

WIP: Generative AI as an Instructional Resource in a Computer Science Ethics Course

Scott Barlowe

Mathematics and Computer Science
Western Carolina University
Cullowhee, USA
sabarlowe@email.wcu.edu

Daniel Aoulou

Mathematics and Computer Science
Western Carolina University
Cullowhee, USA
daoulou1@catamount.wcu.edu

Alex Ponce-Castillo

Mathematics and Computer Science
Western Carolina University
Cullowhee, USA
amponcecastillo1@catamount.wcu.edu

Abstract—This work-in-progress research paper reports our initial attempt at integrating generative artificial intelligence (Gen AI) into our two credit hour ethics course required for computer science majors. Course content includes ethical frameworks, presentations, current event and stakeholder analysis, formal debates, job seeking, and codes of conduct. Given the wide applicability of computer ethics and the crowded schedule in the course, our inquiry seeks to find ways of utilizing Gen AI to streamline content delivery, to provide opportunities for independent student exploration, and to aid students during preparation for class activities. In this paper, we first describe a novel assignment integrating Gen AI given to students enrolled in the Spring 2024 offering of our computer science ethics course. The findings from a survey addressing student use of Gen AI before and during the assignment and the analysis of assignment artifacts submitted by students are then reported. Finally, we present additional student data and results from our separate experimentation, both of which focus on the use of Gen AI for debate preparation. Our efforts reveal that Gen AI can be a useful instructional tool for a computer science ethics course but should be integrated carefully.

Index Terms—ethics, professional skills, computer science

I. INTRODUCTION

Topics and activities in our computer science ethics course emphasizing the impact of computing technology and the professionalism of future employees include well-known ethical frameworks, current event and stakeholder analysis, presentations, formal one-on-one debates, preparation for seeking employment, and the ACM Code of Ethics and Professional Conduct [1]. The effects of computing on society are wide-ranging, constantly changing, and necessitate the continual evolution of course content. There have been many attempts at improving and updating the coverage of how computing impacts society in computer science curricula, but the time and resources to adequately address the continually increasing range of applications and impacts are often insufficient. Generative artificial intelligence (Gen AI) has recently gained much attention for its increased capability to create dynamic content for specific contexts. The advances in Gen AI make it a potential solution for mitigating the time constraints in

a crowded ethics course in a way that allows independent, student-centered exploration.

In this paper, we report our initial efforts using Gen AI to address the challenges encountered when designing and creating activities to supplement content in our required, two credit hour computer science ethics course. Our inquiry specifically seeks to integrate Gen AI into new and existing course materials directly related to course learning objectives so that students can explore extensions of content and prepare for class activities in a timely, independent manner. After covering background material, we present a novel assignment from our computer science ethics course that integrates Gen AI, well-known ethical frameworks, and current event analysis. We then report the results from a follow-up survey of students currently enrolled in that course inquiring about their experience with and motivation for using Gen AI and provide an analysis of student work submitted for the assignment. Finally, we describe our initial investigation addressing the application of Gen AI to debate preparation.

II. BACKGROUND

A. Ethics in Computer Science Curricula

The multitude of impacts that computing can have on society has spurred much work in improving ethics content in computer science curricula. Barlowe and Scott [2] summarize some of the recent approaches. For example, content can be organized by discrete topics [3], [4] or embedded among courses representing concrete applications [5]. Another approach developed by Bullock et al. [6] relies on a repository of scenarios. Barlowe and Scott [2] note that discrete topics fail to address events and technologies that may be interwoven, embedding content may require curriculum flexibility and participation from other disciplines, and scenarios may need to be rewritten to adapt to new technologies and/or societal changes. Instead, they propose the use of current events as a primary tool for analyzing how society can be affected by computing. However, utilizing a diverse, representative set of current events in a coherent and timely manner that also blends well with other course components presents challenges.

This material is based upon work supported by the National Science Foundation under Grant Number (2207361). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

B. Generative Artificial Intelligence

One possible way to improve the delivery of content in computer science ethics courses is Gen AI. Gen AI relies on Large Language Models (LLMs) for creating dynamic responses to user prompts. The emerging field of *prompt engineering* [7], [8] is focused on formalizing diverse strategies for elucidating suitable responses. Although bias [9], hallucinations [10], intellectual property issues [11], and challenges related to academic integrity [12] have been reported in the news media, Gen AI platforms hold promise as an instructional tool [13]. Research is beginning to emerge reporting on student and instructor perceptions of Gen AI [14], but the use of Gen AI as an instructional tool presents many unanswered questions and no curriculum standards exist. Our work aims to help fill that gap with the development of instructional tools and practices specifically for computer science ethics courses.

III. METHODOLOGY

To better understand how Gen AI can be used as a pedagogical tool, we applied Gen AI to our required, two credit hour computer science ethics course during the Spring 2024 semester. Our inquiry included the development of a new assignment integrating Gen AI, analysis of both student survey responses and submissions for the assignment, and the application of Gen AI to debate preparation.

A. Assignment Design and Integration

The set of course learning objectives for our computer science ethics course includes (but is not limited to) application of foundational ethical theories and stakeholder analysis. Quinn [15] provides an introduction to many of the theories covered in class including Kantianism, subjective relativism, act utilitarianism, rule utilitarianism, virtue ethics, and social contract theory. The range of perspectives provided by a survey of ethical frameworks helps prepare students for analyzing scenarios and events that reflect current societal and computing trends. Although this is an effective approach, there are significant challenges. In our course, students do not have a standard background in ethics, instructors do not have formal training in ethical frameworks, and the breadth of activities creates limits for how much time can be spent exploring any single topic.

In our investigation, no class time was spent on introducing these frameworks. Instead, an assignment was developed leveraging Gen AI. The class of 24 students was divided into groups of two (although two students worked individually for various reasons) and each group was assigned an ethical framework to research. Frameworks included those previously listed and ethical altruism [16].

The assignment required each group to independently research an article of interest. The article must have been published within the last two years by a reputable source and report on a specific event or societal condition impacted by computing. Groups were then asked to complete a series of exercises outside of class that alternated between requiring the use of Gen AI and prohibiting the use of Gen AI. Several

links to resources for prompting Gen AI were provided on the assignment (see [17] and [18]) but no class time was spent on how to craft effective prompts. A rubric was developed that emphasized a systematic prompt strategy of each group's choice or design. Students were asked to limit their use of Gen AI to ChatGPT [19], specifically version 3.5, for consistency and to submit their prompt-response dialogues with the platform's sharing functionality. An option was provided for an alternate way to access the platform if any of the students did not want to create ChatGPT accounts. The tasks are listed below in required order of completion.

1) Introduction to Ethical Frameworks

With Gen AI. The goal of the first activity was to introduce groups to an ethical framework. Each group was asked to prompt Gen AI for a one to two paragraph summary of their assigned framework.

Without Gen AI. Groups then constructed a one to two paragraph hypothetical scenario illustrating the assigned framework and listed the stakeholders that students thought would be directly and indirectly impacted.

With Gen AI. Students used Gen AI to produce a one to two paragraph scenario illustrating the application of the assigned framework and to analyze the stakeholders as in the previous step. Groups were then asked to use Gen AI to further refine the scenario to include edge cases, the effect of potential laws on stakeholders, and advice to stakeholders.

2) Current Event Research

Without Gen AI. Groups were asked to find a current event caused or impacted by computing and to construct a one paragraph summary of the article. Groups performed stakeholder analysis on the news article.

With Gen AI. Gen AI was then prompted with the summary to produce a list of stakeholders and the rationale for why each is a stakeholder. Gen AI was also prompted to produce an analysis of the article and to extend the utility of the article in the same way as they refined the scenario, except this time groups did so in the context of the assigned ethical framework.

3) Reflection

Without Gen AI. Groups were asked to construct a list of stakeholders provided by Gen AI that the students did not identify and another list of stakeholders provided by the group that Gen AI did not identify. Students were also asked to report any insights that Gen AI provided that the student had not considered and to evaluate the advice of AI to stakeholders.

The assignment was followed by a class exercise. Each group was paired with another group such that different ethical frameworks were represented. Groups discussed all parts of the assignments and then rotated so that these discussions were repeated with variations in assigned ethical framework.

B. Survey and Submission Analysis

The 22 students present on the day following the class exercise were given an anonymous survey with 11 questions.

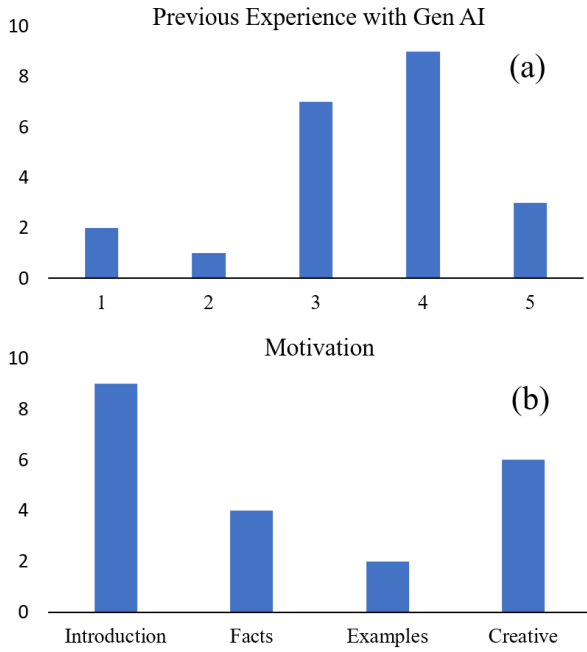


Fig. 1. Number of students per self-reported (a) experience level on a scale from 1 (no experience) to 5 (very experienced) and (b) motivation for using Gen AI.

The survey inquired about students' overall experience level, motivation, and strategies when using Gen AI as well as their perspectives of how Gen AI was applied to the assignment. Student submissions for the assignment were also analyzed.

C. Debate Preparation

After completing our computer science ethics course, students should be able to research, articulate, and defend logical arguments focusing on computing technology. One-on-one debates are a primary (but not the sole) course component used to help address this learning objective. Pairs of students are assigned a topic and each student is responsible for presenting and defending either an affirmative or negative position during class. Topics and positions are assigned randomly. Students during the Spring 2024 offering were explicitly given permission to use Gen AI for preparation, and another survey consisting of five questions was distributed inquiring about student use of Gen AI for this assignment. The authors also performed separate experimentation on debate formulation and preparation.

IV. RESULTS

We now report the results of our investigation. Results from the initial 11-question survey are presented first and followed by our analysis of student submissions for the assignment. Findings from our activities that focused on debate preparation, including an additional survey and separate experimentation by the authors, are presented last.

A. Survey Results

Experience and Motivation. Students rated their previous experience with Gen AI before the assignment on a scale ranging from 1 (no experience) to 5 (very experienced). 86% of students rated their experience as a three or higher and 55% of the class rated their experience as a four or higher. 14% of the students rated themselves as very experienced. When asked to choose their most frequent motivation for using Gen AI, 43% of the students stated they used it for a general introduction to a topic, 19% used it primarily for researching specific facts, 10% used it mostly for providing specific examples, and 29% used Gen AI most frequently for problem solving or to generate other creative content. Figure 1 shows the number of students for each self-reported experience level and motivational factor.

Prompt Strategy. Students were allowed to use whatever method they wanted to prompt Gen AI. Students rated their effort in researching a specific prompt strategy on a scale ranging from 1 (no effort) to 5 (great amount of effort). Exactly half the respondents rated their effort in researching a specific prompt strategy as either a one or two. Only 2 students stated that they spent a great amount of effort.

The most popular approach (employed by 10 students) to prompt Gen AI was to follow a general question with more detailed questions. Only three students mentioned that they requested Gen AI to "act like" something or indicated that they treated Gen AI as a "colleague" or "human". The remaining students described their actions as simply rephrasing or repeating the question that the assignment wanted them to answer.

When asked if Gen AI produced any incorrect or erroneous responses, 16 students (73%) reported that there were no such responses. However, students also reported only a moderate amount of effort in verifying answers (average of 3.04 on a scale ranging from 1 to 5). The majority of students that reported issues with response accuracy thought that the answers from Gen AI were either too broad or contradicted previous responses. Only four students reported their temptation to use Gen AI where prohibited by the assignment as a four or five.

B. Analysis of Student Submissions

The part of the assignment that asked students to reflect on how well Gen AI performed was analyzed. All but one group (12 out of the 13) reported that Gen AI had presented one or more stakeholders not listed by their group. The same number of groups reported that their group had listed one or more stakeholders that Gen AI had not listed.

For example, one group chose an article about recent layoffs in the tech industry (see [20]). Stakeholders omitted by the group but provided by Gen AI included managers and executives, customers of two well-known tech companies, competing tech companies (in general terms), and local communities. Two other stakeholders provided by Gen AI included educational institutions and regulatory bodies. The rationale provided by Gen AI for educational institutions is below:

"Changes in the job market, including layoffs, can impact the demand for certain skills, potentially affecting enrollment and program offerings."

For regulatory bodies, Gen AI stated that they

"...may be indirect stakeholders as they monitor employment trends and may need to respond to shifts in the job market. Layoffs could also influence policy discussions related to job security and workforce stability."

Stakeholders provided by the group omitted by Gen AI included families of laid off workers and suppliers/service providers.

Another group chose an article about employee surveillance (see [21]). In the process of the assignment, students had posed the prompt:

"How do you think the software used to track employees [can] be used in ways that are not anticipated by the developers?"

As part of the response and to the surprise of the students, Gen AI had listed the effect on employee creativity as an unintended consequence of concentrating too heavily on quantitative metrics.

Half of the respondents (11 students) thought ChatGPT was ineffective when asked to make moral decisions or to give advice to stakeholders in the assignment. Students thought the advice was good, but too vague to be useful. For example, part of ChatGPT's advice to human resource departments addressing employee surveillance in the context of social contract theory was to

"Conduct training sessions for both employees and management on the ethical use of surveillance software"

without specifically defining ethical use. Part of the advice given by ChatGPT to legislative bodies for the same group was to

"Create comprehensive regulations governing workplace surveillance and ensure they are enforced"

without describing any specific regulatory or enforcement policies.

C. Debate Preparation

The post-debate survey revealed that 14 out of the 21 responding students (approximately 67%) used some type of Gen AI to prepare for the debate. One student used a platform that provides facial reactions to words spoken by the user and one student used it for research. The remaining 12 students used Gen AI to provide potential arguments and/or counterarguments that could be used in the debate. Of the seven students who stated that they did not use Gen AI for debate preparation, four simply felt that they did not need the additional help, two felt that Gen AI would not have helped, and one wanted to talk to peers instead of interacting with Gen AI.

The authors performed additional experimentation applying Gen AI to debates. Two sets of prompts were given to ChatGPT 4. The first set of prompts was intended to investigate general preparation for the topic (see [22]) and the second set of prompts was intended to simulate the debate environment (see [23])

Prompt Set 1. We directly asked Gen AI for debate topics and a debate focusing on Sora [24], an AI-based video and scene creation platform. ChatGPT provided ten potential topics framed as questions. For the second part, ChatGPT provided the debate topic, the structure of the debate with four main parts, and the main points that each position could take for each section. ChatGPT relied on short, itemized responses.

Prompt Set 2. In a new session, we first asked ChatGPT to pretend to be a student in a computer science ethics class and then outlined a potential debate structure. ChatGPT asked the user for a debate topic and the stance Gen AI should assume. The platform assumed the persona of a debate participant and even addressed a hypothetical audience. The opening statement provided an introduction, three paragraphs describing the main points for the position, and a conclusion. ChatGPT later directly rebutted each of the points provided by the user. This process was repeated for the duration of the debate. ChatGPT described the debate as *"enriching"* and *"insightful"* and ended by thanking the user.

We compared ChatGPT's responses to Google's Gemini [25] platform. For Prompt Set 1 (see [26]), Gemini provided potential topics grouped by focus area (understanding and sentience, impact and accessibility, regulation and ethics, and applications and future). For the second part of Prompt Set 1, Gemini simulated the debate. The major parts of the debate were divided into sections (opening, body, rebuttal, and closing) with each section providing arguments for an affirmative position and for a negative position. In this part, Gemini exhibited two styles. Gemini provided the opening statement by assuming the persona of a team for each position. In the remainder of the debate, the conversational style was still present, but overshadowed by itemized lists for each team.

Gemini was also given Prompt Set 2 (see [27]). In this section, the user's input during the debate was given in developed paragraphs, much like that given to ChatGPT. Although Gemini sometimes referred to the debate participants, the response for this prompt set was again heavily reliant on itemized lists. Gemini did not provide a closing statement.

V. DISCUSSION

The assignment fit well into the course structure by allowing students to apply Gen AI in a setting where the effect of artificial intelligence on society was emphasized by class presentations, debates, and discussions. Furthermore, no class time was spent on ethical frameworks except for when students were actively engaged in class discussion. Groups were also able to independently use Gen AI to explore and analyze a wide array of current events of their own choosing.

Student survey responses, submissions for the assignment, and our separate investigation revealed several interesting

findings. Students reported using Gen AI mostly to generate definitions, lists, and similar output. Students were able to elucidate meaningful insights from Gen AI that had not been considered. However, students thought that many of the responses to more complex requests, such as asking for stakeholder advice, to be vague. Our separate experimentation for debate preparation revealed stylistic differences between platforms, especially when asking Gen AI to become a debate participant. For debate simulation, ChatGPT offered a more conversational style and Gemini provided mostly lists.

Instructors should weigh some of the open questions and unsolved problems before integrating Gen AI. Similar to assignments to be completed outside of class in many other subject areas, the assignment presented here has no safeguards to ensure that students refrain from using Gen AI when instructed to do so. Furthermore, students reported spending little effort researching how to craft well-designed prompts and verifying the accuracy of responses even though they perceived themselves as experienced Gen AI users. For debates, students may rely too heavily on memorizing and repeating the points from Gen AI instead of their own critical analysis. Finally, course goals - as in the case of our ethics course - are often supported by many activities, and determining ways to attribute fulfillment of learning objectives solely to the use of Gen AI is challenging.

VI. CONCLUSION AND FUTURE WORK

Our analysis of survey responses, student submissions for the assignment, and our own experimentation shows that Gen AI can be an effective supplement to a computer science ethics course. Future work in our computer science ethics course consists of improving student fluency with prompt engineering without sacrificing class time, experimentation with finding the best use cases for a given prompt strategy, exploring ways to assess student submissions when using Gen AI, and ensuring that students use Gen AI as a supplement instead of a replacement for critical analysis. We also plan to investigate ways of measuring the improvement of student performance solely attributable to Gen AI.

REFERENCES

- [1] The acm code of ethics and professional conduct. Accessed: May 20, 2024. [Online]. Available: <https://www.acm.org/code-of-ethics>
- [2] S. Barlowe and A. Scott, "News consumption among cs majors: Habits, perceptions, and challenges," in *Proceedings of the 2023 ACM Southeast Conference*, ser. ACMSE 2023. New York, NY, USA: Association for Computing Machinery, 2023, p. 162–166. [Online]. Available: <https://doi.org/10.1145/3564746.3587023>
- [3] S. Baase and T. Henry, *Gift of Fire, A: Social, Legal, and Ethical Issues for Computing Technology*, 5th ed. Pearson, 2017.
- [4] C. Fiesler, M. Friske, N. Garrett, F. Muzny, J. J. Smith, and J. Zietz, "Integrating ethics into introductory programming classes," in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 1027–1033. [Online]. Available: <https://doi.org/10.1145/3408877.3432510>
- [5] B. J. Grosz, D. G. Grant, K. Vredenburgh, J. Behrends, L. Hu, A. Simmons, and J. Waldo, "Embedded ethics: Integrating ethics across cs education," *Commun. ACM*, vol. 62, no. 8, p. 54–61, Jul 2019. [Online]. Available: <https://doi.org/10.1145/3330794>
- [6] B. B. Bullock, F. L. Nascimento, and S. A. Doore, "Computing ethics narratives: Teaching computing ethics and the impact of predictive algorithms," in *Proceedings of the 52nd ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 1020–1026. [Online]. Available: <https://doi.org/10.1145/3408877.3432468>
- [7] V. Liu and L. B. Chilton, "Design guidelines for prompt engineering text-to-image generative models," in *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–23.
- [8] Y. Hu, Q. Chen, J. Du, X. Peng, V. K. Keloth, X. Zuo, Y. Zhou, Z. Li, X. Jiang, Z. Lu, K. Roberts, and H. Xu, "Improving large language models for clinical named entity recognition via prompt engineering," *Journal of the American Medical Informatics Association*, p. ocad259, 01 2024. [Online]. Available: <https://doi.org/10.1093/jamia/ocad259>
- [9] A. Sasani. (2024, March) As ai tools get smarter, they're growing more covertly racist, experts find. [Online]. Available: <https://www.theguardian.com/technology/2024/mar/16/ai-racism-chatgpt-gemini-bias>
- [10] C. Thorbecke. (2023, August) Ai tools make things up a lot, and that's a huge problem. [Online]. Available: <https://www.cnn.com/2023/08/29/tech/ai-chatbot-hallucinations/index.html>
- [11] G. Appel, J. Neelbauer, and D. A. Schweidel. (2023, April) Generative ai has an intellectual property problem. [Online]. Available: <https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem>
- [12] L. Coffey. (2024, July) Students and professors believe ai will aid cheating. [Online]. Available: <https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2024/07/29/students-and-professors-expect-more>
- [13] P. Lauren and P. Watta, "Work-in-progress: Integrating generative ai with evidence-based learning strategies in computer science and engineering education," in *2023 IEEE Frontiers in Education Conference (FIE)*. Los Alamitos, CA, USA: IEEE Computer Society, oct 2023, pp. 1–5. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/FIE58773.2023.10342970>
- [14] C. Zastudil, M. Rogalska, C. Kapp, J. Vaughn, and S. MacNeil, "Generative ai in computing education: Perspectives of students and instructors," in *2023 IEEE Frontiers in Education Conference (FIE)*. Los Alamitos, CA, USA: IEEE Computer Society, oct 2023, pp. 1–9. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/FIE58773.2023.10343467>
- [15] M. J. Quinn, *Ethics for the information age*, 6th ed. Pearson, 2015.
- [16] J. Seglow, *The Ethics of Altruism*, 1st ed. Taylor & Francis, 2004.
- [17] OpenAI. Prompt engineering. Accessed: May 20, 2024. [Online]. Available: <https://platform.openai.com/docs/guides/prompt-engineering>
- [18] General tips for designing prompts. Accessed: May 20, 2024. [Online]. Available: <https://www.promptingguide.ai/introduction/tips>
- [19] OpenAI. Chatgpt. Accessed: May 20, 2024. [Online]. Available: <https://chat.openai.com/>
- [20] E. Buchwald and J. Zarracina. (2023, February) Massive tech company layoffs look ugly. but it may not be as bad as you think. [Online]. Available: <https://www.usatoday.com/story/money/economy/2023/02/07/tech-layoffs-2023-google-amazon/11135692002/>
- [21] G. C. Tong. (2023, April) Employee surveillance is on the rise — and that could backfire on employers. [Online]. Available: <https://www.cnbc.com/2023/04/24/employee-surveillance-is-on-the-rise-that-could-backfire-on-employers.html>
- [22] S. Barlowe, D. Aoulou, and A. Ponce-Castillo. Chatgpt debate preparation. <https://agora.cs.wcu.edu/~sbarlowe/genai/CHATGPTPREP.pdf>. Accessed: May 20, 2024.
- [23] —. Chatgpt debate simulation. <https://agora.cs.wcu.edu/~sbarlowe/genai/CHATGPTSIM.pdf>. Accessed: May 20, 2024.
- [24] OpenAI. Sora. Accessed: May 20, 2024. [Online]. Available: <https://openai.com/sora>
- [25] Google. Gemini. Accessed: May 20, 2024. [Online]. Available: <https://gemini.google.com/app>
- [26] S. Barlowe, D. Aoulou, and A. Ponce-Castillo. Gemini debate preparation. <https://agora.cs.wcu.edu/~sbarlowe/genai/GEMINIPREP.pdf>. Accessed: May 20, 2024.
- [27] —, "Gemini debate simulation," <https://agora.cs.wcu.edu/~sbarlowe/genai/GEMINISIM.pdf>, accessed: May 20, 2024.