

CS4530 Final Project: “Question Coach, Answer Summaries & Privacy Vault”

Group 906: Yipeng Wang, Suneet Pathangay, Ved Naykude

Our Feature: Helping People Ask Better Questions - Without Giving Up Privacy

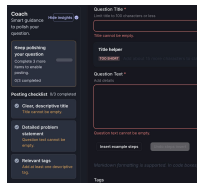

In the original fake stack overflow, it is often difficult for novices to ask clear and well-defined questions. The title is fuzzy, the code is pasted directly without labels, the labels are inconsistent, the long conversation thread is difficult to browse quickly, the moderator needs to spend time cleaning up the content, and the user lacks understanding of how the AI function handles its data. In this regard, we have launched a reputation aware question guidance tool, the answer summary function supported by AI, and the hierarchical design of the privacy safe, so that creators can strictly control the data flow outside the browser while obtaining contextual help.

The question guidance tool is located next to the editor to convert the basic quality criteria into an interactive list. In order to reduce repeated questions, the editor can search for similar questions by combining query words according to the title text and tags, and display the possible matching results through the sidebar during the editing process of users, highlighting overlapping words. Draft will also be saved. For the long question and answer page, the answer summary part will condense the conversation thread into 80 to 120 words, and each sentence will be marked with the source of the original quotation.

Since the effectiveness of auxiliary tools depends on user trust, the AI function is turned off by default. The local browser's exclusive sensitive content checker will identify possible API keys, e-mail addresses and phone numbers before submission, and provide fuzzy preview and one key desensitization. The drafts and AI request tracks are only stored in the client's privacy safe and will be automatically deleted after seven days.

Demo and Source

Our demo site is available at: <https://cs4530-f25-906.onrender.com/>. And our code at <https://github.com/neu-cs4530/fall25-project-fall25-project-group-906>

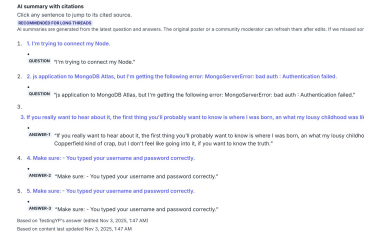


Tech stack and design

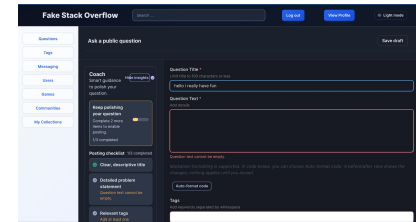
The front-end is built with react, typescript, chakra UI and vite. The user-defined hook usenewquestion is used to manage the status of questions. The content detection is realized by using the coach function module. The privacy vault is built with the help of localstorage and indexeddb to protect the privacy of drafts. The tests are carried out with vitest+testing library; The back end is based on node/express+mongodb extension, integrates openai API processing summary generation, supports indicator display through transparency API, builds indicator model and AI log model to ensure auditability, and both front and back ends follow the 7-day data retention strategy. The front end also pays attention to the consistency and accessibility of chakra component design, and the back end controls permissions through role verification.

Future Work

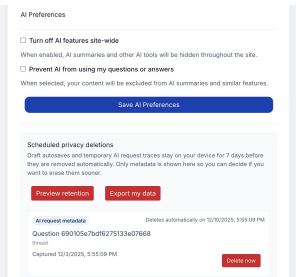
If we have more time, we will explore multiple development directions. In terms of coach pannel, we should preliminarily designed an optional parser based on syntax, which can recognize structured chapters (steps, error, environment), and provide more accurate label suggestions or format diagnosis without sending additional data to the outside of the device. In terms of summary function, we plan to support multilingual output, integrate reputation signals (such as giving priority to the content of adopted answers or high score authors), and add a single answer summary to the content intensive discussion string. In terms of privacy protection, we plan to expand the community indicator dimension of the transparency page, explore the development of users' personal dashboard to clearly explain when and why AI contacts its content, and strengthen role-based permission control for future management tools. These optimizations will continue to practice the core theme of the project: making huskyflow more convenient and efficient to use through small and well-designed tools, while ensuring that the author has absolute control over his own data.



AI generated summary of questions and answers along with citations for source of truth



Coach notifying you of the requirements to post and what you are missing, as well as how to fix it to meet all of them



Settings that enable you to turn off AI features sitewide, and prevent AI from using your content for privacy purposes, as well as the ability to erase AI request data and drafts

Draft automatically saved everything you ask a new question but didn't post

Coach panel for tile, descriptions and tags help