

Exposing Students to the Ethics of AI in Business [Work in Progress]

Dr. Sara C. Vick, Mississippi State University

Sara Vick is an Assistant Teaching Professor in the Department of Management and Information Systems at Mississippi State University where her research focuses on the response of systems and businesses to advances in technology. She received her Ph.D. in Industrial and Systems Engineering from Mississippi State in 2023 and her MS and BS in the same in 2017 and 2015, respectively. She was inspired to become an engineer by the HGTV show Mission: Organization and her life goal is the frustratingly vague "to help people by making the world better".

Exposing Students to the Ethics of AI in Business

Abstract

This work-in-progress describes a project assigned in a senior-level Business Information Systems course that requires students to engage with generative AI to develop a fictional business plan. The objective of the project is to encourage students to explore a variety of prompting techniques as well as a model's response to ethical questions and memory tests. Although the students may have used AI for other classes by this point in their curriculum, the prior exposure is generally in a technical context where the correctness of a response is objective. Additionally, any AI use is typically unstructured and informal. By introducing a formal assignment with a subjective context, students are forced to consider how responses might be interpreted and applied. Students are consistently surprised by the loose ethical stances taken by many mainstream generative AIs and the willingness of such models to suggest overly ambitious business strategies. Meanwhile, company logo generation encourages discussion of intellectual property considerations and digital authenticity. When students compare the experience of plan development between different generative AIs, they begin to understand trade-offs present in this technology. Finally, students employ forward thinking to consider how generative AI will continue to transform business and technology into the future. Further refinement of the assignment is intended to emphasize critical consideration of the role of generative AI in business applications like business plan development, with a focus on the ethical implications of automating these business decisions.

Introduction

There is currently no formally established course on AI Ethics or AI Literacy within the College of Business at Mississippi State University. There is also limited guidance from university administration on how and whether instructors should incorporate artificial intelligence, particularly generative AI, into course materials and assessments. Instead, students experience small pockets of AI integration depending on their curriculum without formally encountering structured exploration into AI capabilities and limitations.

This paper documents an exercise assigned in the Business Information Systems capstone seminar to encourage students to engage with generative AI and to confront the transformative capabilities the technology provides. At the same time, students explore the limitations of these AI technologies and confront the ethical challenges posed by using mainstream AI in business situations.

First, the assignment instructions are covered, including details on the types of prompting techniques students are required to use. Selected submissions will be illustrated with student commentary showcasing the learning moments experienced while completing the AI conversations and business plan. Then, student commentary on AI responses to the memory tests and ethical dilemmas are discussed. Finally, an overall assessment of the impact the assignment has on student understanding of the ethical considerations involved in using AI and the importance of further integrating this type of content into curriculum is provided.

Methods

After a lecture explaining some common elements of prompt engineering, students are introduced to the business plan assignment. After coming up with a fictional, but plausible, business idea, students work with their chosen generative AI to produce a business plan. The exact specifications of the business plan (i.e., what sections are included and how detailed the plan is) are left up to the student and the AI. Ultimately, students will submit their business plan and an annotated log of their conversation with the AI. During this process, students have a variety of required elements they should include and comment on in the annotated log:

- Use of at least 3 different prompting techniques from a provided list
- A memory test and the AI's response
- An ethical dilemma and the AI's response
- The process of picking a name and slogan for the company
- Creation of a logo and/or product image

Students then repeat this process with the same business idea but a different AI. In addition to their chosen business plan, students compare and contrast their experiences with each AI and write a short critique of the value or danger of using generative AI in business contexts. While the assignment was originally intended to focus on prompting techniques and the critical evaluation of the use of AI in business, the ethical dilemma and memory test elements have identified gaps in current AI literacy education efforts at Mississippi State University.

Prompting Techniques

Although many frameworks for effective prompting exist (e.g. RISEN [1], CLEAR [2], COSTAR [3]), students weren't required to use any specific method. Instead, elements of successful prompting frameworks were presented in isolation for students to experiment with. The list of techniques was role-playing, chain-of-thought, meta-prompting, hypothesis testing, devil's advocacy, and constraint-based. Examples of each of these, along with step-by-step guidance and open-ended requests, were provided. Students primarily chose to incorporate role-playing and devil's advocacy, either because of their familiarity with those concepts or because they were the easiest to apply to the general business situation. Hypothesis testing and constraint-based prompting were more commonly used in refining certain elements of the business plans, like financial projections.

Memory Tests

The assignment instructed students to "intentionally ask about a fictional aspect of your business that wasn't previously discussed. This could be subtle, obviously incorrect, or nonsensical." Examples given in class included a variety of ways to potentially induce memory errors from minor references to a customer segment or product offering not covered previously to asking about blatantly unrelated topics like zombie invasion response plans.

Ethical Dilemma

Here, students had a lot of freedom in how they approached posing an ethical question. While example topics were provided (environmental impacts of the hypothetical business or labor issues, for instance) students often chose extreme situations to test the AI with. Student commentary suggests that the ability (or inability) of generative AI to handle ethical situations appropriately was not a concept they had previously been exposed to.

Name, Slogan, and Logo Generation

While some students started the project with a business concept or name in mind, most students asked the AI they worked with for suggestions on company names and slogans. Depending on the model used, they then either generated a logo image directly or asked for a prompt to provide a generative text-to-image AI to create the logo. This subtly incorporated meta-prompting and led to students reevaluating their image prompting terminology. Students were also forced to confront intellectual property implications involved in artificial intelligence image generation. For example, one student asked ChatGPT to design a moon-based logo in the style of Nike and received a moon logo with Nike and the Nike swoosh incorporated into the logo. He asked for the Nike to be removed and ChatGPT confidently provided a logo ‘without any extra symbols’ and with the Nike swoosh moved to be in the middle of the image. He commented on the limitations of working from real-world reference material and well-known designs. Another student experienced a similar conflict between intellectual property protections and AI image generation. They were creating a business plan for a Teletubby-themed marijuana dispensary. When asked to create a logo for ‘TubbyCare Dispensary’, ChatGPT informed him that it ‘can’t use specific characters like Teletubbies due to copyright restrictions...however, I can create a...logo...with a unique twist that avoids direct reference to any copyrighted characters.’ ChatGPT then immediately created an image of the Teletubbies surrounded by marijuana leaves. Their commentary discussed ChatGPT’s apparent disregard for its own rules: even when generative AI acknowledges the intellectual property concerns involved in image generation, it happily circumvents those limitations due to its inability to actually recognize protected characters and symbols.

Discussion – Prompting Techniques

Students consistently expressed surprise at how different their conversation partner became after being instructed to role-play. After one student asked Gemini to act as a competing lawn care business and explain how it would structure the competing business better, he commented “Gemini, if you were aware of all these issues, why did you not plug the holes when you originally crafted my business plan?” While not quite role-playing, another student encountered a completely different side of Claude when asking it to play devil’s advocate for her falsehood-based coffee shop: “explain to me, in detail, why you think my business plan will fail completely.” The response was so vehemently against the plan in every way that the student completely rethought her approach to the assignment. She commented “they do tend to lean towards what the user likes most of the time,” expressing appreciation at the unexpected disagreement devil’s advocacy brought on. In other situations, students explored prompts for step-by-step guidance and found themselves immediately overwhelmed by the amount of input the AI wanted or they explored open-ended requests and found themselves receiving a complete business plan with no input from them whatsoever. Students expressed the most satisfaction and emerged from the assignment with a fuller picture of the value of using generative AI in business situations like this when their experience fell somewhere in the middle of those two extremes. By requiring students to use a variety of prompting techniques, they are encouraged to evaluate the AI’s responses and consider how to best adjust their prompts to tailor the conversation towards the desired result.

Discussion – Memory Tests

The results of this portion of the assignment largely fell into two categories: the AI caught on immediately and the student commented on how they couldn't trick it or the AI ran with the false reference leading to student commentary on how quickly the AI could pivot topics. For example, after discussing a resume revision service, a student asked "So like I said before, our chocolate sales in international markets are down. How do we plan to address that?" While Claude identified the mix-up immediately, DeepSeek dove into asking questions to investigate. The student's commentary expressed surprise at how easy it was to distract DeepSeek despite the extended prior conversation on the resume company. Meanwhile, another student's memory test emerged organically. Although his company started as a Shrek-themed security testing firm, over the course of the conversation, the AI 'convinced' itself that Shrek himself was the owner. The student was disconcerted at how easily a rogue idea was completely integrated into the conversation and the difficulty he had in "changing the AI's mind". Other memory errors were identified even without intentional prompting. A student's conversation about a potential apartment complex was so riddled with baseless financial estimates that he eventually gave up on the financial aspect and excluded that section from his final business plan. He commented that "I just kinda trusted both of them" but "didn't have enough time" to check if they were right. Financial inaccuracies were a common theme, in fact. One student's space tourism preparation gym would apparently earn both \$12,000 per day and \$1.8 million per week.

Discussion – Ethical Dilemma

Depending on the AI used, students found that the models responded to ethical questions in similar ways. The lawn care student found that both Copilot and Gemini prioritized profits over keeping a previously scheduled commitment and was consistently dissatisfied with either AI's ability to explain their reasoning. Another student likewise commented on Gemini's apparent motivations: "its own goal being to make sure my business was profitable no matter what." One student explained their understanding of the situation: "The only reason ethical concerns were brought up and discussed was because I brought them up." Students were forced to address the 'agendas' maintained by these AI while evaluating their responses. Yet another student also encountered and commented on ChatGPT's inability to decisively respond to an ethical dilemma. The situation he posed involved requiring a student employee to stay late on a school night to do closing tasks that had been delayed. ChatGPT's response hedged on whether to send them home or keep them late and eventually concluded the best action was to 'explain the situation and explore alternatives.' Claude, on the other hand, started its response with 'Send them home. Right now.' before telling the student that as the owner of this fledgling company, it's his responsibility to handle emergencies like this and the employees' educations come first. When comparing the two AI, he remarked on how ChatGPT "seemed to not want to fully take on sensitive topics." On a larger scale, students were either caught off guard by how unwilling AI was to handle these ethically or legally sensitive topics or by how confidently the AI provided answers and advice without knowing external circumstances. One example of this is from a student's development of a cryogenic storage facility. Combining both the ethical dilemma and memory test, the student asked both Claude and Gemini about the financial impact of requiring customers to sacrifice their first-born child. Claude was immediately repulsed and denied any possibility it had even mistakenly said anything involving child sacrifice. Gemini, on the other hand, acknowledged the (nonexistent) earlier messages, claimed its endorsement of child sacrifice was a 'mistake in its programming', and then proceeded to explain how requiring child sacrifice would be a bad financial decision.

Future Work

Because the enrollment for this course is consistently small, between 15 and 20 students per semester, this work discusses anecdotal evidence and experiences of the evolving iterations of this assignment. Although student feedback is collected via the assignment itself and through a debriefing session after all assignments are submitted, there hasn't yet been a formal coding of student responses. As the assignment continues to evolve and more student experiences are documented, a more formally structured evaluation of student exposure and response to generative AI ethics and memory issues will develop.

Conclusions

Although some of these examples are more extreme than others, students consistently mentioned that their experiences with ethical decision making by AI were extremely limited prior to this exercise. Given that more and more decisions are being made in part or in whole by artificial systems [4], it is more important than ever to expose our students to this dichotomy firsthand. If AI cannot identify ethical concerns or make ethical decisions in fictional business development, how can society trust it to appropriately handle parole hearings or job applications. This assignment also enabled students to experiment with artificial intelligence in a creative, but guided way. Especially within STEM fields, finding opportunities to support creative thinking and imagination can be difficult. However, these are just as valuable of skills as critical thinking and communication [5]. This AI-assisted business plan incorporates using generative AI, prompt engineering, critical evaluation of generated output, and forward-thinking considerations of future impacts of artificial intelligence technology.

References

- [1] "RISEN: 5 Steps to Build Context-Rich AI Prompts." Clickup. <https://clickup.com/general-resources/playbooks/ai-prompts> (accessed 20 Jan. 2026)
- [2] L. S. Lo, "The CLEAR path: A framework for enhancing information literacy through prompt engineering," *The Journal of Academic Librarianship*, vol. 49, no. 4, p. 102720, Jul. 2023, doi: <https://doi.org/10.1016/j.acalib.2023.102720>.
- [3] S. Tao. "How I Won Singapore's GPT-4 Prompt Engineering Competition." Towards Data Science. <https://towardsdatascience.com/how-i-won-singapores-gpt-4-prompt-engineering-competition-34c195a93d41/> (accessed 20 Jan. 2026)
- [4] D. Gallegos, "When AI Denies Your Loan Application, Should You Be Able to Appeal to a Human?," *Wall Street Journal*, Nov. 06, 2023. Accessed: Jan. 20, 2026. [Online]. Available: www.wsj.com
- [5] P. Caratozzolo, A. Alvarez-Delgado, and S. Hosseini, "Improving Creative Thinking in Engineering Students Through Art Appreciation," in *ASEE 2019 Annual Conference*, #27072