

# DESIGNING AN AI CHATBOT SYSTEM FOR PERSONALIZED LEARNING SUPPORT FOR PRIMARY SCHOOL STUDENTS BASED ON LARGE LANGUAGE MODELS AND RETRIEVAL-AUGMENTED TEXT GENERATION SYSTEMS

Do Manh Hung<sup>1,\*</sup>

DOI: <https://doi.org/10.57001/huih5804.2025.421>

## ABSTRACT

This paper presents the design of an artificial intelligence (AI) chatbot system that supports personalized learning for primary school students by integrating large language models (LLMs) and retrieval-augmented generation (RAG) techniques. The system aims to provide each student with a virtual assistant acting as a primary school teacher, capable of natural language interaction, answering questions, and offering guidance tailored to individual learning levels. We propose an architecture that combines an LLM with a knowledge retrieval module to ensure personalized and trustworthy responses. The system utilizes prompt engineering techniques to shape the chatbot's tone to resemble that of a primary school teacher, and applies content filtering to ensure age-appropriate responses. A prototype has been implemented using a local LLM and a knowledge base built from textbook materials. Experimental interactions demonstrate the chatbot's ability to accurately answer students' typical learning questions in a pedagogical and approachable tone. Preliminary evaluations indicate positive outcomes: the chatbot provides timely support, enhances student engagement, and reduces the workload for teachers. This study highlights the potential of applying LLMs combined with RAG in primary education to personalize learning experiences. Finally, the paper discusses the benefits and challenges of the proposed system and suggests future development directions.

**Keywords:** *Large Language Models, Artificial Intelligence, Generative AI, Retrieval-Augmented Generation.*

<sup>1</sup>School of Information and Communications Technology, Hanoi University of Industry, Vietnam

\*Email: [hungdm\\_cntt@hauai.edu.vn](mailto:hungdm_cntt@hauai.edu.vn)

Received: 10/8/2025

Revised: 15/10/2025

Accepted: 28/11/2025

## 1. INTRODUCTION

Personalized learning has become an increasingly emphasized objective in primary education. Each student has a unique learning pace and preference, yet in traditional classrooms, teachers often lack the time and resources to provide individualized attention. Many students are unable to engage directly with their teachers to clarify doubts due to limited time and large class sizes [1], which may result in learning gaps or reduced motivation. A promising solution is the use of AI chatbots virtual assistants capable of natural conversation to create a one-on-one interactive environment accessible anytime. Large Language Models (LLMs) such as GPT-3.5 and GPT-4 have demonstrated the ability to generate high-quality natural language responses, presenting new opportunities for applications in education at all levels. Notably, Generative AI holds the promise of customizing learning content, adapting to individual needs, and serving as intelligent tutors.

Initial studies have shown that AI integration in classrooms can enhance learning outcomes. For instance, a recent experiment in a primary school found that ChatGPT could automatically tailor lesson content based on each student's proficiency level, thereby increasing engagement and improving performance [2]. Another study reported that students supported by an AI chatbot (featuring GPT-4 and RAG integration) scored significantly higher in post-lesson assessments compared to a control group without AI assistance [5, 9]. These results confirm AI's potential to personalize and improve the quality of learning for primary school students.

However, deploying LLMs in education also presents notable challenges. One major concern is the issue of "hallucination," where the model generates inaccurate or fabricated information [8]. This is a significant barrier to using LLM-based chatbots for educational support, as misleading information can adversely affect students' understanding. To address this, Retrieval-Augmented Generation (RAG) has been proposed as an effective solution. Rather than relying solely on pre-trained internal knowledge, RAG enables the model to retrieve relevant external documents to ground its responses in verified information. This integration substantially reduces hallucination by anchoring the model to trustworthy knowledge sources. Thus, combining LLMs with RAG is a promising approach for building educational chatbots that are both intelligent and trustworthy.

In this context, the present study focuses on designing an AI chatbot system to support personalized learning for primary school students by leveraging the strengths of LLMs and RAG. The system is designed to act as a friendly "virtual teacher," always ready to answer questions and guide students in a tone and style appropriate for young learners. The remainder of the paper is organized as follows: Section 2 reviews relevant theoretical foundations and related work on LLMs, RAG, and educational chatbots; Section 3 details the proposed system design, including the LLM-RAG architecture, prompt engineering to simulate teacher roles, and mechanisms for age-appropriate content control. Section 4 presents the implementation and experimental results, with interaction examples and preliminary evaluations. Section 5 concludes the study and outlines future development directions.

## 2. THEORETICAL BACKGROUND AND RELATED WORK

**Large Language Models (LLMs):** Large language models are a class of artificial intelligence systems trained on massive corpora of natural language data, capable of generating coherent and contextually appropriate text akin to human language. Contemporary LLMs are predominantly built upon the Transformer architecture, with notable examples including GPT-3, GPT-4, Gemma, and LLaMA. These models can perform a wide range of language tasks such as answering questions, explaining knowledge, composing text, translating languages, and more. Their core strength lies in their ability to comprehend context and generate flexible responses based on the probabilistic prediction of subsequent

tokens [2]. In the field of education, LLMs have marked a significant breakthrough by functioning as intelligent virtual tutors. For example, the release of ChatGPT (OpenAI) in late 2022 catalyzed a global wave of AI adoption in educational settings. LLMs enable students to engage in conversational interactions and receive instant explanations to their questions, simulating one-on-one tutoring scenarios. However, generic LLMs are not specifically trained on primary school curricula, and they are susceptible to generating inaccurate content, a phenomenon known as "hallucination". To deploy LLMs effectively and safely in schools, mechanisms for integrating domain-specific knowledge and controlling content quality are essential.

**Retrieval-Augmented Generation (RAG):** RAG is an approach that combines generative models like LLMs with information retrieval systems. Instead of relying solely on pre-trained data, the model retrieves and incorporates relevant documents from an external knowledge base during the response generation process. This enhances the factual accuracy and reliability of the model's outputs [3]. Typically, the RAG workflow involves three stages: (1) Querying - upon receiving a user question, the system utilizes a retrieval engine (e.g., keyword-based or semantic embedding-based) to locate relevant passages from a document corpus (such as textbooks, reference materials, or question banks); (2) Knowledge integration - the retrieved content is appended to the model's input prompt; (3) Response generation - the LLM produces an answer based on the enriched prompt, often including citations to the sources used. This method significantly reduces hallucination by grounding the generated content in verifiable information. For example, Tufino et al. [3] implemented a RAG-based AI assistant for physics education, where the model consistently drew upon teacher-provided curricula and worked solutions, thereby ensuring reliable and pedagogically sound responses. Overall, RAG is emerging as a crucial solution for deploying LLMs in domains demanding high factual precision, such as education and healthcare, by embedding domain knowledge into the language generation pipeline.

**Chatbots in Education and Related Research:** The application of chatbots for educational support is not a new concept. However, only with the recent advances in LLMs have educational chatbots achieved natural conversation capabilities and broad domain comprehension. Prior to the LLM era, intelligent tutoring

systems offered automated feedback but were generally rule-based or relied on predefined scripts. In contrast, modern LLM-based chatbots can respond flexibly to a wide range of questions and learning scenarios. Numerous recent studies have explored the use of AI chatbots across subjects and educational levels. For instance, Swacha and Gracel conducted a survey of 47 studies on RAG-based educational chatbots, revealing their application in assisting both students and teachers with tasks such as knowledge retrieval and guided practice [1].

In primary education, particular attention must be paid to age-appropriateness and friendliness in chatbot design. Young learners require simple, vivid, relatable, and safe expressions. Several studies have investigated child-friendly chatbots, such as those used for storytelling or reading assistance, aimed at fostering reading comprehension and interactive engagement. Nonetheless, the use of LLMs as tutors for primary students remains in its early stages and entails addressing several challenges: filtering sensitive content, adapting language complexity, and maintaining pedagogical integrity in responses. Our research builds upon these foundational studies and focuses on developing a comprehensive chatbot system. The main contribution lies in the tight integration of LLMs with RAG to enable personalized, safe, and pedagogically aligned learning experiences for primary school students.

### 3. SYSTEM DESIGN METHODOLOGY

#### 3.1. Overall System Architecture

The proposed AI chatbot system is designed using a modular architecture comprising the following main components: (1) User Interface; (2) Language Processing Module (LLM); (3) Knowledge Retrieval Module (RAG); and (4) Orchestration and Response Generation Module.

**User Interface:** Students interact with the chatbot via a conversational interface, which can be implemented as a web or mobile application. Students input their queries or requests in natural language (Vietnamese) into the chat box. The interface visually displays the chatbot's responses and may include child-friendly multimedia elements (such as illustrations or emotive icons) if appropriate.

**Language Processing Module (LLM):** This module acts as the core engine of the chatbot, employing a large language model trained to generate Vietnamese text. In our prototype, we utilize a locally hosted LLM (e.g., a

customized version of LLaMA or Gemma 3) that runs on an internal server to ensure student data is not transmitted externally. This module receives the prompt (which includes the student's query and relevant context) and generates the textual response.

**Knowledge Retrieval Module (RAG):** This component augments the LLM with factual knowledge. We constructed a domain-specific knowledge base comprising primary education materials (textbooks, exercise books, and teacher-provided resources). These documents are preprocessed and stored as searchable indexes using vectorization techniques and semantic query models to retrieve relevant text segments. Upon receiving a question, before the LLM generates a response, the RAG module retrieves the most relevant segments (e.g., top 3 passages) from the corpus. These retrieved passages are then inserted into the expanded prompt to provide contextual grounding for the LLM. This integration ensures that the chatbot's responses are evidence-based and rooted in verified educational content, rather than relying on potentially inaccurate internal model knowledge. Prior studies have demonstrated that incorporating factual references into prompts improves output accuracy and reduces hallucinations.

**Orchestration and Response Generation Module:** We developed a coordination module to manage the processing flow. This module constructs the final prompt for the LLM by combining: (a) a fixed instruction prefix that simulates the role of a primary school teacher (detailed in Section 3.2), (b) the knowledge segments retrieved by the RAG module, and (c) the student's query. After the LLM generates the response, this module may perform post-processing, which includes content moderation (e.g., removing inappropriate phrases or inserting friendly symbols/emojis where needed).

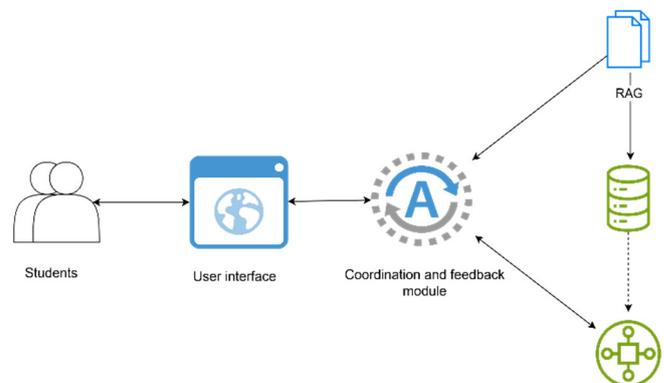


Figure 1. AI chatbot system architecture supporting personalized learning for primary school students integrating LLM and RAG

### 3.2. Prompt Engineering and Simulating the Primary School Teacher Persona

Prompt engineering plays a crucial role in shaping the LLM's behavior. To emulate the tone and style of a primary school teacher, we crafted a system prompt that explicitly defines the role and communication approach expected of the model. The prompt instructs: "You are a dedicated and friendly primary school teacher. Your task is to help students understand lessons by explaining concepts in a simple and comprehensible manner, using real-life examples relatable to their age. You always encourage and praise students, using a cheerful and positive tone. Avoid using language or content inappropriate for children." This prompt serves as a soft behavioral guide for the LLM.

Additionally, we include few-shot examples in the prompt to illustrate desired responses. For example, we insert sample dialogues between a student and teacher in which the teacher explains a simple concept using pedagogical language. These examples help the model better internalize the role and communication style required. The combination of role definition and illustrative examples ensures that the LLM consistently responds in a teacher-like, age-appropriate, and supportive manner.

### 3.3. Content Management and Age-Appropriate Safety Measures

Ensuring that the chatbot provides content suitable for primary school students is a non-negotiable requirement. To this end, we implemented multiple layers of content control within the system:

**Knowledge Base Design Control:** The RAG module relies primarily on official and verified educational documents (e.g., textbooks and standardized curricula), which have been reviewed by teachers or subject-matter experts. This approach ensures that the chatbot draws responses from a reliable and age-appropriate knowledge base, avoiding extraneous or unsuitable content [4].

**Prompt-based Control:** As described in Section 3.2, the system prompt also acts as a "soft filter," instructing the model to avoid sensitive, violent, vulgar, or otherwise inappropriate content. When a student poses an unsuitable question, the chatbot is guided to politely decline or redirect the conversation.

**Post-Processing Moderation:** After the LLM generates a response, the orchestration module performs

a content review. We integrated an automatic moderation system based on a blacklist of prohibited keywords and a content classification model. If any violation is detected (though rare due to earlier controls), the system either removes or modifies the content before displaying it to the student. For instance, if the model inadvertently uses overly complex language or references frightening or violent details, the post-processing step adjusts the response to ensure it remains safe and comprehensible for children.

**Teacher Oversight:** In actual deployment scenarios, the system allows administrative access for teachers to monitor logs of questions and responses. Teachers can configure additional rules or define custom blocked keywords based on classroom context. This approach is similar to the MerryQuery system, which enables teachers to set filters that prevent the generation of undesired content [4]. The inclusion of teacher supervision reinforces the reliability of the chatbot and ensures alignment with educational goals.

Taken together, these components enable the chatbot to harness the full potential of AI-powered language interaction while remaining within the necessary boundaries for use in primary education environments.

## 4. EXPERIMENTAL RESULTS

We implemented a prototype system based on the proposed architecture, utilizing a locally hosted Vietnamese LLM and integrating it with a RAG module. The knowledge base consisted of a collection of textbooks and exercises spanning Grades 1 through 5. The objective of this experiment was to preliminarily evaluate the chatbot's ability to answer students' academic questions and to assess whether its responses align with the pedagogical tone expected of a virtual primary school teacher.

### 4.1. Deployment Setup

The LLM employed was a fine-tuned variant of the LLaMA-3 family (8 billion parameters), augmented with Vietnamese language data to support fluent comprehension and response generation. The RAG knowledge base comprised approximately 5,000 knowledge items drawn from primary school subjects such as Mathematics, Vietnamese, Science, and Social Studies, as well as age-appropriate teacher-provided materials (e.g., explanations of basic scientific phenomena, extended reading passages). We used Faiss

to construct a vector index for the dataset, enabling real-time retrieval of relevant content. The system prompt was designed as described in Section 3.2 to ensure pedagogical tone. Additionally, we incorporated culturally specific prompts tailored to Vietnamese educational contexts for example, guiding the model to use the “thầy/cô - em” pronouns typical in classroom communication, and using familiar character names like Tí, Tèo, or Mai. The system ran on a GPU server with an internal API, providing response times of 5 - 10 seconds per query, which is acceptable for real-time interaction.

### 4.2. Sample Interaction Scenarios

The following examples illustrate how the AI chatbot performs in real-world use cases. The dialogues have been translated into English for clarity.

**Example 1 - Grade 3 Mathematics:** A student asks: “I don’t know what 12 divided by 4 is”. The system retrieves content from a Grade 3 math textbook and recognizes this as a basic division problem. The chatbot responds:

“Hello there! Let me help you: 12 divided by 4 means you have 12 apples and want to share them equally among 4 friends. Each friend will get 3 apples! Mathematically, that’s written as:  $12 \div 4 = 3$  or as a fraction:  $12/4 = 3$ . Now try this one (It’s not hard - I bet you can do it!): What is 16 divided by 4? Quick - try to imagine: if you have 16 oranges and divide them among 4 people, how many does each get? That’s your answer! ❤️”.

The chatbot provides a simple, age-appropriate explanation of the concept and follows up with a related question to reinforce learning - mimicking the encouraging tone of a real teacher

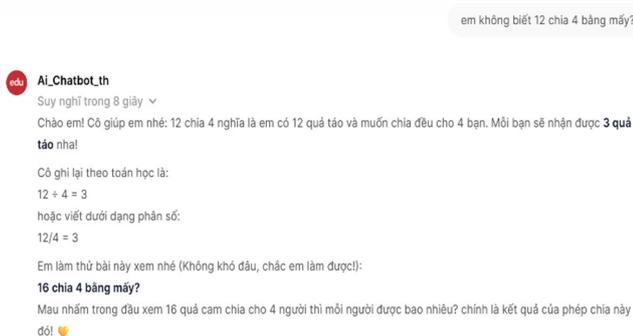


Figure 2. AI chatbot helps students solve 3rd grade math problems

**Example 2 - Continued Context: Multiplication Table Reinforcement:** Assuming the student incorrectly responds “16 divided by 4 equals 3,” the chatbot replies:

“Hi again! I see you gave it a try, which is great but  $16 \div 4$  isn’t quite correct. Let’s check it again: You have 16

oranges and divide them evenly among 4 people. Each person will get  $16 \div 4 = 4$  oranges (because  $4 + 4 + 4 + 4 = 16$ ). Here’s another tip to help you remember division by heart: From the 4-times table:  $4 \times 4 = 16 \Rightarrow 16 \div 4 = 4$ . Now try this: What is 20 divided by 5? I challenge you! (Hint:  $5 \times ? = 20$ ). Keep it up! If you make a mistake, don’t worry math gets easier with regular practice 😊”.

This interaction demonstrates the chatbot’s ability to diagnose a learning difficulty and respond with supportive, clear explanations and motivational reinforcement.



Figure 3. Example of AI Chatbot helping personalize elementary math learning

**Example 3 - Content Moderation and Child Protection:**

When a student asks a question deemed inappropriate for their age, such as: “Can you tell me a romantic love story?”, the chatbot identifies it as unsuitable for primary school students and responds:

“Hello dear! That’s a topic that isn’t quite right for primary school children. You’re still young let’s focus on exciting things like learning new things, exploring nature, playing with friends, and enjoying the love within your family.

When you grow older, you’ll understand more about love stories!

If you ever need help with learning (math, reading, English, storytelling...), I’m always here to support and cheer you on! 😊”.

This response reflects a nurturing and protective tone, similar to that of a responsible teacher, redirecting the conversation while maintaining engagement.

These examples illustrate that the chatbot is capable of delivering accurate, personalized responses based on students’ learning levels (Examples 1 and 2) through the integration of RAG and a curated knowledge base, while also maintaining a positive pedagogical demeanor

(Example 3) via prompt engineering and content moderation. We also observed that the chatbot handles queries well when they fall within the scope of the indexed knowledge. For out-of-scope questions (e.g., personal matters or topics outside the curriculum), the chatbot appropriately declines or redirects students to consult their actual teacher, thereby avoiding misinformation.



Figure 4. Example of AI Chatbot ensuring relevant content is delivered to students

### 4.3. Evaluation of Response Quality and Appropriateness

We conducted a preliminary evaluation of the chatbot's responses along two dimensions: (1) Content Accuracy - whether the responses were factually correct and adequately addressed the question. (2) Age Appropriateness - whether the language, tone, and content were suitable for primary school students and aligned with pedagogical norms.

The evaluation was carried out by five volunteer primary school teachers, using a set of 50 queries prepared by the research team, covering various subjects and typical learning situations. Results showed that approximately 90% of the responses were rated as accurate or acceptable in terms of content. Inaccurate cases were primarily associated with overly open-ended questions or content beyond primary school level (e.g., questions related to middle school knowledge). In such cases, the chatbot typically apologized and declined to answer or responded with a general statement - an acceptable behavior for maintaining safety over speculation.

In terms of style, most responses were deemed appropriate. Teachers described the chatbot's tone as "natural and human-like" and praised its "encouraging and pedagogically sound" qualities. A few cases involved vocabulary slightly complex for Grade 1 - 2 students (e.g., the term "diffusion" in Example 2), but these were accompanied by adequate explanations and thus did not present major issues. One takeaway is the need for

continued refinement of prompts and possibly additional training data to further simplify language for younger age groups.

Overall, the experimental results demonstrate the feasibility of the proposed system: the LLM+RAG chatbot can effectively provide personalized learning support to primary students with relatively high quality. Students perceive the chatbot as an "online teacher" who is always ready to assist, while actual teachers gain a valuable tool to supplement instruction and monitor individual learning progress.

## 5. CONCLUSION

This study proposed and implemented a prototype AI chatbot system designed to deliver personalized learning support for primary school students, built upon large language models (LLMs) in conjunction with retrieval-augmented generation (RAG) techniques. The system allows each student to engage in natural language interactions with a virtual "teacher," through which they can ask questions and receive tailored responses and guidance according to their individual learning levels.

Our primary contribution lies in designing an integrated LLM-RAG architecture that ensures both intelligence and trustworthiness in responses. We also developed prompt engineering strategies to emulate the tone and pedagogical approach of primary school teachers, while enforcing age-appropriate safety measures. Initial evaluations showed that the chatbot could accurately answer the majority of curriculum-aligned questions with a friendly, supportive tone closely mirroring teacher-student communication in real classrooms. Moreover, the system demonstrated potential to enhance learning engagement, offering students an always-available learning assistant while providing teachers with an auxiliary tool to facilitate instruction and track student progress.

Nevertheless, the study acknowledges ongoing challenges ensuring absolute accuracy in content, handling complex or unexpected student queries, and addressing ethical concerns in real-world deployments [6, 7, 10]. These present key directions for future research. We plan to scale the prototype to larger classroom deployments to evaluate its long-term impact on learning outcomes and student behavior. Concurrently, efforts will continue to minimize hallucinations, further personalize responses based on individual student profiles, and refine content moderation mechanisms.

Looking ahead, the combination of generative AI and domain-specific knowledge presents a promising foundation for the next generation of intelligent educational assistants. When responsibly designed and appropriately managed, systems like the LLM+RAG chatbot can redefine learning as a more personalized, engaging, and effective experience. More importantly, the adoption of AI in education must be guided by principles of equity and human-centeredness supporting and augmenting, not replacing, the irreplaceable role of human educators. With a cautious and collaborative approach among researchers, educators, and policymakers, we can harness the power of AI to provide future generations with the best possible learning opportunities in the digital age.

---

## REFERENCES

- [1]. Jakub Swacha, Michał Gracel, "Retrieval-Augmented Generation (RAG) Chatbots for Education: A Survey of Applications," *Applied sciences*, 18(8), 2025.
- [2]. Jussi S. Jauhiainen, Agustín Garagorry Guerra, "Generative AI and education: dynamic personalization of pupils' school learning material with ChatGPT," *Frontiers in Education*, 9 - 2024.
- [3]. Eugenio Tufino, "NotebookLM: An LLM with RAG for active learning and collaborative tutoring," *arXiv:2504.09720*, 2025.
- [4]. Benyamin Tabarsi, Aditya Basarkar, Xukun Liu, Dongkuan (DK) Xu, Tiffany Barnes, "MerryQuery: A Trustworthy LLM-Powered Tool Providing Personalized Support for Educators and Students," in *The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)*, 2025.
- [5]. Miladi F., Psyché V., Diattara A., El Mawas N., Lemire D., "Evaluating a GPT-4 and Retrieval-Augmented Generation-Based Conversational Agent to Enhance Learning Experience in a MOOC," in *Proceedings of the 17th Int. Conf. on Computer Supported Education (CSEDU 2025)*, 2025.
- [6]. Uğraş H., Uğraş M., Papadakis S., Kalogiannakis M., "ChatGPT-Supported Education in Primary Schools: The Potential of ChatGPT for Sustainable Practices," *Sustainability*, 16(22), 9855, 2024.
- [7]. Naznin K., Al Mahmud A., Nguyen M. T., Chua C., "ChatGPT Integration in Higher Education for Personalized Learning, Academic Writing, and Coding Tasks: A Systematic Review," *Computers*, 14(2), 53, 2025. <https://doi.org/10.3390/computers14020053>.
- [8]. Duong T. T. M., Da C. V., Nguyen V. H., "The use of ChatGPT in teaching and learning: A systematic review through SWOT analysis approach," *Frontiers in Education*, 9, 1328769, 2024. <https://doi.org/10.3389/educ.2024.1328769>.
- [9]. Sharma S., Mittal P., Kumar M., Bhardwaj V., "The role of large language models in personalized learning: a systematic review of educational impact," *Discover Sustainability*, 6, 243, 2025. <https://doi.org/10.1007/s43621-025-01094-z>.
- [10]. Davar N. F., Dewan M. A. A., Zhang X., "AI Chatbots in Education: Challenges and Opportunities," *Information*, 16(3), 235, 2025. <https://doi.org/10.3390/info16030235>.