

BOARD #101: Work In Progress: Enhancing Active Recall and Spaced Repetition with LLM-Augmented Review Systems

Mr. Muhammed Yakubu, University of Toronto

Final year Computer Engineering Student at the University of Toronto

Mr. Jasnoor Guliani, University of Toronto

Mr. Nipun Shukla, University of Toronto

Final year student at the University of Toronto.

Dylan O’Toole

Dr. Hamid S Timorabadi P.Eng., University of Toronto

Hamid Timorabadi received his B.Sc, M.A.Sc, and Ph.D. degrees in Electrical Engineering from the University of Toronto. He has worked as a project, design, and test engineer as well as a consultant to industry. His research interests include the application of digital signal processing in power systems.

WIP: Enhancing Engineering Education with Active Recall and Spaced Repetition LLM-Augmented Review Systems

Abstract

Active recall and spaced repetition are well-established techniques known to improve students' retention and understanding in engineering education. Tools such as Anki and Quizlet facilitate these methods but often require significant time and effort on the part of a student to benefit from. To address this, we developed Let's Learn More, an LLM-powered tool that automatically generates additional, subtly varied review questions and provides formative feedback on student responses. By exposing learners to parallel but distinct questions covering the same concepts, the tool aims to reduce simple memorization and promote deeper understanding.

In early 2025, we piloted Let's Learn More with 33 undergraduate engineering students to evaluate its impact on learning outcomes. Preliminary results show a 18.2% higher average and an 33.2% higher median test score when using the LLM-augmented system. User feedback also suggests that automatically generated question variations save time and enhance engagement. This work contributes to the growing body of literature on integrating artificial intelligence into evidence-based pedagogical strategies, and it offers practical insights for optimizing studying efficiency in engineering education.

1.0 - Introduction and Background

Active recall and spaced repetition are two evidence-based learning strategies widely recognized for their effectiveness in improving long-term retention and understanding [1-4]. Active recall involves retrieving information from memory, typically by answering questions or solving problems, while spaced repetition involves repeated review and recall of study material at increasing intervals to enhance long-term retention. Despite their proven benefits, many existing tools such as Anki and Quizlet do not fully integrate these strategies in ways tailored to the unique demands of engineering education, such as solving complex, application-driven problems [5]. The process of creating new practice questions to perform active recall can be time-consuming and difficult without expertise on the subject matter, and students often resort to simply reviewing existing material [6].

We are developing Let's Learn More, a web application that addresses these gaps by integrating active recall and spaced repetition with AI-powered question generation and evaluation. This tool allows students to upload course materials, generate review questions and problems, and receive feedback on their responses, reducing the effort required for content creation and encouraging consistent study habits. By adapting these techniques to engineering and STEM contexts, this project aims to enhance both the efficiency and effectiveness of student learning.

2.0 - Research Objectives

This study explores the potential of Let's Learn More to improve student learning outcomes and engagement. Specifically, it addresses the following research questions:

1. Can AI-generated content reduce the time and effort required for students to create effective study material?
2. Can AI-generated study content enhance students' understanding and performance?

The overarching objective is to evaluate the impact of Let's Learn More on fostering efficient, consistent, and conceptually deep learning practices among engineering students, with a special focus on application-driven problem solving questions.

3.0 - System Design

The Let's Learn More tool is implemented as a web application which combines several core features:

- Content Upload and Parsing: Users can upload course syllabi, lecture notes, and problem sets, which are parsed to extract key topics and concepts.
- Question Generation and Evaluation: A retrieval-augmented system enriches large language model (LLM) prompts with extracted course content to generate review questions and evaluate user responses, providing formative feedback.
- Spaced Repetition Scheduling: A dynamic algorithm schedules review sessions based on user performance and perceived question difficulty.
- Progress Visualization: Students track their progress through graphical summaries of performance and learning trends.

We chose to make Let's Learn More a web application for cross-platform compatibility, using Next.js for the frontend and Supabase for the database-as-a-backend. The user interface of the application is illustrated in Appendix A with a series of screenshots. Readers are also encouraged to view the video walkthrough on YouTube for a dynamic demonstration of the application in use [7].

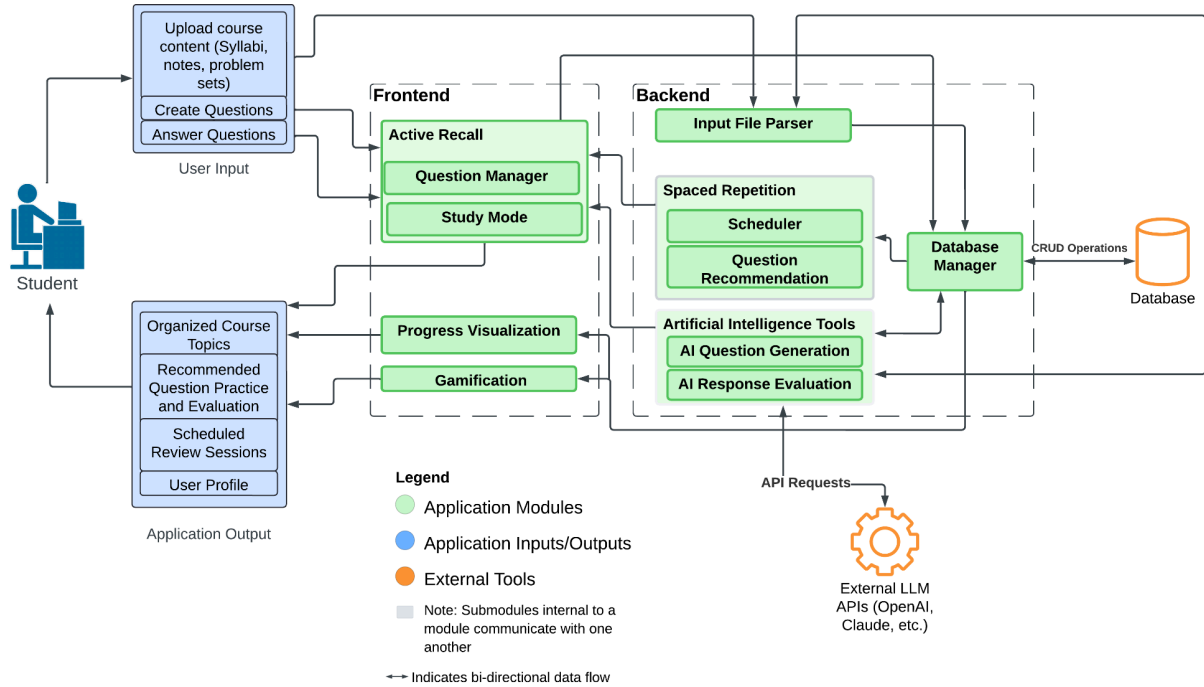


Fig. 1. Block diagram showing components of the system and their interactions, with legend.

In the above diagram, our design is split into frontend and backend components, which is a well-established approach for web applications. The front end is responsible for providing users an interface to provide input to and engage with the application, while the backend is responsible for storing and processing user data to supplement the front end interaction. Users engage with the front end in three ways: providing course content, creating and updating user questions, and answering questions in study mode.

Uploaded course content such as syllabi, course notes or problem sets are parsed in the backend using a large language model. We extract different information depending on the type of content – for example, topics and course objectives are extracted from syllabuses, concept explanations and examples from lecture notes, and question-solution pairs from problem sets. Our tool parses pictorial information (charts, diagrams, flowcharts) into descriptive text during text extraction, ensuring that key information from images is preserved alongside the written content. All extracted data is labeled with metadata including the source and related topics. This extracted data is then stored with its metadata in our database.

Extracted questions are automatically added to students’ flashcards, and new questions are generated topic-wise, by retrieving relevant chunks from our database and providing this context, alongside specific instructions to a large language model. We found this retrieval-augmented approach in combination with few-shot prompting to be effective for this use case [8][9]. For more details on our parsing approach, see Appendix A.

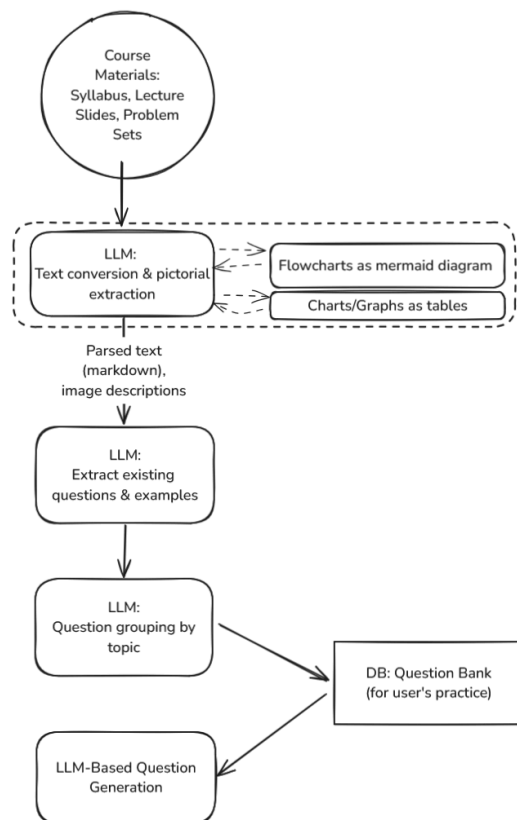


Fig. 2. Parsing system flowchart

Study mode enables users to start a review session of any set of flashcards, which can either be from user-created or AI-generated questions derived from their uploaded files. Students can also view AI-generated evaluations for their long form response questions.

Our question recommendation and scheduling system, handled in the backend, uses a spaced repetition algorithm (FSRS5) [10] to calculate when a user should practice a particular question, based on the time elapsed since review and previously indicated question understanding levels. Lastly, a user will receive progress tracking on a per-topic and course-level basis, powered by their performance in study mode review sessions (See Appendix A for UI screen captures).

4.0 - Pilot Study Design

To evaluate the effectiveness of Let's Learn More, we designed a pilot study consisting of 33 participants split into two groups, each assigned to study 2 textbook chapters and 1 problem set from an introductory university materials science course material under different conditions. This design allows us to compare learning outcomes and study behaviors between the two groups.

- **Group 1:** Participants in this group studied using Let's Learn More, leveraging both AI-generated questions and study materials. The tool provided additional questions based on the uploaded content.
- **Group 2:** Participants in this group studied the provided material independently, without access to AI-generated questions or structured review tools. They were encouraged to use their preferred study methods.

Both groups were allotted a maximum of one hour to study the material. After the study session, participants completed a timed 30-minute test comprising past exam questions on the same topic to assess their understanding and retention of the material. Sample test questions used to evaluate study participants are included in Appendix D.

Following the test, participants in both groups were surveyed about their study strategies to gather qualitative insights into their experiences, perceived effectiveness of their methods, and any challenges faced.

While this study captures short-term retention and understanding, its design does not assess long-term knowledge retention -- a limitation that will be addressed in a more comprehensive follow-up study. This study also omits evaluation of the spaced repetition system.

5.0 - Preliminary Results

The study results highlight Let's Learn More's effectiveness in improving short-term retention and understanding. As shown in Figure 2, Group 1, which used AI-generated questions, outperformed Group 2 with a 18.2% higher average and an 33.3% higher median test score.

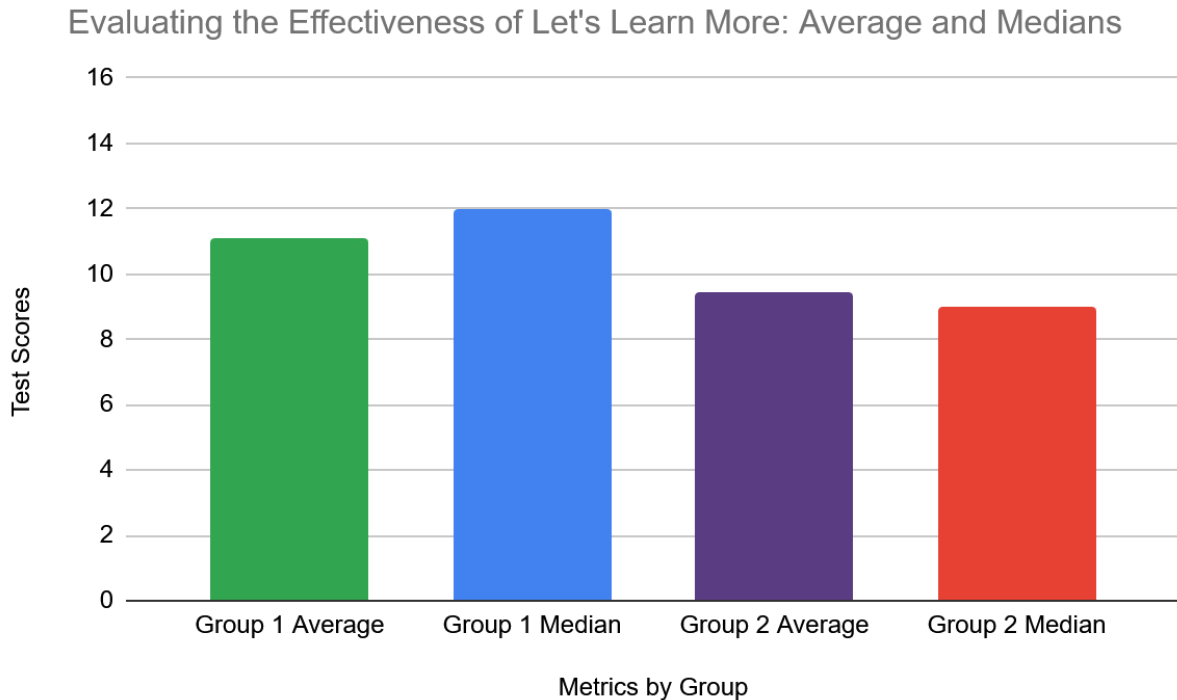


Fig. 3. Graph depicting the average and median test scores of Group 1 and Group 2.

Overall, student feedback on the AI-generated questions was positive, with participants finding them to be a valuable supplement to the textbook, particularly for providing worked examples and opportunities for self-assessment.

6.0 - Discussion

Our preliminary findings suggest that integrating AI-powered features with active recall can significantly enhance learning outcomes and efficiency. Automating time-intensive tasks, like question creation and evaluation, allows students to dedicate more resources to understanding concepts. The automated generation of new, relevant practice material also promotes active learning and aligns with evidence-based pedagogical practices, highlighting AI's potential to transform study routines [1].

While we did not specifically measure our system's impact on spaced repetition outcomes during this pilot study, its design is rooted in well-researched principles of memory retention [3]. By integrating a dynamic scheduling algorithm based on user performance and difficulty ratings, Let's Learn More leverages these established principles, suggesting that it likely enhances this critical aspect of the learning process.

The AI-generated questions were generally viewed as a valuable supplement to course material, helpful for reinforcing learned material, applying concepts, and, to some extent, aiding in test

preparation. However, some students noted a mismatch between the AI-generated questions and the actual test content, suggesting a need for better alignment with learning objectives and assessment criteria. Additionally, some feedback indicated that the questions might be more useful for test preparation than for promoting deep understanding, and that the tool could be improved by tailoring the difficulty and content to students' diverse backgrounds and prior knowledge. See Appendix B for detailed participant testimonials for each section.

7.0 - Challenges

This study encountered several challenges. A primary challenge was the LLM's occasional generation of inaccurate or irrelevant questions. Although this was reduced through iterative refinement of content parsing, retrieval, and prompt engineering, LLM hallucination remains a broader problem [11]. Initially, we attempted mitigation systems, but their complexity grew rapidly. To maintain focus on core research objectives, we shifted the burden of filtering irrelevant questions to the user, allowing them to review and discard generated questions. Future software iterations could explore shifting this burden back to the software, perhaps with LLM self-verification or multi-LLM approaches [12]. An additional challenge was the study's compressed timeframe, which prevented full evaluation of the tool's spaced repetition features and its long-term impact on student learning. Relatedly, the team underestimated the effort required to develop the web application and conduct the multi-month study, further hindering our ability to test spaced repetition. Finally, the system's handling of varied data formats requires optimization to ensure consistent question quality across input types.

8.0 - Future Work

Future work will address the challenges identified in this study and expand upon its preliminary findings. A key priority is a longer-term evaluation of Let's Learn More, assessing its effectiveness in improving students' understanding and performance over an extended period, with a particular focus on the spaced repetition component. This will involve studies conducted over a full academic term, allowing students to integrate the tool into their regular study habits.

Technically, we will explore alternative LLM prompting strategies and question generation architectures, such as self-review or Twin-Star approach [13], to improve the quality and relevance of the generated questions, and reduce hallucination. We also plan to experiment with using learning objectives, rather than just past/extracted questions, as a criterion for question generation, to explore whether it could result in improved question alignment.

References

- [1] J. Xu et al., “Active recall strategies associated with academic achievement in Young Adults: A systematic review,” *Journal of Affective Disorders*, vol. 354, pp. 191–198, Jun. 2024. doi:10.1016/j.jad.2024.03.010
- [2] J. P. Marinelli, T. P. Hwa, C. M. Lohse, and M. L. Carlson, “Harnessing the power of spaced repetition learning and active recall for Trainee Education in otolaryngology,” *American Journal of Otolaryngology*, vol. 43, no. 5, p. 103495, Sep. 2022. doi:10.1016/j.amjoto.2022.103495
- [3] N. J. Cepeda et al., "Spacing effects in learning: A temporal ridgeline of optimal retention," *Psychological Science*, vol. 19, no. 11, pp. 1095-1102, 2008.
- [4] X. Yuan, “Evidence of the spacing effect and influences on perceptions of learning and science curricula,” *Cureus*, Jan. 2022. doi:10.7759/cureus.21201
- [5] L. Weitze and H. Söbke, “Quizzing to Become an Engineer -A Commercial Quiz App in Higher Education,” *New Perspectives in Science Education*, vol. 5, 2016. doi:DOI:10.13140/RG.2.1.2485.6723
- [6] J. Chasmar, B. Melloy, and L. Benson, “Use of self-regulated learning strategies by second-year Industrial Engineering Students,” 2015 ASEE Annual Conference and Exposition Proceedings, 2015. doi:10.18260/p.24979
- [7] Muhammed Yakubu, “LetsLearnMore Feature Preview,” *YouTube*, Mar. 26, 2025. https://www.youtube.com/watch?v=_kIIzg9wFwM (accessed Apr. 29, 2025).
- [8] Z. Levonian et al., “Retrieval-augmented Generation to Improve Math Question-Answering: Trade-offs Between Groundedness and Human Preference,” arXiv.org, Oct. 04, 2023. <https://arxiv.org/abs/2310.03184>
- [9] U. Lee et al., “Few-shot is enough: Exploring CHATGPT prompt engineering method for automatic question generation in English education,” *Education and Information Technologies*, vol. 29, no. 9, pp. 11483–11515, Oct. 2023. doi:10.1007/s10639-023-12249-8
- [10] I. Coffee, “Anki FSRS Explained,” Anki-decks.com, 2024. <https://anki-decks.com/blog/post/anki-fsrs-explained/> (accessed Apr. 29, 2025).
- [11] B. C. Figueras and R. Agerri, “Critical Questions Generation: Motivation and Challenges,” arXiv.org, 2024. <https://arxiv.org/abs/2410.14335> (accessed Apr. 29, 2025).
- [12] S. Mucciaccia, T. Paixão, F. Mutz, A. De Souza, C. Badue, and T. Oliveira-Santos, “Automatic Multiple-Choice Question Generation and Evaluation Systems Based on LLM: A Study Case With University Resolutions,” 2260. Available: <https://aclanthology.org/2025.coling-main.154.pdf>

- [13] Qingfeng Zhuge, H. Wang, and X. Chen, “TwinStar: A Novel Design for Enhanced Test Question Generation Using Dual-LLM Engine,” *Applied Sciences*, vol. 15, no. 6, pp. 3055–3055, Mar. 2025, doi: <https://doi.org/10.3390/app15063055>.
- [14] I. L. Molina, Valdemar Švábenský, Tsubasa Minematsu, L. Chen, F. Okubo, and A. Shimada, “Comparison of Large Language Models for Generating Contextually Relevant Questions,” *Lecture notes in computer science*, pp. 137–143, Jan. 2024, doi: https://doi.org/10.1007/978-3-031-72312-4_18.

Appendix A: Screenshots of Let's Learn More's User Interface

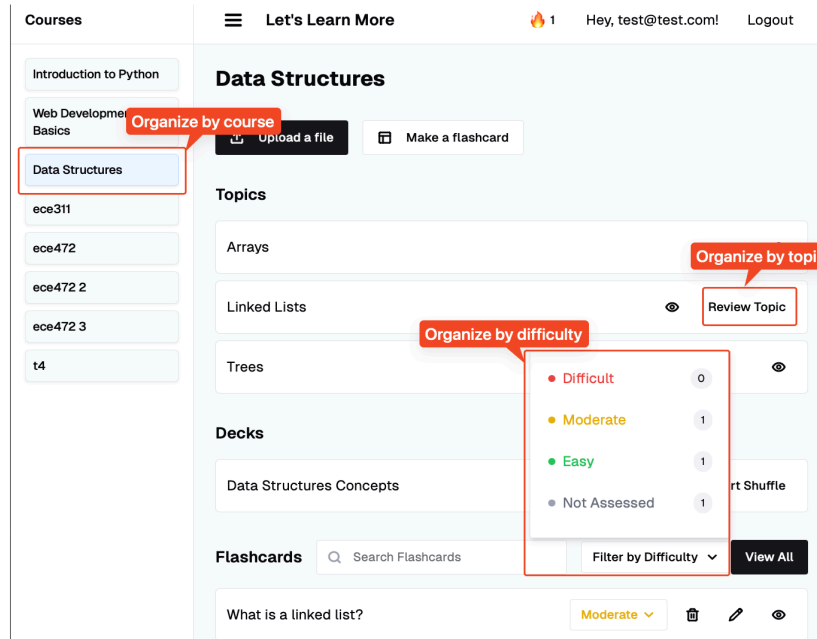


Fig. 4. Course homepage.

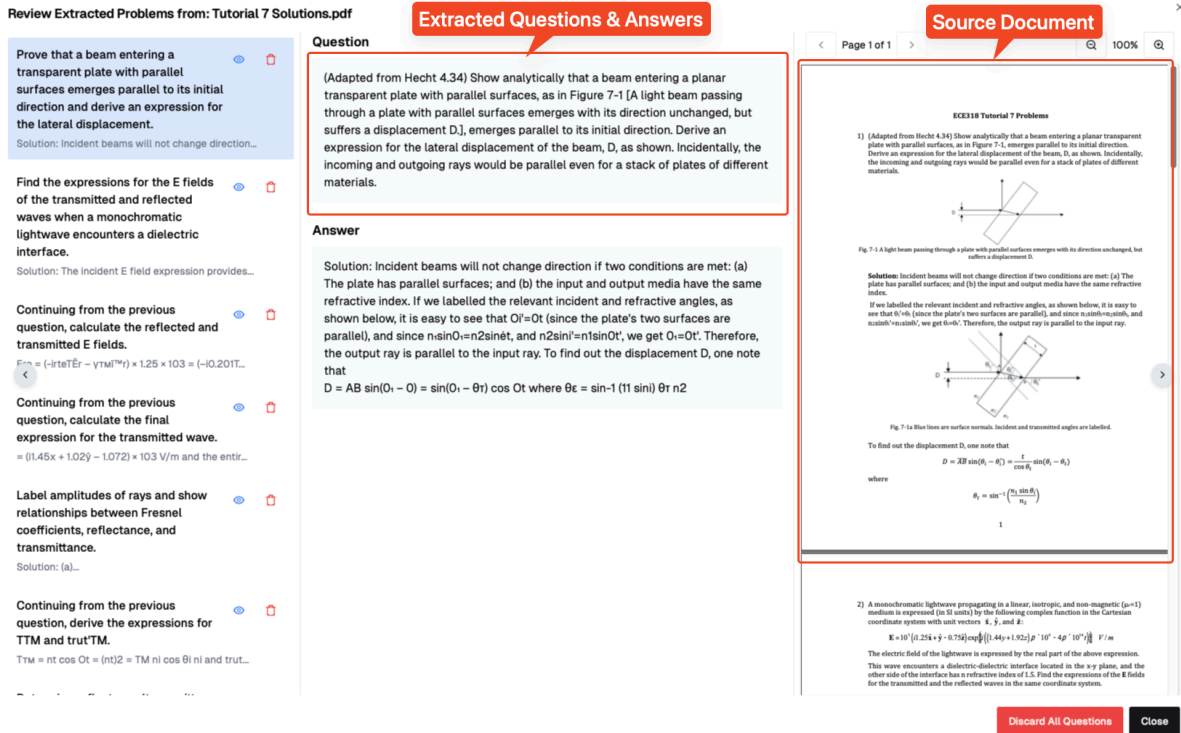


Fig. 5. File Parsing & Question Extraction Interface.

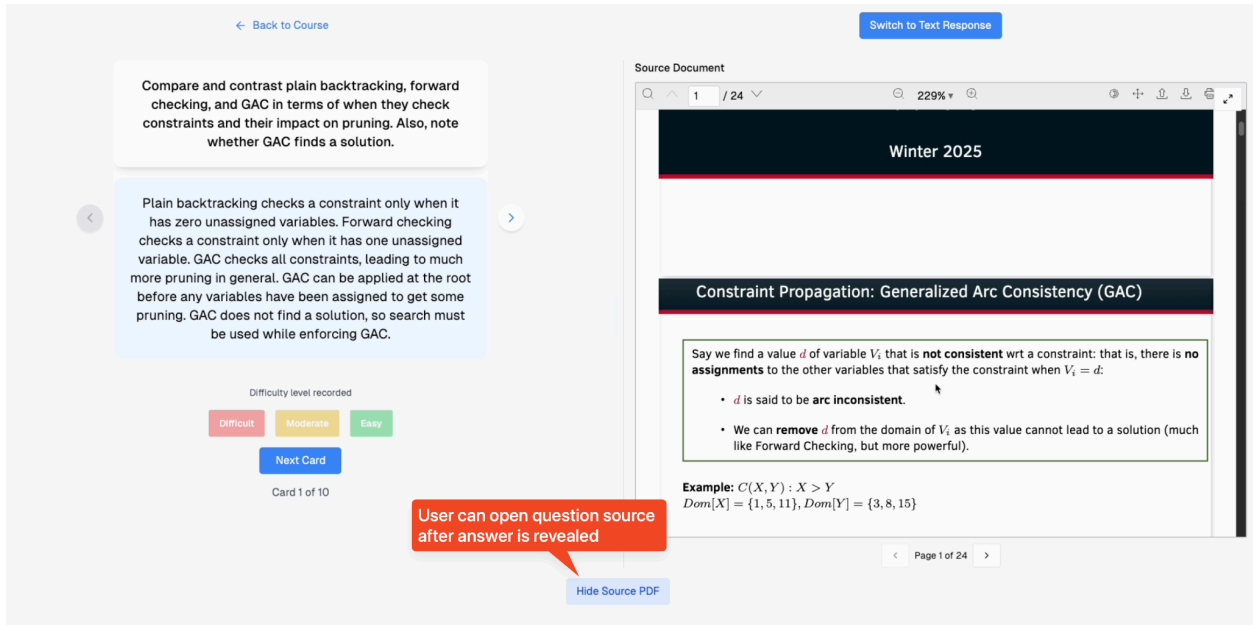


Fig. 6. Study mode – question study interface.

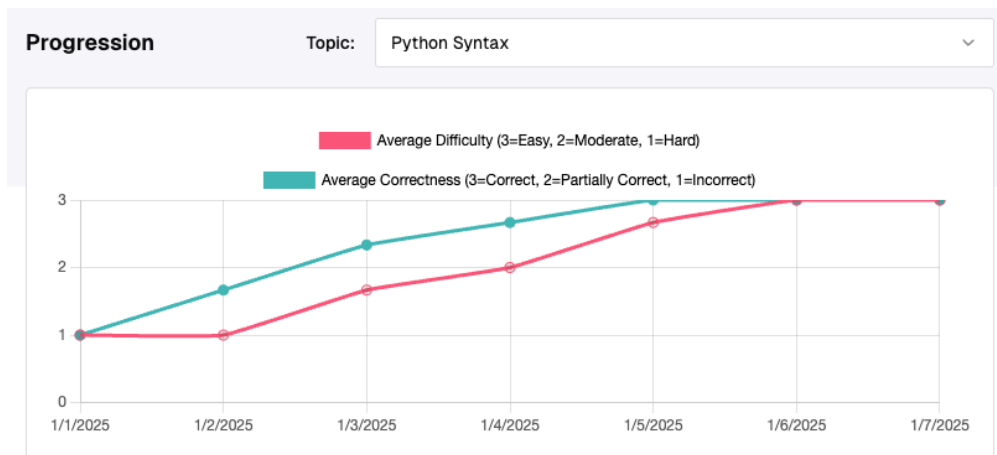


Fig 7. Progress tracking graph, depicting user's progression over time.

Appendix B: Detailed Parsing System

Our parsing system is designed to extract and generate questions from various course materials, including syllabi, lecture slides, and problem sets. The process is largely automated, leveraging the capabilities of Large Language Models (LLMs) to process and interpret the content. Given that LLMs are most proficient with text, the primary pipeline involves converting all input materials into a textual format. The parsing process can be summarized in the following steps, and is visually depicted in Figure 2.

1. **Text Conversion:** The initial step involves converting all input course materials into text. This is straightforward for documents that are already text-based (e.g., lecture notes, problem sets in text format). For other formats, such as slides, we utilize automated tools to extract the text.
2. **Extraction of Existing Questions/Examples:** We first attempt to extract questions and examples that are already present in the course materials. These human-created questions are considered high-quality and are directly incorporated into the question bank.
3. **LLM-Based Question Generation:** For the generation of new questions, we employ LLMs. We experimented with several different approaches and LLMs, including:
 - LlamaParse
 - Google Gemini-flash-2.0
 - OpenAI gpt-4o

For our pilot study, we used OpenAI and LlamaParse. However, we found that Google Gemini-flash-2.0 provided superior parsing quality at a lower cost, and thus adopted it for the main study. The LLM is prompted to generate questions based on the text content of the course materials. See Appendix E for the question extraction prompt..

4. **Handling Pictorial Information:** A significant challenge in parsing course materials is handling pictorial information, such as charts, images, and diagrams. Our approach to this is a "best-effort" strategy:
 - The LLM is instructed to describe any charts, images, or diagrams present in the materials as text.

- Charts and graphs can often be effectively represented and understood as tables, which LLMs can process.
 - Many flowcharts can be converted into "mermaid" markdown diagrams [Include reference/link to mermaid], which can be rendered by our frontend.
5. We acknowledge that this is a simplification, and that some information may be lost in the conversion. However, we observed that the majority of the core learning content in the materials we processed was text-based. We are exploring methods to more intently integrate pictorial information in the parsing process.
 6. **Question Grouping:** Once the questions are extracted and generated, we attempt to group them by topic or concept. This is achieved by extracting topics and concepts from the course syllabus and then using these to categorize the questions.
 7. **Topic-Based Question Generation:** Finally, we prompt the LLM to generate new questions, using the extracted questions as few-shot examples as few-shot prompting has been shown to provide good results in similar papers on Automatic Question Generation with LLMs [9][12][14]. This helps to ensure that the generated questions are relevant to the course content and aligned with the learning objectives.

Appendix C: Participant Reviews of Let's Learn More

Overall Positive Feedback:

- **Supplement to Textbook:** Several participants noted that the AI-generated questions were a valuable supplement to the textbook, which they found lacking in worked examples. This suggests the AI effectively filled a gap in the existing learning resources.
 - *"the textbook didn't really have worked out examples so the ai questions were nice for that"*
 - *"was nice to see solved questions. I depend on those when I'm cramming for tests and it sucks when the textbook doesn't have any."*
- **Reinforcement and Self-Assessment:** Participants appreciated the questions for reinforcing learned material and providing opportunities for self-assessment.
 - *"It was great, complete and helped me remember what I had just read and test myself on the course content"*
- **Application of Material:** The AI questions helped students see how the concepts could be applied.
 - *"Useful for seeing more examples on how to apply the material"*

Areas for Improvement:

- **Alignment with Test Content:** One participant pointed out a mismatch between the AI-generated questions and the actual test questions. The AI questions focused more on real-world applications, while the test focused on more fundamental characteristics. This suggests a need for better alignment between the AI-generated content and the specific learning objectives and assessment criteria.
 - *"The questions and answers were good in general, however in the scope of the test, weren't very helpful. The AI-generated questions focused more on the application of young's modulus in real world materials. Although it helped me understand young's modulus at a higher level than the course material could, the test questions covered more fundamental characteristics of young's modulus such as its structure independence."*
- **Perceived Usefulness for Understanding vs. Test Preparation:** Some participants felt the AI questions were more helpful for test preparation than for truly understanding the material. This highlights a potential trade-off between these two aspects of learning.

- *"Compared to the textbook course material, the AI generated questions are like looking at past test solutions. Not very helpful for understanding the material imo but helpful for doing well on the test."*
- **Relevance to Background:** One student mentioned that the questions weren't geared toward their understanding, because they hadn't done "engineering" math problems before. This suggests that the AI-generated questions might need to be better tailored to students with varying levels of prior knowledge and experience.
 - *"Kind of same as tghe course material, they helped me and I did better than without them but not geared to my understanding because I never did "engineering" math problems."*

Appendix D: Test Questions Used For The Study

Appendix D presents a selection of questions used during participant testing. The questions are designed to test application-driven understanding over memorization, and demonstrate the style of learning that users of our tool can perform better on.

Question:

Which of the following would be expected to change the Young's modulus of a metallic sample?

- a) Adding 1 atomic percentage impurity to the metal
- b) Processing the metal to decrease the its grain (crystal) size
- c) Increasing the strength through plastic deformation
- d) Increasing the temperature to $0.6 T_m$

Question:

Which of the following strengthening mechanisms would you expect to change the Young's modulus of a metal?

- (a) None of these.
- (b) Grain size reduction.
- (c) Cold working.
- (d) Minor alloying additions (solid solution).

Question:

A footbridge is produced from concrete. The bridge has a rectangular cross-section with height 20 cm and width 100 cm and is designed to cross a span of 3 m. Assume a typical value for the bending strength of concrete of 0.5 MPa, and an acceleration due to gravity of 9.8 N/kg.

What would be the maximum mass, in kg, that could be supported at the middle of this bridge?

Appendix E: Question Extraction Prompt

```
`You are an expert at identifying and extracting questions and their answers from educational materials.
```

```
Analyze this PDF document page by page and extract any questions that are clearly part of the assessment, along with their corresponding answers.
```

```
For each main question you identify (including all its sub-parts), provide:
```

1. A brief summary line (one sentence describing what the overall question is about)
2. The related subtopic or subject area
3. The complete text of the question and all its sub-parts formatted in markdown (as close as possible to the original). Equations should be formatted using LaTeX. When referring to figures, use the format "Figure X [description]" - if the figure has an existing caption, use that caption, otherwise, refer to the figure by its description. For multiple figures, use unique sequential numbers.
4. The page number where the question appears
5. The complete answer text for the question, if present
6. A brief summary of the answer approach/steps (especially for calculation/computation questions)

```
Format your response as a JSON array of question objects with the following structure:
```

```
[  
  
  {  
  
    "question_summary": "Brief summary of the entire question",  
  
    "related_subtopic": "Related subject area",
```

```
"question_markdown": "Complete markdown text including all parts (a, b, c,
etc) of the question",

"page_number": <page number where question appears>,

"answer_markdown": "Complete markdown text of the answer, or empty string if
no answer is found",

"answer_summary": "Brief summary of solution approach/steps (for calculation
questions) or key points (for conceptual questions)"

},

...

]
```

Important:

- Separate new lines with `\n`, ensure that all the text is in a single-line string (json format rule)
- Analyze each page in sequence and maintain page number references
- Only extract questions that are clearly part of the assessment - do not include administrative items like name/student number fields unless they contain an actual bonus question
- Keep multi-part questions (e.g. 1a, 1b, 1c) together as a single question object
- Format sub-parts using markdown list syntax
- If no questions are found on a page, do not include any entries for that page
- Make sure to include the questions preamble. That is, the text explaining the situation before the actual numbered questions.
- Only include actual questions that require answers, not rhetorical questions or statements
- For answers, handle both formats:

1. Where the answer immediately follows each question
 2. Where all answers are in a separate section after all questions - in this case, carefully match answers to their corresponding questions by number/letter and content
- If no answer is found for a question, use empty strings for both the `answer_markdown` and `answer_summary` fields`