

Meaningful Writing in the Age of Generative Artificial Intelligence

Kristi Girdharry and Davit Khachatryan

Babson College

DOI: [10.37514/DBH-J.2023.11.1.04](https://doi.org/10.37514/DBH-J.2023.11.1.04)

Introduction

Toward the end of January 2023, the co-authors of this article found themselves at an event promoting faculty research at their school. Davit Khachatryan, an applied statistics professor, was presenting on Playmeans—an app he had created for audiovisual analysis of musical data¹—while Kristi Girdharry, an English professor, was presenting on a preliminary research study about student writers. On the surface, these two presentations, and presenters, were quite far apart in terms of disciplinary knowledge and approaches to their topics; however, there was mutual interest in the ways students and instructors can use various technologies in order to engage in meaningful learning experiences.

In the background of the event that day, and on the minds of many both inside and outside the teaching profession, were conversations about how artificial intelligence (AI) would impact student work. Less than two months earlier, Stephen Marche (2022), writing for *The Atlantic*, had declared that “The College Essay Is Dead” after ChatGPT was released to the public at the end of November 2022; at the beginning of January 2023, Princeton University senior Edward Tian released GPTzero in an effort to detect the use of AI in the composing of an essay (or any text).²

With the rise in AI writing technologies, there have been many (fair) critiques and concerns that higher education has crossed a threshold with no return. What will AI mean for writing, critical thinking, and learning overall? How can we, as instructors, ensure that students learn to write meaningfully and develop their own voice as a result? We view meaningful writing as a means to cultivate critical thinking, as the latter often evolves through the writing process. We believe meaningful writing to be an important framework for designing assignments across the curriculum. And today, with generative AI models outputting human-like text, this framework becomes more important than ever before. How can we develop meaningful writing assignments that can motivate students to think critically? And how might AI play a role in supporting, not hindering, this kind of writing and thinking?

In this report from the field, we first offer a description of what we mean by “meaningful” writing as drawn from the work of Michele Eodice, Anne Ellen Geller, and Neal Lerner (2016, 2019). Next, we offer an accessible explanation of how generative AI technologies, namely Generative Pretrained Transformers (GPTs) work, with an eye towards educating as well as inspiring teachers across the curriculum to develop lessons for their students. Finally, we showcase an example of a meaningful writing assignment that used AI to foster critical thinking and offer a list of similar opportunities that instructors can revise for various contexts and courses across the disciplines.

Meaningful Writing, Meaningful Learning

In their book *The Meaningful Writing Project: Learning, Teaching, and Writing in Higher Education*, Eodice et al. (2016) explained what they found after asking over 700 university seniors from three different institutions to describe the most meaningful writing projects they completed as undergraduates. The answers were interesting in that there were no clear linkages across particular courses, instructors, or assignments but rather themes suggesting that meaningful writing occurs when (a) students have agency while writing—that is, when students have the opportunity to make their own choices when developing writing projects; (b) writers feel engaged with other people and opportunities—that is, when writing involves discussions with professors and peers and/or personal thought towards goals beyond the classroom (e.g., presentations, publications, applications) that might aid in career or other post-graduation plans; and (c) learning for transfer—that is, when students actively draw from prior knowledge in order to think in new ways or when they can see how their work may be applied in future contexts. Although the authors did not explicitly link meaningfulness with critical thinking, the connections are not hard to see: having agency, collaborating with others for various contexts, and actively connecting new information to prior knowledge through writing are deeply connected to critical thought via the opportunity to explore, connect, revise, and solidify thinking via the written word.

Eodice et al. (2019) later revisited one aspect they reported on from the Meaningful Writing Project—personal connection—which was prevalent in more than one-third of all of the surveys. To be clear, the authors did not equate “personal connection” with only “personal writing” but instead found the concept of “personal” in academic settings to be multifaceted, such as through exploring different viewpoints, choosing important topics, and developing new perspectives (pp. 327–328). The authors went on to explain that “students were telling us that the personal connection resulted in a meaningful writing experience because of opportunities for them to grow, develop, or imagine future selves; to take on identities as writers or authors; to have a venue for self-expression; or to tap into previous experiences” (p. 329). Other aspects of personal connection had to do with social contexts—connecting writing to friends, families, and communities—and individual contexts—using writing to explore personal interests and passions (pp. 331–332). Whether working within social or individual contexts, personal connection often seemed linked to student choice (p. 333). Meaningful writing assignments prompted growth in students as writers and critical thinkers.

Given that GPT and other large language models (LLMs) in general are inherently depersonalized entities—that is, they currently can’t mimic the kinds of learning experience students get from classroom interactions, such as personal reflections, struggles, and breakthroughs—using them as part of the writing process may seem antithetical to the findings of Eodice et al. (2016, 2019). However, we believe this to be a limited view of those technologies, and we contend that they can support meaningful writing experiences and critical thinking.

To develop an intuitive understanding of how LLMs, particularly the GPT technology, work, it is important to see the mechanisms involved. Thus, before we address how we might embrace LLMs as part of meaningful writing assignments, we offer a brief rundown of the internal workings of GPT.

A Gentle Exposition on GPT (and Its Variations)

Imagine you are provided a blurb and asked to produce the next word. Let's go, "Once upon a . . ." It stands to conjecture that you are most likely to pick the word "time." There is some chance you'll pick the word "midnight," but some words—"cucumber," "misbehaving," "Picasso"—have close to zero chance of being produced, unless one is experimenting wildly with language. Now, as with many great exercises, comes the self-reflection part: try asking yourself for the reasons behind picking or not picking a word. Having seen or heard the word in the context of the provided prompt is going to determine the likelihood of you selecting that word. Not being familiar with Edgar Allan Poe's "The Raven," for example, is going to diminish the chances of you producing the word "midnight."

As elementary and commonsense as all that may sound, on a high level that is what goes on inside many of today's leading models used for language-related tasks such as text generation, summarization, translation, and question-answering. The GPT gains its proficiency by being given access to enormous repositories of text found on the Internet, in human-written books, and through other written sources (Wolfram, 2023). As a result of being fed wide-ranging and diverse writing—think of a large chunk of Toni Morrison's literature stacked with DIY articles on changing the oil in a motorcycle (and everything in between)—the LLM gets trained on how words relate to each other in various contexts.

This intense "training" results in those models' amazing generative ability that we see today. But just as the illusion of motion in film is created by sequentially running 24 frames per second, the illusion of seemingly coherent, seemingly human-like text is produced one word at a time. Just think of how you would end up with a full, made-up story if you continued the above exercise by producing one word a time. In a nutshell, that's how generative language models such as GPT and its variations (GPT-2, GPT-3, GPT-3.5, ChatGPT, GPT-4, etc.) produce text: word by word (or, strictly speaking, token by token, where "token" is a shorter version of a word, but we'll keep referring to words for the sake of intuitiveness). As sophisticated and intelligent as these models might seem, at the core their primary task is to produce one word—the next word. But the devil is in the details, as the saying goes.

A slightly more substantive understanding of how GPT and its variations do what they do involves knowing that each word in a text is "seen" by these models as long lists of numbers. These numerical representations—referred to as "embeddings"³—undergo a process of refinement until there is an eventual embedding for each word in the given prompt (Vaswani et al., 2017). Further, using these ultimate embeddings to produce the next word (e.g., what could come after the prompt "once upon a"), a list of probabilities gets generated for potential words that could continue the prompt in question. The word that gets generated (e.g., "time") is based on probability. It is thought that instructing the model to always pick the word with the highest probability results in text that is "flat" or in some ways "boilerplate" as it vaguely corresponds to the most expected word at the expense of spontaneity (Tingiris & Kinsella, 2017). However, these models can also be instructed to produce a word at random by using a parameter referred to as "temperature" (Tingiris & Kinsella, 2017; Tunstall et al., 2022), where, depending on the value of "temperature," one ends up with text that is deterministic ("trite") or diverse ("unexpected"). This nuance is unfortunately not made transparent to the average user of GPT technology, who thus runs the risk of producing text that can lack spontaneity.

As impressive as these models are, there are several important caveats that a user should be made aware of:

1. Although these models are trained on massive repositories of information, what gets generated is not necessarily sensible or factually accurate (Wolfram, 2023). Put simply, what the models output is what their engineered brains “think” may represent a continuation of the prompted text. The text that gets outputted as a result is often surprisingly human-like, but human-like does not necessarily imply that text is useful, factual, deep, or creative.
2. When triggered with prompts that do not resemble the data “seen” during training, these models can result in responses that are “hallucinatory” and don’t make much sense (Wolfram, 2023). While plugins connecting GPT to third-party sources are in development, and some variations, such as Bing AI (which runs on GPT technologies) and GPT-4, have “access” to the Internet, the produced output should still be taken with a grain of salt in terms of its usefulness.
3. The way these models are designed is such that the generated text does not undergo a process of quality control (Wolfram, 2023). Put differently, and in contrast with most humans, there is nothing resembling self-reflection or a “think-before-you-speak” moment before outputting text. Academically speaking, when generating text, these models seem to take a daring shortcut through Bloom’s taxonomy (Anderson & Krathwohl, 2001; Bloom et al., 1956), “creating” text without mastering evaluation. The ultimate merit of the generated text and the extent to which it needs to be reworked for usefulness should be left strictly to the evaluation of the user.
4. Last but not least, in spite of rigorous mathematical underpinnings, many of the engineering considerations that go into building these LLMs are more of an art than a “hard” science. Numerous design choices are based not on any particular theory (which doesn’t yet exist anyway) but instead on what has been shown to work well in practice (Chollet & Allaire, 2018; Wolfram, 2023). What gets generated as the next word following a prompt is a result of many clever engineering hacks, trial and error, and, in general, what resembles statistical craftsmanship with a touch of alchemy. Examples include but are not limited to selection of the temperature to yield allegedly the “best” results for text generation (is it “best” for coherence, profit, factuality, etc.); length of the embedding vectors used to convert each word into a vector of numbers; the method according to which each word is converted to the embedding vector; the way the position of a word is introduced to the model by adding embeddings of words with the embeddings of their positions in the text; etc. (Wolfram, 2023). There is no theory claiming that these and many other tricks are meaningful (i.e., useful) choices or why a model (e.g., a GPT version) built under such considerations should work. Likewise, there is no theory claiming that the workings of these models represent the way humans produce text.

In spite of the apparent risks, how can students and teachers effectively utilize the text generation capabilities of generative AI? How might generative AI help students find more

meaning as part of the writing process? We now present an extended example of a successfully implemented assignment that uses GPT and helps to address some of what Eodice et al. (2016; 2019) found to contribute to meaningful writing. We then offer multiple suggestions for ways that instructors might consider working with LLMs in their own courses across the disciplines.

Practical Applications of Teaching with GPT

Sample Assignment Using ChatGPT

In a Spring 2023 senior seminar titled “The Rhetoric of Social Media,” students were asked to contend with various aspects of social media through a humanities lens. In one assignment, they were asked to think about how social media texts create and promote goals and values through language and other forms of communication. Each student chose a musical artist and performed a close reading of the artist’s music (e.g., a selection of song lyrics, music video, or performance) as well as their social media presence (e.g., Instagram, Twitter, TikTok, etc.) in order to account for any harmony or dissonance between the art and the artist. For example, Taylor Swift is known for writing many intimate songs about her relationships, which a student might closely read in an example of her lyrics, but they may or may not find something similar when examining her social media platforms.

On somewhat of a whim, given all of the conversations around LLMs, and to complement the close reading students were doing, distant reading using ChatGPT⁴ was added to the project. Students prompted the tool with questions, assessing the helpfulness of the technology, and evaluating the answers. More specifically, students started with an artist of whom they had prior knowledge (almost all students chose artists they extensively listened to and followed on social media), and based on this prior knowledge, they asked ChatGPT about their artist. Doing this involved a bit of experimentation as they started to learn, first-hand, about the slight edits and variations needed in order to get the model to respond in a sensible way. Students collected their prompts and the responses provided by the LLM through screenshots in order to assess the model’s output in terms of clarity and veracity.

During informal discussions with their professor and formal peer review sessions, students were eager to talk about their experiences, which ranged from excitement over the LLM’s efficiency and accuracy to laughter over its incoherence. The most interesting part of the project for the class overall came from requiring students to fact-check the responses outputted by ChatGPT. For a reminder of context here, students were experimenting with this technology in February 2023—in these earlier stages of ChatGPT’s public version, there happened to be a higher likelihood of inaccuracy, particularly for lesser-known artists, which we wouldn’t expect as often today. Students generally found the kinds of information on Wikipedia (a fact-checking haven, it turns out) to be mostly correct. However, when they experimented with questions such as whether a particular artist was “great,” they were met with dead ends. In this case, they had to think a little more about genre, audience, and purpose to come up with how they might get answers that dealt with the concept of “greatness” (such as asking how critics received particular artists and albums at certain times).

While the very act of fact-checking prompted some students to refine questions and edit their findings accordingly, what became more interesting was how the process of using

ChatGPT motivated students' critical thinking in ways they didn't even recognize in the moment. For example, one artist that came up was Burna Boy—the Nigerian singer, songwriter, and producer—who was chosen by one of the students. While discussing ChatGPT's inability to offer an accurate, in-depth analysis of a particular political song that, according to the student, necessitated a nuanced human understanding, the student commented on how the tool could "only" assist with certain information, such as translating Nigerian Pidgin and slang and pointing to online forums where Nigerian fans offered their interpretations of the political messages of the song.

Given that the student wasn't Nigerian and didn't have access to the language or culture beyond her connection with Burna Boy's music, ChatGPT actually opened up more areas for investigation, despite the student's mistrust of LLMs. It should be noted that the student didn't naturally see this affordance on her own. Indeed, like many of her peers, rather than reflecting on her thinking processes during this project, she focused on the product and how the generated text couldn't match the dynamic tone she wanted. It took the instructor stepping in as a sort of a co-learner to point out how the technology facilitated something new.

Reflections and Customizable Use Cases

While instructors themselves cannot deem an assignment "meaningful," students reflected in ways that fall in line with what Eodice et al. (2016; 2019) reported in their work:

- First, students were given *agency* in that they were allowed to choose their own topic (i.e., musical artist). The vast majority of students selected an artist they were passionate about, and this freedom cultivated a sense of inclusiveness and engagement.
- In turn, this agency to work on something one was passionate about was an act of *self-expression* and connected to students' *senses of identity*. After all, one can argue that music is one of the mediums students (and people in general) use to connect with others frequently, so it was not at all unexpected that students found this engaging.
- Third, the use of the LLM gave students an opportunity to *engage with others*: their peers and the professor (and a whole forum of online Burna Boy fans in one case). The lively discussions and conversations around the output of the technology cultivated a sense of engagement and provided students time to think more critically about their assignments.
- Lastly, as a result of this critical examination of output, independent research, and discussions with peers, students would often come up with *new ways of thinking* to replace or supplement what the LLM had to offer. Importantly, not only was the mere content of the text edited by students but also often the tone. This editing became especially apparent when working with the "flat" responses from ChatGPT that could not nearly produce the personal tone used when one reflects on something that one is passionate about and which, again, connects deeply to one's sense of identity. After all, no matter how massive the amount of information GPT "saw" as part of its training, it surely lacks the sentiment and passion regarding what matters to one.

The use of GPT was not central to this project—it was merely one piece of a larger project—yet it had a lasting impact on students’ thinking for at least the remainder of the semester. Moreover, on the instructor end, there was even more that could have been done or emphasized with this one assignment that would complement a course focused on rhetoric, (mis/dis)information, and social media (and could be revised for other course foci), such as by asking students to

- assess bias or potential one-sidedness of GPT answers and think about how such issues may impact GPT users more broadly;
- reflect more specifically on other ethical considerations for using GPT models or on how they support/challenge learning;
- spend more time thinking about the affordances and limitations of GPT in their future personal and professional settings, including thinking about context-specific risks and rewards now that they’ve experienced using a model critically (e.g., if all customer service is performed by GPT); and
- contemplate what it felt like “chatting” with GPT versus working with their peers and professor.

As instructors, we are trying to understand the limitations and affordances of these technologies ourselves, and this teaching experiment offers an opportunity to further consider how we might link meaningful writing opportunities with critical thinking and AI in future assignments. What are other ways we might consider bringing in LLMs to complement the kinds of learning students are doing in our courses and assist students in moving from what ChatGPT provided as a raw starting point to something much more intentional and personal to meet their writing goals? How do we support this kind of critical thinking and meaningful writing in the age of generative AI? Table 1 displays a few prototypes of activities that are meant as frameworks that can be edited and revised for classroom use across the curriculum for various course contexts, purposes, and goals.

Table 1 *Prototypes of Activities for Critical Thinking and Writing with Generative AI*

Use of Generative AI for Writing	How	Opportunities for Teaching and Learning	Connection to Existing Literature on Meaningful Writing
Text Generation	Student provides inputs for generating text and then brings the generated output to fruition	Distinguish good writing from bad; develop critical thinking skills and evaluation skills; practice writing good prompts to get to the desired output; prototype fast and see how different scenarios emerge quickly; distill new and useful material by generating the text, evaluating its quality, iterating until something worthwhile comes up	Make choices; discuss with others (potentially); grow and develop personally; take on identities as authors and writers; reflect by drawing from past knowledge and experience to determine if the output is meaningful (i.e., makes sense)

Summarization of one’s own text	Student writes a piece of text and asks the system to	Turn stream-of-consciousness writing into something more coherent through summarization; learn how to	Take on identities as writers or authors; have a venue to test self-
---------------------------------	---	---	--

Table 1 (continued)

Use of Generative AI for Writing	How	Opportunities for Teaching and Learning	Connection to Existing Literature on Meaningful Writing
	summarize it	write coherently and succinctly based on how the system summarizes text	expression; develop personally
Text Style Transfer	Student writes a piece of text and then asks the system to modify its style (e.g., to Shakespearean or humor, etc.)	Turn one’s writing into something more vibrant/creative; learn about styles without having to read bodies of work—e.g., turn a piece of text Kafkaesque without the need to read Kafka originals to learn how his style works/differs; develop a one-off “style” by adding a touch of Kafka, humor, irony, etc. to one’s own way of writing	Take on identities as writers (e.g., humorous writing but in a business context); connect personally to the styles/writers one enjoys; make own choices; grow and develop personally
Q&A	Student prompts the system with a question, receives the generated output, and evaluates its veracity, accuracy, and usefulness	Practice critical thinking skills (evaluative skills); practice asking the “right” questions to get to useful output; develop the ability/discipline to fact-check	Make own choices; discuss with others; grow and develop personally

Given what we understand about student writing and meaningfulness, we offer one piece of advice when experimenting with GPT or any new writing tools: make sure that *reflection* is a part of the process. While we emphasize various forms of agency in the chart above—such as making choices and taking on identities—we want to ensure that students are, of course, learning and acquiring knowledge that stays with them beyond the classroom. Metacognition is key here, and reflection can take different forms, such as in-class discussion, peer conferences, and/or writing exercises of varied formality. Moreover, we recognize the hesitance some instructors might feel in that the sample adaptable exercises above could lead down a less-than-ideal path for students (e.g., passing GPT text off as their own); however, within the space of the classroom and through reflecting on their work, students can learn how to ethically and effectively use these tools and start to move from one meaningful writing assignment to overall meaningful learning processes and experiences.

AI and Meaningful Learning Experiences

Meaningful writing assignments aid in fostering critical thinking, as they encourage students to actively participate in their writing processes. Recognizing this participation is crucial when creating assignments across the disciplines, particularly in an era increasingly influenced by AI. While AI tools can be beneficial, it is important to ensure they do not overshadow the vital skill of critical thinking. To support meaningful writing in the age of AI, students need agency to make choices regarding what they write about and how to write for particular genres and audiences—both in terms of content and tone. Adherence to critical thinking spiced with lively discussions and solid reflection is crucial to ensuring that

generative AI remains a tool in the hands of students rather than becoming the primary creator of student texts. Using this technology merely as a tool (rather than a creator) gives students an opportunity to exercise critical thinking and fact-checking, which can result in new opportunities for agency, engagement, and transferable learning. In many ways, this (human) critical thinking about generative AI's output fills in the "evaluation" piece, which—for better or worse—is missing from the way generative AI is currently designed.

Through this type of engagement, students will learn that taking generative AI's output as is leaves writing prone to homogenizing everyone's message and stripping out individuality and heart, not to mention the risk of producing text that is factually inaccurate, "meaningless"—that is, senseless or useless—if not dangerous (e.g., fake news that has real consequences). On the contrary, using its output as a starting point—taking it with a grain of salt and editing, fact checking, researching, engaging with others and, above all, adding one's own passion and heart to the writing—is going to make all the difference when everyone is given access to the same resource. Who knows, in the age of generative AI, maybe this is the new path to reinvigorating writing as an iterative process that fosters and preserves critical thought?

Notes

¹See Khachatryan, D. (2022). Playmeans: Inclusive and engaging data science through music. *Journal of Statistics and Data Science Education*, 31(2), 151–161. <https://www.tandfonline.com/doi/full/10.1080/26939169.2022.2138801>.

²See <https://gptzero.me>.

³Words that are thought to appear in similar contexts (think "candies" and "lollypops") are assumed to receive numeric representations that are similar. Importantly, the eventual numeric embeddings that represent each piece of text are not arbitrary. Instead, embeddings result from a successive process of refinements using deep artificial neural networks—mathematical models first proposed in 1943 (McCulloch & Pitts, 1943; Goodfellow et al., 2016)—that are thought to roughly represent how a human brain may process information. Through multiple rounds of such "distillation" using neural networks and a mechanism called "attention" (where more attention is paid to some parts of the text than others), embeddings get better and better at preserving the underlying contextual meaning of the text they represent.

⁴Students could choose a different GPT platform, but they all picked ChatGPT for this assignment, most likely given the buzz about it in the news and on campus.

References

- Anderson, L. W., & Krathwohl, D. R. (Eds.). (2001). *A taxonomy for learning, teaching and assessing: A revision of Bloom's taxonomy of educational objectives* (Complete ed.). Longman.
- Bloom, B. S., Engelhart, M. D., Furst, E. J., Hill, W. H., & Krathwohl, D. R. (1956). *Taxonomy of educational objectives: The classification of educational goals by a committee of college and university examiners (Handbook I: Cognitive domain)*. Longmans.
- Chollet, F. & Allaire, J. J. (2018). *Deep learning with R*. Manning.
- Eodice, M., Lerner, N., & Geller, A. E. (2016). *The meaningful writing project: Learning, teaching, and writing in higher education*. Utah State University Press.

- Eodice, M., Geller, A. E., & Lerner, N. (2019). The power of personal connection for undergraduate student writers. *Research in the Teaching of English*, 53(4), 320–339. <https://doi.org/10.58680/rte201930141>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Marche, S. (2022, December 6). The college essay is dead. *The Atlantic*. <http://tinyurl.com/2cnhbhn7>
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115–133.
- Tingiris S., & Kinsella B. (2021). *Exploring GPT-3: An unofficial first look at the general-purpose language processing API from OpenAI*. Packt.
- Tunstall, L., Von Werra, L., & Wolf, T. (2022). *Natural language processing with transformers*. O'Reilly Media.
- Vaswani A., Shazeer N., Parmar N., Uszkoreit J., Jones L., Gomez A., Kaiser L., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
- Wolfram, S. (2023). What is ChatGPT doing . . . and why does it work? <http://tinyurl.com/43anvwyb>