
Voice, Authorship, and Generative AI

Open Access Teaching Case Developed for the Tech for Humanity Pathways Minor

Funded by the Andrew Mellon Foundation

Developed by Sarah Cozort

Background

A recent issue of *The Chronicle of Higher Education* announced that the study of literature in the academy is no longer in crisis but in collapse, an assertion supported by the closure of a number of literature programs across the US, as well as a shift in hiring practices that reflects a preference for specialists in rhetoric and technical communications over literature and creative writing. A decline in the perceived value of writing (and therefore writers and writing instructors) is frequently attributed to the rise of generative AI. Not only are some students less inclined to develop their own skills as writers when programs are available that can “write for them,” many writing professionals are rightfully concerned about issues of copyright infringement and their own potential obsolescence. Perhaps a greater contributing factor to this sense of crisis than generative AI, however, is the fact that the value and limitations of these technologies are not widely or fully understood in the English department and by professional writers.

Here, I suggest one way for writing practitioners to get a firmer grip of generative AI’s capacities and limitations is to revisit questions about the necessary elements of authorial voice, first posed by a handful of pioneering rhetoric and composition scholars in the 1960s and ‘70s. Analyzing the process through which AI generates material “in the voice of” an author will allow us to approach some of these questions about voice anew and potentially empower writing practitioners within the academy to discuss the relationship between generative AI and the production of writing, i.e., authorship, in a way that is factually grounded and compelling to both students and administrators.

A recent study on authorial voice and AI begins by suggesting that, “Recent developments in AI content generation software have rendered human-authored texts indiscernible from AI-generated ones” (Nañola et al., 2025). While this is frequently the case in both advanced scholarly writing and student writing, the claim brings up questions about the features of authorial voice that distinguish AI-generated writing from human writing. Keeping in mind the fact that undergraduate composition courses are intended to prepare students for writing in the workplace, as well as potentially writing in an advanced research capacity, insights gathered about both authorial voice in advanced research and in undergraduate writing are equally pertinent to discussion of undergraduate writing. Looking more deeply at the claim that AI-generated texts are indiscernible from human generated texts, Nañola et al. (2025) found several pronounced distinctions between the two, especially when comparing undergraduate student writing to AI-generated texts. These distinguishing features are as follows:

1. First-person (“I” voice) is not the consistent dominant textual identity for AI-generated texts.
2. AI-generated text is closer to an expert’s writing whereas student-written work is closer to a novice’s.
3. The writing style of AI-generated texts leans toward predictability.

The following case study explores the meaning of these three distinctions, as well as their indications for how writing instructors are to help students think about the relationship between AI and their own writing.

Nañola et al. (2025) begin by recognizing the ongoing relationship between technological development and the history of writing, pointing specifically to the rapid innovations in composition and communications technology that are characteristic of the 20th and 21st century. The recognition of this history serves as a reminder that, while technological developments such as the typewriter, mimeograph machine, word processor, and early predictive text models were initially met with skepticism, these technologies have proven to shape writing (and communication more widely speaking) in unexpected and productive ways. It is with this cautious optimism that Nañola et al. undertake their comparison of the authorial voice in AI-generated texts and undergraduate student writing, recognizing that “AI integration has opened avenues for increased variations in language learning opportunities through assistive features such as translation, revisions, paraphrasing, textual explanations, artificial

conversations, and even content generation.” Despite these affordances, the authors do recognize that considerable evidence suggests that concerns over the use of generative AI by students to commit plagiarism and engage in other forms of academic dishonesty are well-founded. Further, because AI models are trained on a large sample of pre-existing texts, older strategies for the detection of academic dishonesty are not as effective when detecting whether a piece of writing turned in by a student has been largely (or wholly) generated via AI. Thus, it is suggested that one alternative strategy for detection of this form of academic dishonesty is through a closer look at the distinguishing features of voice.

I-voice is not the consistent dominant textual identity for AI-generated texts

Many students today have worked with an instructor or writing tutor who has discouraged the use of first person singular or plural in academic writing, despite the fact that in recent years the American Psychological Association (APA), the Modern Language Association (MLA), and Institute of Electrical and Electronics Engineers (IEEE), have all taken the position that using the first person singular or plural (“I voice” or “we voice”) is acceptable (and even desirable) in some cases. This stylistic shift in academic writing is tied to the recognition that the long-held standard of the use of passive voice in scientific writing suggests an overly confident claim to “objective” truth, which is additionally dismissive of worldviews that value pluralism and collectivity. A classic example of the difference between passive and active voice in scientific writing is the following:

Passive Voice: “The experiment was conducted.”

Active Voice: “The researchers conducted the experiment.”

We can see in this case how use of the passive voice removes the researchers from the scenario in which the experiment took place. This suggests that, regardless of who these researchers were or their relationship to the subject of study, they have conducted the experiment as any other researcher would have conducted it. We know, however, that many factors particular to the individual researcher or group of researchers can influence how they conduct research. These factors could include a researcher’s unique training, personal ethics, or cultural background, among others. We can imagine, for instance, that a researcher who is testing the medicinal properties of a specific herb based on its use within an indigenous community would be sensitive to how that herb is treated within that community. In such a case, it may be appropriate for said researcher to acknowledge their own relationship to the

community whose use of the herb prompted the study, as well as their personal reflection on these cultural practices shaped their understanding of respectful and ethical methods of inquiry.

AI-generated text is closer to an expert's writing whereas student-written work is closer to a novice's

In recent years, reflections on the appropriate use of I-voice in academic writing have gained traction in the teaching of undergraduate writing. This has had some influence on the increased use of first-person “I” and plural “we” in undergraduate writing (and more recently in the writing of professional researchers). Students tend to use first person “I” in their college writing courses unless instructed otherwise, and because of these shifts in attitude toward the use of I-voice in academic writing they are discouraged less often from it. Despite this shift, generative AI models are trained on corpuses of scientific texts that largely favor the use of passive voice. Thus one distinguishing feature between AI-generated texts and texts generated by students is a pronounced difference in first-person “I” and plural “we.” This observation is not only handy as a method for the detection of academic dishonesty in student writing, explanation of this distinction in the teaching of academic writing may serve as a deterrent for the overreliance on generative AI on the part of students. Further, this distinction highlights the fact that generative AI is unable to discern when it is appropriate to use first-person “I” for ethical or culturally sensitive reasons, whereas a human being with a sense of the cultural context in which an experiment is being conducted and a text is being composed is more likely to use I-voice when appropriate.

The writing style of AI-generated texts leans toward predictability

In their comparison of 12 AI generated texts to 12 student-written texts, Nañola et al. (2025) coded for specific differences in language use. This means that they looked for the use of *kinds* of language, specifically boosters and hedges, among others. They then documented each occurrence of these kinds of language use in the two corpuses before reflecting on patterns of use. A “booster” is a word which intensifies a claim. For instance, when we add a booster to the sentence, “the rain in Spain falls on the plain,” it becomes “the rain in Spain falls *mainly* on the plain.” “Mainly” here intensifies the claim that rain falls on the Spanish plain in the sense that it conveys a matter of degree. This is referred to as an adverbial booster because it modifies the verb “falling,” i.e., “mainly” tells us that the rain is not only falling, but it is falling on the plain more than it does on other places. Adverbial boosters, or boosters of degree, were found to occur more in the AI generated samples than in the student samples, the latter of which tended

to use boosters such as “every,” “most,” “just,” and “should” (19). These boosters observed frequently in student writing are called “determiners,” meaning that they suggest totalizing/binary distinctions, i.e., “the rain falls here but not there,” rather than “the rain mostly falls here but also some falls elsewhere.” A similar pattern emerged in the comparison between the use of hedges in AI-and student-generated texts, suggesting that undergraduate students tend not to have acquired the genre features characteristic of advanced scholarly writing to the degree found in the repository on which the language learning model generates texts. In addition, despite the fact that professional scholars often use more complex sentence structures than most undergraduate students, the study found that the language learning model tends to homogenize sentence structure based on the samples in the repository, where as, surprisingly, there tends to be more variation in the structure of student writing at the sentence level than in the AI-generated texts.

Case Study

Dr. Philippe Migers*¹ is currently working on lesson plans for an upcoming semester of English 102, a sophomore-level research and composition course. Due to the preponderance of students who use AI to assist in the composition of their assignments, the department has advised instructors to create a lesson plan considering the use of AI in research and composition. Although Dr. Migers teaches undergraduate students, she is at a prestigious R1 university, known for producing a high number of undergraduate students who go into careers as researchers in STEM. There are also many opportunities for upperclassmen to engage in important research alongside graduate students and faculty. Thus, while Dr. Migers understands that her course will serve as an introduction to topics in research practice and design suitable to students who will be going directly into the workforce after graduation, she also understands that there may be a number of folks in her course that she will be preparing for life as a professional researcher and scholarly writer. This makes designing her unit on generative AI feel especially high stakes. Having recently read a study on the differences between authorial voice in AI-generated texts and texts generated by students (Nañola et al., 2025), Dr. Migers decides to build a lesson plan around some of the observations of the study. Specifically she wants to draw attention to the following three observations about key differences in texts generated by AI and those written by students: 1) use of first person “I”; 2) the fact that AI generated texts trained on repositories of academic writing tend to mimic professional scholarly

¹ The scenario and rhetorician described in this case are fictional and meant only to facilitate reflection and conversation.

writing rather than “novice” or student writing; and, 3) the fact that AI generated texts tend to “lean toward predictability.”

Processing Questions

- 1) With the knowledge you are privy to from the case study here, imagine that you find yourself in the position of Dr. Miger. Create a lesson plan that draws attention to the primary differences between AI-generated writing and student writing. A lesson plan should include ten or so minutes of “lecture” material, where the problem and the key findings of the study are laid out. This should be followed up with some kind of activity that allows students to find evidence for – or “make meaning” out of – these key points. For instance, you might give students a page from an AI-generated “scholarly” research report and a page from an undergraduate research report and ask for them to read through each text and highlight lines/passages where they observe evidence for the three key differences highlighted in the study. You might also ask them to highlight evidence that refutes these key findings. This is just one example of a way that some of these key differences might be demonstrated to and tested by students. Can you think of another activity that might help students explore these differences?
- 2) How might you use the three findings highlighted by the study on which Dr. Miger’s is basing her lesson plan to demonstrate or emphasize lessons that are already built into the English 101 curriculum? You might want to think back to your own experience of Freshman Writing. One lesson that is often part of Freshman Writing is about varying one’s sentences in length and structure in order to hold the interest of the audience. One strategy could be to ask students to look at their own writing, sentence by sentence, to examine how varied their own sentences are. Then, perhaps after some revision, students might be asked to compare the sentences in a passage from their own writing to a passage generated by AI. Some questions for analysis: Do you notice the homogeneity observed by researchers in AI-generated material in the passage you have examined? Are there notable differences between your own writing and the AI-generated passage? Do these differences support or diverge from the findings of the researchers?

Thematic Questions

Ethics

1. One motive for studying the differences between AI-generated and student texts is to potentially dissuade the unethical use of AI by students. First, given what you know

about your own use of AI-generated material and also how your peers use it, define ethical use of AI when conducting scholarly research and composing scholarly reports of said research.

2. Then, after you have a good working definition of the ethical use of generative AI in the academic context, reflect on how you think the ethical use of AI by students might best be addressed in the classroom. Is there an approach to the topic that you feel would be particularly persuasive to students? Do you think that students who are preparing for careers as researchers should receive different or more in-depth training on the ethical use of AI in a research context? If so, what would general training look like and how would specialized training for those headed into research-oriented careers look different?

References

- Nañola, Elisha L., RG L. Arroyo, Nicole Jazz T. Hermosura, Mark Ragil, Janen Nicole U. Sabanal, Henelsie B. Mendoza. (2025). Recognizing the artificial: A comparative voice analysis of AI-Generated and L2 undergraduate student-authored academic essays. *Systems* (130). <https://doi.org/10.1016/j.system.2025.103611>
- Watson, S., Brezovec, E. & Romic, J. The role of generative AI in academic and scientific authorship: an autopoietic perspective. *AI & Soc.* <https://doi.org/10.1007/s00146-024-02174-w>
- Yasuda, Sachiko. (2022). Natural scientists' perceptions of authorial voice in scientific writing: The influence of disciplinary expertise on revoicing processes. *English for Specific Purposes*. 67. 31-45. 10.1016/j.esp.2022.03.001.