Generative Artificial Intelligence and Writing Assignments

This July, Harvard undergraduate Maya Bodnick published a piece in *The Chronicle of Higher Education* under the splashy headline "GPT-4 Can Already Pass Freshman Year at Harvard." In it, Bodnick chronicled how she had fed the essay prompts from a number of popular humanities and social science courses—including EC 10, GOV 1295, GENED 1033, and a First Year Seminar on Proust—into GPT-4 and submitted the A.I.-generated papers *verbatim* to their respective instructors, all of whom apparently found them indistinguishable from A- and B-level work produced by real Harvard undergraduates. Though her article was more stunt than study, arguably Bodnick succeeded in proving at least the most moderate form of her argument: i.e., that Harvard instructors teaching writing-based humanities and social science courses "need to adapt to their students' new reality—fast."

But how? It's far from clear what it means in practice to "adapt" to the rise of generative A.I., and much of the advice that's been made available—including Bodnick's (viz. that instructors move most, if not all, of their assessment into the classroom, swapping all of their paper assignments for modalities akin to blue book exams)—is unrealistic, pedagogically dubious, or both. In this document, we try to provide you with some guidance about how you might revise or rethink your writing assignments to account for the capacities and affordances of generative A.I., with the understanding that individual instructors are best positioned to think through the details of exactly how this guidance can be best applied to their own courses.

If you would like to talk through your plans or get feedback on a proposed assessment, please email bokcenter@fas.harvard.edu to book an appointment with a member of our senior staff.

1. Three High-Level Principles

In the next section of this document, below, we offer specific guidance and examples as to how you might adapt some of the most commonly-assigned genres of writing assignments to preserve the learning objectives that you likely had in mind when first developing them. First, however, we suggest three near-"universal" approaches that any instructor assigning writing may find useful in discouraging students from outsourcing their thinking and writing to A.I. in ways that contravene the purposes of your assignments.

1. **Talk directly and specifically with students about how your assignments are meant to work.** Our students are not, by and large, looking for opportunities to cheat or take shortcuts. The *vast* majority, in fact, are just as concerned to determine the ethical and responsible use of A.I. as are their instructors. The primary challenge posed by generative A.I. is not that, in making cheating easy it will, therefore, make it rampant, but rather that its utility will blur the lines for even our most scrupulous students between "seeking help" or "brainstorming ideas," on the one hand, and "soliciting an unacceptable degree of assistance," on the other. Talking to students about what you want them to learn, how completing your assignments on their own will help them practice and learn it, and whether or not employing various kinds of assistance from A.I. is consonant with those goals, will go a long way toward narrowing the problematic use cases you might be worried about.

2. **Disaggregate process from product, and render it visible.** Even prior to the advent of generative A.I., most of the writing that students did in our courses was, unfortunately, a black box—instructors

distributed prompts, and a few weeks later, students submitted final products, with little to no evidence of how they arrived at their conclusions. Now more than ever, we would encourage instructors to open up that black box by asking students to share early stages of their research and writing, in the form of preliminary assignments like project proposals, lists of analytical questions, annotated bibliographies, brief source analysis exercises, draft introductions, etc. Asking students to share their work in progress makes it considerably harder, not to mention less appealing, for students to outsource their thinking and writing to a large language model, as it would require them to forge, convincingly, not one but multiple phases of thinking and drafting.
    - This approach may come with a bonus: Whatever it costs in labor to review so many preliminary submissions—e.g. In the sense that you may not be able to assign as many final products—will probably be more than compensated by the improved coaching that students receive on their preliminary work.

3. **Create opportunities for students to reflect on/talk about their work.** So long as students imagine that they are submitting their final written work to a single reader (i.e. the instructor), *and* that said reader will never ask them to elaborate on, defend, or recapitulate their ideas in further conversation, leaning on generative A.I. might seem like a relatively safe (even victimless) indiscretion. If, however, students realize that they may have *many* readers—and, moreover, that those readers will ask them many questions about their writing (including how they arrived at their conclusions, how they would respond to various counterarguments, what they would have included had they had more space or time, …)—the value proposition of outsourcing all of those decisions to a large language model that won't be able to help them respond to their readers in the moment becomes much less appealing. Students in courses that employ peer feedback opportunities (like Expos, and many department tutorials) consistently report that reading and interrogating other students' writing is one of the most rewarding parts of a Harvard education. This seems like an ideal time to lean into that feedback, and ask students to think of the writing that they submit to their courses as the first, rather than the last, word on their thinking.

2. Specific Approaches Tailored to your Objectives

Jump to … Response paper | Single source analysis | Article summary | Annotated bibliography | Research paper | Final exam | Inhabit a voice | Creative writing

Response paper

Objective	Identify a key issue in a reading
How to incorporate A.I.	Use a large language model (LLM, like ChatGPT) to develop the ideas that emerge from the student's initial insight. In the sense that this kind of response paper is often used to gauge students' intuitions about a reading or measure comprehension, ChatGPT's role might be to put that intuition into dialogue before class discussion. For example, asking students to list the 3–5 main issues of a reading and rank them in importance before asking an AI to do the same could allow for both authentic "prediction" by the student along with a corroborating or competing interpretive response that could be discussed further in class. Students could perhaps ask the AI what kinds of evidence or what kinds of scholars typically address these issues and/or what some of the pressing debates around those issues are or have been.

Other ways to assess	**In-class, small group discussion.** Have students create the same ranked list of 3–5 issues and then work in small groups to compare their responses. This kind of small-group work requires facilitation by instructors to make sure students are on the right track (in the same way that facilitating input from an AI would be necessary).
Scaffolding required	Students need to understand how the "prediction effect" works, namely, that there's inherent value to guessing "the answer" to a problem on one's own before being given the answer or exploring it too much further through experimentation or research. This initial, flexible "guess" or hypothesis both keeps us invested in learning and makes us reflective about where we were at the start of that learning process. If students are motivated in this way, the role of AI or its absence is less material to the learning experience than the students' investment in what that experience is helping them learn.

Objective	Pose an analytical question
How to incorporate A.I.	Perhaps students could develop the question and then pose it to the AI assistant to see if it yields an interesting response.
Other ways to assess	**In-class, small group discussion.** Have students create a handful of analytical questions and then use a rubric to workshop them together in class. This kind of small-group work requires facilitation by instructors to make sure students are on the right track (in the same way that facilitating input from an AI would be necessary).
Scaffolding required	**Framing and models**. Students need to understand the role of asking questions—certain kinds of questions—in moving from readings and lectures to developing a thesis and working with evidence. This kind of framing, along with models of more/less productive analytical questions will help students here, whether or not an AI is part of the activity.

Objective	Prepare for discussion
How to incorporate A.I.	Students might ask an LLM to generate a handful of questions that might provoke insightful academic discussions of a given text, then practice responding to these as preparation for class.
Other ways to assess	**Class discussion.** The quality of the class discussion is the best evidence of the students' preparation, but if one intends to put more emphasis on the quality of classroom discussion this year, then transparency about evaluation criteria will be essential. **LEARN MORE**
Scaffolding required	If greater emphasis is put on the quality of the discussion itself (rather than on documents like response papers), giving the students both practice opportunities and feedback early and often will be essential.

Single source analysis*

Objective	Close read a text for language / formal aspects
How to incorporate A.I.	While the ultimate perspective, interpretation or argument should be the student's own, the student could use the LLM as an assistant in analyzing the text, potentially by getting the LLM to analyze short passages of the text, then working to critique and synthesize the ideas the A.I. offers with the student's own.
Other ways to assess	**Annotation.** Ask students to show you their close reading skills in action by having them annotate a passage. You can have them do this by hand, by distributing a printed passage with wide margins, or let students do it online (either in class or at home) by uploading your readings into a Harvard-supported online annotation tool like Perusall. LEARN MORE
Scaffolding required	Students would benefit from practicing close reading in class, with instructor feedback, and/or from watching how the instructor performs a close reading.

Objective	Close read an object for form / function / conservation
How to incorporate A.I.	In cases where students are developing the ability to describe objects with precision, they may practice doing this "for" an LLM in order to test their skills. Early in the term, they could describe the object but not name it, then see if LLM can guess the object.
Other ways to assess	**Show and tell.** An in-class (or video-recorded) "show-and-tell" project, in which students need to give the audience a "guided tour" of the object could be a way to ensure that they grasp the ways the ideas they've developed relate to various features of the object.
Scaffolding required	Students would benefit from an opportunity to perform an in-class object-analysis in oral or written form (and perhaps the requirement that they build on this in-person thinking as they develop their paper).

Objective	Make field / ethnographic observations
How to incorporate A.I.	After students collect their observations, you might ask them to input them into a LLM and prompt it to sift through fragmentary notes in search of patterns on which the student researcher may wish to reflect.
Other ways to assess	**Multimedia observations.** Ask students to make and compile their observations in forms that students cannot outsource to an LLM (e.g. annotated video observations, annotated images, photo essays).
Scaffolding required	Students would benefit from some low stakes "show and tell" assignments that ask them to capture and share raw audiovisual data in class in order to practice analyzing it.

Article summary

Objective	Parse another writer's logic / argument
How to incorporate A.I.	You might ask students to begin by producing their own summary, after which they would be tasked with prompting an LLM to do the same. Students could share both their own summary and the LLM's, and spend time in class comparing their results with the goal of developing intuitions not only about the particular article assigned, but also more generally about what makes a "summary" more or less insightful / useful.
Other ways to assess	**Concept mapping.** Students produce a visual diagram of the ways in which an author's premises, evidence, and arguments connect and flow. **LEARN MORE**
Scaffolding required	Students would benefit from some in-class time devoted to introducing concept mapping and its conventions.

Objective	Characterize another writer's use / interpretation of evidence
How to incorporate A.I.	Depending on the nature of the evidence cited in the article, students might be able to ask an LLM to generate feedback on how the author interpreted the evidence (e.g. "Does [the cited source] support the claim that [the article's author makes]?") or to find out whether other researchers have ever cited the same or similar evidence (and to what purpose).
Other ways to assess	**Scavenger hunt.** Rather than asking students to produce a prose summary, have them spend some time in the library drilling down on one (or a handful) of the author's footnotes. You could prepare a "scavenger hunt" worksheet that students complete before class, where they are asked to ferret out information about the material cited in the footnote—who wrote it? For what purpose? Does the cited passage in fact support the claim that the article's author is making about it? Etc.
Scaffolding required	Students would benefit from a session with a research librarian to orient them to library catalogues, databases, etc. that can help them complete the "scavenger hunt."

Annotated bibliography

Objective	Form a picture of the literature within a field
How to incorporate A.I.	In addition to getting help with the low-level formatting technicalities that AI tools have mastered, students may get feedback on the sampling of titles that they have put together (that they must of course evaluate critically).
Other ways to assess	**Source map.** Instead of a list of sources, students produce a diagram of the ways that various sources relate to each other and to the primary work the student is analyzing

	(some might function as comparisons, some as contestable claims, some as applicable theories, etc).
Scaffolding required	Students would benefit from framing that motivates (the otherwise potentially tedious seeming) annotated bibliographies, including the elements of summary and citation. Students will also likely need models for summary and guidance on how to locate and evaluate the kinds of sources that are best suited to the annotated bibliography at hand.

Objective	Assess the utility of various sources
How to incorporate A.I.	Practice by having the students critique an LLM's own attempts to create a bibliography of sources in response to an analytical question or topic.
Other ways to assess	**In-class lightning round.** In some cases you can achieve the same learning objective in a quicker, lower stakes way than through academic writing. For instance, students could simply rank five sources in terms of utility, then be called on in class to defend their rankings or grades.
Scaffolding required	Students would benefit from guidance on how scholars assess and articulate the quality/veracity of sources.

Research paper*

Objective	Define a research question
How to incorporate A.I.	Ask a LLM (much in the same vein as the "identify a key issue" activity above) to formulate a handful of questions based on a series of themes/key terms and a set of criteria for the assignment. Then, ask students to evaluate the strengths and/or feasibility of the responses.
Other ways to assess	**Question drafting peer workshop.** Have students identify the themes/key terms they're most interested in from the course and then use the criteria for the assignment to practice drafting research questions (perhaps by swapping questions with peers and offering each other feedback on those questions' alignment with the assignment prompt).
Scaffolding required	As with all assignments—but especially ones that include a synthesis of several smaller steps and a have longer-term focus on process—going over (and coming back to) the prompt will help clarify for students the role of any given step and how it relates to other steps. In the case of developing research ideas, this is likely to include helping students develop research questions (which is a specific instance of formulating analytical questions en route to a provisional thesis. In any writing context, including the use of A.I., it is important to make sure students know that a provisional thesis isn't a commitment: it's a heuristic for framing the early stages of writing and thinking about structure, and/but it's likely to change as the process of

	inquiry unfolds and evidence + counterevidence get worked through. That is: throughout the writing process, one's thesis and one's evidence/analysis mutually shape each other.

Objective	Identify relevant research material
How to incorporate A.I.	You might ask students to use an LLM to help them generate lists of sources, and take note of where it succeeds and fails in order to see whether this yields any insights into how well different subfields are represented in the datasets on which the LLM was trained. You might also encourage students to use an LLM to generate lists of most frequently cited materials in a subset of literature within a field.
Other ways to assess	You might try a variation on the **scavenger hunt** described above, asking students to use HOLLIS and other library databases (i.e. rather than A.I.) to assemble a collective course bibliography of sources during a class field trip to the library.
Scaffolding required	Students would benefit from a session with a research librarian to orient them to library catalogues, databases, etc. that can help them complete the "scavenger hunt."

Objective	Enter an existing conversation in a field
How to incorporate A.I.	After a student has identified the main "schools," perspectives, or approaches that characterize their field of research, have them ask an LLM to roleplay one (or more) of them in a conversation. You might ask the student to submit a screen recording of their conversation.
Other ways to assess	**Role play.** Students could be assigned to embody the various schools or approaches they have identified in the field, and to perform a live discussion or debate in class with instructor or peer feedback. **Card game.** Building on the "dinner party" framework of putting scholars in conversation with each other, turn it into a card game: have each writer / reading / perspective discussed in class on a separate index card. Divide class into small groups and distribute cards (randomly or have them pick their favorites). The first player plays a card to start a conversation; the next student plays a card adding a new perspective or argument to the conversation, using their card/idea to build on the conversation that the previous card/idea started. This can be expanded by incorporating students' own ideas or arguments (in a "they say / I say" framework, as mentioned below).
Scaffolding required	Students would benefit from instruction in the key rhetorical moves of academic argument, as in the book *They Say, I Say*, as well as some practice in public speaking (which can be arranged through the Bok Center).

Final exam*

Objective	Synthesize course readings
How to incorporate A.I.	While it's unlikely that you would invite students to use A.I. during a final exam, they might be encouraged to use it to generate practice essay or short answer prompts that you could practice in a review session.
Other ways to assess	N/A?
Scaffolding required	N/A?

Objective	Apply course concepts to new problems / scenarios
How to incorporate A.I.	See above.
Other ways to assess	N/A?
Scaffolding required	N/A?

Inhabit a voice

Objective	Identify and be able to reproduce key stylistic patterns in a primary source or key conceptual features of a thinker's perspective
How to incorporate A.I.	After drafting their own piece of writing (e.g. a letter from one historical figure to another), students could upload it to a LLM and ask the LLM to respond (e.g. as the imaginary letter's recipient), to which the student might then draft a second response.
Other ways to assess	**Briefing book.** Students may be reluctant to attempt to inhabit the voice of a person whose identity they do not share; in these circumstances, outsourcing the writing to a LLM may be particularly tempting. Instead of asking students to inhabit a different voice, you might ask them to compile a briefing book of the background research that would go into *preparing* to inhabit that other voice.
Scaffolding required	Students would benefit from sessions with librarians and curators, as relevant, as well as opportunities to read as many ego documents as possible produced by the kind of person they are meant to channel.

Objective	Cultivate empathy for a figure / perspective distant in time, space, or experience
How to	See above.

incorporate A.I.	
Other ways to assess	See above.
Scaffolding required	See above.

Creative writing

Objective	Produce a piece of original fiction (or creative nonfiction)
How to incorporate A.I.	Students could ask an LLM to produce "background" materials that flesh out the life and experiences of their fictional characters (e.g. a diary kept by someone matching the description of their short story's protagonist). They could then read these for inspiration when drafting their story.
Other ways to assess	**N/A.** LLMs remain largely ineffective at producing narrative fiction, and the workshop format of most creative writing courses would make it difficult/infeasible to employ A.I. illicitly.
Scaffolding required	N/A.

3. A Final Note About Generative A.I. and Equity

Just as we used to assume (incorrectly) that our students were all "digital natives," and would all, uniformly, understand how to do things online better than we would, we might be tempted to assume that all of our students will know more about how to access and use generative A.I. than we could ever teach them. That may be true of some of our students—but it's likely not true of all of them, meaning that encouraging them to use A.I. as part of your graded assignments may introduce unforeseen inequities into your assessment scheme. Insofar as you want to allow, or even encourage, your students to make use of A.I. to enhance their written work, you'll want to make sure that they all have equal access to the most useful platform(s), and a fair chance to develop the requisite amount of proficiency in what is now called "prompt engineering" (i.e. the ability to compose queries that will produce the richest/highest quality output from an A.I. tool).

This may mean holding a course "hackathon" where students all receive some instruction in effective use of A.I. and opportunities to practice with more expert users on hand; it could also mean encouraging students to go through a free online "prompting course" like this one. At the very least it means including explicit instruction about how you want students to use generative A.I. in the high level principles articulated at the start of this document.

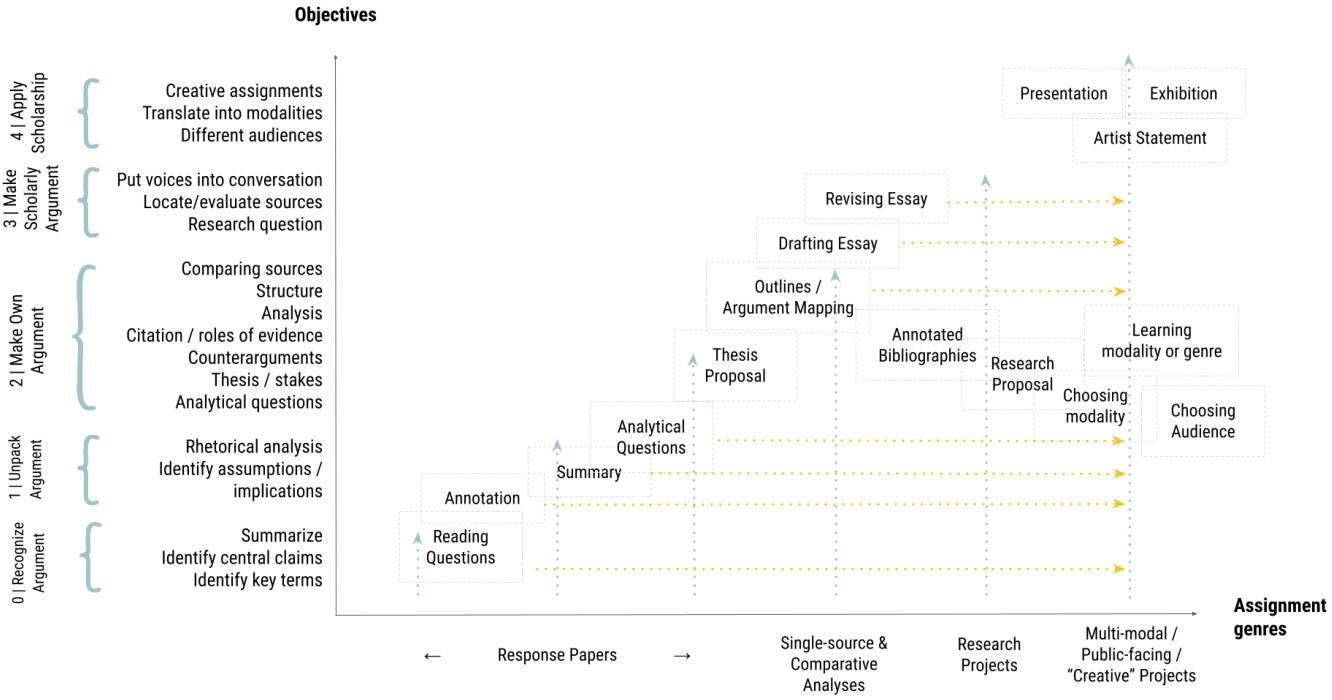

The figure above presents a typology of common writing assignment sequences. On the far left (in vertical text), the labels for the **y-axis** indicate a progression of general learning objectives, moving from "show you understand an author's argument" to the more complex operations of "formulating your own argument" and putting that argument into dialogue with other scholars or different audiences. Closer to the y-axis itself are the groupings of specific learning objectives that tend to belong to each stage of this progression.

The **x-axis** breaks down the most common genres of writing assignments, and on the plane of the graph itself are common kinds of assignments, along with blue and yellow arrows. The blue arrows suggest the most complex objectives for which a kind of assignment is likely to serve as evidence, while the yellow arrows indicate the ongoing role of less complex assignments as scaffolding for increasingly complex ones.

The figure is not exhaustive, and different instructors or disciplines will have different intuitions about the exact order of some elements of the graph. That being said, **the figure's argument** is relatively simple: assignments should be sequenced and scaffolded in alignment with their objectives, and learning experiences tend to be more positive when the sequencing and scaffolding is both intentional and transparent.

Considerations for the Use of Generative AI

Two Possible Roles

Two possible roles for AI in these kinds of assignments are:

- **Role 1.** The AI “does” part(s) of an assignment for a student, and the student then evaluates the AI’s work and gives it feedback (before making use of the output)
- **Role 2.** The student does part(s) of an assignment, and the AI then evaluates the student’s work and gives it feedback (before making use of the output). In this case, the AI arguably also “does” part(s) of the assignment for the instructor.