

Generative AI Tools Experiences in Computer Science Capstone
January 8, 2024
Michael Heroux
Scientist in Residence
Department of Computer Science
College of Saint Benedict & Saint John's University

During the Fall 2023 semester I again integrated generative artificial intelligence (AI) tools into the capstone course for Computer Science (CS) students starting with a strategy outlined in [this writeup](#). This semester, we focused on using ChatGPT 4.0. Many other tools are available, but none are as readily usable or helpful, in my experience.

The primary direct goal of the CS capstone course is producing a comprehensive research paper and presentation on a relevant topic of the student's choice and producing a working demonstration of what was learned in the form of a prototype software product or something similar. The primary learning goal of the course is to build skills for life-long self-education.

During the Fall 2023 semester, we used ChatGPT in many ways, including:

1. **Topic selection:** Using personas, contexts, and scenarios, students described their relevant personal experiences, education, and interests. They further provided a description of what was required for a successful capstone project (as described in the course handbook) and requested a list of ideas for what could be pursued as a good capstone topic. This activity, completed during the first class, was extremely valuable. It also led to some challenges we discuss below.
2. **Key concepts identification and explanation:** Once the students identified the topics they would pursue, ChatGPT was also very useful in identifying the next level of detailed understanding. Specifically, students could ask what key concepts underpinned their topic and could get explanations of these concepts using an intuitive question-and-answer dialogue.
3. **Concept reinforcement via quizzing:** One of the more interesting and useful exercises throughout the semester was prompting ChatGPT to present quiz questions on concepts. Students used this technique to privately assess and improve their understanding of their topic.
4. **Draft text generation:** Many, but not all, students used ChatGPT to generate draft text. Primarily, they generated a paragraph or two at a time and then edited the content to make it consistent and aligned with the larger narrative of their paper. This activity was useful but also problematic, as described below.
5. **Virtual course instructor:** Midway through the semester, OpenAI (the company that produces ChatGPT) enabled a new feature: The ability to generate a special-purpose version of ChatGPT (called a GPT) that was trained with additional information and guidance. I used this feature to create my own GPT by uploading the capstone course handbook along with some guidance on responding to prompts in ways that were helpful to the students. In effect, I had created a "virtual instructor." I published this GPT so students could ask course-specific questions and get feedback on draft papers. It

became a very useful resource. The GPT could advise students on how their paper addressed the writing rules of the course and how their paper could be improved.

6. **Demonstration prototype creation:** Every student used ChatGPT to assist in creating a demonstration of their knowledge in the form of a prototype software product. This activity was helpful but also presented challenges discussed below.

Observations

All the observations I discussed last semester (see [this writeup](#)) continued to be true this semester. Briefly, these were:

1. Assessing level of expertise via dialogue (Q&A) promotes the use of generative AI tools to improve student understanding.
2. Increasing the frequency of draft submissions appears to reduce incentives to submit content that students do not understand well.
3. Student ability to determine the topic and scope of their research improved substantially.
4. Students submitted papers with near-perfect grammar and spelling.
5. Students who did not acquire a deep level of expertise as evidenced by Q&A assessment wrote papers that were technically correct but very hard to read.
6. Students who used generative AI tools effectively were able to learn and produce more about their topic, more quickly, and with less effort than before these tools were available.

Further observations:

1. **Using ChatGPT to assist in identifying possible topics, learning about a topic, and reinforcing knowledge worked very well:** Every student benefited from the ability to quickly identify an interesting topic that matched their background and interests. In the past, I often had one or more students each semester who struggled to find a good and accessible topic. This challenge no longer exists.
2. **Some students struggled to build depth of understanding:** Because ChatGPT made topic selection so easy, some students did not fully understand their topic until late in the semester. Every semester, I have students who struggle to understand the deeper technical concepts of their topic. This year, I had quite a few students with this issue. I believe this happened for two primary reasons:
 - a. **Using ChatGPT to select a topic did not necessarily require students to build depth of understanding:** Prior to using ChatGPT, building depth of understanding occurred naturally as part of the topic selection process. Students would implicitly build a deep understanding of their topic as part of picking the topic.
 - b. **The first major paper draft was due before some students were ready to submit it:** I was unaware that some students had acquired only a surface understanding of their topic until their first paper was submitted. After discovering the shallow treatment of some papers, I asked students to explicitly state the technical theme of their topic and build a depth of understanding to address this theme.
3. **Students who relied too heavily on draft text generation created papers that were wordy and repetitive:** Most students used text generation modestly to generate a

paragraph from a bulleted list, polish content, or summarize a paragraph. In these cases, the text was eventually revised manually to be legitimately owned by the student. However, there were instances where students relied too heavily on text generation.

- 4. A focus on assessing expertise through interviews was essential and effective:** This semester, 40% of a student's grade was based on an oral assessment of their level of expertise. Several times during the semester, I would interview each student and provide a tentative level of expertise score from 1 to 10. A final level of expertise individual interview was the last course activity. Because level of expertise interviews were such an important part of the grade, students had very little incentive to use generative AI tools as a substitute for learning.
- 5. Use of generative AI tools dramatically reduced the time and effort to build expertise in a topic and to write, present, and effectively use this expertise:** We were able to produce better papers, talks, and prototypes more quickly and with less effort. We were able to use the remaining course time to explore new topics and reduce the workload at the end of the semester.
- 6. Students produced the same papers, presentations, and prototypes as in the past, but the processes we use are dramatically different than a year ago:** As generative AI tools and our abilities to use them continue to improve, I believe the economy of academic life will change dramatically. How much time we and our students spend on some (but not all) activities will dramatically decrease. In other cases, our role will shift from primarily generating content to primarily reviewing and revising it.
- 7. Thinking of generative AI tools as intellectual companions and administrative assistants provides inspiration for creative new uses:** During the past semester, we have discovered a variety of new uses for generative AI tools. Almost always, inspiration comes from considering how these tools could act as intellectual companions or administrative assistants. As a companion, we can converse with these tools in brainstorming sessions, ask wild and crazy ideas, look for connections between topics, and build our understanding of new concepts and terms. As an assistant, we can ask these tools for an analysis of a data set, reformatting of a table, converting bibliography entries to a new format, improving the tone of a paragraph, and more.

Future Plans

This semester I plan to actively address the depth of understanding challenge by asking students to clearly identify the specific computer science elements of their research as soon as it is feasible, to ensure they are building depth of understanding as an explicit goal.

We will also explore expanding course expectations. Last semester we were able to reduce time and effort to accomplish the same course objectives as in the past. This semester I want to explore what more we can do to improve the outcomes of the course. Some ideas include incorporating cross-discipline content that intersects with a student's primary research topic, to take advantage of the liberal arts courses each student has taken.