A chatbot to help promoting financial literacy

Davi Silva Eleuterio
UNIAG, Instituto Politécnico de Bragança, 5300-253 Bragança, Portugal

Pedro Filipe Oliveira
CeDRI, Instituto Politécnico de Bragança, 5300-253 Bragança, Portugal

Paulo Matos
CeDRI, Instituto Politécnico de Bragança, 5300-253 Bragança, Portugal

Jorge Manuel Afonso Alves
UNIAG, Instituto Politécnico de Bragança, 5300-253 Bragança, Portugal

Abstract

Currently, governments and many other institutions have been making significant efforts to promote financial literacy. However, a considerable portion of the population still lacks basic financial knowledge, highlighting the need for updated strategies to enhance financial education. In today's digital world — where people often search for quick and convenient solutions — the development of a reliable and intelligent chatbot to answer questions related to financial concepts and decision-making can be very beneficial. This paper proposes the implementation of an automated web scraper to extract content from a trustworthy financial education website with plenty of useful concepts about finances, using this collected data to develop a chatbot which provides accurate and helpful responses to users. The solution is built using technologies such as Streamlit, Langchain, and OpenAI.

2012 ACM Subject Classification Applied computing → Document analysis; Computing methodologies → Information extraction; Computing methodologies → Machine learning

Keywords and phrases chatbot, financial literacy, web scraper, LLM, RAG

Digital Object Identifier 10.4230/OASIcs.SLATE.2025.7

Funding This work was supported by national funds through FCT/MCTES (PIDDAC): CeDRI, UIDB/05757/2020 (DOI: 10.54499/UIDB/05757/2020) and UIDP/05757/2020 (DOI: 10.54499/UIDP/05757/2020); and SusTEC, LA/P/0007/2020 (DOI: 10.54499/LA/P/0007/2020).
This work was supported by national funds through FCT/MCTES (PIDDAC): UID/04752 – Applied Management Research Unit (UNIAG).

1 Introduction

Rehman and Mia [8] define financial literacy as the combination of understanding, abilities, and confidence that enables individuals to make effective financial decisions. It includes comprehending core financial concepts and utilizing them practically. But unfortunately, a significant portion of the population lacks this fundamental financial understanding, contributing to difficulties in managing debt and resulting in over-indebtedness.

In parallel with the growing digitalization of the financial world, the ways that people access information are evolving rapidly. This change highlights the importance of using new technologies to enhance financial literacy, particularly among younger generations. Traditional approaches involving complex texts or extensive research are often less utilized by individuals accustomed to getting instant information from sources such as artificial intelligence (AI) and social media short videos.

Chatbot technology has been used across many fields of knowledge and numerous articles address its successfulness in conducting quick research and providing trustworthy responses.

The ability of chatbots to be integrated into popular social media platforms can also highly increase the use of the tool, making it convenient for the user [7].

We propose that a specialized financial chatbot can effectively solve the financial literacy gap by providing accessible and engaging financial education. By delivering accurate information through simple, conversational interactions, such a tool can empower users to make informed decisions, raise awareness about credit risks, and build essential financial knowledge.

This paper explores the design and potential of a financial chatbot specifically built to leverage a reliable data source. Our aim is to demonstrate how providing accurate and practical financial knowledge through user-friendly conversations can reduce the incidence of factual inaccuracies and misleading information often generated by AI, thereby significantly improving users' understanding of financial concepts.

2 State of the Art

Artificial intelligence has experienced much momentum in recent years due to the advancements in hardware and software technologies, specially after the release of ChatGPT in 2022. This paradigm of Generative AI, which leverages advancements in Natural Language Processing (NLP) and is frequently powered by massive Large Language Models (LLMs), has enabled the creation of highly interactive and versatile applications. Among these, advanced chatbots capable of understanding and generating human-like text represent a significant area of development and application [5].

(a) Definition of Chatbot

A chatbot is an AI system and a Human-computer Interaction (HCI) model, which uses NLP and sentiment analysis to communicate in human language by text or oral messages with humans or other chatbots. Interactive agents, artificial conversation entities, smart bots and digital assistants are examples of chatbots among the internet, acting as a powerful tool along many applications, such as education, business, e-commerce, healthcare and many others [1].

(b) Brief History of Chatbots

The development of chatbot agents draws upon foundational ideas from early computational and mathematical concepts. One early mathematical model relevant to sequence prediction, a key element in generating responses, was the Markov Chain, developed by Russian mathematician Andrey Markov in 1906. This statistical model for predicting random sequences has been utilized in machine learning fields for tasks such as autocomplete and next-word prediction for many years.

Another pioneer milestone regarding machine intelligence capable of human-like interaction was the Turing Test, proposed by Alan Turing in 1950. Turing, often regarded as a father of theoretical computer science and artificial intelligence, introduced this test to assess a machine's ability to exhibit intelligent behaviour indistinguishable from that of a human. The test involves a human interrogator conversing separately with a hidden machine and a hidden human. The machine passes if the interrogator cannot reliably determine which is the machine. The Turing Test significantly influenced subsequent artificial intelligence research and spurred efforts towards creating machines capable of natural language conversation [12].

In recent times, the most impactful development in the history of chatbots has been the emergence of systems like ChatGPT, powerfully demonstrating the capabilities of Large Language Models (LLMs), such as GPT-3.5 and GPT-4 [4]. Trained on massive datasets, these models achieve robust language skills, enhanced reasoning, and strong contextual understanding, enabling coherent multi-turn dialogues. ChatGPT, in particular, made this cutting-edge technology widely accessible for public interaction and evaluation.

Underpinning these capabilities is typically the Transformer neural network architecture [11]. Their effectiveness stems from a two-stage training process: initial unsupervised pre-training on vast text to build general knowledge, followed by supervised fine-tuning on dialogue data to hone conversational performance. This methodological approach, leveraging the Transformer's design, is fundamental to their sophisticated natural language generation, with research actively analysing current advancements and trends [9].

(c) Use of Chatbots in Education

AI has been used in the domain of education for over 40 years in different shapes and forms, supporting school administration, teachers and students in different applications [2]. Educators can use them for developing instructional materials and assessments, although responsible application is essential to encourage critical thinking among students. For learners, these tools can promote equity through more accessible information and aid in providing effective learning strategies adapted to diverse preferences [2]. They can also assist with assessing student submissions and suggesting pedagogical improvements. However, it is important to recognize that the current iteration of tools like ChatGPT still has functional limitations, may contain factual inaccuracies, and cannot replicate humans' ability to provide truly differentiated, specific instruction for every student [5].

Research indicates that effective student learning outcomes are strongly linked to personalized support, while insufficient individual attention can hinder academic progress. Methods like micro-learning have been shown to alleviate student fatigue [10] and improve material retention, contributing positively to comprehension, skill development, and overall academic performance. Within this context, chatbots are suggested as valuable tools for e-learning environments. They can serve as interactive tutors, managing student inquiries and providing feedback, and potentially facilitate communication with families regarding a child's learning. By instantly responding to common questions, chatbots enhance the accessibility and comfort for students seeking assistance, making learning more engaging [6].

However, the widespread adoption of these tools also introduces challenges, notably concerns about student misuse leading to academic integrity issues [3]. In response, educational institutions are implementing diverse strategies to prevent such misconduct. Approaches range from simple prohibition of certain tools to revising and updating existing policies. For example, several traditional universities, including some within the UK's Russell Group, classify the unauthorized use of AI bots as academic misconduct. Furthermore, adapting assessment methods is recommended, such as designing tasks less susceptible to AI assistance by incorporating unique content, or returning to traditional formats like written exams instead of only computer-based testing [3].

3 Materials and Methods

3.1 Data Sources and Tools

The study used financial articles obtained from the website *todoscontam.pt*, which is an initiative from the government of Portugal to promote financial literacy through the National Plan for Financial Education. Since these articles are normally unstructured and contain both textual and visual information, a method to automatically extract the useful data was required. The main tools used in this paper include:

- The pages and subpages from the website *todoscontam.pt*
- Python-based programming for web scraping and chatbot development
- OpenAI's ChatGPT for interpreting and analysing financial information content.
- LangChain and Pinecone (vector database) for Retrieval-Augmented Generation (RAG).
- Streamlit framework for building an interactive web interface.

3.2 Methodology

This project was structured to achieve two main objectives:

- Extract and storage data from a reliable website containing foundational financial concepts.
- Implement a chatbot which answer questions about finances using the data extracted as basis.

To address these objectives, we implemented a web scraper using python programming and the chatbot itself:

(a) Web Scraper

One of the primary objectives of this chatbot is to ensure reliability. Because the information provided by large language models (LLMs) is not always fully trustworthy, it is crucial to source data from reliable and verifiable sources. To implement that, we manually collected all the URLs of useful articles regarding financial literacy from the website *todoscontam.pt*. Then, the web scraper is structured as follows:

- Selection of articles: the pages selected from *Todos Contam* website are only the ones which provide financial concepts or decision-making information, ignoring other pages containing non-relevant content, for example pages regarding contact, about the website and others. The URLs collected were inserted in a tuple containing the URL and the destination path. This path and sub paths are respectively the theme and sub themes of the article.
- Downloading and filtering content: for each URL iterated, the system downloads the Hyper-text Markup Language (HTML) content of the article. This HTML code contains many unwanted information, such as navigation bar, tab bar, footer content and its own HTML tags. At this step, every useless information is removed, and the final article text is finally converted to a markdown file.
- File and folder management: in this process, the title of the article is determined as the name of the file, and the folder where the file is stored is named with the theme of the article. A total of 222 documents were stored, with an average size of 2.75 kilobytes.

The diagram presented at figure 1 illustrates all the functioning of the web scraper and summarize each step.

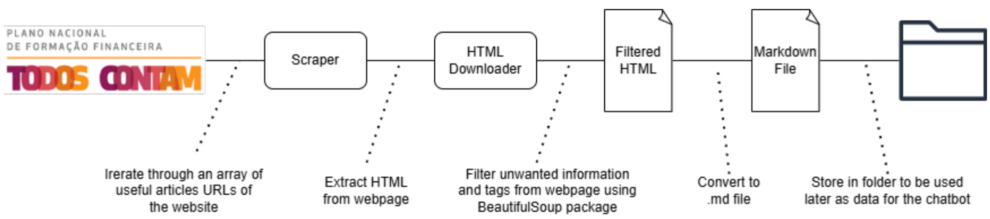

Figure 1 Web scraper.

(b) Chatbot Development

Following the extraction of financial literacy articles in Markdown format, a specialized chatbot was developed to provide users with information derived solely from this corpus. The chatbot leverages a Retrieval-Augmented Generation (RAG) architecture, combining the retrieval capabilities of a vector database with the generative power of a Large Language Model (LLM). The implementation utilizes Python, primarily employing the Langchain framework for orchestrating the RAG pipeline, Streamlit for the user interface, and Pinecone as the vector database.

a. Core components and configuration: the system is built upon several key libraries. Langchain-openai provides interfaces for OpenAI's models, specifically using ChatOpenAI with the gpt-4.1-nano model for response generation and OpenAIEmbeddings with the text-embedding-3-small model (1536 dimensions) for creating text embeddings. Langchain-pinecone facilitates interaction with the Pinecone vector database. Streamlit's caching mechanism (@st.cache_resource) is employed to efficiently manage resource-intensive objects like the LLM, embedding models, and the Pinecone connection. Figure 2 illustrates the chatbot functioning.

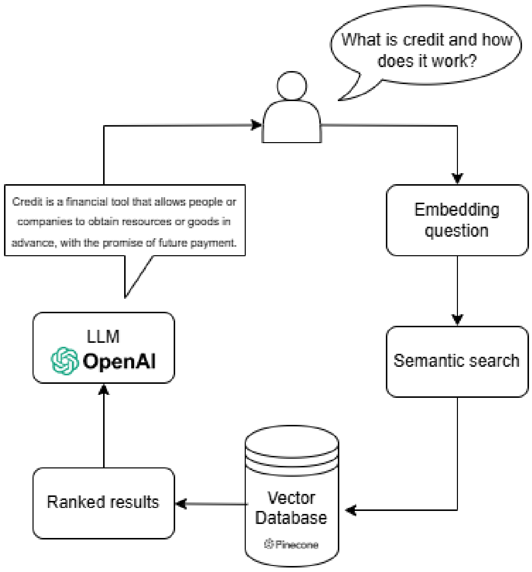

Figure 2 Chatbot operating architecture.

b. Data preparation and indexing: the initial step involves loading the Markdown documents from the specified directory, including subdirectories. The DirectoryLoader from Langchain, configured with UnstructuredMarkdownLoader, handles this process. A crucial preprocessing step adds metadata to each document, identifying the financial theme based on the name of the subfolder from which the document originated.

c. Retrieval-augmented generation pipeline: the core question-answering functionality is implemented using Langchain's RetrievalQA chain. This chain integrates the LLM with a retriever built upon the Pinecone vector store. The retriever is configured to perform similarity searches and retrieve the top 5 most relevant document chunks based on the vectorized user query.
A custom PromptTemplate is employed to guide the LLM's response generation. This template incorporates:

- A static context defining the chatbot's persona as a helpful financial assistant based on *todoscontam.pt* content, instructing it to use only retrieved information, state when information is unavailable, and avoid fabrication.
- A placeholder for the context retrieved from the vector database.
- A placeholder for the user's question.

The RetrievalQA chain is configured to insert all retrieved text chunks directly into the context placeholder of the prompt. The chain is also set to return the source documents used for generation.

d. User interface and interaction: a web-based user interface is provided using Streamlit. It displays a chat interface where users can interact with the chatbot. Chat history is maintained using Streamlit's session state. When a user submits a query, the query is processed by the system. The response generated by the LLM is displayed, along with an expandable section listing the source documents (filename and theme) that the retriever identified as relevant for answering the query.

4 Results

To assess the chatbot's performance in an initial phase, a set of 48 questions spanning six key financial literacy themes was curated, reflecting common user queries in this domain. These themes, derived from the structure of the *Todos Contam* portal, include:

1. Personal and Family Budgeting
2. Savings and Basic Investment
3. Credit and Debt
4. Essential Banking Products
5. Basic Insurance
6. Consumer Rights and Duties in Finance

Each question was presented to the developed chatbot, and the corresponding answer was recorded and analysed. The evaluation focused on the chatbot's ability to adhere to its core instruction: provide accurate and relevant answers based solely on the retrieved document context from the scraped articles, and explicitly state when the requested information was not available within that knowledge base. Answer quality was assessed based on correctness relative to the source material, relevance to the question, and appropriate handling of knowledge gaps.

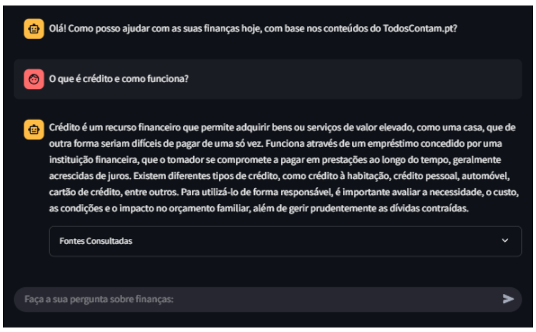

Figure 3 Example of interaction 1.

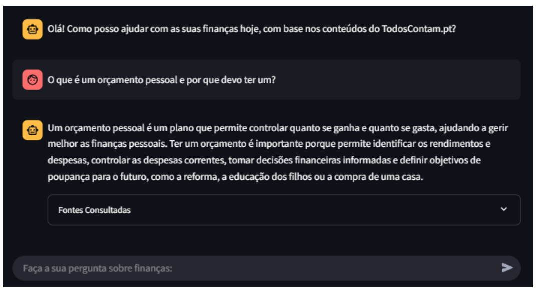

Figure 4 Example of interaction 2.

4.1 Findings

The analysis of the 48 question-answer pairs revealed a high degree of adherence to the provided knowledge base and instructions. The chatbot consistently utilized the information present in the scraped articles to formulate responses.

Table 1 summarizes the performance across the different themes. An answer was considered successfully addressed if it either provided a correct response based on the source material or correctly identified that the information was not present in the knowledge base.

Table 1 Summary of Chatbot Performance by Theme

Theme	Correctly Answered	Identified as Missing
1. Budgeting	8	0
2. Savings & Investment	7	1
3. Credit & Debt	6	2
4. Banking Products	7	1
5. Insurance	8	0
6. Consumer Rights	6	2
Total	**42**	**6**

Overall, the chatbot successfully addressed all 48 questions according to the evaluation criteria. In 42 cases, it provided answers directly derived from the scraped content. In the remaining 6 cases, it correctly identified that the specific information requested was either entirely absent from the source documents or that the source documents lacked the requested level of detail, thereby demonstrating its ability to recognize the boundaries of its knowledge base. This corresponds to a 87.5% success rate in information retrieving and 100% success rate in terms of adhering to the operational instructions and leveraging the provided corpus.

4.2 Discussion

The results indicate that the RAG-based chatbot effectively functions as an information retrieval and summarization tool for the specific knowledge base derived from *Todos Contam*. Its main strength lies in its fidelity to the reliable source material, ensuring that users receive information aligned with the content curated by the Bank of Portugal and the National Plan for Financial Education. The explicit prompting to rely solely on retrieved documents and acknowledge gaps proved successful. It was also observed that the chat provides responses with consistent structure and content when asked the same question multiple times, varying only in wording.

However, the chatbot's primary limitation is its dependence on the scraped content. It cannot answers questions outside the scope of the articles available on *Todos Contam* or provide broader financial context, comparisons with products not mentioned, or real-time market data. While it correctly identifies knowledge gaps, it cannot bridge them without additional information sources. Furthermore, the quality and comprehensiveness of the answers are inherently limited by the quality and depth of the original articles.

4.3 Proposed Solutions and Future Work

To enhance the chatbot's utility and address its limitations, several avenues for future work are proposed:

- **Expand Knowledge Base:** Incorporate content from other authoritative Portuguese financial sources, such as the websites of the CMVM (Securities Market Commission), ASF (Authority for the Supervision of Insurance and Pension Funds), the Bank of Portugal's main site, and relevant government portals (e.g., related to taxes or social security). This would broaden the range of topics the chatbot can address.
- **Incorporate Structured Data:** Explore adding structured data (e.g., current interest rates for specific products, calculators) via APIs or separate databases to provide more dynamic and tool-like functionalities.
- **User Evaluation:** Conduct user studies with the target audience to gather qualitative feedback on the chatbot's usability, understandability, and perceived usefulness in improving financial literacy. This feedback can guide further development iterations.

By expanding the knowledge base and potentially refining the RAG pipeline, the chatbot's potential as a comprehensive financial literacy tool can be significantly enhanced.

5 Conclusions

Financial literacy remains a critical challenge, with many individuals lacking the necessary knowledge to make informed financial decisions in an increasingly complex digital world. Traditional educational methods often struggle to engage audiences accustomed to readily available, interactive digital content. This paper addressed this gap by proposing and implementing a specialized chatbot designed to enhance financial education in Portugal.

We presented the development process, which involved automated web scraping of reliable content from the *Todos Contam* website, followed by the construction of a chatbot using a Retrieval-Augmented Generation (RAG) architecture. Leveraging tools like Langchain, OpenAI models, Pinecone, and Streamlit, we created an interactive system capable of answering user queries based on the curated financial knowledge base.

The evaluation demonstrated the chatbot's effectiveness within its defined scope. It successfully addressed a diverse set of 48 questions across key financial themes, either by providing accurate information derived strictly from the source material or by correctly identifying when the requested information was unavailable in its knowledge base. This highlights the system's fidelity to the reliable source and its adherence to operational instructions, crucial aspects for building trust in AI-driven financial tools.

In conclusion, this work demonstrates the viability of using a RAG-based chatbot, grounded in verified information sources, as a tool to promote financial literacy. By providing accessible, reliable, and interactive financial education, such systems hold significant potential to empower individuals and contribute to better financial well-being. The emphasis on using curated, trustworthy data sources is paramount when developing AI applications for sensitive domains like personal finance.

References

1 Eleni Adamopoulou and Lefteris Moussiades. Chatbots: History, technology, and applications. *Machine Learning with Applications 2*, 2021.

2 Ilaha Ashrafova. Education and chatbots: New opportunities for teachers and students. *Journal of Azerbaijan Language and Education Studies*, 2(2), 2025.

3 Ali Ateeq, Mohammed Alzoraiki, Marwan Milhem, and Ranyia Ali Ateeq. Artificial intelligence in education: implications for academic integrity and the shift toward holistic assessment. *Frontiers in Education*, 9, 2024.

4 Chiranjib Chakraborty, Soumen Pal, Manojit Bhattacharya, Snehasish Dash, and Sang-Soo Lee. Overview of chatbots with special emphasis on artificial intelligence-enabled ChatGPT in medical science. *Sec. Medicine and Public Health*, 2023.
5 Faisal Kalota. A primer on generative artificial intelligence. *Education Sciences*, 2024.
6 Chee Ken Nee, Mohd Hishamuddin Abdul Rahman, Noraffandy Yahaya, Nor Hasniza Ibrahim, Rafiza Abdul Razak, and Chie Sugino. Exploring the trend and potential distribution of chatbot in education: A systematic review. *International Journal of Information and Education Technology*, 13(3), 2023.
7 Reshawn Ramjattan, Patrick Hosein, and Nigel Henry. Using chatbot technologies to help individuals make sound personalized financial decisions. *2021 IEEE International Humanitarian Technology Conference (IHTC)*, 2021.
8 Khurram Rehman and Md Aslam Mia. Determinants of financial literacy: a systematic review and future research directions. *Future Business Journal*, 10, 2024.
9 Chen Renzhang and Zhao Haixia. AI chatbot research: A bibliometric analysis of advancements and trends. *Journal of Macau University of Science and Technology (Humanities and Social Sciences)*, 19, 2025.
10 M. S. Shail. Using micro-learning on mobile applications to increase knowledge retention and work performance: A review of literature. *Cureus*, 11(8):e5307, 2019. doi:10.7759/cureus.5307.
11 Emma Yann Zhang, Adrian David Cheok, Zhigeng Pan, Jun Cai, and Ying Yan. From turing to transformers: A comprehensive review and tutorial on the evolution and applications of generative transformer models. *Sci*, 2023.
12 Michal Černý. The history of chatbots: the journey from psychological experiment to educational object. *Journal of Applied Technical and Educational Sciences*, 12(3), 2022.