

# AI and new forms of knowledge work: How writing education can help equip students for an uncertain future<sup>1</sup>

Chad Nilep  
*Nagoya University*

---

One way in which disruptive technologies have historically resulted in social and economic change is by creating new jobs while displacing or reducing the value of existing jobs. While it is easy to predict that current developments under the label of “artificial intelligence” will transform work, it is exceedingly difficult to predict just what skills will best serve in the future. Perhaps the best tools that educators can provide today come from classical liberal arts, especially logic, rhetoric, and critical thinking. Questions of how to improve job skills, in addition to being difficult to answer, may be counter-productive to ask. As Ekbia and Nardi (2017) note, 21<sup>st</sup> century innovations from social media to self-service grocery checkout so far have served not to reduce human labor but to transfer its value to the corporations that control the technology. Thus, important questions arise regarding how best to govern and to organize society in ways that allow for continued gains without sacrificing equity. Obviously, writing education will not answer these questions on its own, but our students will need to grapple with them. The changes wrought by earlier disruptive communication technologies took more than a decade to resolve and are in many ways still being worked out. Therefore, the most beneficial contributions of writing education at present are to develop skills in critical inquiry, rhetorical reading and writing, and logical argumentation with which today’s learners can face tomorrow’s crises.

---

## 1. Disruption, winners, and losers

Generative AI is described as a “disruptive technology”, one that promises to remake the economy, labor, and society. Historically, disruptive technologies from steam power to the microprocessor have enriched societies by reducing labor costs, increasing productivity, and sometimes introducing new categories of products. At the same time, however, these technologies have tended to displace some workers or to reduce the value of their existing jobs. Industrial and social progress usually – perhaps inevitably – produces winners and losers. Some individuals gain a great deal from changes, while others gain only marginally. Likewise, some individuals stand to lose a great deal more than they gain, while for others such progress is indeed a boon.

One familiar story about the Industrial Revolution of the 19th century in England involves the Luddites, individuals who feared that the new machine age would somehow destroy the world. According to one version of the story, the Luddites resisted progress by smashing engines and burning down factories in a futile attempt to stop the arrival of the industrial future in which we would all enjoy a higher standard of living through modern technology. The epithet *luddite*, still used to disparage technophobes whose excessive conservatism is thought to slow

---

<sup>1</sup> This manuscript is based on “Equipping students for an uncertain future”, an oral presentation at the 5<sup>th</sup> International Symposium on Academic Writing and Critical Thinking, held at Nagoya University 16 February 2024, and responses from participants in the symposium.

progress, comes from this telling. This version of the story is, however, wrong in most of its details. In fact, the Friends of Ned Ludd, as they styled themselves, were not a band of English rustics but rather a labor movement with specific concerns about how their work would be valued in the machine age. They were mainly highly skilled, highly paid workers in the textile industry, such as carters, weavers, and knitters. While the Luddites did smash factory machines, among other forms of labor action, this was not because they believed that the machines would destroy the world in some unforeseeable manner.<sup>2</sup> Rather, it was a rational fear of how coming changes in the market would affect their own standard of living. Skilled artisans worried that the availability of inexpensive, factory-made cloth would reduce the need for their productive labor and cost them their high standard of living. With the coming of industrialization, this is indeed what happened. By 1850 hand-made cloth was essentially gone from the English marketplace, and with it the high wages of the people who had made it. While the Industrial Revolution created new jobs and a generally rising standard of living, it destroyed the standard of living of some artisans. Average pay for members of the working class was essentially flat through most of the 19<sup>th</sup> century (Kestenbaum and Goldstein 2015). This average, however, masks the inequality among those who gained and those who lost income during the period. Nicolas Crafts (2010), synthesizing work by Robert Allen (2009) and Joel Mokyr (2009), suggests that less than five percent of Britain's population accounted for about one third of the gains during the Industrial Revolution. Incomes for the Friends of Ned Ludd dropped sharply and remained lower for the next generation of workers than had been the case for these highly paid artisans. It would not be until their grandchildren's generation that factory workers' standards of living rose to the level that had been enjoyed by the Friends of Ned Ludd (Kestenbaum and Goldstein 2015).

Similar to the ways that factory machines and fossil fuels disrupted laborers in manufacturing, some contemporary commentators predict that generative AI will remake knowledge work. In fields from education to law, from journalism to management and far beyond, generative AI may replace many of the tasks today performed by educated workers. It is difficult to predict whether these changes will reduce the value (and the compensation) of human workers, free them to work in more productive ways, or increase demand for their labor, perhaps through the introduction of new categories of product or services. What we can expect, though, is that there will likely be changes in how we labor in the future, and that this will produce economic winners and losers.

The anthropologist Ilana Gershon compared her recent studies of AI in the workplace with her earlier work on New Media (e.g. Gershon 2011). Gershon (2023) said that when she started studying New Media in the 2000s she was "constantly surprised" during fieldwork. Of course, it is common to be surprised during fieldwork; that surprise is a major reason to do it. Scholars go to the field to discover new things and then analyze them. But what was unusual about studying New Media, Gershon suggests, was the *constant* surprise. Every time she returned to a field site or re-interviewed subjects, she would find people using internet or cellular communication in new and unexpected ways. This constant surprise continued for about ten years until people gradually settled on new norms for communication in the workplace and their private lives. She says that recent swings in public discourse over AI, from techno-

---

<sup>2</sup> Given what we now know about links between fossil fuel emissions and climate change, such a prediction, had it been made, might seem prescient now.

optimism to moral panic to resignation, remind her of that constant surprise. And like norms around New Media, it may take a decade or so to settle on norms for the use of AI at work, school, and anywhere else we end up using it.

Since it is currently difficult to predict what knowledge work will look like in just a few years, educators today struggle to decide how to deal with AI in education. What should we be teaching students about the nature and usefulness of current and coming applications? How and to what degree should we guard against students' use of AI as a form of academic dishonesty, for example by submitting computer-generated term papers? Should we, as educators, embrace it as a useful tool for our own work, or should we fear it as a potential threat to our own positions? My own struggles with these questions may serve to illustrate some of the challenges involved, as well as my reason for looking past specific answers and focusing instead on rhetoric and logic in writing education.

In April 2023, the start of the first school term after the release of GPT-3.5 and ChatGPT (OpenAI 2022), I spent a considerable amount of time developing and explaining an AI policy for my undergraduate writing courses. I illustrated ways in which AI can produce grammatically correct texts that were nonsensical or factually wrong. During the first class meeting we watched excerpts from a YouTube video (Cocktail Chemistry 2023) in which a bartender asks ChatGPT to invent a new cocktail recipe, then mixes and tastes the cocktail. One image captured from the video shows a quite plausible-looking recipe in the ChatGPT display screen. The recipe shows measurements of ingredients, instructions for mixing them appropriately, and a description of the cocktail. Later in the video, the bartender samples the cocktail that he mixed using the recipe and declares it too sweet; apparently the algorithm did not know that both grapes and syrup are sweeteners. I concluded the lesson by pointing out that algorithms do not think in the same way that humans do or know things that are common knowledge among people. Therefore, the texts that they produce tend to be grammatically correct but not necessarily meaningful. Instead of expecting AI output to make sense, I cautioned the students, they need to bring meaning to their writing and rely on the algorithms solely for assistance.

My course policy in April 2023 regarding the use of AI in class had four points. First, students were encouraged to use AI or other software to check for errors of grammar or spelling. Second, students were allowed (though not necessarily encouraged) to use some amount of translation, especially at the level of individual sentences or phrases. I warned them, however, not to translate entire paragraphs or essays, as the results tend to sound stilted and are sometimes nonsensical. Third, I warned students that they should not write paragraphs or essays in Japanese and use machine translation to render it into English text. As described above, the results of such practice are less than ideal. Moreover, the point of the course was to learn methods for communicating to various audiences using appropriate styles and patterns (read: rhetoric, argumentation, and grammar) in English. Finally, I declared that students must never submit work produced by AI without explaining what tools were used and how. Students who violate this rule would receive an F in the course.

Notwithstanding the prohibition on AI-produced work, I received some suspicious papers. In each of these cases I put a sample of the paper into GPTZero (Tian, Cui, and Mimi 2023) and received a numerical estimate of the likelihood that GPT or similar algorithms had produced it. I then contacted the student, confronted them with my suspicion and the numerical evidence, and asked them to provide evidence that they had drafted earlier versions of the paper.

In every case, the students I approached admitted that they had submitted a paper produced by ChatGPT or a similar product.<sup>3</sup>

In contrast to the detailed but not completely effective policy described in April 2023, my courses in October 2023 used a much less comprehensive policy. I did not dedicate the first course meeting to illustrating AI's foibles. I merely told students that such tools are useful but fallible, and that I would not accept assignments that had been created with generative AI that was not disclosed. As in the spring term, I received a small number of suspicious papers. Again, I confronted students with my suspicions and with numerical evidence suggesting that the work was computer-generated, and again each student confessed to the violation. In both cases – the detailed policy description of the spring courses and the brief mention in the autumn term – I received about the same number of suspicious papers, and in each case the number of such papers was small – somewhat less than 5% of all papers received.<sup>4</sup> That suggests either that most students do not attempt to cheat, or (less probably) that most cheating is not detected.

Studies of plagiarism (e.g. Comas and Surenda 2010, Park 2010) suggest that cheating in writing courses is most common when students lack academic experience or feel alienated from the course content and practices. The writing courses described above are required courses for first-year undergraduates. Students therefore must enroll and earn a passing grade even if they have no interest in the topics. The relatively low rate of attempted cheating in courses where it may be expected suggest (albeit anecdotally) that education policy and teachers' attention may be better focused elsewhere. Furthermore, given that the students who feel tempted to submit AI-generated papers are usually the least confident English users, it is not hard to guess that their essays filled with complex grammar and obscure vocabulary are not their own work. In such courses, carefully crafted and enforced AI policies may not be strictly necessary.

Even if AI is not a problem to be defended against, however, questions of how best to structure writing courses for a future in which its use is commonplace are not easy to answer. Experiences with then-new information and communication technologies at the beginning of this century suggest that answers may not be forthcoming for several years. But a first step in adapting writing education should involve thinking about the nature of AI and thinking about the nature and goals of writing education.

## **2. On the nature of AI**

In the early 2000s I worked briefly with interdisciplinary teams in natural language processing. Our goals at that time included making systems to read texts and parse their syntactic structure; to transcribe speech into written words; and to understand and respond appropriately to spoken requests over the telephone, among others. Today, computer systems can do most of these tasks – not perfectly, but far beyond what was possible a generation ago. There was a common joke in interdisciplinary labs at that time: “The more linguists that we fire from the team, the better the algorithm works.” The first time I, as a linguist, heard this joke, I understood it to be a form of good-natured ribbing at my expense. In addition, though,

---

<sup>3</sup> Despite the course policy stating that such behavior would result in the student failing the course, I instead gave no credit for that single paper and required the offending student to file a report detailing what they had done. Future violations of the university's AI policy, published subsequently, might result in loss of course credit or expulsion.

<sup>4</sup> This figure is avowedly inexact. To protect students' privacy, precise data was not collected or analyzed. Figures in this section should be understood merely as general descriptions and not the result of any quantitative analysis.

the joke provides a useful window into how the teams approached these tasks, and into the nature of AI systems. Linguists are interested in – among other things – how the human mind processes language. As a result, linguists on the teams tended to pursue computational models that work the way we think that the human mind works. Engineers and computer scientists, on the other hand, were interested in designing systems that give the correct output. This can be a much simpler problem. “Firing linguists”, therefore, may simplify the process by concentrating on the output without excessive concern for how it comes about.

The output of language-based generative AI systems resembles human communication, but it may not be correct to think of these as the same thing. Elena Esposito (2017, 2022) reframes natural language processing systems not as “artificial intelligence” but as “artificial communication”. When human beings interact with AI-based systems, we experience a form of communication that is different from interaction with other humans – perhaps so different that we should not regard the algorithms as participants in communication. Although the algorithm produces text or sound that appears to be the product of intelligence, it is difficult to say that the algorithm has knowledge or communicative intent in the way that a human being would in the same situation. It may be nearer the mark to say that the actual intelligence present is in the human being’s interpretation of the text more than the algorithm’s production of it.

Christoph Purschke and colleagues (2023) further limit communication between humans and algorithms from “artificial communication” to “artefactual communication”. Large language models are artefacts, built from the texts they are trained on. And these texts are themselves artefacts of human intelligence. Thus, what we call AI – large, pre-trained models built from corpora of texts, images, music, or what have you – are not *intelligent* in the same way that humans or animals may be. They are good at producing output that resembles the kinds of material they were trained on, but the resemblance to the human intelligence that produced that training input is at best imperfect, and perhaps of a different type.

### **3. Looking back to move forward**

How should we think about the relationship between writing education now and the knowledge work of the future? As described in section 1, new technologies tend to have winners and losers. The Industrial Revolution displaced artisans, but it created factory jobs – and for those who controlled the new means of production, it created the modern industrialist. New media displaced corporate in-house graphics departments but created web developers and search engine optimizers. It is a relatively safe bet that generative AI and the applications that come from it will create new ways of working and will displace certain tasks from the workplace. It is, at the same time, extremely difficult to predict exactly what those changes will look like. Although I will not here venture specific predictions, I do want to reflect on worries about what may result from currently ongoing disruptions.

Hamid Ekbia and Bonnie Nardi (2017) note that work in which human labor combines with artificial systems is not just coming: it is already here. They call this computer-mediated work *heteromation*. Unlike automation, in which industrial machines do work in place of human laborers, heteromation relies on both human and computer inputs. This includes creative work, such as social media “influencers” or YouTubers who create works which computer algorithms then deliver to audiences; service providers, such as ride-sharing drivers whose work is sold via apps; and self-service systems such as grocery check-out, in which customers or users are themselves the laborers. In heteromated systems, Ekbia and Nardi point out, human labor is

essential, yet the bulk of the monetary value – in some cases, its entirety – goes to the corporations or businesses that control the computer systems. In a system where human user-laborers contribute most of the value, they argue, those users should share more fully in decision-making and economic gains.

We are beginning to grapple with new technology that will affect – and is affecting – how we work and live. Some changes will have major implications, and some only minor consequences. Society will need to balance various demands such as productivity versus equality, and the needs of workers and producers against the desires of consumers and users and the hopes of owners and investors. These changes are new in their particulars (heteromation of human and machine labor; novel products and jobs), but they are familiar in their broad contours. There will be winners and losers, and groups and individuals will try to do what is best for themselves, their group, and the world more broadly. In the face of this specifically uncertain but broadly predictable future, the best thing that I can do for my students and the institutions I consider my in-group is to pattern my teaching on classical liberal arts.

From Europe's Classical period, training in logic and rhetoric has been considered fundamental to liberal education. In medieval Europe these subjects, along with grammar, were identified as the Trivium, the elementary subjects of liberal arts education. Looking back to these long-practiced subjects may be the best way – or at a minimum, one way to prepare today's students for tomorrow's challenges. In my own teaching I emphasize logic and rhetoric more than grammar, for two reasons.

The first reason that I do not feel the need to emphasize grammar in my writing courses is that the students I mainly work with, scholars writing in a language that is not their dominant spoken language, are well versed and sometimes overly focused on grammar. It takes a lot of time and effort to master the form of a foreign language, and in that process we convince ourselves that the form of that language is what is most important. Of course grammar is important, but in concentrating too much on linguistic form we may forget that it is fundamentally a tool for the communication of thought. Language – grammar, along with rhetoric and logic – allows us to analyze ideas and to communicate our own ideas to other people. It is those ideas, rather than the tools used for analyzing and communicating them, that are central to the process.

The second reason to de-emphasize it, is that AI powered applications are already relatively good at producing good grammar. In fact, I encourage my students to use grammar checking applications to help them spot errors or suggest appropriate vocabulary. But I caution them not to trust what the applications suggest. While AI is relatively good at producing sentences like the ones it was trained on, it does not understand the ideas that those sentences represent. As a result, AI can “hallucinate” untrue ideas. Grammar checking systems, likewise, sometimes suggest changes in the name of eliminating grammatical errors that have the unwanted effect of changing the meaning that the writer intended. Therefore, I advise students to use software solutions for the tasks they are good at, such as spotting errors or providing first-pass translations of text from one language to another, but I caution them to be wary of the output and to take care that the corrected or translated texts accurately communicate the writer's ideas.

Relying on medieval and classical education may seem perverse as a response to new technologies. Perhaps, like the historical Luddites, I seem to be rejecting new technology in order to preserve an occupation that is becoming outmoded. I do not believe that is the case,

though. The situation we find ourselves in today around emerging technologies built on generative AI is new in its specific details. In its broad outlines, however, the situation is familiar and at least somewhat predictable. New styles of communication, new jobs and work styles, and new social norms are emerging all the time. What is true now and will be true in the future is that people will need to agree on ways to integrate these novelties into their lives and societies. Rhetoric, logic, and the forms of critical thinking that flow from them offer the most effective tools for people to effect such change.

Logical argument – informal as well as formal argument forms – allow us to test our own knowledge and intuition, and to try to approach understandings of “truth” to the greatest extent that we can. It is important not only to recognize valid and invalid types of arguments, but also to be able to make valid arguments in order to produce clear and convincing texts. Rhetorical appeals, likewise, allow us to share our understanding with others, to convince them of the value of particular contributions to knowledge, and in so doing to build social structures and coalitions to act in common. Not only is the crafting of rhetorical appeals necessary, but the ability to read rhetorically (Peterson 2017) also allows us to respond most effectively, resisting shallow or manipulative appeals and embracing and joining those that further our best interests and what we believe to be best for society.

Generative artificial intelligence, like other new technologies, promises to disrupt labor, education, and society over the next several years. Like most disruptive moments, this one is likely to produce both positive and negative effects, which will almost certainly be unevenly distributed among people and groups. Although the specific changes are difficult to predict, it is easy to foresee that today’s students will need appropriate tools and training to deal fairly and effectively with whatever comes. The greatest contributions that writing education can make toward this future do not depend on embracing or resisting the technology, but on keeping our and our students’ attention on the fundamentals of communication, argument, and rhetoric, and the social goods that can flow from them.

### References

- Allen, Robert C. (2009) *The British Industrial Revolution in Global Perspective*. Cambridge: Cambridge University Press.
- Crafts, Nicholas (2010) “Explaining the first Industrial Revolution: two views”. *European Review of Economic History* 15(1):153-168. doi:10.1017/S1361491610000201
- Cocktail Chemistry (2023 January 1) “Can ChatGPT Make a Delicious Cocktail? I Put It to the Test”. YouTube. <https://youtu.be/ld7iF9zradQ>
- Comas-Forgas, Ruben, and Jaume Surenda-Negre (2010) “Academic plagiarism: Explanatory factors from students’ perspective”. *Journal of Academic Ethics* 8(3):217-232. doi: 10.1007/s10805-010-9121-0
- Ekbia, Hamid R., and Bonnie A. Nardi (2017) *Heteromation, and Other Stories of Computing and Capitalism*. MIT Press.
- Esposito, Elena (2017) “Artificial communication? The production of contingency by algorithms”. *Zeitschrift für Soziologie* 46(4): 249-265. doi:10.1515/zfsoz-2017-1014
- Esposito, Elena (2022) *Artificial Communication: How Algorithms Produce Social Intelligence*. MIT Press.
- Gershon, Ilana (2010) *The Break-up 2.0: Disconnecting Over New Media*. Cornell University Press.

- Gershon, Ilana (2023 May 26) “ChatGPT: Genre, scale, animacy”. Keynote address, *Technolinguistics in Practice: Socially Situating Language in AI Systems*. University of Siegen.
- Kestenbaum, David, and Jacob Goldstein (2015 May 6) “Episode 621: When Luddites Attack”. Planet Money. <https://www.npr.org/transcripts/404701816>
- Mokyr, Jerome (2009) *The Enlightened Economy: An Economic History of Britain 1700–1850*. New Haven: Yale University Press.
- OpenAI (2022 November 30) “Introducing ChatGPT”. Blog. <https://openai.com/blog/chatgpt>
- Park, Chris (2010) “In other (people’s) words: Plagiarism by university students – literature and lessons”. *Assessment and Evaluation in Higher Education* 28:471-488.
- Peterson, John (2017 February 18) “Reading Rhetorically to Support Writing Instruction”. *Third Symposium on Academic Writing and Critical Thinking*. Nagoya University.
- Purschke, Christoph, Alistair Plum, and Catherine Tebaldi (2023 May 25) “From ‘natural’ to ‘culturally grounded’ and ‘socially anchored’: examining the notion of language in NLP.” *Technolinguistics in Practice: Socially Situating Language in AI Systems*. University of Siegen.
- Tian, Edward, Alex Cui, and Yazan Mimi (2023 January 3) *GPTZero*. <https://gptzero.cc/>