Yale JACKSON SCHOOL OF GLOBAL AFFAIRS

Designing Writing Assignments in the Era of AI
David Morse, Writing Program Director

Large language models (LLMs) such as ChatGPT present a new challenge to educators. Almost instantaneously, LLMs are able to complete a variety of writing assignments with a middling to moderate degree of success. Students can therefore avail themselves of LLMs to quickly complete work that might have otherwise taken them a dozen hours or more to complete.

Unfortunately, using a program like Turnitin or ZeroGPT to detect student use of AI is problematic for several reasons. Therefore, enforcing a ban on LLMs is not currently feasible. Nor is such a ban necessarily desirable. LLMs will inevitably be a part of students' professional lives; it may therefore behoove them to become familiar with LLMs.

Jackson faculty should therefore determine the extent to which they will permit the use of LLMs in their classroom. Faculty could plan to help familiarize students with ethical (which is to say, limited) use of LLMs to improve students' skills using the tool. And/or faculty could attempt to minimize or limit student use of LLMs. Either way, faculty need to think carefully about how they will craft assignments and exams with these goals in mind.

Regardless of your stance on this issue, it is important that you <u>explicitly state your policy on the use of LLMs in your syllabus</u> and discuss this policy in class. This document is one of many resources available to help faculty form and articulate their LLM policies; for further reference, please see in particular the excellent guidance from Yale's Poorvu Center for Teaching and Learning: https://poorvucenter.yale.edu/AIguidance#3%20ChatGPT%20on%20Syllabus. The website includes a thorough overview of LLMs, guidance for assignment design, helpful links, and suggestions for the language that instructors might use to articulate LLM policies on their syllabi.

This is new territory for all of us. Please experiment with the suggested methods below and try methods of your own. And do note which approaches seem to work well, which work less well, and why—and report back to the Jackson community (and the Writing Center specifically). We're all learning here.

THREE APPROACHES

1. Develop students' LLM literacy
By introducing a critical, targeted approach regarding the use of LLMs , instructors reinforce to students the inherent value of producing their own written work, while helping students become familiar with what can be a useful learning tool. Here are a few ways that instructors might incorporate LLMs into a suite of assessments.

<u>Critique LLMs' work</u>

You might ask students to first complete a writing assignment on their own, then have them submit the prompt to an LLM and compare and contrast their draft with that of the LLM. You could also skip this first step and simply have students critique LLMs' responses to a given prompt. In either case, ask students to critique the LLM's work in terms of logical reasoning, accuracy, structure, and harmful stereotypes. At this time, LLMs' primary strength, from the point of view of Jackson students, is likely its ability to almost instantaneously survey a vast amount of information and consolidate that information into a coherent document. However, at least for the time being, because LLMs are essentially content aggregators, the writing produced by LLMs tends to be generic and formulaic. LLMs struggle to put forward an original or creative argument. LLMs can also struggle to put forward convincing, detailed evidence to support claims.

According to Anna Mills at the College of Marin, "Showcasing AI failings has the added benefit of highlighting students' own reading, writing and thinking capacities. We can remind them that they are learning to understand and to express themselves with purpose, things a language model cannot do. Draw attention to the virtues of human-written prose and prompt students to reflect on how their own cognitive processes surpass AI."[1]

Create a bibliography

Earlier iterations of ChatGPT were notorious for creating false source materials. This problem seems to have been mostly minimized if not eliminated. Students could be encouraged to use an LLM to suggest a reading list on a given topic, thus potentially saving them hours of work and providing them access to a variety of sources they might not have been able to discover otherwise. Of course, we don't want students to rely exclusively on LLMs to find their sources for them: the hard work of research, which requires perseverance and creativity, is a skill that it behooves students to learn. If you do permit students to use LLMs to help create a bibliography, be sure to remind them of LLMs' potential for error: all suggested sources must be confirmed.

Brainstorm ideas

Students often struggle to transition from broad ideas for a topic to something specific that might be suitable for an in-depth analysis. LLMs can help with that. Students could be encouraged to have a conversation with an LLM, prompting the LLM for various ideas related to a more general topic. If the LLM suggests a topic that the student finds intriguing, the student could use the LLM to explore the topic further.

Solicit feedback

Students benefit from writing and revising multiple drafts—but it can be a challenge, given time and logistical constraints, to provide them with the feedback they need at each step of the writing process. Students could be asked to write an initial draft of a paper, submit the paper to an LLM for critique, then revise the paper in response to the critiques that the students found most relevant or helpful. To ensure that LLMs are not leading students

---

[1] https://libguides.asu.edu/c.php?g=1311696&p=9743203

astray, you could ask students to submit their original draft, the LLM's criticism, and students' revision.[2]

2. Minimize students' incentives to cheat

According to Yale's Poorvu Center, "students are less likely to cheat when they:

- Are pursuing questions they feel personally connected to
- Understand how the assignment will support their longer-term learning goals
- Have produced preliminary work before the deadline
- Have discussed their preliminary work with others
- Make their writing process visible in the completed assignment."[3]

One of the best ways to meet many of these criteria is to break down a larger assignment (for example, an assignment due at or near the end of the semester) into smaller component parts, each part building on the previous part. For example, a decision memo might be broken down into a backgrounder and/or analysis of key findings, options analysis, and recommendations. Students would receive critical feedback either from their instructor or their peers (or an LLM) with each part of the whole. At the end of the term, students would incorporate the various parts into a coherent final document, responding to previous criticism to improve upon their past work. Sequencing assignments in this way can make a larger assignment seem less daunting, help make the writing process more visible, and give the students a chance to acquire confidence as they learn from each other and the instructor—all of which reduces their tendency toward feeling overwhelmed and thus resorting to LLMs.

3. Limit LLMs' ability to help

LLMs generate text by anticipating what the next word in a sequence will be based on its scan of vast amounts of data. This method has several limitations. First, the most popular version of an LLM, ChatGPT, has been trained on data sets that have a cut-off date of 2021, which means that students cannot rely on it to provide information or written work that is concerned with more immediately contemporary events. Second, if the information is not on the Internet, then LLMs probably know nothing or very little about it. Assignment prompts can be devised with these weaknesses in mind.

Devise writing prompts that exceed LLMs' knowledge base

Given the fact that LLMs don't have access to online materials after 2021, instructors can consider asking students to write about events that have transpired after that date—or books/articles that have been published after that date. Also, instructors might ask students, in a given assignment, to be as granular or as highly specific as possible. Consider having students write about local events that may be less well known on the internet.

---

[2] https://thehill.com/policy/technology/3905672-chatgpt-sends-shockwaves-across-college-campuses/

[3] https://poorvucenter.yale.edu/AIguidance

Chances are that some or possibly the majority of your course materials are not yet available on the Internet. If that is the case, then assignments that ask students to grapple with these texts will be difficult for LLMs to complete. Ask students to compare and contrast passages or analytically relate the texts to the larger concerns of the course. You might also ask students to incorporate into their work themes or topics from class discussions.

You might also ask students, if appropriate, to write about their own experience or that of their personal or professional connections. You might ask them, for example, to analyze a past professional experience and demonstrate how it confirms or undermines certain themes or arguments from the course.

Require tasks that exceed LLMs' abilities

LLMs struggle to put forward original or analytical arguments. While LLMs are able to compare and contrast various sets of data or events or policies, this skill remains (at least for now) a relative weakness. Simple prompts (e.g., asking students to summarize information) are relatively easy for LLMs. The more in-depth the prompt—the more critical thinking that the prompt requires—the harder it is for LLMs to produce a quality answer.

Simple prompt: What are the causes of inflation?
Better prompt: Compare and contrast present-day US inflation with the 1970s.[4]

Besides challenging students to analyze complicated phenomena and put forward original arguments, you might also introduce some alternative means of assessments that require creativity and that cannot easily be accomplished by LLMs, e.g., podcasts, annotated maps, videos, diagrams, or storyboards.

Incorporate assessment tasks that require students to assimilate information, then process, analyze, and reformulate that information on the fly

Examples include:
- in-class debates between individuals or teams
- Q & A sessions between student presenters and the class
- role-playing exercises
- in-class writing assignments in which students are not allowed access to the Internet/LLMs.

Any of these assessments could be used as the starting point for students to then embark on their own longer writing assignment. If students have already gotten started in a meaningful way on an assignment by doing their own work, they may be more likely to finish the task without additional assistance from LLMs.[5]

---

[4] https://academictech.uchicago.edu/2023/01/23/combating-academic-dishonesty-part-6-chatgpt-ai-and-academic-integrity/

[5] https://criticalai.org/2023/01/17/critical-ai-adapting-college-writing-for-the-age-of-large-language-models-such-as-chatgpt-some-next-steps-for-educators/

Note that while these tasks could complement an assessment portfolio, there is still great value in asking students to complete longer writing assignments that require days or weeks to complete. Writing tasks prioritizing critical thinking remain a superior way for students to synthesize ideas, develop their own ideas, and remember these ideas for longer.[6]

<u>Introduce friction to signal how seriously you take academic integrity and to make cheating more difficult</u>

Ask students to compose their writing assignment on a Google doc. A Google doc preserves all the changes made to the document over time, so that, by clicking "version history," an instructor is able to glance quickly at previous versions of the document to see how the work has grown as students have increased their learning and improved their thinking. To access the version history, click on the clock icon in the upper right corner.

Or you might ask students to write a reflection explaining their research methods and how their understanding of the topic evolved over the course of their research. Where did they begin? What wrong turns did they take and how did they alter their course? Where/when/how did they achieve a breakthrough?

<u>Use LLMs to determine the quality of your prompt</u>

To determine whether students might be able to use LLMs effectively to complete a prompt, enter that prompt into one or more LLMs and evaluate the results. If the results are quite strong, consider revising the prompt.

---

[6] https://news.yale.edu/2023/03/03/chatgpt-campus-assessing-its-effects-college-writing-and-teaching