covering academic programs, administrative procedures, regulations, student services, and institutional policies, users are required to manually navigate multiple sections to locate specific answers. This process is inefficient, particularly when users are unfamiliar with the website structure or require immediate responses to precise questions.

From a technical perspective, conventional website search mechanisms are insufficient for addressing this problem. Keyword-based search functions do not capture semantic meaning and often fail when users phrase queries differently from the terminology used on the website. As a result, relevant information may remain undiscovered even when it exists within the system. In practice, this limitation leads students to repeatedly contact administrative offices for routine inquiries, increasing response delays and operational workload.

Existing chatbot solutions applied in university environments only partially address these challenges. Rule-based and FAQ-driven chatbots are constrained by manually defined patterns and limited datasets, which restrict their coverage and scalability. While more advanced language model-based chatbots offer improved conversational capabilities, they introduce the risk of generating unverified or hallucinated information when not grounded in authoritative data sources. In an academic environment, where institutional accuracy and trustworthiness are essential, such limitations make unconstrained generative systems unsuitable as standalone solutions.

These challenges demonstrate the need for an intelligent system that combines semantic understanding, factual grounding, and scalability. A Retrieval-Augmented Generation (RAG) approach satisfies these requirements by integrating semantic retrieval over verified documents with controlled language model-based answer generation. By retrieving relevant content directly from official university webpages and using it as context for response generation, the system ensures that answers remain accurate, transparent, and institutionally verified.

Based on this motivation, this study proposes the design and implementation of an intelligent chatbot trained exclusively on official International Burch University website data. The system employs automated web scraping, semantic text embeddings, similarity-based retrieval, and large language model-based response generation to provide accurate and context-aware answers to user queries. The following sections describe the data collection and preprocessing pipeline, system architecture, employed methodologies, and technical implementation details of the proposed chatbot solution.

A. Data and Findings

This study required a reliable textual dataset representing academic, administrative, and informational content from IBU. Because no structured dataset of this type existed publicly, a complete pipeline has been developed to collect and prepare the required data. The resulting corpus forms the foundation for our NLP project, including retrieval augmented generation, semantic search, and linguistic analysis

1) Data Collecting

The dataset was collected through automated web scraping of the official IBU website. The process began by extracting all valid URLs from the publicly available XML sitemap located at *https://www.ibu.edu.ba/sitemap-0.xml* [6]. From this sitemap, every page belonging to the domain was parsed and the links that contained relevant academic or institutional content were filtered. To ensure proper loading of dynamic components, the scraper relied on Selenium [7] WebDriver with headless Google Chrome. This configuration allowed the system to handle Java Script rendered elements, delayed content, and occasional Cloudflare verification prompts. For each visited URL, the full visible text was extracted after removing header and footer elements, producing a clean body of textual content associated with each page.

2) Data Cleaning

After extraction, the raw text required several cleaning steps to ensure consistency across the dataset. HTML markup, scripts, navigation items, and repetitive template elements were removed using BeautifulSoup parsing and custom filtering rules. Pages containing fewer than 100 characters of meaningful text were discarded to avoid noise. Additional normalization steps included whitespace correction, removal of content related to interface, and verification that each document contained only the main page body. The cleaned documents were stored in JSON format, with each entry containing the extracted text and its corresponding source URL.

3) Dataset Summary

The final dataset consists of the scraped webpage texts, their cleaned versions, the chunked representations, the semantic embeddings, and a FAISS similarity index designed for retrieval tasks. Together, these components form a coherent and structured corpus that accurately reflects the informational content of International Burch University's website. The dataset captures details about academic programs, university administration, campus life, student services, and institutional operations. This curated and preprocessed dataset supports all downstream NLP experiments conducted in this research.

B. Methodology

This section will outline the approach taken to develop the IBU chatbot using RAG. The methodology includes data preprocessing, vector embedding creation, semantic retrieval mechanisms, and response generation. The overall workflow follows a pipeline from taking raw data to response generation, using both traditional NLP techniques and modern transformer models. The implementation uses Sentence Transformers for semantic understanding, FAISS for fast similarity search, and pickle serialization for efficient storage of embeddings and document chunks, creating a robust question answering chatbot based on the university's official documentation.

1) Data Preprocessing Pipeline

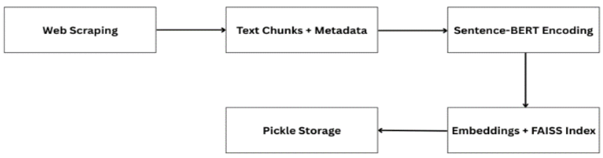

Fig. 1. RAG Embedding pipeline

a) Text Chunking: Raw scraped text from over 470 university web pages was segmented using LangChain's [8] RecursiveCharacterTextSplitter with chunk size being 1000 characters where the overlap was 200 characters. The splitting was prioritized at paragraph boundaries ('\n\n', '\n'), then sentences ('. '), then words (' '), and finally characters.

This chunking strategy addresses two key challenges, maintaining semantic coherence by respecting natural text boundaries, and ensuring contextual continuity through overlapping segments. Each chunk retains metadata linking it to its source URL and position within the original document (chunk_id), enabling source attribution in chatbot responses. Pages containing fewer than 100 characters were filtered out to remove navigation elements and error pages that lack useful information.

b) Text Preservation Approach: Unlike traditional information retrieval systems that employ aggressive preprocessing (stop word removal, lemmatization, stemming), this implementation preserves the original natural language text. Modern transformer-based embedding models capture semantic meaning from complete sentences, including stop words and grammatical structure. Only minimal cleaning was applied to remove HTML artifacts while maintaining the text's natural readability for the language model's generation phase.

2) Vector Embedding and Indexing

a) Embedding Generation: The retrieval system was implemented using SentenceTransformers with efficient serialization via Python's pickle module. Document embeddings were generated using the 'all-MiniLM-L6-v2' SBERTs model, which is 5 times faster than the 'all-mpnet-base-v2' model but still offers good quality [9]. This model produces 384-dimensional dense vector representations. The resulting All preprocessed text chunks were encoded into dense vector representations using the selected embedding model. These embeddings capture the semantic meaning of the text, enabling similarity-based retrieval in which relevant content can be identified even when exact keyword matches are absent. This representation allows the system to retrieve semantically related documents that align with the intent of user queries rather than relying solely on lexical similarity.

b) FAISS Indexing: To enable efficient similarity search across thousands of chunk embeddings, the system uses FAISS (Facebook AI Similarity Search) with an 'IndexFlatL2' index. FAISS organizes embeddings in an optimized data structure that enables rapid nearest-neighbor search using L2 (Euclidean) distance as the similarity metric. The complete system state (preprocessed chunks, embeddings, and FAISS index) is serialized using Python's pickle module and stored persistently, eliminating the need to regenerate embeddings for each session. The storage files include 'chunks.pkl' - Text chunks with metadata (source URLs, chunk IDs), 'embeddings.pkl' - Numpy array of 384-dimensional vectors, and 'faiss_index.index' - FAISS similarity search index.

3) Retrieval Mechanism

Query processing is done in a way that when a user submits a question, the system first encodes it using the same SBERT model used for document chunks. Next, FAISS performs efficient nearest neighbor search to identify the top-k most relevant chunks based on vector similarity. Finally, the original natural language text of the most similar chunks is retrieved for context formation.

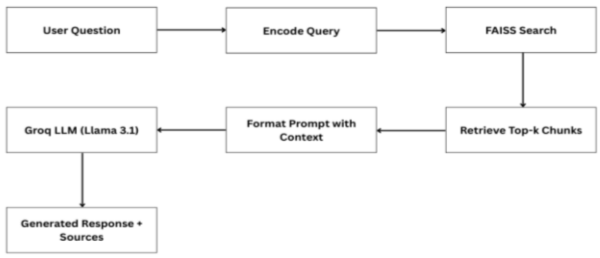

Fig. 2. System Workflow

4) Response Generation Pipeline

a) Context Preparation: Retrieved chunks are concatenated to form a context window that provides the language model with relevant information to answer the user's question. Each chunk's source URL is preserved for citation purposes.

b) Prompt Engineering: A structured prompting strategy was used to ensure that the language model generates accurate and institutionally grounded responses. The system prompt defines the model's role as an assistant for International Burch University and restricts answer generation exclusively to the retrieved contextual information from official university webpages. When the required information is not present in the provided context, the model is instructed to explicitly state this limitation. User queries are combined with the retrieved context and supplied to the language model as a single prompt. This approach reduces hallucination risks and ensures that generated responses remain aligned with verified university documentation.

c) Language Model Selection: Response generation uses Groq's Llama 3.1 8B Instant model via API. This lightweight model is well suited for RAG applications where the primary task is synthesizing information from provided context rather than relying on extensive world knowledge. The model was configured with 0.3 temperature, favoring factual accuracy over creative elaboration, and 500 max tokens, which is sufficient for comprehensive answers.

IV. Results

After the text edit has been completed, the paper is ready for the template. This section presents the evaluation of the IBU chatbot across two key dimensions: retrieval quality and answer generation quality. The evaluation employed both automatic metrics computed across 20 test questions and manual human evaluation of answer usefulness. The results demonstrate the chatbot's effectiveness in retrieving relevant information and generating helpful responses for university related queries.

A. Evaluation Metrics

The chatbot was evaluated using the following metrics:

Hit Rate@k: Binary metric indicating whether at least one relevant source URL appears in the top-k retrieved results. A score of 1 indicates successful retrieval, while 0 indicates failure. The average across all queries measures overall retrieval success rate.

Precision@k: Measures the proportion of retrieved source URLs (from top-k results) that are relevant to the query. Calculated as the number of relevant source URLs divided by k, then averaged across all queries.

F1 Score: Harmonic mean of precision and recall at the token level, measuring lexical overlap between the generated answer and reference answer. Scores range from 0 (no overlap) to 1 (perfect overlap). This metric provides partial credit for answers that capture key information even when not identical to the reference.

$$\text{Precision} = \frac{(\text{common tokens})}{(\text{tokens in predicted answer})} \quad (1)$$

$$\text{Recall} = \frac{(\text{common tokens})}{(\text{tokens in reference answer})} \quad (2)$$

$$\text{F1} = 2 \times \frac{(\text{Precision} \times \text{Recall})}{(\text{Precision} + \text{Recall})} \quad (3)$$

Usefulness: Manual rating of answer quality on a 3-point scale where 2 represents an answer that is clear, directly answers the question, and helpful for students, 1 represents a partially helpful answer which might be missing details or is unclear, while 0 represents an answer that is not helpful or is off-topic. This metric evaluates whether generated answers would be genuinely helpful to students in practice, capturing qualitative aspects like clarity and completeness that automatic metrics based on word overlap cannot measure.

B. Retrieval Performance

The retrieval system's ability to identify relevant source URLs was evaluated across 20 test questions. Fig. 3 presents the retrieval performance for different values of k (number of retrieved chunks). Hit Rate increases with more chunks retrieved, while Precision decreases as expected due to the inclusion of less relevant results.

The results demonstrate strong retrieval performance with Hit Rate@5 reaching 100%, indicating that for all the queries, at least one relevant source URL was successfully retrieved within the top 5 results. The Precision@5 of 40% suggests that two fifths of the retrieved source URLs were relevant.

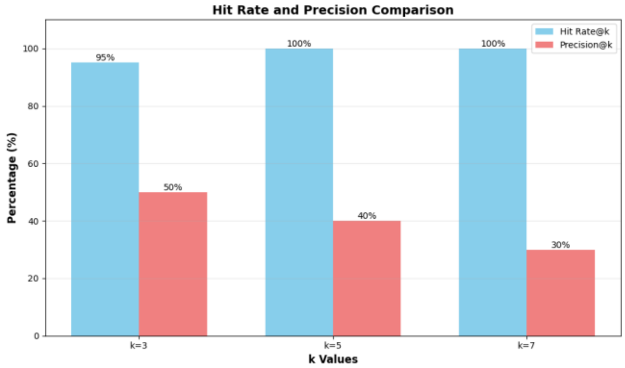

Fig. 3. Retrieval Quality Metrics

C. Answer Generation Performance

The quality of generated answers was assessed using automatic F1 scoring against reference answers and manual usefulness ratings. Results are presented in Table 1.

TABLE I. Answer Quality Evaluation

Metric	Score	Description
Average F1 Score	0.35 (35%)	Token-level overlap with reference answers
Average Usefulness	0.60 (60%)	Manual rating of answer usefulness (scale: 0-2)
Questions Evaluated (Automatic)	20	Full evaluation set
Questions Evaluated (Manual)	20	Full evaluation set for usefulness

Table 1 presents answer quality metrics, showing low lexical overlap (F1=0.35) and moderate usefulness ratings (1.2/2.0 or 60%) from manual evaluation. The F1 score of 0.35 indicates a lacking overlap between generated and reference answers, suggesting that either the system does not capture key information from the retrieved context or that the retrieved chunk from the correct source URL is wrong. However, the moderate usefulness score of 0.60 reveals room for improvement in answer completeness and clarity.

D. Error Analysis and Limitations

Manual inspection of the 20 evaluation questions revealed several retrieval and generation failure patterns. The most common failure pattern involved correctly identifying the relevant source page but retrieving the chunks that did not contain the specific requested information. This occurred in 25% of the evaluated questions. Analysis revealed that important information was sometimes split across chunk boundaries. For a tuition question, the system retrieved chunks from the tuition-and-fees page but got introductory text rather than the actual fee information. Additionally, long pages with multiple topics (e.g., department page containing program descriptions, staff listings, and requirements) resulted in retrieving general information rather than specific details requested.

An example of this error type occurs with the question "What are the working hours of the student affairs office?" Although the retrieved source is correct, the Student Affairs Office page, the selected chunk discusses the office's purpose and services rather than the working hours. Additionally, some retrieved chunks come from entirely irrelevant sources. The problem arises due to the specific "8:00 AM to 4:00 PM" information that was in a different chunk that ranked lower in similarity and was therefore not selected.

V. Conclusion

The aim of this study was to investigate whether a RAG based chatbot can improve students' access to accurate and relevant information on the IBU website. The sample for this research consisted of a dataset made out of data scraped from more than 470 official IBU webpages, representing academic programs, administrative procedures, student services, campus resources, and institutional policies. Data was collected through automated web scraping using Selenium and processed through a multi stage NLP pipeline that included text cleaning, chunking with LangChain, semantic embedding using the SBERT all-MiniLM-L6-v2 model, and FAISS similarity indexing. The processed data was analyzed using retrieval performance metrics (Hit Rate@k, Precision@k), answer generation metrics (F1 score), and human usefulness ratings to assess the effectiveness of the developed system.

Findings showed that the chatbot achieved strong retrieval performance, with Hit Rate@5 reaching 100%, demonstrating its ability to consistently locate relevant university documents. However, answer generation produced a moderate F1 score of 0.35, indicating limited lexical overlap with reference answers, while human evaluated usefulness received a 0.60 rating, reflecting partial but meaningful practical value for students. These results highlight that while the system reliably retrieves information, improvements in response synthesis and clarity are still needed for optimal user experience. Therefore, it can be concluded that the implemented RAG based chatbot successfully addresses the core accessibility challenges present in the current IBU website, offering faster and more direct access to university verified information. The system provides a functional foundation for improving digital student support and sets the groundwork for future improvements in answer quality and user interaction.

The findings of this study provide clear answers to the established research questions. Regarding the first research question, the results demonstrate that the developed chatbot significantly improves information accessibility for students by allowing them to obtain relevant and accurate answers without the need to manually navigate multiple webpages. The strong retrieval performance indicates that the system consistently identifies the correct and most relevant institutional sources, effectively addressing the problem of hard finding of information on the IBU website. Concerning the second research question, the implemented RAG framework proved to be a reliable approach for generating grounded and institution verified responses. Because the chatbot depends exclusively on data scraped from official university webpages, it minimizes the probability of misinformation and ensures higher correctness. This reduces the students' dependence on administrative offices for routine inquiries and provides a more efficient and automated method for accessing university related information.

This study will bring benefits to several key groups. For the university, the chatbot reduces administrative workload, provides 24/7 student support, and improves the overall usability of the institutional website. It also helps standardize information delivery, ensuring students receive consistent and accurate guidance. For the scholarly community, this research contributes empirical evidence on the effectiveness of RAG systems in educational environments, introduces a scalable, fully automated data collection pipeline, and demonstrates the existence of transformer based semantic retrieval for university knowledge systems. Future research can extend this work by improving answer generation models, integrating multilingual support, conducting large scale usability studies, and evaluating long term adoption trends to further advance intelligent campus information systems.

References

[1] M. Dibitonto, K. Leszczynska, F. Tazzi, and C. M. Medaglia, "Chatbot in a campus environment: Design of LiSA, a virtual assistant to help students in their university life," in HCI 2018: Lecture Notes in Computer Science, vol. 10903, pp. 103–116, 2018. doi: 10.1007/978-3-319-91250-9_9. [Online]. Available: https://repository.unilink.it/chatbot-in-a-campus-environment-design-of-lisa-a-virtual-assistant-to-help-students-in-their-university-life/

[2] M.-T. Nguyen, M. Tran-Tien, V. Phan, H.-T. Vu, et al., "Building a chatbot for supporting the admission of universities," in Proc. 13th IEEE Int. Conf. Knowledge and Systems Engineering (KSE), 2021. [Online]. Available: https://www.researchgate.net/publication/357082032_Building_a_Chatbot_for_Supporting_the_Admission_of_Universities

[3] S. Neupane, E. Hossain, J. Keith, H. Tripathi, F. Ghiasi, N. Amiri Golilarz, A. Amirlatifi, S. Mittal, and S. Rahimi, "From questions to insightful answers: Building an informed chatbot for university resources (BarkPlug v.2)," 2024. [Online]. Available: https://arxiv.org/abs/2405.08120

[4] B. R. Ranoliya, N. Raghuwanshi, and S. Singh, "Chatbot for university-related FAQs using AIML + LSA," in Proc. Int. Conf. Advances in Computing, Communications and Informatics (ICACCI), 2017. [Online]. Available: https://researcher.manipal.edu/en/publications/chatbot-for-university-related-faqs/

[5] K. M. Soe and Soe, "University chatbot using Artificial Intelligence Markup Language (AIML) in Myanmar language for University of Computer Studies, Yangon," in Proc. ICCA, 2020. [Online]. Available: https://www.academia.edu/78186369/University_Chatbot_using_Artificial_Intelligence_Markup_Language

[6] International Burch University, "Sitemap," IBU Official Website, [Online]. Available: https://www.ibu.edu.ba/sitemap-0.xml

[7] Selenium Project Contributors, "Selenium WebDriver Documentation," SeleniumHQ, 2024. [Online]. Available: https://www.selenium.dev/documentation/webdriver/

[8] R. Taylor and LangChain Team, "LangChain: Building applications with large language models," 2024. [Online]. Available: https://python.langchain.com

[9] UKPLab, "Pretrained Models - Sentence Transformers Documentation," SBERT.net. [Online]. Available: https://www.sbert.net/docs/sentence_transformer/pretrained_models.html