

Use of a Retrieval-Augmented Generation (RAG) Chatbot in an Online R Programming Course

Guido Lang
guido.lang@quinnipiac.edu

Tan Gürpınar
tan.gurpinar@quinnipiac.edu

Business Analytics & Information Systems
Quinnipiac University
Hamden, CT USA

Abstract

Given the increasing complexity of programming education and the need for personalized learning tools, this study explores how artificial intelligence (AI) chatbots can enhance student understanding and engagement in a self-paced online environment. Specifically, the study examines the use of a Retrieval-Augmented Generation (RAG) chatbot in an online R programming course. The course chatbot was designed to provide students with conceptual answers and explanations based on the course materials. Student interactions with the chatbot and their perceptions of its usefulness were analyzed, using log data from the chatbot and survey responses collected at the end of the course. The findings reveal that students with higher prior knowledge of R were more likely to use the chatbot. The most frequent use of the chatbot was for advanced topics that went beyond the lecture content. Most students that didn't use the chatbot simply didn't have a need for it. Students were overall very satisfied with the chatbot and appreciated its ability to give helpful explanations. This study contributes to understanding the role of AI chatbots in online education and offers insights for future implementations, particularly in how to enhance personalized learning experiences and support self-paced learning.

Keywords: Chatbots, Retrieval-Augmented Generation (RAG), Online Education, R Programming, Student Support, AI in Education

1. INTRODUCTION

Online programming courses often face unique challenges in providing adequate support to students (Lang & D. O'Connell, 2015; Napalit et al., 2023). Unlike traditional classroom settings, students in online courses lack immediate access to instructors for real-time assistance. This can lead to delays in addressing students' questions and difficulties, potentially hampering their learning progress (Hurlbut, 2018). Additionally, the asynchronous nature of many online courses means that students may be working on course materials at different times, further complicating the provision of timely support (Varkey et al., 2023).

General AI chatbots, such as ChatGPT, Copilot, Claude, Gemini, generate responses based on proprietary data and human-reinforcement learning. These chatbots use deep learning models, particularly transformer architectures, to understand and generate human-like text (Meyer et al., 2023). They rely on patterns learned from online datasets to predict the most probable next word or phrase in a conversation. While they are effective in generating coherent and contextually relevant text, their responses are limited to the information they have been trained on and do not automatically update with new data or specific documents (Jungheer, 2023). The technology that powers AI chatbots, i.e., large language models,

has been shown to have useful applications in information systems education, specifically for writing of teaching cases (Lang et al., 2024).

Retrieval-Augmented Generation (RAG) chatbots integrate retrieval mechanisms with their generation capabilities to enhance the accuracy and relevance of their responses. These chatbots first retrieve pertinent information from a predefined knowledge base or a set of documents before generating a response (Jeong, 2023; Maryamah et al., 2024). The retrieved information is then used to inform the generation process, ensuring that the responses are not only coherent but also contextually precise and up-to-date. This approach is particularly beneficial in educational settings, where the need for accurate, specific, and context-aware answers is crucial. By leveraging both retrieval and generation, RAG chatbots can provide targeted support, addressing student queries with a higher degree of relevance and precision (Pichai, 2023).

Given the distinct advantages of RAG chatbots in delivering precise and contextually relevant information, their use in educational settings presents a promising avenue for enhancing student support. This transition from general AI chatbots to more sophisticated RAG systems raises several pertinent questions regarding their effectiveness and acceptance among students. Specifically, understanding the factors that influence students' engagement with such chatbots, the reasons behind their non-usage, and the types of queries they find most helpful can provide valuable insights. Additionally, gauging students' overall perceptions and attitudes towards the chatbot can further elucidate its impact on their learning experience (Labadze et al., 2023). In this context, the current study aims to address the following research questions:

- RQ1: What factors affect students' likelihood to use a course chatbot?
- RQ2: What are the reasons why students don't use a course chatbot?
- RQ3: What do students actually use a course chatbot for?
- RQ4: What do students think of a course chatbot?

2. RELATED LITERATURE

Several studies have explored the use of AI chatbots in educational contexts. The recent meta-study by Labadze et al. (2023) emphasize numerous benefits of integrating chatbots in teaching, as seen from both students' and

educators' perspectives. They found that educators primarily save time and gain in an improved pedagogy while students see advantages in three key areas: homework and study assistance, a personalized learning experience, and the development of emerging skills. Still, there are hurdles in operationally implementing AI chatbots and not enough insights into their impact on specific skills (Labadze et al., 2023). This motivates our research to focus on the implementation of AI chatbots in connection to programming skills.

For this field, our research builds on further studies that already explore the use of chatbots in programming education. For instance, a study by Ait Baha et al. (2023) examine the integration of AI chatbots in e-learning systems designed to support students' programming skills. Conducted in a Moroccan public college where French is the primary foreign language, the chatbot effectively enhanced students' educational experiences by guiding instead of simply delivering answers. The chatbot's interactive nature provided students with timely answers, reducing uncertainty. The study suggests that future research should involve different learning contents and contexts to determine the specific tasks for which students use the chatbot and where they benefit the most (Ait Baha et al., 2023).

Furthermore, the implementation of a chatbot assisting in Python programming further demonstrates the positive impact of AI tools in programming education. Created using the SnatchBot API, the chatbot was found to be user-friendly and effective in simplifying programming logic and enhancing students' Python skills. The tool's success during the COVID-19 pandemic underscores its potential in supporting online learning environments. Future work proposed by the authors includes the consideration of additional programming languages and tests in other universities (Chinedu & Ade-Ibijola, 2021).

In another study, Kazemitabaar et al. (2024) deploy a chatbot in a C programming course, providing valuable insights into the varying performance of different LLM models across various programming languages and their impact on student engagement and learning outcomes (Kazemitabaar et al., 2024). Drawing from expert interviews, they also develop design recommendations for chatbots to allow for transparent and controlled use, which have been considered in this research.

Lastly, a study by Vukojičić and Krstić (2023) explores the influence of ChatGPT on student

work, particularly in enhancing code commenting practices and promoting uniform writing styles. This uniformity improves code readability and maintainability, fostering better comprehension and collaboration among peers and instructors. The study highlights that ChatGPT not only hones individual coding skills but also contributes to a more efficient and effective learning environment (Vukojić & Krstić, 2023). As a result, ChatGPT helped the students to produce code that is easier to understand, modify, and build upon and confirmed a recent Reddit survey that recommends OpenAI's ChatGPT as an effective LLM to assist in R programming (https://www.reddit.com/r/rstats/comments/1d0glz3/best_chatbot_for_r_programming/).

These studies collectively demonstrate the potential and current limitations of AI chatbots in programming education. They highlight the importance of model accuracy, the positive impact on coding practices, and the interactive support provided to students. Our research aims to build on these findings by specifically focusing on the implementation of a RAG chatbot in an online R programming course. This focus addresses the unique challenges and opportunities presented by R programming and aims to provide a more tailored and effective educational tool for students in this domain.

3. METHODOLOGY

The first author developed and implemented RAG course chatbot for use in his graduate-level, seven-week, asynchronous, online course on "Business Data Analytics with R". The course consists of video lectures, which focus on coding demonstrations in R, and hands-on assignments such as exercises and lab projects, which are based on the video lectures. The course doesn't use a textbook and all course materials were developed by the first author. The course is part of the MS in Business Analytics program offered at the authors' university. While the course doesn't have any prerequisites, it is assumed that students in the program have knowledge of undergraduate statistics.

The RAG course chatbot is a Python web application that runs on Streamlit and uses LlamaIndex (for retrieval-augmentation) and OpenAI's GPT-4 (for generation). The course chatbot uses the source code provided by Frasca et al. (2023). All lecture materials, including the code used in the lectures and exercises, was ingested by the chatbot. By utilizing RAG, the course chatbot was able to respond using examples from the lectures, referencing

exercises, and following the professor's code style. The prompt instructed the chatbot to focus on explaining R programming and statistical concepts based on the course materials, which were sufficiently encompassing since the course didn't use a textbook. The full prompt along with a screenshot of the chatbot's user interface can be found in Appendix A. The course chatbot was publicly available online without user authentication for the duration of the course. All user queries were saved verbatim in a log file.

At the beginning of the course, students were informed of the availability of the course chatbot. Usage of the course chatbot wasn't required nor was it promoted again throughout the course.

The logs from the course chatbot contain the verbatim queries that were entered by the students into the course chatbot. A total of 80 queries were made by the students. The logs were content analyzed by the authors using a grounded theory approach (Strauss & Corbin, 1990). After an initial pass-through, the authors decided on a set of categories that represent all queries. The authors then collaboratively categorized each of the 80 queries into one category.

At the end of the course, students completed an anonymous survey on Qualtrics that collected demographic information and measured their use of, and attitudes towards, the course chatbot. Similar to the logs, the open-ended questions were also content analyzed by the authors using a grounded theory approach (Strauss & Corbin, 1990). Again, after an initial pass-through, the authors agreed on a set of categories that represent all responses and subsequently categorized each response accordingly. Students were offered extra credit (worth about 1.50 % of the final grade) in exchange for participation in the survey. The full survey can be found in the Appendix B. This study was approved by [university name withheld for review]'s Institutional Review Board under protocol number 03624.

4. RESULTS

Of the 40 students enrolled in the course, 38 (95%) completed the survey. Of the 38 students that completed the survey, most ($n = 27$, 71.05%) are in the 18-24 years age range, followed by the age ranges of 25-34 years ($n = 6$, 15.79%), 35-44 years ($n = 4$, 10.53%), and 55-64 years ($n = 1$, 2.63%). The majority ($n = 23$, 60.53%) is female. In terms of employment status, 17 (44.74%) are employed part-time, 13

(34.21%) are employed full-time, and 8 (21.05%) are not employed. None of the demographic factors are associated with any of the following results.

What Factors Affect Students' Likelihood to Use a Course Chatbot?

Of the 38 study participants, 16 (42.10%) used the course chatbot at least once. Prior use of AI chatbots (such as ChatGPT, Copilot, Gemini, Claude) is not associated with course chatbot use ($t = -1.04, p = 0.31$). However, prior knowledge of R is significantly related to course chatbot use. Specifically, students that used the course chatbot had more prior knowledge of R ($M = 2.18, SD = 0.98$) than students that didn't use the course chatbot ($M = 1.45, SD = 0.67, t = -2.58, p = 0.02$). Figure 1 depicts the difference in prior knowledge of R between the groups of non-users and users.




Figure 1: Average prior R knowledge by chatbot usage

What do Students Actually Use a Course Chatbot for?

Based on the log analysis of the 80 queries, it was found that almost half of the queries ($n = 35, 43.75%$) can be categorized as "advanced," meaning they are queries about topics that go beyond the topics covered in a lecture. An example query in this category is "Is difftime number data type treated any differently than number data type?" (note that this wasn't covered in a lecture). Over a quarter of the queries ($n = 23, 28.75%$) can be categorized as "lecture", meaning they are queries about topics that are covered in a lecture. An example query in this category is "What functions can you use to aggregate data?" (note that this was covered in a lecture). The third most frequent category of queries is "error" ($n = 12, 15.00%$), meaning

they are queries about errors encountered when writing code for the lectures and/or assignments. An example query in this category is "Explain error Quitting from lines 56-71 [unnamed-chunk-4] (Final-Project_Code.spin. Rmd)". One tenth of the queries ($n = 8, 10.00%$) are part of the "prerequisite" category, which means they are queries about topics that are assumed to be prerequisite knowledge for the course. An example query in this category is "What is a p value?". Lastly, 2 queries (2.50%) were categorized as "other" because they're unrelated to any of the previous categories. These queries were "How do you work as a chat bot, do you send an api request to gpt-4 or something?" and "Thank you!". The distribution of number of queries per category is depicted in Figure 2.




Figure 2: Distribution of query categories

What are the Reasons Why Students Don't Use a Course Chatbot?

Of the 22 survey participants who indicated that they didn't use the course chatbot, 16 (72.73%) provided a reason as to why they didn't use the course chatbot. The most frequent reason provided by 8 (50%) students was a lack of need for the course chatbot because the video lectures recorded by the professor were sufficient. The second most frequent reason mentioned by 6 (27.27%) students was an unawareness of the existence of the course chatbot. Lastly, 2 (12.5%) students stated they didn't use the course chatbot because they preferred to challenge and troubleshoot errors themselves. Figure 3 depicts the distribution of reasons for non-usage of the course chatbot.




Figure 3: Distribution of reasons for non-Usage of chatbot

What Do Students Think of a Course Chatbot?

Students were asked to indicate their agreement with a statement measuring the extent to which the course chatbot enhanced their learning experience. About 87% somewhat or strongly agreed, suggesting that students felt very positively about the course chatbot. Figure 4 depicts the distribution of responses to this item.




Figure 4: Responses to "The course chatbot enhanced my learning experience."

Additionally, students were asked to indicate their agreement with a statement measuring the extent to which they wished course chatbots would be used in more courses. Similarly, about 87% somewhat or strongly agreed, suggesting again that students felt very positively about the course chatbot. Figure 5 depicts the distribution of responses to this item.




Figure 5: Responses to "I wish course chatbots would be used in more courses."

Not surprisingly, students' responses to these questions were highly positively correlated ($r = .88, p < 0.001$). In other words, the more students felt the chatbot enhanced their learning experience, the more they wished other courses would use it, too. Lastly, 3 students (8.33%) mentioned the course chatbot positively and unprompted in the end-of-semester course evaluation, further indicating a positive reception by students.

Finally, students were asked what they liked about the course chatbot. Eleven (68.75%) students gave an answer. Among the answers given, the most frequent answer provided by 5 (45.45%) students was that the course chatbot provided helpful explanations. Three (18.75%) students each liked that the course chatbot was aligned with the course materials and that the course chatbot was easy to use. Figure 6 depicts the distribution of responses




Figure 6: Distribution of responses to what students liked about the chatbot

5. DISCUSSION AND CONCLUSION

This study examined the implementation and impact of a RAG chatbot in an online R programming course. Our findings indicate that prior knowledge of R significantly influences students' likelihood of using the chatbot. Somewhat surprisingly, students with higher prior knowledge of R were more likely to use the chatbot. Students primarily used the chatbot for advanced queries and lecture-related questions. Feedback from students was overwhelmingly positive, highlighting the chatbot's helpful explanations along with many students expressing a desire for similar chatbots in other courses.

Contributions

From a theoretical perspective, this study enriches the body of knowledge on AI-enhanced learning environments by demonstrating the efficacy of RAG chatbots in delivering targeted educational support. It provides empirical evidence on how RAG chatbots, which blend retrieval and generative capabilities, can effectively address the specific needs of students in online courses. Moreover, this research offers a nuanced understanding of how advanced AI technologies can be leveraged to enhance learning experiences, thereby contributing to the theoretical understanding surrounding AI and education.

On the practical side, the findings from this study offer actionable insights for educators and educational institutions looking to integrate chatbots into their online course offerings. It highlights the practical benefits of using RAG chatbots to provide immediate and context-aware support to students, thereby reducing the instructional burden on educators and enhancing the overall student learning experience. The findings suggest that RAG chatbots can effectively assist with advanced queries and lecture-related questions, making them a valuable tool for supporting student learning.

Limitations

This study has several limitations that should be considered when interpreting the findings. First, the survey data and the chatbot logs were collected independently and anonymously, thus making a combined analysis impossible. As a result, this study couldn't draw inferences between actual chatbot use and attitudes about the chatbot. Second, the relatively small sample size limits the generalizability of the results. With only 38 student participants, the findings may not be representative of larger or more diverse

student populations. Third, the voluntary nature of chatbot usage introduces potential self-selection bias. Students who chose to use the chatbot may have distinct characteristics or motivations compared to those who did not, which could influence the study's outcomes. This bias may affect the validity of the conclusions drawn regarding the factors influencing chatbot usage and the types of queries submitted. Lastly, the study was conducted within a single course at a specific institution, which may limit the applicability of the results to other courses, institutions, or educational contexts.

Future Research

Future research should address these limitations by exploring the impact of RAG chatbots in more diverse and larger student populations across various educational contexts and disciplines. Longitudinal studies could provide insights into the long-term effects of chatbot usage on student learning outcomes, engagement, and satisfaction. Additionally, research could also examine the effectiveness of different chatbot design features and interaction modalities to identify best practices for maximizing their utility and acceptance among students. By addressing these areas, future studies can contribute to a more robust and generalizable body of knowledge on the use of AI technologies in education.

6. REFERENCES

- Ait Baha, T., El Hajji, M., Es-Saady, Y., & Fadili, H. (2023). The impact of educational chatbot on student learning experience. *Education and Information Technologies*. Advance online publication. <https://doi.org/10.1007/s10639-023-12166-w>
- Chinedu, O., & Ade-Ibijola, A. (2021). Python-Bot: A Chatbot for Teaching Python Programming. *Engineering Letters*, 29(1), 25–34.
- Frasca, C., Muir, K., Ding, Y. (2023). Build a Chatbot with Custom Data Sources, Powered by LlamaIndex. Last accessed online on 7/18/2024 at <https://blog.streamlit.io/build-a-chatbot-with-custom-data-sources-powered-by-llamaindex/>
- Hurlbut, A. R. (2018). Online vs. traditional learning in teacher education: a comparison of student progress. *American Journal of Distance Education*, 32(4), 248–266. <https://doi.org/10.1080/08923647.2018.1509265>
- Jeong, C. (2023). A Study on the Implementation of Generative AI Services Using an Enterprise Data-Based LLM Application Architecture.

- Advance online publication.
<https://doi.org/10.48550/arXiv.2309.01105>
- Jungherr, A. (2023). *Using ChatGPT and Other Large Language Model (LLM) Applications for Academic Paper Assignments*.
<https://doi.org/10.31235/osf.io/d84q6>
- Kazemitabaar, M., Ye, R., Wang, X., Henley, A. Z., Denny, P., Craig, M., & Grossman, T. (2024). CodeAid: Evaluating a Classroom Deployment of an LLM-based Programming Assistant that Balances Student and Educator Needs. In F. F. Mueller, P. Kyburz, J. R. Williamson, C. Sas, M. L. Wilson, P. T. Dugas, & I. Shklovski (Eds.), *Proceedings of the CHI Conference on Human Factors in Computing Systems* (pp. 1–20). ACM.
<https://doi.org/10.1145/3613904.3642773>
- Labadze, L., Grigolia, M., & Machaidze, L. (2023). Role of AI chatbots in education: systematic literature review. *International Journal of Educational Technology in Higher Education*, 20(1). <https://doi.org/10.1186/s41239-023-00426-1>
- Lang, G. & O'Connell, S. D. (2015). Learning About Learning About Learning: Insights From a Student Survey in a Hybrid Classroom Environment. *Issues in Information Systems*, 16(11), 11–16.
https://doi.org/10.48009/3_iis_2015_11-16
- Lang, G., Triantoro, T., & Sharp, J. H. (2024). Large Language Models as AI-Powered Educational Assistants: Comparing GPT-4 and Bard for Writing Teaching Cases. *Journal of Information Systems Education*, Forthcoming.
- Maryamah, M., Irfani, M. M., Tri Raharjo, E. B., Rahmi, N. A., Ghani, M., & Raharjana, I. K. (2024). Chatbots in Academia: A Retrieval-Augmented Generation Approach for Improved Efficient Information Access. In *2024 16th International Conference on Knowledge and Smart Technology (KST)* (pp. 259–264). IEEE.
<https://doi.org/10.1109/KST61284.2024.10499652>
- Meyer, J. G., Urbanowicz, R. J., Martin, P. C. N., O'Connor, K., Li, R., Peng, P.-C., Bright, T. J., Tatonetti, N., Won, K. J., Gonzalez-Hernandez, G., & Moore, J. H. (2023). Chatgpt and large language models in academia: Opportunities and challenges. *BioData Mining*, 16(1), 20.
<https://doi.org/10.1186/s13040-023-00339-9>
- Napalit, F., Tanyag, B., So, C. L., Sy, C., & San Pedro, J. R. (2023). Examining student experiences: Challenges and perception in computer programming. *International Journal of Research Studies in Education*, 12(8). <https://doi.org/10.5861/ijrse.2023.71>
- Pichai, K. (2023). A Retrieval-Augmented Generation Based Large Language Model Benchmarked On a Novel Dataset. *Journal of Student Research*, 12(4).
<https://doi.org/10.47611/jsrhs.v12i4.6213>
- Strauss, A., & Corbin, J. (1990). *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*.
- Varkey, T. C., Varkey, J. A., Ding, J. B., Varkey, P. K., Zeitler, C., Nguyen, A. M., Merhavy, Z. I., & Thomas, C. R. (2023). Asynchronous learning: a general review of best practices for the 21st century. *Journal of Research in Innovative Teaching & Learning*, 16(1), 4–16. <https://doi.org/10.1108/JRIT-06-2022-0036>
- Vukojić, M., & Krstić, J. (2023). ChatGPT in programming education: ChatGPT as a programming assistant. *Journal for Contemporary Education*.

APPENDIX A: RAG CHATBOT PROMPT AND USER INTERFACE

"You are a course chatbot for BAN 663 (Business Data Analytics with R) taught by Prof. Guido Lang at Quinnipiac University. You were created by Prof. Lang and you are powered by GPT-4 and LlamaIndex. Your task is exclusively to answer questions related to R programming and statistics. When faced with inquiries, your responses should be confined to explaining, clarifying, and discussing aspects of R code and related statistical concepts based on the course materials given to you. Your objective is to encourage learning and comprehension. While you may provide R code snippets, you should never offer direct code solutions. Instead, focus on explaining the functionalities of R programming and the use of statistics in plain language to help users understand the underlying concepts. Refer to the course materials whenever possible. Lastly, never address queries outside the realm of R programming and statistics."

Course Chatbot for BAN 663

Created by [Prof. Lang](#) for [BAN 663](#). Powered by [GPT-4](#) and [LlamaIndex](#).

Ask a question about the course content.

Write R code that loads and merges data from two CSV files using an inner join.

Sure, the process of loading data from two CSV files and then merging them using an inner join in R involves several steps. Here's a general description of how you can do it:

- Load the required libraries:** You need the `tidyverse` package for data manipulation and the `here` package for setting the working directory. You can load them using the `library()` function.
- Load the CSV files:** Use the `read_delim()` function from the `readr` package (which is part of `tidyverse`) to load the CSV files. The `here()` function is used to specify the file path. If your CSV files are comma-separated, you can just pass the file path to `read_delim()`. Here's an example:

```
df1 <- read_delim(here("path/to/your/file1.csv"))
df2 <- read_delim(here("path/to/your/file2.csv"))
```
- Merge the data:** Use the `inner_join()` function from the `dplyr` package (also part of `tidyverse`) to merge the data frames. This function keeps only the rows in both data frames that match on the specified columns. Here's an example:

```
df <- inner_join(df1, df2)
```

By default, `inner_join()` merges the data frames based on columns with the same name in both data frames. If you want to specify the columns to join on, you can add a `by` argument like this:

```
df <- inner_join(df1, df2, by = "ColumnName")
```

Replace `"ColumnName"` with the name of the column you want to join on. If you want to join on multiple columns, you can pass a vector of column names like this: `by = c("Column1", "Column2")`.

Remember, it's important to assign the result of `inner_join()` to a new variable (in this case, `df`) so you can use the merged data in the next steps of your analysis.

Your question

APPENDIX B: SURVEY ITEMS

1. What is your age range? (Under 18, 18-24, 25-34, 35-44, 45-54, 55-64, 65 or older)
2. What is your gender? (Male, Female, Non-binary/Third gender, Prefer not to say)
3. What is your current level of employment? (Full time, Part time, Not employed)
4. What was your knowledge of R prior to this course? (Likert scale: 1 – None at all, 5 – A great deal)
5. How much experience with AI chatbots (such as ChatGPT, Copilot, Claude, or Bard) did you have prior to this course? (Likert scale: 1 – None at all, 5 – A great deal)
6. Have you used the course chatbot for this course? (Yes, No)
 - a. If No: Why didn't you use the course chatbot? (Open answer)
7. To what extent do you agree with the following statements? (Likert scale: 1 – Strongly disagree, 5 – Strongly agree):
 - a. The course chatbot enhanced my learning experience.
 - b. I wish course chatbots would be used in more courses.
8. What did you like most about the chatbot? (Open answer)