

Knowledge Compass: A Question Answering System Guiding Students with Follow-Up Question Recommendations

Rui Sheng
rshengac@connect.ust.hk
The Hong Kong University of Science
and Technology
Hong Kong, China

Yan Luo
luoyan.me@gmail.com
The Hong Kong University of Science
and Technology
Hong Kong, China

David Gotz
gotz@unc.edu
The University of North Carolina at
Chapel Hill
Chapel Hill, United States

Leni Yang
lyangbb@connect.ust.hk
The Hong Kong University of Science
and Technology
Hong Kong, China

Ziyang Xu
xuziyang@whu.edu.cn
Wuhan University
Wuhan, China

Huamin Qu
huamin@cse.ust.hk
The Hong Kong University of Science
and Technology
Hong Kong, China

Haotian Li
haotian.li@connect.ust.hk
The Hong Kong University of Science
and Technology
Hong Kong, China

Zhilan Zhou
zsl@live.unc.edu
The University of North Carolina at
Chapel Hill
Chapel Hill, United States

ABSTRACT

Pedagogical question-answering (QA) systems have been utilized for providing individual support in online learning courses. However, existing systems often neglect the education practice of guiding and encouraging students to think of relevant questions for deeper and more comprehensive learning. To address this gap, we introduce Knowledge Compass, an interactive QA system. The system can recommend follow-up questions that provide potential further explorations of the topics students ask about. Additionally, the system applies a course outline visualization and a set of interactive features for students to track the relationship between their questions and the course content.

CCS CONCEPTS

• **Information systems** → **Query suggestion**; • **Social and professional topics** → **Informal education**.

KEYWORDS

Online Learning, E-learning, Question Answering, Follow-up Question, Question Recommendation

ACM Reference Format:

Rui Sheng, Leni Yang, Haotian Li, Yan Luo, Ziyang Xu, Zhilan Zhou, David Gotz, and Huamin Qu. 2023. Knowledge Compass: A Question Answering System Guiding Students with Follow-Up Question Recommendations. In *The 36th Annual ACM Symposium on User Interface Software and Technology*

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

UIST '23 Adjunct, October 29–November 1, 2023, San Francisco, CA, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0096-5/23/10.

<https://doi.org/10.1145/3586182.3615785>

(UIST '23 Adjunct), October 29–November 1, 2023, San Francisco, CA, USA.
ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3586182.3615785>

1 INTRODUCTION

In online learning courses, students have limited one-to-one interaction with teachers. The lack of individual support can negatively impact the effectiveness of learning and cause dissatisfaction and high dropout rates [9]. To address this issue, research has explored automatic question-answering (QA) systems for helping students understand a course [10]. Most of them primarily focus on improving the capability of QA systems to understand and clarify questions [8] and generate precise [5, 6, 11], personalized responses [7].

However, the role of teachers in question-answering conversations with students extends beyond passively responding to questions. To facilitate deep and comprehensive learning, teachers should guide and encourage students to think of related questions [1, 3, 4]. Therefore, a question-answering system that can recommend relevant questions based on students' previous questions is desirable for achieving better learning results. To the best of our knowledge, there is no pedagogical QA system that has achieved this goal.

We present Knowledge Compass, an interactive QA system based on GPT-3.5¹, designed to answer students' questions and recommend follow-up questions, with the ultimate goal of facilitating deep and comprehensive learning. Focusing on the scenario of online learning, the system utilizes video recordings as the primary source materials for answering questions. The follow-up questions and previous questions are from the same or related topics that are identified from course outlines. We design the prompts for GPT to systematically generate different types of recommended questions, adhering to question taxonomies from previous literature [2]. Additionally, the system outlines the topics of a course through visualization, along with a set of interactive features that

¹<https://platform.openai.com/docs/models/gpt-3-5>

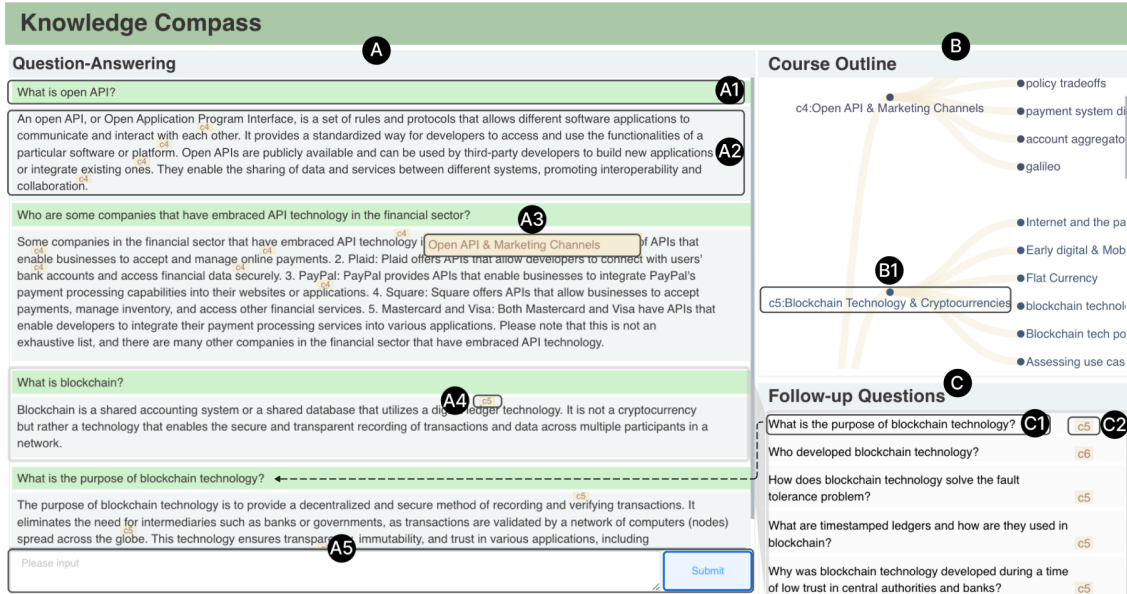


Figure 1: Knowledge Compass has three views. (A) The Question-Answering View allows students to ask questions about the course content; (B) The Course Outline View presents an outline of the course; (C) The Follow-up Question View enables students to explore recommended follow-up questions.

allow students to track the relationship between their questions and the topics. It can further help students have a comprehensive understanding of the course.

2 KNOWLEDGE COMPASS

The system consists of three views: Question-Answering View (Figure 1-A), Course Outline View (Figure 1-B), and Follow-up Question View (Figure 1-C). The Question-Answering View allows students to interact with our system to ask questions. The Course Outline View outlines the topics of a course, allowing students to grasp the connections between different topics. The Follow-up Question View lists recommended follow-up questions. In the following sections, we will introduce the interface and interaction of each view, along with the underlying algorithm that powers it.

Question-Answering View. Our system can answer questions (Figure 1-A1, A2) based on the course materials we provide. Given the extensive length of course materials, we leverage LangChain², a framework that enables GPT to analyze lengthy inputs. We extracted the corresponding lesson topics of each sentence, which are displayed in the form of a corner label at the top of each sentence (Figure 1-A4), such that students can track back where the answer is from. Furthermore, students can simply hover over a corner label to view detailed information (Figure 1-A3) about the corresponding lesson topic. Besides, clicking the corner label will move the corresponding lesson topic node to the center of the Course Outline View, giving students a contextual overview of the topic.

Course Outline View. The outline of a course provided by the course website is displayed through a tree diagram (Figure 1-B). The course name serves as the root node in a tree, while the lesson topics are represented as the second layer nodes. The individual

video topics, which form the building blocks of each lesson, are depicted as the leaf nodes of the tree.

Follow-up Question View. Several questions (Figure 1-C1) related to a particular conversation round are recommended. For each question, the corresponding lesson topics are extracted (Figure 1-C2) in a similar manner as in the Question-Answering View. There are mainly three steps for generating follow-up questions. First, we make sure the follow-up questions are from the same topics in the previous question or related topics. Specifically, we define related topics as child and sibling nodes (i.e., lesson topics) in the course outline visualization. After identifying related topics, we use GPT to recommend questions based on corresponding course materials. To overcome GPT’s content length limitation, we segment the material into smaller chunks and summarize them using GPT, which produces a concise version (up to 2000 words). Using this summary, GPT can suggest questions related to previous ones, and these questions can also be answered by utilizing the course materials. Second, we design the prompt to generate questions that start with “what”, “how”, “why”, and so on based on the question taxonomy concluded by [2]. Finally, our system can return the top 5 recommended questions by topic relevance and diversity.

3 PRELIMINARY EVALUATION

To gain early feedback from potential users, we conducted a preliminary evaluation of Knowledge Compass.

Dataset. We selected an open course from MIT OpenCourseWare, “FinTech: Shaping The Financial World”³ for our preliminary evaluation with users. We also prepared another two courses “Introduction To Algorithms”⁴, and “Innovation Systems For Science,

²<https://github.com/hwchase17/langchain>

³<https://ocw.mit.edu/courses/15-s08-fintech-shaping-the-financial-world-spring-2020/>

⁴<https://ocw.mit.edu/courses/sts-081-innovation-systems-for-science-technology-energy-manufacturing-and-health-spring-2017/>

Table 1: This table presents examples of follow-up questions generated by Knowledge Compass. The generated questions are diverse and demonstrate their depth and relevance regarding the previous question.

Course	Question	Follow-up Questions
FinTech: Shaping The Financial World	Why banks need open banking to share the data with non-banks?	<p>What is open banking and how does it enable financial institutions to share data with non-banks?</p> <p>How do open APIs eliminate the need for manual processes like screen scraping and robotic process automation?</p> <p>Why are web-based APIs considered protocols that multiple parties can use to access services?</p>
Introduction To Algorithms	What is the difference between Static arrays and Linked lists?	<p>What are the advantages of using static arrays for inserting and deleting elements?</p> <p>How does inserting at the beginning of a static array affect its invariant?</p> <p>Why does inserting and deleting from static arrays require $O(n)$ time?</p>
Innovation Systems For Science, Technology, Energy, Manufacturing, And Health	How does classical economic theory affect innovation?	<p>How does government intervention in driving economic growth and increasing economic well-being relate to the potential for innovation?</p> <p>What are the implications of different waves of innovation in the US economy on society and the government's role in balancing winners and losers?</p>

Technology, Energy, Manufacturing, And Health”⁵ that will be used for demonstrating our system.

Setup. We invited four Ph.D. students (two males and two females, denoted as *P1-P4*) to study the fourth section (i.e., Open API & Marketing Channels) of the FinTech course. None of them had a background in FinTech before taking part in our evaluation. They were instructed first to study the section using around one hour, similar to the length of the course video. To mimic a real-world scenario of studying new content, we provided all related course material on the homepage of that section⁶. They could learn with any materials provided by the course instructor, including course videos, slides, and reading materials. During the 1-hour study, the participants were encouraged to use our system to receive help when they met any questions. After studying the section, the participants were invited to recall their learning experiences and comment on our system through a semi-structured interview.

Experiment Result. Through the preliminary evaluation, we found that each participant spent approximately 20 minutes interacting with Knowledge Compass during a one-hour learning session. Each participant's number of submitted questions to Knowledge Compass ranged from three to six, with an average of five. Follow-up questions recommended by our system accounted for a roughly average of 50% of these queries, ranging from 33% to 66%. We discovered that the four participants tended to initially learn the section with the course slides. Then they would turn to Knowledge Compass for assistance when encountering a topic that lacks a clear introduction. The majority of their inquiries began with “what”, with “how” and “why” following closely behind. We inquired the participants about in what situations they found the recommended follow-up questions useful. We concluded two major benefits of follow-up questions. First, *P1*, *P3*, and *P4* reported that follow-up questions were useful for understanding what to focus on next through follow-up questions and eventually establishing a general understanding of the course. Besides, follow-up questions sometimes helped clarify the question asked by participants. For

instance, *P1* found that our system initially struggled to provide a satisfactory response to his question (i.e., what is the relationship between open API and open banking). However, he was pleasantly surprised when a follow-up question specified his query (i.e., how does open API facilitate open banking). With the new question, *P1* obtained a satisfactory answer.

Nonetheless, the participants pointed out several potential improvements in the follow-up question recommendation. *P1* and *P3* suggested a way to personalize the system by not recommending follow-up questions that users have previously indicated as uninteresting (e.g., not selecting before). *P4* hoped that the strategy for the recommending could be progressive. The system can focus on the breadth of knowledge to help students understand the framework of the content at the beginning. Then the system can shift to recommending more questions about in-depth content.

Follow-up Question Demonstration. We showcased some follow-up questions from the user experiment mentioned above and collected several examples from another two courses (Table 1). The results demonstrated the effectiveness of our proposed prompt to generate context-specific and diverse follow-up questions.

4 CONCLUSION AND FUTURE WORK

We introduced Knowledge Compass, an interactive QA system designed to address students' questions and recommend follow-up questions for promoting deep and comprehensive learning. Our evaluation affirmed the system's efficacy in assisting students in achieving a comprehensive understanding of the course material. Our future work includes: 1) conducting a larger scale user study to evaluate the system thoroughly 2) delving deeper into theories and algorithms for follow-up question recommendations, including leveraging epistemology and learning theories in the recommendation process, and 3) incorporating personalized responses and follow-up questions based on users' historical interaction records.

ACKNOWLEDGMENTS

This project is partially supported by HK ITF grant PRP/017/22FX and HK RGC GRF grant 16210722.

⁵<https://ocw.mit.edu/courses/6-006-introduction-to-algorithms-spring-2020/>

⁶<https://ocw.mit.edu/courses/15-s08-fintech-shaping-the-financial-world-spring-2020/pages/class-4-open-apis-marketing-channels/>

REFERENCES

- [1] Christine Chin. 2001. Learning in Science: What Do Students' Questions Tell Us about Their Thinking? *Education Journal* 29, 2 (2001), 85–103.
- [2] Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. QuAC: Question Answering in Context. In *Proceedings of the Empirical Methods in Natural Language Processing*. 2174–2184.
- [3] Marylou Dantonio. 1987. Develop Concepts, Question by Question. *Science Teacher* 54, 5 (1987), 46–49.
- [4] Meredith D Gall. 1970. The Use of Questions in Teaching. *Review of Educational Research* 40, 5 (1970), 707–721.
- [5] Brandon Jaipersaud, Paul Zhang, Jimmy Ba, Andrew Petersen, Lisa Zhang, and Michael R Zhang. 2023. Decomposed Prompting to Answer Questions on a Course Discussion Board. In *Proceedings of the International Conference on Artificial Intelligence in Education*. 218–223.
- [6] Srikrishna Raamadhurai, Ryan Baker, and Vikraman Poduval. 2019. Curio SmartChat: A System for Natural Language Question Answering for Self-paced K-12 Learning. In *Proceedings of the Workshop on Innovative Use of NLP for Building Educational Applications*. 336–342.
- [7] Anuschka Schmitt, Thiemo Wambsganss, and Jan Marco Leimeister. 2022. Conversational Agents for Information Retrieval in the Education Domain: A User-Centered Design Investigation. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 22.
- [8] Svitlana Vakulenko, Shayne Longpre, Zhucheng Tu, and Raviteja Anantha. 2021. Question Rewriting for Conversational Question Answering. In *Proceedings of the ACM International Conference on Web Search and Data Mining*. 355–363.
- [9] Thiemo Wambsganß, Leonie Haas, and Matthias Söllner. 2021. Towards the Design of a Student-Centered Question-Answering System in Educational Settings. In *Proceedings of the European Conference on Information Systems*.
- [10] Florian Weber, Thiemo Wambsganss, Dominic Rüttimann, and Matthias Söllner. 2021. Pedagogical Agents for Interactive Learning: A Taxonomy of Conversational Agents in Education. In *Proceedings of the International Conference on Information Systems*. 1–17.
- [11] Brian Zylich, Adam Viola, Brokk Toggerson, Lara Al-Hariri, and Andrew Lan. 2020. Exploring Automated Question Answering Methods for Teaching Assistance. *Artificial Intelligence in Education* 12163, 610–622.