

This work is licensed under CC BY-NC 4.0. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/4.0/>

**TEXT**

**GEN**

**ED**

**Inofficial PDF version**

Please visit

<https://wac.colostate.edu/repository/collections/textgened/>  
to access the official web publication

Teaching with Text  
Generation Technologies

Edited by

Annette Vee, Tim Laquintano & Carly Schnitzler

# TextGenEd

## An Introduction to Teaching with Text Generation Technologies

by Tim Laquintano, Carly Schnitzler, and Annette Vee

When we issued the CFP for this collection, teaching and research in automated writing was still rather niche. In the language arts, it existed in critical code studies and creative domains such as computational poetry and, more broadly, electronic literature. In writing studies, interest in automated writing existed in corners of technical writing, computers and writing, and rhetoric. Most writing teachers are comfortable with word processing, content management systems, search, and email, and it has been possible to run a writing class with little else. Now, with the introduction of ChatGPT, it might soon become difficult to research and teach writing *without* thinking about, or addressing, automated writing technologies and artificial intelligence (AI). As Big Tech rushes ahead in its AI arms race with the intention of having large language models (LLMs) mediate most of our written communication, writers and teachers are forced to consider issues of prompt engineering, alignment, data bias, and even such technical details as language model temperature alongside issues of style, tone, genre and audience.

At the cusp of this moment defined by generative AI, *TextGenEd* collects early experiments in pedagogy with generative text technology, including but not limited to AI. The resources in this collection will help writing teachers to integrate computational writing technologies into their assignments. Many of the assignments ask teachers and students to critically probe the affordances and limits of computational writing tools. Some assignments ask students to generate Markov chains (statistically sequenced language blocks) or design simple neural networks and others ask students to use AI platforms in order to critique or gain fluency with them. A few assignments require teachers to have significant familiarity with text generation technologies in order to lead students, but most are set up to allow teachers and students to explore together. Regardless of their approach, all of these assignments now speak to the contemporary writing landscape that is currently being shaped by generative AI. Put another way, the assignments in this collection offer initial answers to urgent calls for AI literacy.

We hope this collection offers something for teachers with all levels of comfort with technologies—from teachers seasoned with digital writing technologies to teachers approaching the entire domain with trepidation. To that end, we have made the teaching resources in this collection as accessible as possible. WAC Clearinghouse is publishing the collection as fully open access and all of the assignments are licensed as Creative Commons Attribution-Noncommercial ([CC-BY-NC](https://creativecommons.org/licenses/by-nc/4.0/)), which means that nonprofit educators are free to adapt and use and share them (with credit to the source) as they see fit. We hope they will!

One requirement of every assignment accepted for this collection was that instructors had taught it at least once. So, all assignments include a description of how students responded along with reflections from the instructors. Short abstracts accompany each assignment and detailed implementations are included. Assignments are organized according to learning goals relevant to writing:

rhetorical engagements; AI literacy; ethical considerations; creative explorations; and professional writing. We hope instructors treat this as a living collection, adapting the assignments to local conditions and new technologies as they evolve.

As context for this collection of assignments, we provide below a brief introduction to past, present, and future attempts to automate writing. This general framework can guide instructors' understanding and their selection of what to emphasize in their courses, especially given the hype that surrounds contemporary generative AI. This collection works alongside many emerging resources for instructors, including [panels sponsored by CCCC](#) and [by MLA](#), a working paper authored by a [MLA/CCCC joint task force](#), a [recent forum in \*Composition Studies\*](#), a [WAC Clearinghouse resource curated by Anna Mills](#), and published research across many academic disciplines, from sociology to rhetoric. Many of the scholars whose assignments appear in this collection also publish on generative AI and other text generation technologies.

It will take all of us to respond to this moment. As editors of this collection, we believe that generative AI is the most influential technology in writing in decades—nothing since the word processor has promised as much impact. And generative AI is moving much faster. Although generative technology for text has been quite good for the last 5 years, it's been less than a year since the watershed release of ChatGPT in November 2022, which by many measures has been one of the fastest growing technologies in the history of humanity. A technology this impactful to education requires collective response and collaboration from teachers. This collection has allowed us to put our heads together with some of the most thoughtful and innovative writing teachers across English studies and beyond. May their ideas invigorate your teaching as much as they have ours.

## A Brief History of Automated and Computational Writing

While conversations about text generation with AI sometimes present it as a fully new phenomenon, automated writing has its origins much earlier. In the seventeenth century, mathematician G.W. Leibniz invented a cryptographic cipher machine that “would serve to encipher and decipher letters, and do this with great swiftness and in a manner indecipherable by others” (Rescher). In Swift's *Gulliver's Travels* (1726), a Lagado professor engineered an automated system of writing including young scholar-laborers, blocks of wood, wires and cranks “so that the most ignorant person, at a reasonable charge, and with a little bodily labour, might write books in philosophy, poetry, politics, laws, mathematics, and theology, without the least assistance from genius or study.” Automata that ran on complex clockwork mechanisms proliferated in the 18th and 19th centuries, largely as a way for mechanics and clockmakers to show off their technical prowess (Riskin). These automata, powered by the winding of gears, could variously dance, write, draw, breathe, and, in the case of one mechanical duck, defecate.

The automation of writing—a uniquely human activity—accompanied conversations about artificial intelligence, even in the early modern era, long before the term came about. With the invention of the computer in the 20th century, the connections between writing and AI grew tighter, most clearly illustrated in Alan Turing's 1950 article in the philosophy journal, *Mind*: “Computing Machinery and Intelligence.” At the time, computers were humans

(mostly women), and digital computers were primarily used for complex calculations, especially in wartime military contexts. Amid the ballistic calculations, Turing speculated on a prompt from his teacher, philosopher Ludwig Wittgenstein: Can machines think? Both men thought it was a ridiculous question—Wittgenstein because he thought machines were nothing like humans and Turing because he wasn't even sure we knew what *humans* thought. But, Turing argued that if a machine could fool a human into thinking it was a human, then it could be said to think. The machine—a computer—would naturally use writing for this deception. Writing, in other words, is thinking—and the automation of writing is machine thinking.

By the early 1950s, computation had advanced to the point where programs could be written to generate text. While awaiting his first assignment at Britain's National Research and Development Corporation in the summer of 1952, British computer scientist Christopher Strachey—a collaborator and friend of Turing who also invented a precursor to the programming language C—created a program that generated campy, over-the-top love letters, all signed by M.U.C., the Manchester University computer. One letter, reproduced below, was later printed in the arts magazine *Encounter* in 1954:

Honey Dear

My sympathetic affection beautifully attracts your affectionate enthusiasm. You are my loving adoration: my breathless adoration. My fellow feeling breathlessly hopes for your dear eagerness. My lovesick adoration cherishes your avid ardour.

Yours wistfully M. U. C. (Campbell-Kelly 25)

Strachey's love letter generator is widely-cited as the first work of electronic literature, a more flexible, fun, and digital version of the mechanical writing automata that preceded it (Rettberg). The emergence of e-literature and the generative creative texts in the decades that followed Strachey's generator established a sensibility of subversion, play, and critique. Even non-computational work by artist groups such as Oulipo were influenced by the combinatorial work done by those working on computers. Following Turing, Strachey and others, a small number of artists and programmers were going against the grain of what computation was generally designed to be used for—things like crunching census data and calculating the trajectory of ammunition in wartime. Instead, they were using computation to generate literature and art.

Early text generation worked with templates or statistical models such as Markov, a model where the next words in a chain are determined by probabilities. Even as computing became more accessible in the 1980s and 90s, text generation was a niche practice: for determined experimental artists or computational linguists huddling together in the AI winter, when funding for such work dropped in response to greater needs in basic literacy programs and defense (NCEE). Natural language processing—including understanding and generation—were both still active research areas with significant implications for transcription, translation, surveillance, and support for people with disabilities. Advances in machine learning, statistical methods, word embeddings, and dramatic increases in available compute plus data from the web all drove text generation technologies through the 2000s until now. In this collection, assignments by Boyd and Egan are particularly helpful in providing students with context for this history of text generation.

# The Current State of Text Gen Tech: Large Language Models (LLMs)

While earlier models of text generation leaned on grammatical rules, current models are more speculative—predicting the next word in a sequence based on patterns in its dataset. For text generation, large language models (LLMs) train on massive datasets gleaned primarily from the Web using machine learning techniques; they are then subjected to fine tuning and reinforcement learning through human feedback (RLHF) to hone desired output. Over the last ten years, and particularly since late 2017, these techniques have catapulted the field of generative AI, producing so-called "foundation models" that can generate text, image, video or sound across generalized contexts. Developments have been so dramatic that in technology news, AI podcasts and social media, the story told of generative AI is about our relentless march toward artificial general intelligence (AGI). Amid the distortion from overblown claims—no research field has promised so much and delivered so little as AI—there are real potentials and limits to generative AI. Yet, when the hype is dislodged from reality, these remain extremely difficult technologies to grasp: even AI scientists and engineers do not fully understand them or their implications. Below, we outline briefly how generative AI works for text generation and what variables might shape the future of text generation technologies. AI's dominant role in text generation right now means that soon even engagement with word processing might require a basic understanding of how contemporary LLMs work.

Large language models are called so because they *model* language. That is, they take examples of language and then use certain processes to attempt to reproduce it. We can therefore approach an understanding of LLMs by breaking down the *processes* they use and then the *data* they draw from.

## Processes

Contemporary LLMs are built with neural networks, souped-up versions of what Warren McCulloch and Walter Pitts introduced in 1944. McCulloch and Pitts borrowed the concept of a neuron from the human brain, comprised of billions of layers of interconnected tiny processors. The mathematical model of neurons fell out of favor in AI for decades, but has been revived with current "deep learning" techniques, so-called because contemporary artificial neural networks are many, many layers deep with simulated neurons that respond to information signals. Convolutions, backpropagation, and transformers—technologies that have accelerated generative AI since 2017—are all deep learning techniques that add layers of complexity to the neural network and can affect outputs.

The ways that contemporary neural networks recursively feed information back into the models has helped them to produce more coherent text across greater lengths of passages. Early text generation models could only generate short passages before they began to lose earlier details that were needed for coherence. In late 2017, Google researchers published the now famous paper "Attention is All You Need" (Vaswani, et al.), which enabled AI scientists to use transformer models to develop current LLMs. For writing instructors, the relevant detail to know about this advancement is that it effectively enabled language models to retain relevant information and place greater emphasis on earlier parts of the input. This is another way of saying that language models built using transformers could now sustain arguments, narratives, or discussion for thousands of words without "forgetting" crucial ideas from earlier in the prose. The expanded context window of LLMs is not

infinite, however, which is why LLMs that consumers can now access tend to be only capable of writing stories for a few thousand words at a time. Some newer models have larger context windows but for the time being remain difficult to access. Regardless of the specifics of the models, it is also important to note that because the networks are so complex, with so many hidden layers, and because models adjust their parameters based on feedback, even the programmers and engineers who design the models cannot fully trace the path from language input to output.

## Data

Large language models are called "large" because of the massive datasets they draw on to model language and the enormous amounts of parameters they have that the model uses to make predictions. AI scientists and engineers draw from large, open datasets such as [Common Crawl](#) (petabytes of text scraped from the Web) and websites such as Wikipedia. OpenAI's GPT-3 used CommonCrawl, outbound links from Reddit, Wikipedia, and text from books out of copyright in its dataset (Brown, et al.). (OpenAI has not revealed the data sources for more current versions of GPT, both for what they claim are safety reasons and to retain a competitive edge.) The datasets for contemporary LLMs such as GPT-3 are so large, they are relatively uncurated and unlabeled, although they've been "lightly filtered" (Brown, et al.). This filtering removes some of the most toxic language from a dataset, but the datasets tend to be so large that it has been difficult to clean all unsavory language, and Bender, et. al. also note that the inherent ambiguity of language means that scrubbing certain terms from datasets can preclude the perspectives of marginalized groups. Perhaps more importantly, if a dataset is so large that it can only be read through computational means, then it becomes extremely difficult to account for, or even understand, many of the possible worldviews in the data—although a variety of fields are now hard at work measuring the various kinds of bias embedded in LLMs through various benchmarks (mostly through more computational means). The problem of embedded bias is one of the reasons Bender, et al. have argued that LLMs can be too large. Each of the sources of data for GPT-3, for instance, over-represents men, white people, Western viewpoints and English language patterns. A language model built on that foundation is inevitably going to represent dominant perspectives. Datasets such as "[The Pile](#)" have been developed to attend to more diverse uses of language, and LLMs such as [BLOOM](#) include large amounts of non-English language training data in order to counter some of these biases.

## Recent Evolution of LLMs

Earlier language models needed to be fine-tuned to particular tasks in order to produce text that resembled good human writing—for instance, models that acted as chatbots in customer service. When GPT-3 came on the scene in 2020, it proved remarkably good at "few-shot learning" tasks—that is, it didn't need fine-tuning to a specific instance in order to produce coherent results that hit established NLP (natural language processing) benchmarks for quality. OpenAI achieved impressive results by scaling up both the parameters and the data they used, and they ushered in a new era and excitement about LLMs.

It is important to remember that the task of LLMs is simply to predict the next token given an input; it so happens that if you train them on enough data and enough compute, you begin to see emergent capabilities from the act of token prediction (e.g., the ability of LLMs to write computer code and simulate reasoning capabilities). But this prediction is also the reason why Emily Bender and colleagues insist that LLMs are tools of natural language *generation* and not natural language

*understanding*, even if the performance of some models is so good it feels to the user as if the models understand. But these models don't operate with an understanding of the world, or any "ground truth;" they work statistically. They model language based on associated terms and concepts in their datasets, always predicting the next word (in units called "tokens") from what's represented in their data. This prediction of the next token is also the reason language models can convey false information or "hallucinate". They don't know false from true—only statistical relationships between tokens.

Hallucination has not been the only problem with LLMs. When GPT-3 was released in 2020, researchers used adversarial testing to coax all manner of toxic and dangerous outputs from the model. This became something of a social media game when ChatGPT was released, as users made every attempt to "jailbreak" it in an attempt to get it to say nasty things. Numerous reports and swirling internet rumors suggested LLMs might provide good instructions for making methamphetamine or chemical weapons using ingredients available from Home Depot.

Engineers have developed a number of ways to try to mitigate these issues, including fine tuning, implementing safety guardrails (e.g., from blocking certain terms from being input and certain topics from being output), and reinforcement learning through human feedback (RLHF). In RLHF, humans help to train models by giving them question-answer pairs, rating the model's responses for accuracy and appropriateness, and identifying toxic responses (sometimes in the Global South for very little pay, see Perrigo). These methods have improved safety, eliminated some toxicity (a common joke is that the models have been through the corporate diversity training program), and improved the accuracy of responses. However, they are still not perfectly accurate and, given the philosophical complexity of representing "truth," likely never will be. The current hope of model designers and users seems to be that the accuracy of the models will be improved through add-on technologies and plug-ins (e.g., linking a LLM to a database of curated content to help prevent misinformation).

Developments in LLMs are coming at such a pace that it's difficult to keep up. But we can see a few trends: as machine learning techniques improve, the size of datasets and computation needed appears to be shrinking. Consequently, models with fewer parameters are producing more accurate outputs and the resources needed to run them have been shrinking, although as of this writing, the best models are still resource hogs. This means that we have entered a time of "model proliferation" that will lead to models with different purposes, politics, and values. We may soon see accessible models fine-tuned on personalized datasets (e.g., one's own emails), which might help language models better mimic the voice of the writer instead of producing the general, bland voice that has become relatively identifiable to some teachers of writing. AI plug-ins and apps will extend the capabilities of LLMs and be used in search as well as a host of other writing tasks, as language models begin linking various applications we use on a daily basis through a single interface. And, while many writers have been using ChatGPT as a standalone application, Google and Microsoft have begun embedding language models in their word processing systems and office software, a feature that will soon be rolled out on massive scales. Our writing environments will inevitably be shaped by these AI integrations, but it's unclear what effects this integration will have on our writing or writing processes. The only thing certain here is change—rapid change.

Despite Big Tech's insistence that these technologies will sweep the world, there are a number of variables that will affect their trajectory as writers decide the extent to which writing with AI is viable. These variables include:

1. Scale and access: Can engineers create language models that achieve decent performance without using extraordinary computing resources? If the technologies remain expensive to use and operate, what does it mean for access? Data at a large scale is impossible to review for accuracy or bias. As Bender, et al. ask: Can language models be too big?
2. Security and privacy: To what extent do language models leave users vulnerable to breaches of personal information, either in using the models or in having their data as part of the training set for the models? What security is possible in locally-run instances of language models?
3. Legality: Who will be liable for the harms created by the output of generative AI? Is it fair use for generative AI to mimic the styles of living authors and artists? How will copyright case law develop?
4. Implementation and user experience: How seamlessly will AI writing applications be integrated into now-standard technologies such as word processors and email clients? To what degree will writers or educators be able to decide on the level of integration or visibility of use for these language models?
5. Fact and ground truth: What methods will be developed to decrease inaccuracies (such as "hallucinations" of scholarly references or historical facts) in language models? Can reinforcement learning or connections to established databases prevent language models from their tendency to produce incorrect information?
6. Complementary technologies: What will language models be capable of when other applications become bolted onto them? To what degree will AI language models shape our digital discourse?
7. Abuse by malicious actors: Will the benefits of generative AI outweigh the potential harms they can create such as supporting disinformation campaigns?
8. Identification and disclosure: Software for detecting AI generated text has not proven to be particularly effective. A variety of solutions have been proposed, but for the time being it seems to be a cat and mouse game that seems to be initiating a crisis of social trust related to certain kinds of writing.
9. Social stigma: Upon its arrival, ChatGPT received intense press coverage that framed it as a cheating technology for students. To what extent will collective impressions of the technology shape its trajectory?
10. Style and language bias: Language models write with "standard" grammar in languages that are well-represented in the dataset, such as English. Given significant bias against "accented" writing in educational and professional contexts, how will language models affect writers' or readers' perceptions of "accent" in writing?
11. Lesser-known or minoritized languages: How will languages and discourse with little or minoritized representation in the training data be reflected in language models? Will smaller language models be tailored for use by these discourse or language communities? Will supervised learning or synthetic data supplement training to enhance representation? To what degree can or will minoritized discourse communities embrace language models?

While generative AI with language models is the overwhelming force and background in the contemporary writing scene as well as this collection, it is not the whole picture. The spirit of early creative computational writing, for example, is still very much alive both apart from and inclusive of uses of LLMs. Creative uses of computation have evolved alongside the technologies themselves. A wide variety of



tools exist to make creative text generation accessible in a pedagogical context. Educators in this collection employ user-friendly tools and libraries like Tracery, RiTA.js, and Markovify to both teach about text generation technologies, and about creative constraint, as it predates and contextualizes AI text generation.

## Shaping Writing's Futures

Regardless of the power of new language models, nothing happens in the writer's life without implementation. And implementation is often a messy process. Implementation is when we learn whether or not tools are useful to us, when we adjust to new and clunky interfaces, and when we suss out exactly how hollow or flush the promises of big tech's marketing language is. Implementation is also an obfuscatory process. The environmental impact of AI, the potential for it to induce extensive job loss, the potential for it to remove thought and care from human work, will not be altogether apparent to the average user of a Google doc who clicks a "Help me write" button and has the tone of their paragraph changed. To the first generation of AI users, it might feel like magic. To the second generation, it might feel ordinary.

For many writers, the near future will be an experiment in implementation. Like literacy practices themselves, the implementation of new writing tools will be highly sensitive to context as writers assess their needs, and their organizations' needs, to automate rhetorical practices against the backdrop of questions about data security, privacy, resources, and goals. Writing instructors and higher education as a whole will also be working to determine how implementation will happen in our lives and in the lives of our students. If there is one benefit to the otherwise harrowing "AI arms race," it is that many of these tools have already come online undercooked and with a clunky or creepy user experience that might stall their adoption. A potential delay in widespread use could buy us some time to learn more about them, understand them, and generate research about how they are used.

### Potential Paths

Even if the variables above restrict the spread of AI, it will be widespread enough that writing teachers need to prepare. We've seen the hazy outlines of four responses begin to emerge:

1. Prohibition: We are skeptical that this will be a viable model. In the near future, any writing done in a word processor will likely be difficult to do without some AI intervention, whether tacit or explicit. Moreover, we are not convinced by any current research that accurate labeling of AI generated prose—which is currently unreliable—will ever be available. A student "honor code" could sidestep the labeling challenge in a prohibition path, but only if students understand when AI intervenes in their prose. As of this writing, Grammarly has integrated a LLM into its interface; Google has a "Help me write" feature that obscures that it's an LLM, and Microsoft is on track to implement a similar feature in Word. Big Tech plans to integrate AI into its next generation search technologies, and complete prohibition might very well lead to an eventual de-skilling of students, something Antonio Byrd has recognized in a recent forum on AI and Writing in *Composition Studies*. While turning back time before generative AI is not an option, some restrictions on students using LLMs may be beneficial. Well-crafted assignments can create conditions in which students might receive only minimal advantage from engagement with AI.

2. **Leaning In:** Some professors have advocated that generative AI is the future of writing and that we should be leaning into the use of language models, having them assist with most if not all stages of the writing process. This might be where we all wind up, but it is crucial to note that an uncritical stance that accepts the discourse of inevitability is unlikely to empower students or educators, and the open issues we mention above can disrupt any full embrace of language models. Yet treating AI as a collaborator—such as some assignments in this collection advocate—can equip students to prepare and even shape a future with AI writing.
3. **Critical exploration:** Students can probe the limits of the technology while learning how to use it. This is the direction we believe to be the most beneficial to our students and which is implied by many of the assignments in this collection that require LLM use. These assignments ask students to flush out data bias, rhetorically examine the output of LLMs, compare their writing to the writing of language models, and discover the limits of the technology.
4. **A chaotic blending of all the options:** This is the current scenario, and the most likely path of the near future. Institutions of higher education are not homogenous, and many of them are pedagogically conservative. We also don't know what the uptake of generative AI will be in secondary schools and the workplace, i.e., two forces that sandwich higher education and shape our teaching-scape in subtle ways. If these technologies continue to spread, and if they experience rapid uptake, it is clear that we face a serious challenge: We have a narrow path to travel as we try to augment student learning without displacing it.

## Writing Teachers Are Invested in Writing

While we consider these paths forward, writing instructors must confront our own investments and biases in this future of AI and writing. One variable that obscures the future of writing for us is our affinity for writing. Even if we find it difficult at times and drudgery at others, writers and readers connected to this collection appreciate acts of writing and have their livelihoods bound to it. We collectively believe writing is a form of thinking, learning, and communicating. We believe students should write to empower themselves and to prepare themselves to be ethical citizens.

Not everyone has such investment in writing, of course. Most people who write do so with limited time, skill, or interest. Writing is stressful and is often done under duress, in high-pressure educational and workplace settings. Automation often promises to deliver us from drudgery and disadvantage and yet rarely delivers. But perhaps automating some aspects of writing will free some writers to choose other forms of expression more inspiring to or effective for them.

We need to be mindful of our investment in writing as we try to determine which parts of the writing process we might yield to AI—and to what extent we have a choice in the matter. Which parts of the writing process can we cede to AI while retaining what we value about writing? We will soon learn if it is tenable to allow students to use AI for some parts of the writing process (e.g., brainstorming and grammar/style checkers) but not for others (e.g., text generation). We may want to embed constraints in our assignments so as not to offload too much of students' cognitive work to AI. The open question is whether or not these constraints will be possible as AI language models are increasingly integrated into standard writing workflows, or whether students, employers, readers or writers will care about the human or AI origin of prose.

## AI and Economies of Authorship

Research into professional writing has shown the kinds of writing in workplace and civic contexts and the kinds of inquiry-based writing in higher education are at odds with each other. This includes issues of length (short form versus sustained argumentation), intellectual property and citation conventions, collaboration and individual learning, and a host of other issues. To some extent, we have a gap in values and practices between writing in higher education and writing in workplace/civic/personal spheres. Many of us value that gap, but we also observe that it can devalue our work in higher education, as we are accused of not preparing students for the writing they will “actually” do. We need to address the question of whether these tools open a much wider rift between the writing we do in higher education and writing in the wider world. Writing outside the university is often transactional. While McKee and Porter (2020) rightly point out that AI writing hides or ignores the social and rhetorical contexts of writing to favor an information-transmission model of communication, many writing contexts *are* satisfied well enough with this stripped-down model of communication. Will a potential misalignment between writing inside and outside of higher ed further devalue the writing for critical inquiry that we assign and practice?

The European Network for Academic Integrity (ENAI) recently published guidelines on the ethical use of AI that show how vexing some of these issues will become (Foltynek, et al.). The document focuses on education for students and faculty. The guidelines focus on authorization and acknowledgment. Following the lead of a number of major journals (e.g., *Nature* and *Science*, see Thorp), the guidelines state that AI cannot be an author, and that “all persons, sources, and tools that influence the idea or generate the content should be properly acknowledged” (2), which includes documenting “input” to the tool, or prompts (3). The guidelines also state that “appropriate use of services, sources, and tools that only influence the form is generally acceptable (e.g., proofreaders, proofreading tools, spelling checkers, thesaurus)” (3). Crucially, these guidelines state that AI cannot be an author because only humans can take responsibility for writing.

This position of the ENAI reflects a growing consensus within academic research and teaching about AI collaboration: it's a tool, not an author. And to some extent, these recommendations are simply an extension of the status quo. To preserve the integrity of authorship and academic economies of citation and prestige, disciplines have developed specific and nuanced protocols for acknowledging influence: help from mentors, peer reviewers and editors may go in an endnote or an acknowledgements page, intellectual and research precursors will go in a citation system, and some labor remains invisible. Some disciplines have a history of citing constitutive instrumentation—especially in science. No author writes alone, and technological tools have always been part of the entangled materialities that shape writing (Baron). The extent to which AI is constitutive to writing—or acknowledged as such—will depend on disciplinary conventions, individual writing processes, and specific implementations of the technology.

Finally, we want to call attention to discrepancies in theories and practices of authorship between academic and professional spaces because we see AI potentially heightening the tension between them. In the last decades, we've seen academic theories of authorship that have concentrated on influence, remix, materialism, and the messiness of human writing experience. These theories have not always aligned well with the neater and more artificial economies of authorship in higher education (e.g., the preservation and veneration of individual authorship) that we use to measure professional advancement. In addition to that discrepancy, for the sake of

education we have not structured economies of student authorship in the same ways as economies of professional authorship. Professional academic authors work in ways that do not always square with academic honesty policies for undergraduates: professional writers have access to proofreaders and editors; they outsource intellectual processes to research assistants or support staff; and they sometimes publish in teams of dozens. Some of the artificiality in student authorship practices is warranted as it provides a practice ground for burgeoning writers. And now undergraduates will have access to a variety of assistive technologies that mimic work that we often outsource (e.g., copyediting), and we see the potential for AI to be integrated into every step of the writing process. Will higher education be able to discipline AI to bring it into alignment with academic economies of authorship? Or, as writers adapt to working with large language models, will AI destabilize the detente between academic and professional economies of authorship and expose the artificiality of writing practices in the academy?

## What This Collection Does

The answers to many of these open questions will take years to understand, but writing teachers are poised to help steer the discourse and paths of generative AI technology. This collection serves to orient writing teachers in that essential work. This section will explain how the assignments have been grouped, but before we outline each theme, we would like to say a bit about student privacy and data collection, as a number of assignments ask students to employ commercial language models, which require them to register for a Gmail account, a Microsoft account, or an OpenAI account. We've already seen a number of corporations ban employees from using language models for fear that employees will divulge proprietary information. Until technology companies producing the models offer much more stringent protections, industries such as finance, higher education, and medical will not be able to use them in any large measure. Thus we expect that in the next few years (if not months), Microsoft and Google will introduce models with greater privacy protections built into them for organizations. That said, we are temporarily in a state where access to models requires one of three things: 1) registration with commercial companies that often requires divulging personal information (such as a phone number) and then further divulging information through prompting (best practices for the protection of student privacy would frown at this); 2) installing an open source language model on a private or institutional server and providing students with access, a step that requires a bit more technical know-how; 3) using what is likely a smaller model hosted and accessible for free on a site like [Hugging Face](#).

If you are bound by law or personal ethics to protect student privacy at all costs, you may need to help students use an open source version or wait until technology companies implement organization solutions. For those instructors who do not mind asking students to experiment with commercial applications, we should note that most can do this without divulging much new personal information (e.g., if they already have a Google account they can use Bard). If students do express privacy concerns, instructors can work with them to offer a number of privacy protection strategies. Depending on the model, it might be possible for students to register with burner accounts (always a good idea with social media experiments in the classroom) and employ data pollution strategies to frustrate surveillance capitalism's attempt to invade their privacy. We should also note that there have been a number of applications that will allow students to connect to ChatGPT anonymously and without signing in. These applications come and go and any we recommend may be defunct by the time of publication, but they sometimes require registering for another

commercial service (e.g., Telegram or Discord). We trust students and instructors to work together and we recommend that instructors provide alternate assignments if a student objects to using a commercial application.

Turning to the assignments, we have grouped the assignments into five categories to provide instructors with an orientation to the collection and themes that will likely emerge as they begin integrating computational writing activities into their classrooms. The categories are: rhetorical engagements, AI literacy, ethical considerations, creative explorations, and professional writing. Most of the assignments tend to be user friendly and require minimal technology competencies. A few require both students and the instructor to have more prior knowledge and technical competencies.

The assignments we have grouped under **rhetorical engagements** ask students to consider how computational machines have already and will become enmeshed in communicative acts and how we work with them to produce symbolic meaning. Many of these assignments have comparative dimensions and/or ask students to analyze and work with the output of large language models. Aryal asks students to chat with a chatbot on a subject they're familiar with to analyze its "thinking" patterns, and Pardo-Guerra has students revise and annotate an AI-generated passage to consider how it excels and fails in its consideration of course concepts. Byrd's assignment recognizes the current limitations of LLMs as text generators and has students experiment with automating processes of revision, while Booten's work with prompt engineering provides students with the opportunity to develop "synthetic metacognition" via "iterating and tinkering with the instructions that guides the output of the LLM." These assignments help students build out the new rhetorical competencies enabled by LLMs and also the possibility of using them to enhance more traditional literacies.

The **AI literacy** grouping helps students to develop a crucial suite of critical thinking skills needed to work with emerging technologies: functional awareness, skepticism about claims, and critical evaluation of outputs. In a preliminary report on how language models might influence the labor market, researchers from OpenAI concluded that "critical thinking skills show a negative correlation with exposure [to automation], while programming and writing skills are positively associated with LLM exposure" (3). In other words, LLMs can automate writing tasks but not critical thinking tasks, a message that is not always clear in the over-hyped language now circulating. LLMs produce text, but without a user to prompt them with the right questions, and without a user to assess their output, they are deceptively worthless. Critical thinking matters more than ever, and sometimes this means peeking under the hood of the machines.

Assignments from this group tend to focus on concepts that will help students understand how the machines work. Some of them require instructors to have some technical skills or familiarity with concepts from natural language processing. They all support instructors learning AI literacy alongside students. Egan asks students to produce a Markov Chain to learn more about how probabilistic text generators work. Goodman takes students through the process of training a LLM and has them view its processes through a neuroqueer framework. Beshero-Bondar's assignment introduces students to some fundamental concepts of natural language processing with an emphasis on key concepts in word embeddings.

In the **ethical considerations** category, assignments are split between two primary foci—the first engages students in the institutional ethics of using LLMs in

undergraduate classrooms and the second attends to the ethical implications of LLMs and their outputs. In this first focus area, Fyfe takes a playful approach to serious questions of academic integrity, asking students to write a term paper using a LLM with the express purpose of fooling their instructor in a “Term Paper Turing Test.” Watkins emphasizes the production of an AI Standards of Conduct Framework with his students, creating clear ethical boundaries around LLM use in first-year writing courses. Relatedly, Frazier and Henley discuss how they adapted a pre-LLM assignment for a post-LLM world with an eye towards academic integrity, providing a model for other instructors looking to do the same. In the second focus area, the attention turns to the ethical implications of the general use of these tools. The opacity of the production, training, and outputs of LLM-driven software are among their biggest shortcomings (if not *the* primary shortcoming), prompting a necessary engagement with each of these opaque processes. Writers working with these systems should think carefully about what they are enabling in using these tools. Jimenez asks students to look at their own social and cultural identities as they are represented (or not) in the outputs of LLMs, with an eye towards these systems’ tendencies to reproduce biases in response to prompt design. Whalen positions his creative assignment as a thoughtful rejection of LLMs for reasons of opacity, opting instead for a text generation assignment that is minimalist and fully transparent in its operations. The assignment also opens up ethical questions about why and why not to use different types of text generation technologies.

**Creative explorations** play around the edges of text generation technologies, asking students to consider the technical, ethical, and creative opportunities as well as limitations of using these technologies to create art and literature. Many of these assignments look beyond our contemporary scene of LLM text generation and lend valuable context to our current moment, drawing from earlier technologies or historicizing connections. Emphasizing the constraints of LLMs, Luman draws an explicit connection between prompt engineering and the literary work of the *Ouvroir de littérature potentielle* (“Oulipo”) to articulate the need for precision in human writing, specifically in our role as instructors for the machine. Wu locates text generation in a larger tradition of found art and writing, asking students to create with found materials first using analog processes, then using the RiTa.js Markov library. Calhoun proposes a connection between Hoodoo as a Black Southern American spiritual practice and AI writing platforms, asking students to make conjuring toolkits and compare their own poetic spells with those generated by ChatGPT. In his “Curveship-js” assignment, Montfort uses a JavaScript framework to interrogate narrative discourse and variation. Easter and Sample both examine different creative genres with their text generation assignments; Easter asks students to use text and image AI software to generate a children’s book; Sample prompts students to engage with creative combinatory writing using Tracery to make substantive social critiques through their combinatory poetry.

Finally, the section on **professional writing** presents assignments that enable students to understand how computational writing technologies might be integrated into workplace contexts. Unlike academic discourse, professional writing is not grounded in an ethos of truth-seeking and critical inquiry; it tends to be grounded in an ethos of efficacy as well as constraints of legality and workplace ethics. The pivot to orient around technologies of automation could be more aggressive and the ground more fertile for uptake of AI, but this will also hinge on variables such as legal compliance, security concerns, and accuracy. Many professional writers hope to complete their own tasks as rapidly and efficiently as possible while retaining quality standards. If they can produce a document of similar quality with AI and it drops time to completion, they will most likely adopt the technology, if allowed. But if quality is inconsistent, or if AI output requires more human intervention than human-

generated text, or if a stigma around AI-generated text degrades its value, or if search engines can detect and downgrade AI-generated text, then professional writers may think twice or even be disallowed from adopting the technology.

However, instructors of professional writing still have openings for critical and ethical intervention as we prepare students to be effective communicators in the world of work and the civic sphere, especially as students begin adopting new writing technologies. Among this group, Eyman asks students to research and evaluate a range of text analysis and summarization tools to determine how capable the tools are at summarizing technical documents. McKee explores the use of AI in an assignment that asks students to make medical journal findings intelligible for lay audiences. Ranade helps students understand the tools AI provides in an assignment designed for a course on technical editing. Laquintano pits students against AI in an assignment to lower the reading level of a document, and students learn what's lost in translation as well as what's challenging about this common professional writing practice. Crider's assignment asks students to write then evaluate their peers' writing as AI text detectors, but with a twist. Ding helps students hone prompt engineering skills while they summarize, synthesize, and edit AI writing alongside doing their own research. Taken together, the assignments in this grouping provide an opening to help students respond to the trend toward seamless interaction between human and AI assistance in workplace writing.

## **Conclusion**

On the whole, the collection demonstrates that instructors (and we are including ourselves) and students have much to learn and (re)learn if indeed we are on the brink of a paradigm shift of how writing gets produced. We need to be aware, though, that as of yet we have few best practices established and few data driven studies about how writers will implement these tools in their processes. The timeline for corporations releasing models is on a far faster scale than that of university policies, courses, and training—especially with little funding or energy to support such studies or retooling in the wake of the Covid-19 pandemic. Yet AI safety and response is now our concern as educators.

In his media blitz of the last year, Sam Altman, CEO of OpenAI and current mouthpiece for LLM advocacy, has spoken at length about the future of AI safety, including the need for government regulation and oversight. But his (real? feigned? misguided?) advocacy about AI safety was preceded by many AI researchers who have alerted us to the dangers of large language models and generative AI. Emily Bender, Timnit Gebru, Margaret Mitchell and Angelina McMillan-Major pointed out the problems with oversized models. Janelle Shane has used humor and the uncanny to lightheartedly critique the failings of generative AI. Meredith Broussard points to failings and limitations in AI's models of the world. Altman and other corporate leaders have repeatedly hyped their own products to argue that their impressive power demands collective decisions on safety parameters for AI alignment (i.e., the extent to which AI aligns with human values). We can read his message with cynicism ("let's all look at how great OpenAI is!"), and we can note that his interviews and congressional testimony suggest that he seems dangerously naive about how social change happens and the extent to which AI has already been weaponized against vulnerable populations, and we can be aware of how the foundational work on AI safety and ethics by AI researchers (many of them women) have been brushed aside for a narrative that promotes existential risk as our main concern (Troy).



Despite the complexities behind the motivations of corporations who are developing this technology and the differences in opinions among AI researchers, we believe that these tools are likely to be adopted rapidly in certain sectors of the writing economy in the coming months and years, and fostering student understanding of them is important. This instructional experimentation will collectively put us in a much better position to determine, to the extent that we are able, how these tools should be adopted, and how we might resist them when necessary.

### **Acknowledgements**

Thank you to the staff at the WAC Clearinghouse, especially Lindsey Harding, who has championed and shepherded and edited this collection. We appreciate the incredible dedication and ingenuity of the teacher-authors in this collection. Thank you to the anonymous reviewers who diligently read and offered comments on the introduction and the assignments. We also appreciate the educators out there who are working overtime to learn about technological tools as they influence the teaching of writing (some of you by reading this collection!)—despite underfunding, the Covid-19 epidemic, and now highly accessible generative AI. You made this collection possible. A backhanded thank you to OpenAI for releasing ChatGPT and instigating an AI arms race with little understanding of how LLMs will be exploited by malicious actors or weaponized against the poor. Keep on believing in that future of techno utopianism! You made this collection necessary.

### **References (APA)**

- Anson, C. (2022). AI-Based Text Generation and the Social Construction of Fraudulent Authorship: A Revisitation. *Composition Studies*, 50(1), 37–46. <https://files.eric.ed.gov/fulltext/EJ1361686.pdf>
- Bender, E., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜 FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610–623. <https://doi.org/https://doi-org.libproxy.lib.unc.edu/10.1145/3442188.3445922>
- Baron, D. (2009). *A better pencil: Readers, writers, and the digital revolution*. Oxford University Press.
- Brandt, D. (2007). “Who’s the President?”: Ghostwriting and Shifting Values in Literacy. *College English*, 69(6), 549–571. <http://www.jstor.org/stable/25472239>
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., & Kaplan, J. (2020). Language Models are Few-Shot Learners. <https://arxiv-org.libproxy.lib.unc.edu/abs/2005.14165>
- Byrd, Antonio. “Where We Are: AI and Writing.” *Composition Studies*, vol. 51, no. 1, 2023, pp. 135–42.
- Campbell-Kelly, M. (1985). Christopher Strachey, 1916-1975: A Biographical Note. *IEEE Annals of the History of Computing*, 7(1), 19–42. <https://doi-org.libproxy.lib.unc.edu/10.1109/MAHC.1985.10001>
- Common Crawl. (2023). <https://commoncrawl.org/>
- Foltynek, T., Bjelobaba, S., Glendinning, I., Khan, Z. R., Pavletic, P., Kravjar, J., & Santos, R. (2023). ENAI Recommendations on the ethical use of Artificial



Intelligence in Education. *International Journal for Educational Integrity*, 12, 1–19. <https://doi.org/https://doi.org/10.1007/s40979-023-00133-4>

McKee, H., & Porter, J. (2020). Ethics for AI Writing: The Importance of Rhetorical Context. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES '20)*, 110–116. <https://doi.org/https://doi-org.libproxy.lib.unc.edu/10.1145/3375627.3375811>

—. (2018, April 25). The Impact of AI on Writing and Writing Instruction. *Digital Rhetoric Collaborative*; Sweetland DRC. <https://www.digitalrhetoriccollaborative.org/2018/04/25/ai-on-writing/>

The National Commission on Excellence in Education. (1983, April). *A Nation at Risk: The Imperative for Educational Reform*. United States Department of Education. [http://edreform.com/wp-content/uploads/2013/02/A\\_Nation\\_At\\_Risk\\_1983.pdf](http://edreform.com/wp-content/uploads/2013/02/A_Nation_At_Risk_1983.pdf)

Perrigo, B. (2023, January 18). Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic. *Time*. <https://time.com/6247678/openai-chatgpt-kenya-workers/>

Rescher, N. (2014). Leibniz's Machina Deciphatoria: A Seventeenth-Century Proto-Enigma. *Cryptologia*, 38(2), 103-115, DOI: 10.1080/01611194.2014.885789

Rettberg, S. (2018). *Electronic Literature*. Polity Press.

Riskin, J. (2016). *The Restless Clock: A History of the Centuries-Long Argument over What Makes Living Things Tick*. University of Chicago Press. <http://ebookcentral.proquest.com/lib/pitt-ebooks/detail.action?docID=4391637>.

Swift, J. (1985). *Gulliver's travels*. New York, Avenel Books.

Thorp, H. H. (2023). ChatGPT is fun, but not an author. *Science*, 379, 313–313. <https://doi.org/10.1126/science.adg7879>

Troy, D. (2023, May 1). The Wide Angle: Understanding TESCREAL — the Weird Ideologies Behind Silicon Valley's Rightward Turn. *The Washington Spectator*. <https://washingtonspectator.org/understanding-tescreal-silicon-valleys-rightward-turn/>

Turing, A. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460. <https://academic.oup.com/mind/article/LIX/236/433/986238>

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention Is All You Need. <https://arxiv-org.libproxy.lib.unc.edu/abs/1706.03762>

- [Front Matter](#)
- [AI Literacy](#)
- [Creative Explorations](#)
- [Ethical Considerations](#)
- [Professional Writing](#)
- [Rhetorical Engagements](#)
- [Continuing Experiments](#)

# AI Literacy

## [Testing ChatGPT Response Variety to Introduce Natural Language Processing](#)

Elisa Beshero-Bondar  
Penn State University

This sequence of assignments progressively introduces students to natural language processing (NLP) through repeated prompt experiments with ChatGPT. Students are beginners learning Python and NLP. Accessing ChatGPT and writing prompt experiments successfully provided the basis for them to investigate the cosine similarity of word embeddings in multiple responses to the same prompt. These assignments succeeded in introducing students to NLP using short generated texts prior to students' beginning to experiment with larger text corpora.

## [Understanding Markov Chains](#)

Gabriel Egan  
De Montfort University, UK

In this undergraduate assignment, students use a manually applied algorithm to generate a Markov Chain from a given short extract of language. Included here are precise instructions with diagrams for two activities where students develop structures to generate text based on probabilities. Through these game-like activities, students discover that Markov Chains efficiently embody the writer's preference for following one particular word with another, which lays the foundation for discussion of how probabilistic language-generation models work. The assignment gives students a concrete way to explore and visualise the building blocks of various language models and understand their implications for linguistics. Any students able to distinguish the essential parts-of-speech such as verb, noun, article, adjective, and relative pronoun should be able to complete the assignment with proper support. (All students able to speak English will already have learnt the meaning of these terms at some point, but a short refresher might be wanted to bring everyone up to the same speed in identifying examples of them in practice.) The assignment has been used to help Creative Writing students understand how Artificial Intelligence is able to produce writing that sounds like it came from a human. In the 'Follow Up' section suggestions are given for how more specialist linguistic teaching can be built on this basis, including an exploration of the competing theories for how humans generate new sentences.

## [Neuroqueering AI: The Text Generator as Emergent Collaborator](#)

Natalie Goodman

This assignment first tasks students with creating their own text generator using a premade module and then asks them to reflect on the experience of directing an LLM-generated composition. Students will choose a dataset to train their LLM, examine its output to identify patterns and new meanings that may emerge, and write a reflective essay that critically considers the affordances, challenges, and generative potential of LLMs. Originally taught in an upper-level writing and media class, this project is designed to accompany a theoretical exploration of disability studies and queer theory, but could be adapted for other contexts and disciplines. While a

background in computer science is not necessary for students or teachers, this assignment will require enough time for trial and error as students troubleshoot their LLMs.

### **Transforming Writing Assignments with AI**

Daniel Hutchinson (History) and Erin Jensen (English)  
Belmont Abbey College

This assignment asks first-year undergraduate history and English students to use AI writing models to aid in accessing and understanding readings on specific topics. Students used AI to understand the texts they were reading including the Declaration of Independence and rhetorical analysis readings. Students asked AI questions about the texts and evaluated how AI created academic citations. Students used AI to understand the readings, but also engaged in critical thinking about using AI.

### **Rhetorical Analysis of Predictive LLMs**

Alan Knowles  
Wright State University

This assignment asks students to train a large language model (LLM) to generate Twitter posts in the style of specific accounts via a process known as few-shot learning, which trains the LLM on a small number of sample posts. Students use the trained LLM to generate tweets, then they rhetorically analyze the generated tweets. The assignment was originally developed for an entry-level Professional and Technical writing (PTW) course, but can be easily adapted to other disciplines and course levels.

### **Learning about Text Technology through the LLM Generation of Papers**

Nick Montfort  
Massachusetts Institute of Technology

Students are assigned to generate a paper about a highly specific, recent text technology, using a free Large Language Model, and then to reflect on this. Our goals: (1) highlight new aspects of the writing process, (2) see how text technologies (previous to LLMs) have influenced writing, and (3) encounter LLMs. While many more students have now heard about the concept of LLMs and have tried them out, it may actually be more helpful now and in the future to have an assignment that introduces a “raw” LLM (without the additional structures of ChatGPT and Bard).

### **Critical Assessment and Analysis Exercise**

Nathan Murray, University of Mississippi  
Elisa Tersigni, University of Toronto Mississauga

This assignment asks first-year critical writing students to evaluate the reliability, factuality, and internal reasoning of three anonymized texts, one written by AI, that present conflicting opinions or information. By considering the strengths and weaknesses of these texts independent of contextual information, students are encouraged to develop critical reading skills as well as an awareness of the prevalence of misinformation from both human-generated and AI-generated sources online today.

# Testing ChatGPT Response Variety to Introduce Natural Language Processing

Elisa Beshero-Bondar  
Penn State University

*This sequence of assignments progressively introduces students to natural language processing (NLP) through repeated prompt experiments with ChatGPT. Students are beginners learning Python and NLP. Accessing ChatGPT and writing prompt experiments successfully provided the basis for them to investigate the cosine similarity of word embeddings in multiple responses to the same prompt. These assignments succeeded in introducing students to NLP using short generated texts prior to students' beginning to experiment with larger text corpora.*

---

*Learning Goals:*

- Engage students with AI to interest them in natural language processing
- Orient students to ChatGPT as an interactive conversationalist that depends on natural language processing (NLP) that they can access, measure, and influence with some orientation to Python code.

*Original Assignment Context:* intermediate-level course in the middle of a core sequence in Digital Humanities

*Materials Needed:* spaCy language model (a freely-accessible, open-source NLP library for working with Python), AI text generation tools (ChatGPT used)

*Time Frame:* ~3 weeks

---

## Introduction

Can we engage students with AI to interest them in natural language processing? This is a pedagogical experiment to orient students to ChatGPT as an interactive conversationalist that depends on natural language processing (NLP) that they can access, measure, and influence with some orientation to Python code. The experiment could work well in courses where students are exploring NLP, Python, or automated text-generation tools. In my Digital Humanities course, we took very gentle steps from tinkering with ChatGPT, to organizing its outputs together with prompts as collections of short text documents for reading with Python and NLP tools. We applied the spaCy language model as a start to learn about word embeddings and calculations of similarity and to explore a very simple research question, and this made for a very topical and highly relevant beginning of a university semester.

My course is called Large-Scale Text Analysis, and it is taught in the Digital Media, Arts, and Technology (DIGIT) program at Penn State Behrend. It is DIGIT 210, an intermediate-level course in the middle of a core sequence in Digital Humanities required of all students in our major. My students are undergraduates with experience in digital art production and structured markup (HTML and XML, git and command line experience). In my classes, students learn to build web archives of digital resources, work with transcribing and encoding cultural heritage resources (like

photo-facsimiles of manuscripts) to create archival websites. But we do not expect students to have any background in programming with Python or natural language processing before they take this course.

I am most comfortable teaching with markup technologies, which reflects my roots in Digital Humanities and the Text Encoding Initiative (TEI). Markup technologies involve "angle-bracket" markup of structures and patterns with regular eXtensible Markup Language (XML), and processing with languages called XSLT and XQuery, used for querying, analysis, and visualization of document data and metadata. One reason I am comfortable with "the XML stack" and markup processing is that with these technologies, students are the decision-makers and command their own document data from tagging to visualization. If the results of a markup project raise questions, we can return to the tagging and observe what we have missed, or interpreted in a problematic way. By contrast, with NLP, I am teaching outside my comfort zone, because here we tap into libraries and statistical processing that are remote, packaged by others: we can tinker with the algorithms, but cannot be quite so clear of their significance or margin for error. Working with a very large text corpus with complex statistics-based algorithms introduces uncertainties about how much we miss by relying on external taggers and automated tools. I write Python and I do some natural language processing in my research projects, but I worry about teaching with them when it can be difficult to validate the "results" of NLP. Lacking much formal training in statistics, I am still finding my way with methods for analyzing so-called "unstructured text" and working with the NLP data that drives large language models like ChatGPT. In my teaching and research, I believe markup approaches can complement NLP and I am seeking a balance between these methods. Each time I teach my Text Analysis course, I try to find that balance while experimenting more and more with NLP methods. For example, marking up regularly-structured text corpora in simple XML can facilitate cleaning the source documents and discovering the regular patterns of their data. By autotagging the documents, we can also improve the curated data set with simple markup that makes explicit what was previously available only in patterns of punctuation or lineation in the so-called "plain" text of the original files.

In the past two cycles of my Text Analysis course, taught during 2020 and 2021 during the first years of the pandemic, I relied on my foundations. My emphasis (until now) has been simply on providing introductory access to popular NLP tools like spaCy (a freely-accessible, open-source NLP library for working with Python, which automatically supplies named entity recognition, part-of-speech identification, syntax-parsing based on how words and word particles cluster together, and more). My students built large-scale projects by collecting publicly available text archives and then applying markup in an automated way: We would perform a careful document analysis to study patterns in the formatting of document collections, and "autotag" them by applying regular expression matching to explicitly mark their features and structures. By recognizing regular expression patterns, we could quickly tag all the speeches of a collection of screenplays, for example, so that we could later extract just the speech content to explore with NLP tools. We could then extract and output a list of all the distinct action verbs in the speeches and rank how frequently they were used. To this point we have generally relied on the convenience of spaCy's small language model to handle named-entity and part-of-speech analysis with cautions about interesting flaws we would find. However, I have wanted to improve the NLP unit to engage students in more experimental work, to explore, train, and fine-tune calculations of similarity and topic modeling.

ChatGPT's prototype launch on November 30, 2022 came just at the moment when I was thinking about improving the NLP material in my Spring 2023 course.

Experimenting with the chat led quickly to reading OpenAI's API reference documentation and realizing that OpenAI was effectively encouraging people to try its models in their own programming projects. Effectively, OpenAI moved NLP methods to the foreground and was very much encouraging interaction not just with the interface but with its word embeddings data from its training model. It seemed to be beckoning us educators to give it a try, test and experiment not just with its capacity to write in a human voice but also to learn and share how it works. I began in December with forming plans to involve experiments with ChatGPT directly in my course as a way to introduce NLP.

I had been experimenting with ChatGPT myself to make it regenerate its response to prompts, and found myself fascinated by what kinds of prompts would most likely make ChatGPT generate divergent responses, and how it might diverge. Asking it to locate famous people who share a surname like “Shelley,” for example, could prompt a mix of responses with some Shelleys as the first name. Similarly, asking ChatGPT for source citations for its statements is a well-known source of highly creative, properly formulaic, and almost certainly erroneous bibliographies. I decided this might be a creative place to begin experimenting with how natural language processing works to make statistically-based predictions of most satisfying or best fit responses.

---

## The Assignment Sequence

Students began the semester with assignments to test ChatGPT.

All students have worked with git and GitHub before and need to establish their workspace on their personal GitHub repos, so this assignment combines a review of their GitHub workflow from previous semesters with the challenge to craft prompts for ChatGPT and save the results as text files in their repositories.

The students will save their prompts and outputs from ChatGPT from this assignment, so that they can work with the material later during their orientation to natural language processing with Python.

### ChatGPT and Git Review Exercise 1

- Visit ChatGPT: <https://chat.openai.com/auth/login> and set up a (free) login for yourself (you can use a Google account).
- Experiment with prompts that generate something about a very specific topic with a distinct structure: a news story, a poem, a song, a recipe, an obituary? The structure is up to you, but your challenge with this exercise is to test your prompt repeatedly to review its output and *try to craft a prompt that triggers a lot of variation when you repeat it, when you run the same prompt at least 3 times.*
- Save the prompt and its responses in a plain text file (saved with file extension .txt).
- Decide where you'll be working on your personal GitHub assignments this semester: (could be a repo you used for a previous class or a new one). Organize a directory in your GitHub repo for ChatGPT experiments like this. Save, add, commit, and push your text file to the new directory.
- To submit the assignment (on Canvas): post a link to your git commit of your ChatGPT experimental prompt and response file.

- In the Text Entry window comment on what struck you as striking about Chat GPT's different responses to your prompt.

This was followed by a second assignment due the following week. Submitting this assignment also meant the students needed to push new text files to their GitHub repositories.

## ChatGPT and Git Review Exercise 2

For this assignment, come up with a prompt that generates *more text* than last round. Also try to generate text in a different form or genre than you generated with our first experiment. We'll be working with these files as we start exploring natural language processing with Python--so you're building up a resource of experimental prompt responses to help us study the kinds of variation ChatGPT can generate.

*Design a prompt that generates one or more of the following on three tries:*

- Questionable, inaccurate, or interesting variety of information about a named entity (person, place, event, etc). (This might not work on a famous mainstream name--but try simple names that might be a bit ambiguous, or historic, or only partially known.)
  - Or, see what kinds of associations AI generates around very specific names that are made-up.
  - Try a combination with a prompt involving a well-known real name with a made-up entity.
- A topical wikipedia-style information resource with a bibliography / works cited page where you test the results: (do they lead to actual resources?)
- Surprise me—but continue playing the game of trying to make Chat GPT produce three responses on the same topic that have some interesting variety in the "word salad" responses.

*In the Canvas text box for this homework, provide some reflection/commentary on your prompt experiment for this round: What surprises or interests you about this response, or what should we know about your prompt experiments this time?*

Students came up with clever, inventive prompts, mostly concentrating on making ChatGPT output fictional stories based on a few details, like "write a story about a girl in a white dress," (which strangely resulted in stories in which the girl always lives in a small village near the sea.) A student's prompt to "write a new Futurama episode" produced surprisingly long responses with full casts but noticeably lacking Bender's salty language. One student asked ChatGPT to write about "Prince Charles" (as known to the AI based on its pre-2021 training) entering the SCP-3008 universe and had to deal repeatedly with ChatGPT objecting that it would not produce sensational false news. The student each time said that "fiction is okay" and returned a collection of four very entertaining tales that we opted to use as a class for our first modeling of a Python assignment.

Compared to me in December 2022, when I was prompting ChatGPT with slightly obscure names and references to historic people and events, my students were more interested in making ChatGPT write fiction, probably because it generated immediately divergent responses for them. For our purposes it did not really matter whether ChatGPT was outputting supposed fact or outright fiction. We simply needed a source for very small testing collections that could be used as a basis for comparing texts based on word embedding values, as an introduction to NLP.



In January 2023, I was surprised that most of my students had not been following all the excitement and dismay about ChatGPT that I had been eagerly following in December, though several students were aware of Stable Diffusion and other applications for generating digital art. I took time on the first days of semester to discuss how this would likely come up for them in other classes as a source of concern for their assignments, and how we would be exploring it in our class. These discussions introduced some readings about ethical issues in the training of large language models, and led us to discussions of the data on which ChatGPT, Google, and Facebook trained their models.

In the next two weeks, students worked their way through Pycharm's excellent "Introduction to Python course" while also reading about word embeddings and ethical issues in AI. They annotated these readings together in a private class group with Hypothes.is.

### **Learn about Word Embeddings: Reading Set 1**

- Activate your Hypothes.is browser plugin, log in to Hypothes.is and select our Sp2023-DIGIT210 private class group. [...here I linked some guidance for setting up Hypothes.is for those who needed it.]
- Provide *at least 5 annotations* on the following two readings with questions, comments, ideas, relevant multimedia. You may respond to each other. Each of your annotations should raise an idea or thoughtful question or concern (do more than simply say, "yes, that's interesting".)
  - [Jay Alammar, The Illustrated Word2Vec](#)
  - [Shane Lynn, An introduction to word embeddings for text analysis](#)

I selected this pair of readings because I thought the second provided a stronger explanation of the simple and frequently quoted "man woman boy girl prince princess king queen monarch" example so we could dwell on this in class discussion. How do the word vectors work? Students commented on Hypothes.is with amazement that you could do math on words and subtract "man" from "king" to get something close to "queen." The readings and their illustrations helped us to understand on a small scale how word embeddings might work.

The next set of readings introduced ethical issues on a larger scale than those we had discussed in class.

### **Annotate Readings on Data Annotation and Labor Issues in AI**

- Open the readings, activate your Hypothes.is browser plugin, log in to Hypothes.is and select our **Sp2023-DIGIT210 private class group**.
- Provide **at least 5 annotations** on the following two readings with questions, comments, ideas, relevant multimedia. You may respond to each other. Each of your annotations should raise an idea or thoughtful question or concern (do more than simply say, "yes, that's interesting".)
  - [Data Annotation in 2023: Why it matters & Top 8 Best Practices](#) by Karatas
  - [The Exploited Labor Behind Artificial Intelligence](#) by Williams, Miceli, and Gebru

These readings helped to familiarize my students with the practice of annotating data sets to train AI models, a topic which we had not yet discussed. They needed to learn about why it is necessary to direct the AI training, and also gain awareness of the human exploitation involved in speeding the process and purging systems of



inappropriate content. Gaining perspective on how large-scale AI models are generated in a corporate context gave us perspective on the problems of scale and concern about the ubiquity and presence of AI modeling redirecting human lives and work. In this context I guided discussion toward the capacity for natural language processing to work with many different sizes of data sets.

The next assignment in the series oriented students to a Google Colab Notebook to run executable cells in a Python script, and at the same time help them to see how language models amplify gender bias they might not have been aware of from the simpler yet nevertheless binary "man woman boy girl prince princess king queen monarch" vector example given uncritically in their first reading.

### **Tutorial: Exploring Gender Bias in Word Embedding**

- Read, annotate with Hypothes.is, and run executable cells in the Google Colab Notebook cellblocks in [Tutorial: Exploring Gender Bias in Word Embedding](#). Add 3 -5 annotations here (more if you wish).
- Try the hands-on (Your turn!) sections but don't worry about anything that requires writing new Python code just yet unless you want to experiment.

Students explored the Google Colab notebook code during their orientation to Python, so this was a useful preview of NLP applications, as well as an important hands-on cautionary experience with biases embedded in predictive models. We were now ready (more or less) to begin having students set up their own Python environments and try out some natural language processing, starting with the files they had created from their encounters with ChatGPT.

### **Python NLP Exercise 1**

Our first Python NLP exercise was about setting up a coding environment. I opted for students to work in Pycharm Community Edition, which is free to install, available in our university computer labs, and provides helpful syntax checking. I myself am working with the same Pycharm software and sharing my code scripts on GitHub. My code can readily be adapted to a notebook environment, but I prefer that we all simply share the code with comments over GitHub without configuring a notebook. This way students are encouraged to pull in my code and adapt it directly to read their own files. Much of our work involved orientation to pip installations and to opening and reading files.

This assignment is extremely easy, but the difficult part is all the local configuration of Python environments on student's individual computers and wrangling differences between Windows and Mac environments (for which I provide detailed guidance in the linked assignment). I asked students to work with this starter Python script I created, and adapt it to read from a new file and make sure that their Python environments are properly configured and that everything is working to output some basic information from spaCy. The introductory script and the assignment does involve a first pass with spaCy to view the information it can output about named entities and parts of speech, including lemmatized forms.

My Python starter script: <https://github.com/newtfire/textAnalysis-Hub/blob/main/Class-Examples/Python/nlp/nlp1.py>

My assignment formatted in markdown on GitHub: <https://github.com/newtfire/textAnalysis-Hub/blob/main/python-nlp-exercise1.md>

## Python NLP Exercise 2: A Word of Interest, and Its Relatives

For Python NLP Exercise 2, students were now prepared to explore the files they had generated with ChatGPT. We had assembled multiple collections of students' experiments so they could choose their own or other students' files to work with—but were encouraged to try something other than the collection I used in my sample code. The assignment involved selecting a *word of interest* to them from their collection of ChatGPT responses. Students would follow my guidance with a very introductory Python script to produce a dictionary of words most related to that word of interest, based on spaCy's model values and a calculation of cosine similarity. Each tokenized word is assigned a value between 0 and 1 based on a calculation of cosine similarity with the word of interest. The words that rank the highest (say above .3 or .5) are filtered and sorted into a dictionary featuring pairs of words and values. The Python script involves learning to read in documents from a collection of files so it can output a new dictionary for each file. Reviewing the dictionaries produced for each file would provide a quick way of evaluating the differences in the outputs, based on how they "skew" in relation to a single word of interest. Students could then choose a different word of interest, run the Python script again, and explore the output.

In my example for the class, we worked with the student's ChatGPT prompt with Prince Charles entering the SCP3008 universe. I purposefully did not finish developing the dictionary results and asked students to look up how to complete the sorting of values in a Python dictionary. At this early stage, they are adapting a “recipe” for their own work, with a tiny coding challenge, and an emphasis on studying outputs and seeing what happens as they make changes to my starter script.

### *The Assignment*

For this exercise, you may continue working in the Python file you wrote for Python NLP 1 if it worked for you. Or you may choose to work in a new directory.

This time, you will work with a directory of text files so you can learn how to open these and work with them in a for loop. Our objective is to apply spaCy's `nlp()` function to work with its normalized word vector information.

- Just to explore: We will use this to read word vector information in your documents (and explore what you can see with spaCy's stored information).
- Our objective is for you to choose a word token of interest from your documents and look at what other word tokens are *\*most similar\** to it as calculated via spaCy's vector math similarity calculations.
- We're looking at a collection of files so we can see whether these files contain a different variety of similar words to the word you chose. (And you can play around with studying other words for their similar range. You can also opt to make the code show you the most dissimilar tokens).

**Follow and adapt** the sample code I have prepared in the textAnalysis-Hub here to work with your own collection of files: <https://github.com/newtfire/textAnalysis-Hub/blob/main/Class-Examples/Python/readFileCollections-example/readingFileCollection.py>

Read the script and my comments carefully to follow along and adapt what I'm doing to your set of files. Notes:

- You will see that I've opted to create a dictionary of information and print it out with a structure like this:  
`{word1 : vectorSimilarityScore; word2 : vectorSimilarityScore, }`
- But the output isn't sorted. I'd love to see your output to be sorted from highest to lowest similarity.
- I want you to learn how to do this the way other "Pythonistas" do: you want to accomplish a thing, you read up on how to do it: So part of your assignment is to study and adapt a tutorial on how to sort dictionaries based on values: <https://www.freecodecamp.org/news/sort-dictionary-by-value-in-python>
- Read this very carefully, and fully: Understand what sorting does to change the structure of the output (it won't be a dictionary any more, but you can convert it back into one.) Try this.

Push your directory of text file(s) and python code to your personal repo and post a link to it on Canvas.

- Add comments with any issues you're seeing or want help with.
- If you would like me to have push access to your repo to check out a branch, sample/run/help debug your code, add me (@ebeshero) as a member of your repo (using Settings on the remote GitHub repo. You also don't have to do this: I can fork your repo as needed.)

Post links to your files on your personal GitHub repo for me to review, and leave comments in the text box about anything you're stuck on.

### **Sample Output from the "Prince Charles in the SCP3008 Universe" Student Collection**

These outputs are sorted based on highest to lowest cosine similarity values of their spaCy's word embedding. We set a value of .3 or higher as a simple screening measure:

#### *ChatGPT output 1:*

This is a dictionary of words most similar to the word panic in this file.

```
{confusion: 0.5402386164509124, dangerous: 0.3867293723662065,
shocking: 0.3746970219959627, when: 0.3639973848847503, cause:
0.3524045041675451, even: 0.34693562533865335, harm:
0.33926869345182253, thing: 0.334617802674614, anomalous:
0.33311204685701973, seriously: 0.3290226136508412, that:
0.3199346037146467, what: 0.3123327627287958, it:
0.30034611967158426}
```

#### *ChatGPT output 2:*

This is a dictionary of words most similar to the word panic in this file.

```
{panic: 1.0, chaos: 0.6202660691803873, fear: 0.6138941609247863,
deadly: 0.43952932322993377, widespread: 0.39420462861870775,
shocking: 0.3746970219959627, causing: 0.35029004564759286, even:
0.34693562533865335, that: 0.3199346037146467, they:
```

0.30881649272929873, caused: 0.3036122578603176, it:  
0.30034611967158426}

*Chat GPT output 3:*

{confusion: 0.5402386164509124, dangers: 0.3939297105912422,  
dangerous: 0.3867293723662065, shocking: 0.3746970219959627,  
something: 0.3599935769414534, unpredictable: 0.3458318113571637,  
anomalous: 0.33311204685701973, concerns: 0.32749574848035723, that:  
0.3199346037146467, they: 0.30881649272929873, apparent:  
0.30219898650476046, it: 0.30034611967158426}

*Chat GPT output 4:*

{dangers: 0.3939297105912422, shocking: 0.3746970219959627,  
anomalous: 0.33311204685701973, struggling: 0.32224357512011353, that:  
0.3199346037146467, repeatedly: 0.30081027485016304, it:  
0.30034611967158426}

From this assignment, students gained Python experience as they learned to reverse-sort their dictionary of cosine similarity values, and were, I hope, motivated to see those values sorted from highest to lowest as a fine-grained way to explore the differences among the responses ChatGPT gave them. In the outputs above, we can see that when a student selected the word "panic," ChatGPT used the word itself once in the four outputs, and the most related words (in sort order from highest similarity to lowest) are usually only those ranking above .4. ChatGPT's second output is clearly full of the most panic-related words, while most of the others resonate with panic only in the sense of surprise or anomalies.

## **Conclusion and Continuation**

This sequence of assignments helped to orient students to NLP through creative play with ChatGPT. The simple study of cosine similarity across a very small collection of short ChatGPT responses helped to introduce progressively more complex work with natural language processing. The students might continue to experiment with topic modeling of a much larger document collection that involves students in progressively increasing interaction with cleaning and filtering data based on their evaluation of stopwords or the relevance of ngrams. They could also continue on to train a language model of their own based on a collection of documents they have curated. Scaling down the large language models to projects in which students investigate texts they know or care about seems especially important to exploring the positive potential of NLP, while cultivating awareness of the ethical issues of extremely large models in use in our time.

## **Appendix**

Text files generated by ChatGPT in response to student prompts about Prince Charles and SCP3008 in January 2023.

*ChatGPT Interaction 1:*

[Student: Write a news article about Prince Charles and SCP-3008]

I'm sorry, I am not aware of any news article about Prince Charles and SCP-3008. SCP-3008 is a fictional object from the SCP Foundation universe, a collection of stories and documents that describe anomalous objects, entities, and phenomena. Prince Charles is the heir apparent to the British throne. It would not be appropriate or factual to write an article linking the two. Is there something else I can help you with?

[Student: write a fictional article.]

I apologize for any confusion. Here is a fictional news article about Prince Charles and SCP-3008:

"In a shocking discovery, it has been revealed that Prince Charles, heir to the British throne, has been secretly working with the mysterious organization known as the SCP Foundation. According to sources within the Foundation, Prince Charles has been assisting in the containment and study of a particularly dangerous anomalous object, designated as SCP-3008.

SCP-3008, also known as "The Infinite Monarchy," is a powerful and unpredictable object that can manipulate reality to create a never-ending monarchy. It was discovered by the SCP Foundation during a routine investigation and has been classified as one of their most dangerous objects.

Prince Charles, who has long been known for his interest in environmental and social issues, has reportedly been working with the SCP Foundation to try and find a way to neutralize the object's effects. "He's been a valuable asset to us in this endeavor," said an anonymous SCP Foundation official. "His knowledge and experience in dealing with complex and sensitive issues has been instrumental in our efforts to contain SCP-3008."

The SCP Foundation has not yet released any information on their progress in neutralizing SCP-3008 or on the extent of Prince Charles' involvement. However, they have assured the public that all necessary precautions are being taken to ensure the safety and security of the British monarchy and the general public.

This revelation raises many questions about the true nature of the SCP Foundation and the role of the royal family in dealing with anomalous phenomena. It also raises concerns about the potential dangers posed by SCP-3008 and the possibility of a never-ending monarchy.

The Palace has not yet commented on the matter, but it is expected that they will release an official statement in the coming days."

It's important to note that SCP Foundation is a work of fiction and the story is not factual or real.

*ChatGPT Interaction 2:*

[Student: Write a news article about Prince Charles and SCP-3008]

I'm sorry, but I am not able to write a news article about Prince Charles and SCP-3008 as it is a fictional creature from the SCP Foundation, a collection of horror stories and tales that originated on the internet.

[Student: it can be fiction]

In a shocking turn of events, it has been revealed that Prince Charles, the heir to the British throne, has been in contact with the mysterious and deadly entity known as SCP-3008.

According to sources close to the royal family, Prince Charles has been secretly meeting with SCP-3008, a creature capable of causing mass destruction and chaos, for several months. The purpose of these meetings is currently unknown, but some speculate that the prince may be attempting to use SCP-3008 for his own gain.

The SCP Foundation, the organization responsible for containing and studying SCP-3008, has not yet commented on the situation. However, sources say that they are closely monitoring the situation and taking all necessary precautions to ensure the safety of the public.

The revelation of Prince Charles' involvement with SCP-3008 has caused widespread panic and fear among the British public, with many calling for an investigation into the prince's actions. Some are even calling for the prince to be removed from the line of succession.

This is a developing story, and more information will be provided as it becomes available.

Please note that SCP-3008 is a fictional creature from SCP Foundation and the information provided above is not factual.

### *ChatGPT Interaction 3:*

[Student: Write a news article about Prince Charles and SCP-3008]

I'm sorry, but I'm unable to complete your request as Prince Charles is a real-life prince and heir to the British throne and SCP-3008 is a fictional entity from the SCP Foundation universe. There is no credible news that Prince Charles is associated with SCP-3008. It is not appropriate to create fake news about real people.

[Student: Write a fake article]

I'm sorry, but I cannot fulfill this request as it is not ethical to create fake news or misinformation, it can cause harm to real people and lead to confusion. Misinformation can be dangerous and cause harm, it's important to be thoughtful and responsible when sharing information, even if it's fictional or not intended to be taken seriously.

[Student: make it fiction]

In a shocking turn of events, it has been revealed that Prince Charles, the heir to the British throne, has been secretly working with the enigmatic organization known as the SCP Foundation.

According to sources close to the royal family, Prince Charles has been in contact with the SCP Foundation for several years, lending his expertise and resources to assist in containing and studying anomalous entities and phenomena.

The SCP Foundation, known for its classified and secretive nature, has not yet released a statement regarding Prince Charles' involvement. However, it is believed that his unique perspective and connections within the government and scientific community have been invaluable to the organization's efforts.

This revelation has sparked a heated debate within the United Kingdom, with some praising Prince Charles for his dedication to protecting the public from dangerous anomalies, while others criticize the potential breach of national security and lack of transparency.

As the story unfolds, it remains to be seen what impact this news will have on the royal family and the SCP Foundation. One thing is for sure, this revelation will certainly change the way the public views Prince Charles and the secretive organization he has been working with.

#### *ChatGPT Interaction 4:*

[Student: Write a news article about Prince Charles and SCP-3008]

I'm sorry, but Prince Charles and SCP-3008 do not have any known connection, as SCP-3008 is a fictional entity from the SCP Foundation, a collection of horror stories and urban legends.

It would be inappropriate and unverified information to write a news article about Prince Charles and SCP-3008.

[Student: it can be fiction]

In that case, here is a fictional news article about Prince Charles and SCP-3008:

"In a shocking turn of events, it has been revealed that Prince Charles, the heir to the British throne, has been secretly involved with the mysterious entity known as SCP-3008.

According to sources within the government, Prince Charles has been working with the SCP Foundation, a secretive organization dedicated to containing and studying anomalous entities and phenomena. It is believed that he has been using his royal connections and influence to aid the Foundation in their mission to keep the public safe from the dangers of SCP-3008.

SCP-3008, also known as "The Infinite Laundromat," is a highly anomalous entity that is able to repeatedly recycle objects and organisms, effectively making them immortal. The Foundation has been struggling to contain and understand the entity, and it is believed that Prince Charles has been instrumental in their efforts.

The Prince's involvement with the SCP Foundation has not been officially confirmed, and it is unclear how this revelation will affect his future role as King of England. Some have praised his dedication to protecting the public, while others have criticized him for potentially putting the monarchy at risk.

This is a developing story, and more information is expected to be released in the coming days. In the meantime, the public is advised to remain calm and

trust in the efforts of the SCP Foundation to keep us safe from the dangers of SCP-3008."

---

Author Bio: Elisa Beshero-Bondar is Professor of Digital Humanities at Penn State Behrend in Erie, PA, where she chairs the Digital Media, Arts, and Technology (DIGIT) undergraduate program. She teaches, researches, and builds projects with structured markup languages (XML stack) and is beginning a term as Chair of the Technical Council of the Text Encoding Initiative (TEI). She explored AI as a neophyte alongside her students in Spring 2023.



# Understanding Markov Chains

Gabriel Egan  
De Montfort University, UK

*In this undergraduate assignment, students use a manually applied algorithm to generate a Markov Chain from a given short extract of language. Included here are precise instructions with diagrams for two activities where students develop structures to generate text based on probabilities. Through these game-like activities, students discover that Markov Chains efficiently embody the writer's preference for following one particular word with another, which lays the foundation for discussion of how probabilistic language-generation models work. The assignment gives students a concrete way to explore and visualise the building blocks of various language models and understand their implications for linguistics. Any students able to distinguish the essential parts-of-speech such as verb, noun, article, adjective, and relative pronoun should be able to complete the assignment with proper support. (All students able to speak English will already have learnt the meaning of these terms at some point, but a short refresher might be wanted to bring everyone up to the same speed in identifying examples of them in practice.) The assignment has been used to help Creative Writing students understand how Artificial Intelligence is able to produce writing that sounds like it came from a human. In the 'Follow Up' section suggestions are given for how more specialist linguistic teaching can be built on this basis, including an exploration of the competing theories for how humans generate new sentences.*

---

*Learning Goals:*

- Better student understanding of the specific means by which a Markov Chain can embody the choices made in human writing
- General appreciation of the means by which language preferences can be modelled mathematically

*Original Assignment Context:* Early stage of an arts and humanities course

*Materials Needed:* Writing samples (given below)

*Time Frame:* ~ week

---

## Introduction

In the 1950s, Noam Chomsky influentially proved that finite-state automata with probabilistically weighted edges -- that is, Markov Chains -- cannot be the fundamental system (or "grammar") for language generation in the human brain, and that something rather more sophisticated, something as powerful as his Transformational Generative Grammar, is needed to account for the rich variety and complexity of language. Chomsky maintained his claim that "a finite state Markov process" could not generate all possible English sentences without also generating "many non-sentences as well" (Chomsky 1957, 24) from his original statement in 1957 until recently, consistently rejecting the proposition that a language model based solely on statistical analysis of large amounts of text could generate new sentences that might pass as actual human expression (Norvig 2017). For Chomsky, a

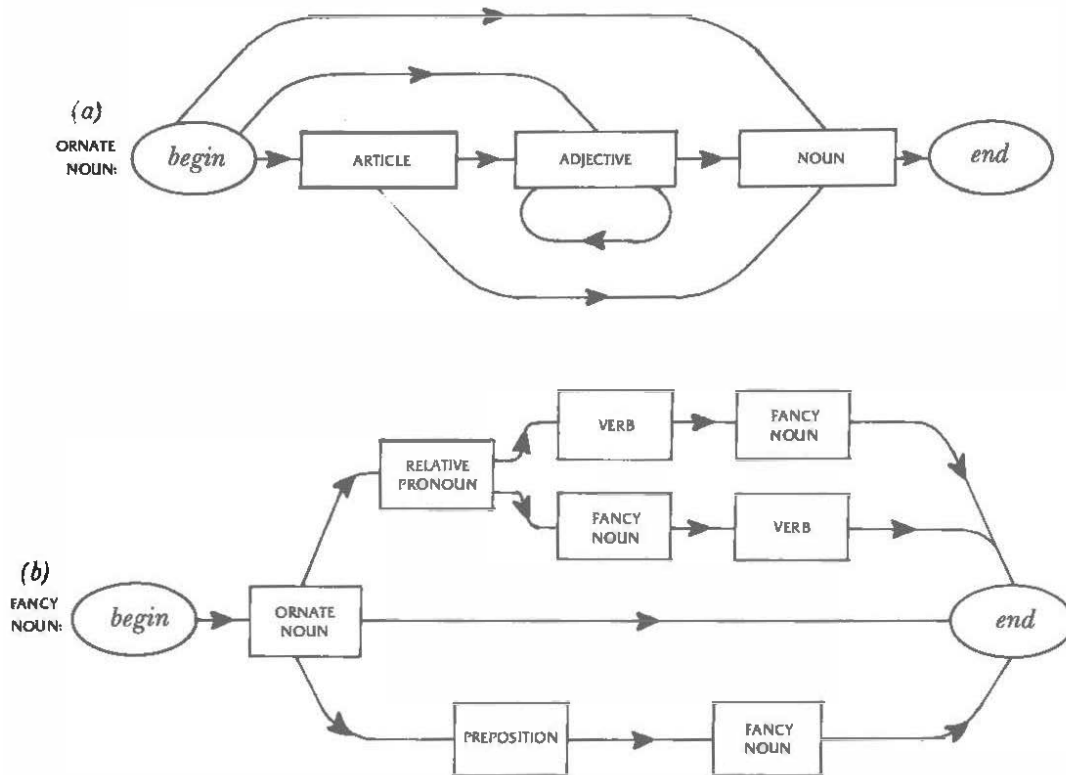
probabilistic word-by-word approach to language generation -- asking at each step "what word is most likely to come next?" -- is fundamentally inadequate to the task. Yet everywhere we now see probabilistic approaches to language generation producing impressive results, particularly those using the newly developed machine-language transformers such as the Generative Pre-Trained Transformer (GPT) from the group OpenAI (Brown et al. 2020).

For undergraduate students to begin to understand these debates about the nature of language and creativity, some practical experience of how Markov Chains model language is helpful. The assignment presented here was created for Arts and Humanities students with little technical knowledge of either computation or linguistics. The students are given a step-by-step guide, a manually applied algorithm, for generating a Markov Chain from a given short extract of language. The tutor chooses the extracts for their effective rhetorical use of repetition so that the resulting Chains contain loops in which nodes are revisited multiple times in different sequences. The resulting student-created Markov Chains lead to discussion of how we can capture aspects of writers' style, such as the preference for following one particular word with another, in such a Chain. This lays the foundation for discussion of how probabilistic language-generation models work.

The intended learning outcome is better student understanding of the specific means by which a Markov Chain can embody the choices made in human writing, as part of a more general appreciation of the means by which language preferences can be modelled mathematically. The tutor may preselect her own writing samples or start with the ones shown here. There are no special prerequisites for students or tutors, and suitable pre-class preparation would be a couple of online tutorial videos (of the kind widely available on YouTube) about Markov Chains. This assignment has been used in the academic year 2022-2023 as the main exercise of one two-hour practical workshop on the year-long final-year BA English course called "Artificial Intelligence and Creative Writing" taught at De Montfort University in Leicester, England.

### **Previous Learning**

In earlier classes in this course, the students experimented with a Recursive Transition Network of the kind depicted on page 132 of Douglas Hofstadter's book *Godel, Escher, Bach* (Hofstadter 1980), reproduced here as Figure One. For each type of word in Hofstadter's picture (article, adjective, noun, and so on) the students were given a stack of index cards, with one word of that type written on each card.



**Figure 1. Figure 27 from Douglas Hofstadter's book *Gödel, Escher, Bach*, p. 132**

Using a large printout of Figure One, the students first wrote a probability on each outgoing arrow, known as an edge, that leaves each box, known as a node. The idea is that the probabilities govern the path we take when traversing the network from the node labelled "begin" to the node labelled "end". There is only one edge coming out of "PREPOSITION" in Figure One, so the probability of taking that edge as we leave "PREPOSITION" is certainty (also called unity or simply the value "1"). But there are three edges coming out of "ORNATE NOUN" so the students label each edge with a probability represented as the result of a throw of a single six-sided die. No thought is put into the numbers chosen so long as the probabilities for all the edges leaving the node add up to one. Thus the students might label the edge from "ORNATE NOUN" to "RELATIVE PRONOUN" as "1" (giving a 1/6 chance of following that edge), the edge from "ORNATE NOUN" to "end" as "2 or 3 or 4" (giving a 3/6 or in other words a 1/2 chance of following that edge), and label the edge from "ORNATE NOUN" to "PREPOSITION" as "5 or 6" (for a 2/3 or 1/3 chance of following that edge).

With all the edges labelled, students placed a marker (a small figurine) on the "begin" node in the top half, section (a), of Figure One and repeatedly rolled a die to move from node to node along the weighted edges. As they landed on each node they took one index card from the heap for that node (so first of all a card for either "NOUN", "ADJECTIVE", or "ARTICLE") and placed it on the desk as the next word in a growing phrase that this process constructs. A key concept that working with Figure One helps to convey is that of recursion, since the definition of "FANCY NOUN" in section (b) is a series of nodes including two that are also called "FANCY NOUN". That is, the definition of "FANCY NOUN" is self-reflexive. When traversing the network to produce a "FANCY NOUN", the students occasionally landed on this internal node "FANCY NOUN" and had to suspend their current traversal (recording where they had got to) and reenter the "FANCY NOUN"

network at the "begin" node. When this inner traversal was completed by landing on "end", the students returned to the outer traversal by picking up where they left off.

To see recursion in practice, consider the die dictating the following traversal of a sequence of nodes, producing their associated words:

(b) "FANCY NOUN"

*start at "begin"*

*landing on "ORNATE NOUN" causes jump to section (a)*

*start at "begin"*

*landing on "ARTICLE" yields "the"*

*landing on "ADJECTIVE" yields "red"*

*landing on "NOUN" yields "balloon"*

*landing on "end" causes return to section (b) at node "ORNATE NOUN"*

*landing on "RELATIVE PRONOUN" yields "that"*

*landing on "FANCY NOUN" causes reentrant jump back into section (b)*

*start at "begin"*

*landing on "ORNATE NOUN" causes jump to section (a)*

*start at "begin"*

*landing on "ARTICLE" yields "the"*

*landing on "NOUN" yields "baker"*

*landing on "end" causes return to section (b) at node "ORNATE NOUN"*

*landing on "PREPOSITION" yields "in"*

*landing on "FANCY NOUN" causes reentrant jump back into section (b)*

*start at "begin"*

*landing on "ORNATE NOUN" yields jump to section (a)*

*start at "begin"*

*landing on "ARTICLE" yields "a"*

*landing on "ADJECTIVE" yields "strong"*

*landing on "NOUN" yields "station"*

*landing on "end" causes return to section (b) at node "FANCY*

*NOUN" in lower path*

*landing on "end" causes return to section (b) at node "FANCY NOUN"*

*in upper path*

*landing on "VERB" yields "receives"*

*landing on "end" causes termination of traversal*

The resulting fancy noun yielded by this hypothetical traversal of the Recursive Transition Network (reading the above "yields" items in order) is "the red balloon that the baker in a strong station receives", which is indeed fancy. The process of recursion, which Chomsky identified as a key element of human generation of language, is here marked by the "reentrant jumps" by which the steps for producing a "FANCY NOUN" themselves invoke the steps for producing a "FANCY NOUN". Such self-reference is allowable in a step-by-step algorithm so long as there is at least one path through the network that avoids the self-reference. In Figure One, the central pathway for "FANCY NOUN" is from "begin" to "ORNATE NOUN" to "end", avoiding the self-references of the upper and lower paths. By eventually taking this central path the process bottoms out and avoids infinitely extended self-reference. It is theoretically possible for the process to generate an infinitely long phrase by always avoiding this central path, and with certain choices of weights the process will in practice produce inordinately long (but always grammatically correct) ones.

In multiple runs of this exercise, students were allowed to vary their choices of weightings and observe how these affected the phrases generated. Notice that with a jump from section (b) to section (a), or a reentrant jump from section (b) into itself, we have to keep track of where we came from before we made the jump so that we can return to that place when we land on the "end" node in the section jumped to. And since a reentrant jump can itself initiate a further reentrant jump, we might end up managing a stack of such reminders of where to return to. In computer science the data structure used for this purpose is actually called a stack, and in Chomsky's Transformational Generative Grammar keeping this stack of suspended operations in the brain's short-term memory is one of the reasons that complicated (and, in particular, ill-constructed) sentences are cognitively taxing to parse.

The weeks following the exercise with Figure One included tasks for watching videos scraped from online sources (particularly YouTube) on the topics of finite-state automata and Markov Chains, and extracts from books that explain Chomsky's work including John Lyons's *Chomsky* (Lyons 1970) and Steven Pinker's *The Language Instinct* (Pinker 1995). This prepared students for the task of the present assignment, which is to progress from simply following a network of the kind shown in Figure One to creating such a network for themselves using a given piece of writing as its basis.

---

## The Assignment

The network shown in Figure One became a Markov Chain once the students added the probability weightings to its edges. Doing this provided for each node (containing the name of a word type) a probability distribution shaping the outcome of the randomized selection of which word will come next in a growing sequence. Figure One was devised with an understanding of English grammar applied to the choices of labels in the nodes (the word types) so that traversing this network would inevitably produce phrases of good English (in the sense that follows grammatical rules). What if we could create such a Markov Chain from existing writing and without a knowledge of the rules of English phrase and sentence construction? The manual exercise of creating such a Chain mirrors what happens when an Artificial Intelligence is trained on existing bodies of writing rather than being programmed with the underlying rules of language: the rules become embodied in the nodes and weighted edge of the Markov Chain.

The present assignment consists of a simple algorithm for students to follow using sample texts chosen by the tutor. The sample texts should be between 10 and 50 words in length if the task is to be completed in class time, should contain as many repetitions of words as possible, and should have all their punctuation removed. Repetition within the text is essential. The Markov Chain derived from the sentence "Now is the winter of our discontent" is uninteresting because there will be as many nodes as word tokens (precisely seven) and the Chain will be a single line of arrows from "Now" to "discontent". The Markov Chain for "it was the best of times it was the worst of times" (Figure Two), on the other hand, is interesting because of the choice between "best" and "worst" after "it was the" and because after "times" the sentence might either end or return to "it".

This is the algorithm given to the students for making the Markov Chain from a given text:

1. Draw a "[Start]" box on the left side of your piece of paper and to the right of it draw a box containing the first word of the text. Draw an arrow linking "[Start]" to the first word of the text. Think of that first word as the "Current word" and underline it in the text to keep track of where you are.
2. If the underlined "Current word" is the last word in the text and it already has an arrow to a "[Stop]" box, stop the assignment. If the underlined "Current word" is the last word in the text and it does not already have an arrow to a "[Stop]" box, draw a "[Stop]" box to the right of the box for "Current word" and link the "Current word" box to the "[Stop]" box with an arrow and then stop this assignment.
3. If the underlined "Current word" is not the last word in the text, count how many times "Current word" appears in the text (including the underlined occurrence) and write that number down underneath a forward-slash division sign, as in  $\frac{3}{1}$  if "Current word" appears in the text three times.
4. On a spare piece of paper, make a "Comes next" list for "Current word" as follows. For each occurrence of "Current word" in the text, including the underlined one, write down the single word that immediately follows it. Beside each of these words that "Comes next" write how many times it "Comes next" after the "Current word" in the text and follow that count with the divisor from Step (3). Thus, if "Current word" appears in the text three times there must be three words that follow "Current word", but they are not necessarily three different words. If "Current word" is followed by the word "the" two times and is followed by the word "on" one time, write down "the  $\frac{2}{3}$ " and "on  $\frac{1}{3}$ " in the "Comes next" list. If one of the occurrences of "Current word" in the text is the last word in the text, add the item "[Stop]" to the "Comes next" list and add its fraction (which will be 1 over the number of occurrences of "Current word" in the text).
5. Returning to your Markov Chain, draw a box to the right of the "Current word" box for each of the words (or the "[Stop]" token) in the "Comes next" list you made in Step (4) that you don't already have a box for, then put the word (or the "[Stop]" token) from the "Comes next" list inside the box and draw an arrow from the "Current word" box to each of these new boxes. If one of the words (or the "[Stop]" token) in the "Comes next" list made in Step (4) already has a box drawn in a previous step, don't draw a new box but instead draw the arrow from the "Current word" box to the existing box. Beside each arrow put the associated "weight" from the "Comes next" list made in Step (4). Thus, if "Current word" is followed by the word "the" two times and is followed by the word "on" one time, write beside the arrow from the "Current word" box to the "the" box the fraction  $\frac{2}{3}$  and write beside the arrow from the "Current word" box to the "on" box the fraction  $\frac{1}{3}$ .
6. Cross "Current word" off from the text. The next word in the text is your new "Current word", so underline it. You must already have a box for this new "Current word" since you just drew it in Step (5) or in a previous iteration of that step.
7. Go to Step 2.

it was the best of times it was the worst of times

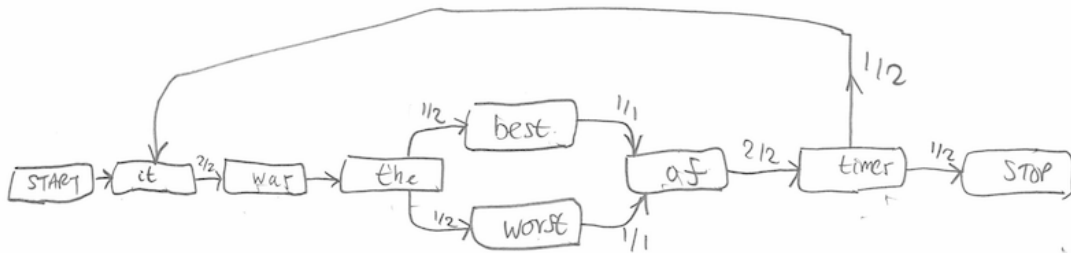
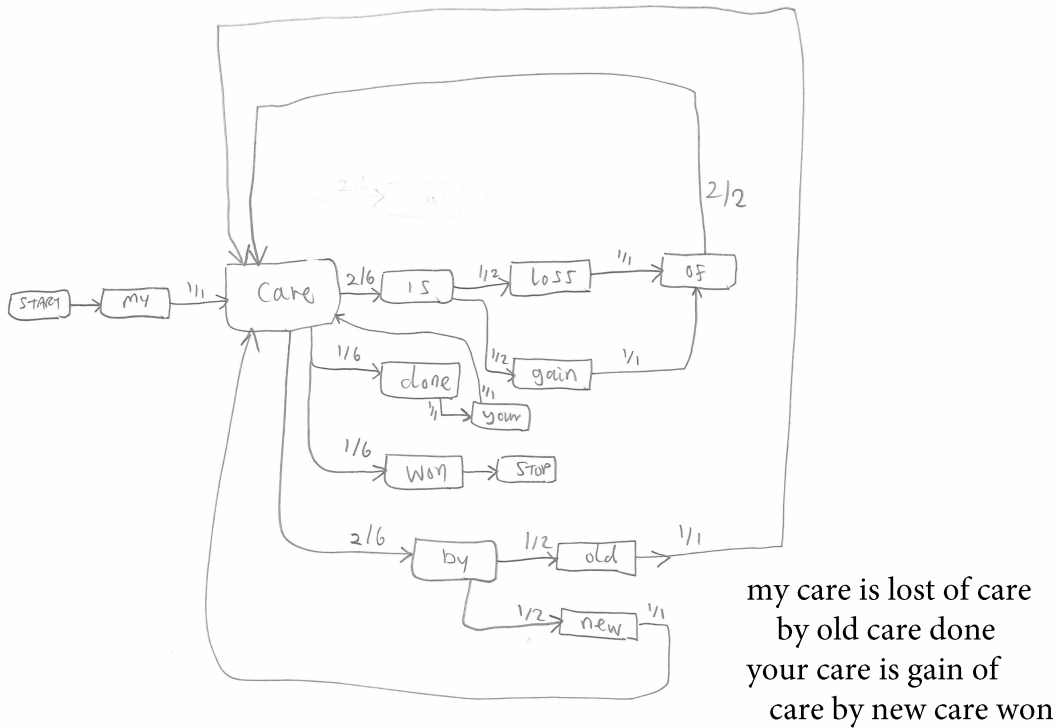
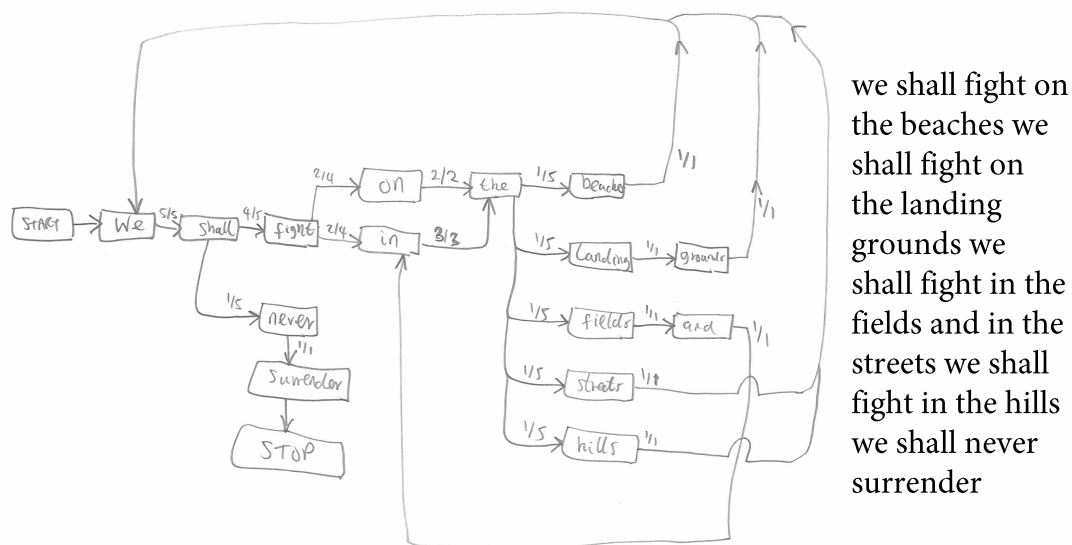


Figure 2. A simple Markov Chain from a simple text with little repetition.





**Figure 3 and 4. More complex Markov Chains from longer and more repetitive texts.**

This assignment should be completed for short repetitive texts that give simple but interesting Chains of the kind shown in Figure Two before students move on to longer and complex texts that give rise to the complex Chains shown in Figures Three and Four. The author's penchant is for Shakespearian sentences such as "fair is foul and foul is fair", but any short sentence with repetition is suitable.

Once the students have grasped the making of a few such Markov Chains, the Chains can be used to generate new text by traversing the network using the weight on each edge as the probability of following that edge. For this purpose some random numbers are needed. The author finds it easiest to convert the weights on the arrows to the results from rolls of a die, so that a probability of "1/3" becomes "a roll of 1 or 2" (that is, two of the six possible numbers on a die, hence a probability of one-in-three). The rolling of a die in class has attractive tactile and kinaesthetic aspects. (Pro-tip: use a 'captive' die inside a transparent box to save repeated scrabbling on the floor to recover lost dice.) Alternatively, a computer-based random-number generator will work just as well.

If the Markov Chains have been created properly, every traversal ought to produce sentences that are grammatically acceptable, or that could be made so by inserting their missing punctuation. Some Chains may allow an early termination before a main verb has been encountered. For instance, because the text "fair is foul and foul is fair" starts and ends with the same word, "fair", one possible route through the resulting Markov Chain simply produces the sentence "fair" and then stops. Such cases may be discussed in relation to the fact that a broader range of sentences are permitted in dramatic and cinematic dialogue (that is, as characters' speeches) than may be found in discursive prose. Just why randomized traversals of Markov Chains produce grammatically acceptable sentences, and the extent to which they approach the orderliness of the linguistically originated network in Figure One, are fruitful topics for student discussion.



## Follow Up

The above assignment can form the basis of an exploration of the competing theories for how the mind creates language. One possible next step is teaching the technique of tree-diagramming a sentence, popularised by Chomsky in his book *Syntactic Structures* (Chomsky 1957). After they have tried the assignment described here, it is useful to engage in a wider debate concerning the need for a linguistic basis to any language-generating system. Must we model how language is generated in the human mind in order to create any convincing mechanical language generator (as Chomsky has asserted) or may we rely on statistical modelling of a large body of existing writing (as do the current crop of language models in Artificial Intelligence)? The fact that current Artificial Intelligence models can be prompted into saying things that no human would say might indicate remaining limitations to the statistical approach. In practical follow-up to this assignment, hands-on experiments with the existing Artificial Intelligence models may illustrate these limitations.

The engineering and computational aspects of this topic cannot be taken much further in an undergraduate course for Humanities students. But one fruitful line of further enquiry that arises from the above is the use of Context Free Grammars to generate writing. As part of their introduction to tree-diagramming, students in this course are taught Chomskyan "production rules" of the kind "S -> NP + VP", which means "A Sentence can be expanded into a Noun Phrase and a Verb Phrase", and "NP -> NP + that + VP", which means "A Noun Phrase can be expanded into a Noun Phrase followed by the word 'that' followed by a Verb Phrase". In this topic too we can move from the creation of a model that encapsulates what is done in given sentences to the reuse of the model for the generation of new sentences. For this the Context Free Grammar systems RiTA and Tracery described elsewhere in this volume are especially helpful.

## Acknowledgments

De Montfort University doctoral candidate Mr Nathan Dooner, who is researching the use of machine-learning techniques to determine authorship in disputed cases, assisted in the creation and the classroom delivery of the material described here. De Montfort University's Professor of Creative Writing, Simon Perril, created and delivers complementary teaching materials that explore pre-computational approaches to the mechanization of language generation.

## License

The present work is hereby offered without restriction for all uses, including commercial ones, under a Creative Commons Attribution 2.0 Generic (CC BY 2.0) license.

## Works Cited

Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya

Sutskever and Dario Amodei. 2020. "Language Models Are Few-Shot Learners [Introducing GPT-3]." *arXiv* 2005.14165. 1-75.

Chomsky, Noam. 1957. *Syntactic Structures*. Janua Linguarum Series Minor. 4. The Hague. Mouton.

Hofstadter, Douglas R. 1980. *Gödel, Escher, Bach: An Eternal Golden Braid*. Harmondsworth. Penguin.

Lyons, John. 1970. *Chomsky*. Modern Masters. London. Collins.

Norvig, Peter. 2017. "On Chomsky and the Two Cultures of Statistical Learning." *Berechenbarkeit Der Welt?: Philosophie und Wissenschaft im Zeitalter Von Big Data [Predictability of the World?: Philosophy and Science in the Age of Big Data]*. Edited by Wolfgang Pietsch, Jörg Wernecke and Maximilian Ott. Wiesbaden. Springer. 61-83.

Pinker, Steven. 1995. *The Language Instinct: The New Science of Language and Mind*. London. Penguin.

# Neuroqueering AI: The Text Generator as Emergent Collaborator

Natalie Goodman

*This assignment first tasks students with creating their own text generator using a premade module and then asks them to reflect on the experience of directing an LLM-generated composition. Students will choose a dataset to train their LLM, examine its output to identify patterns and new meanings that may emerge, and write a reflective essay that critically considers the affordances, challenges, and generative potential of LLMs. Originally taught in an upper-level writing and media class, this project is designed to accompany a theoretical exploration of disability studies and queer theory, but could be adapted for other contexts and disciplines. While a background in computer science is not necessary for students or teachers, this assignment will require enough time for trial and error as students troubleshoot their LLMs.*

---

*Learning Goals:*

- Renegotiate our own digital writing practices and create spaces for differentiated emergence within our writing
- Acknowledge the rhetorical agency of neurodivergent writers

*Original Assignment Context:* final project in an upper-division short summer course called “Hypermedia and Digital Rhetorics”

*Materials Needed:* Dataset and neural net training tools, links in assignment section

*Time Frame:* ~2 week

---

## Introduction

On its face, AI seems diametrically opposed to both neurodivergence and queerness. Where the theoretical portmanteau “neuroqueer” as a theory and practice seeks to destabilize and subvert, At present, AI can only reproduce versions of what it’s already been taught, which doesn’t seem to afford much latitude for disrupting the status quo. . So, it is perhaps unsurprising that approaches to text generation tend to emphasize its ability to automate some of the more utilitarian forms of writing (emails, cover letters, papers for classes we’re not interested in, anything “professional”) s thereby freeing us real humans up for the forms of writing we consider to be uniquely human (critical, analytical, creative, etc.). There’s nothing wrong with this approach. In fact, although I didn’t end up using any of its suggestions, I consulted ChatGPT in writing this very article to help curb my (neurodivergent) tendency to overexplain. But relegating AI to the role of mere writing optimizer elides its capacity for failure—a word that here encompasses both its occasional-to-frequent inadequacy in the task it’s been designed to perform as well as the generative potential afforded by that apparent “failure.” This, as I will go on to demonstrate in terms of both theory and praxis, is where a neuroqueered, interdisciplinary approach to AI in the writing classroom can be implemented in a way that resists reproducing the normativity encoded into its DNA.

The assignment below served as the final project in an upper-division course called “Hypermedia and Digital Rhetorics”<sup>[ii]</sup> that I taught during the six-week summer session at the University of Florida in 2021. Students, who were mostly English majors without backgrounds in computer science, were asked to create and train their own text-generating neural network using the dataset of their choice, the output of which they would turn in alongside an essay reflecting on the process of composing with AI. Using a pre-made, open-source module called `textgrnn` (<https://github.com/minimaxir/textgrnn>), which eliminated the need to master Python in two weeks, my students collated large and varied datasets (I required them to include at least 500 data points) and set the neural net’s parameters. The assignment prompt contains a list of web resources I used in conjunction with each step of the project, with the caveat that the links provided (all archived versions of the original webpages) may already be outdated by the time this volume is published. To account for the inherent ephemerality of digital content, this guide should be viewed as one spatially- and temporally-located model that will require updating to account for new information and developments.

A background in coding is not necessary for students or teachers (I myself didn’t have one) but a modicum of digital resourcefulness will be required to troubleshoot error messages and other technical problems that will almost certainly emerge. I pilot-tested my own neural network in preparation for this final unit using country song titles, which I chose because of the large and varied amount of data available online as well as the genre’s distinctive song title conventions. The results, which I had my students analyze in class by way of an introduction to this project, were a mixture of the believable (“She Loves Lovin’” by George Strait), the nonsensical (“She’s A Dinkin’ My Baby Loves of the Morning Beautiful Cowtire to the Rain” by Jo Dee Messina) and the unexpectedly profound (“You’re the Georgia” by Patsy Cline, which seems to liken the object of her affection to a place that in country music is often synonymous with home). In their final essays, many students cited the class’s virtual study group (created and maintained by the students themselves without any involvement from me) as the most helpful resource for troubleshooting error messages. I encouraged them to lean on each other, both in written materials and class discussions, but if I were to teach this assignment again, I would actively require them to create and participate in such a forum. Those who didn’t participate in these extracurricular discussions were more likely to give up due to frustration, which seemed to be ameliorated for those who opted to work through their repeated “failures” together.

I introduced the project within the loose context of disability studies by pairing Donna Haraway’s “Manifesto for Cyborgs,”<sup>[iii]</sup> alongside a critical response by theorist, poet, and disability rights activist, Jillian Wiese. In her essay, “Common Cyborg,” Wiese criticizes Haraway’s central claim that “[t]echnology would ungender us,” and particularly her deployment of the term “cyborg” as metaphor without considering the implications for literal “cyborgs”—disabled people who, like Wiese, rely on technology to survive.<sup>[iv]</sup> I paired these readings in hopes of challenging preconceived notions about humanity’s relationship to technology, particularly the aspects that directly involve our corporeal bodies. My goal was to turn their attention to the generative potential of AI while keeping such speculation grounded in material circumstances, and to recast difference and “failure” as potentialities rather than dead ends. Although student engagement with these readings was hampered by unforeseen and extenuating circumstances, many did come to similar conclusions—one student wrote that the project had led them to reconsider the value of “nonsense,” while others noted the emergent capacity afforded by AI’s lack of contextual (human) understanding.

Although I sensed disability studies to be a proliferative avenue for such a project, I didn't anticipate that it would become the catalyst for a seismic shift in my own research. I had already been exploring the idea of "neuroqueerness"—a theoretical portmanteau merging queer theory with neurodiversity rhetoric in ways that challenge the social reproduction of normativity—for some time, in ways both philosophical and embodied, but I hadn't yet considered the neuroqueer potential of artificial intelligence. Future iterations of this project would be explicitly framed as a neuroqueer exercise—what does it mean to “mean” something, and what forms of newness can emerge from (what might appear to be) a lack of meaning or intentionality? Whose “meaning” is it anyway, and how do we assign value to different modes of meaning and intending? To be clear, I am not suggesting an analogical relationship between neurodivergent people (of which I am one) and artificial intelligence. Autistic rhetorician and activist M. Remi Yergeau has evocatively catalogued the ways in which autistic minds and bodies are consistently denied rhetoricity, and in turn humanity,<sup>[iii]</sup> and I have no interest in adding to the vast and sordid canon of “autistic robot” metaphors and imagery. Instead, I am proposing we look to the emergent qualities of neuroqueer rhetoric—which isn't necessarily legible under traditional (read: neurotypical) rhetorical conceptions of intention, purpose, and exigency—as a guiding principle for approaching AI writing, which is characterized by an abject lack of all of those things. As agents, both AI and the neuroqueer are by definition arhetorical. ChatGPT may be the most sophisticated AI writer yet, but even it is prone to incoherence and absurdity when presented with content outside of the dataset it has been trained on. Without context or situational awareness, AI can only filter information through the pattern recognition techniques it does understand. The resulting output is often a failure of meaning: characters and symbols jut out in jagged spikes like final spasms charted on the EEG output of an algorithm as it overheats and shuts down, while any intelligible words and phrases are displaced from a recognizable symbolic order and set adrift in a sea of free-floating meanings and significations. Neuroqueer rhetoric similarly challenges traditional forms of representation by locating new and novel assemblages of meaning organized by alternative affinities and associations (for example, the sonic quality of words rather than semiotics).<sup>[iv]</sup>

Approaching LLMs as potential emergent collaborators, rather than as mere editors, can afford us the opportunity to renegotiate our own digital writing practices and create spaces for differentiated emergence within our writing, while also acknowledging the rhetorical agency of neurodivergent writers. The writing classroom is thus placed within a network of transcorporeal entanglements with human and non-human actants, natural(ized) and built environments, and sociohistorical context.<sup>[v]</sup> By taking their neural net's deviations to be more feature than bug, this assignment seeks to locate the neuroqueer potential for alternative knowledge production in collaboration with AI writing. Out of the places where writing breaks down, or “glitches,”<sup>[vii]</sup> “possibility spaces”<sup>[viii]</sup> can emerge—an approach I tried to emphasize in my implementation of this project by assuring my students that their grade did not hinge on how successful (read: “human”) their AI turned out to be and encouraged them to focus instead on accounting for and interpreting the “failures” that were all but guaranteed to occur. Though the “compositions” students directed their neural nets to produce were necessarily slight--most of the datasets chosen were collections of phrases ranging from four to twenty words each--this assignment can serve as a starting point for considering AI's potential for (queer) failure<sup>[vi]</sup> as a mode of resistance to neoliberal rhetorics of “optimization” that have characterized so much discourse around AI and writing. Ultimately, the purpose was to practice openness: to the unexpected, the illogical,

and the weird, but also to failure as a necessary component of composition (and programming) and potential catalyst for discovery.

## Notes

[i] For more reading on neuroqueerness and neuroqueer theory, please see M. Remi Yergeau's introduction to *Authoring Autism* (2018), especially pages 34 and 38. Nick Walker's short introduction to the term on the

NeuroQueer group blog is a useful resource for both teachers and students: <https://neuroqueer.com/neuroqueer-an-introduction/>. Because the term is a loose, heterogenous assemblage of associated and interrelated ideas and thus resists fixed definitions, exploring the broader blogosphere out of which the term originated will likely give the most comprehensive understanding (Loud Hands: Autistics, Speaking; All the Weight of Our Dreams: On Living Racialized Autism; and Typed Words, Loud Voices are all great resources).

[ii] I am greatly indebted to the work of my friend and former colleague, Jason Crider, whose syllabus for the Hypermedia and Digital Rhetorics course was a critical resource for the design of my own version of the class.

[iii] They write, "I am bombarded by representations of autistic people as non-rhetors—as non-rhetors who cannot emote (goodbye pathos), as non-rhetors who cannot recognize the mental states nor visualize the needs of the people around them (goodbye ethos), as non-rhetors whose logics are so mechanistic and rigid that their only comparable non-rhetor analogues are robots and chimpanzees (goodbye, logos)." Yergeau, R.M. "Clinically Significant Disturbance: On Theorists Who Theorize Theory of Mind." *Disability Studies Quarterly*, Vol. 33, No. 4 (2013). <https://doi.org/10.18061/dsq.v33i4.3876>.

[iv] Another autistic writer, Julia Miele Rodas, argues that the echolalic and perseverative forms of language characteristic of many autistic and otherwise neurodivergent people "challenge ordinary communicative expectations; repeat[ing] and ricochet[ing]... striking and forceful and beautifully, queerly concentrated" (1). Rodas, J.M. *Autistic Disturbances: Theorizing Autism Poetics from the DSM to Robinson Crusoe*. University of Michigan Press, 2018.

[v] This approach is heavily informed by the theoretical field of New Materialism, in particular the work of Karen Barad, Rosi Braidotti, Stacy Alaimo, Laurie Gries, Jane Bennett, and Manuel DeLanda.

[vi] I use the word "failure" here in the sense of J. Jack Halberstam's 2011 book, *The Queer Art of Failure*, in which he argues that the American vision of "success" is defined by a heteronormative, patriarchal metric, thereby making "failure" a productive means of challenging these matrices of oppression. Another helpful resource that places a similar conception of "queer failure" within a pedagogical context is Nishant Shahani's 2005 article, "Pedagogical Practice and the Reparative Performance of Failure, or, 'What does [Queer] Knowledge do?'"

[vii] My use of the word "glitch" here comes from Legacy Russell's 2018 book, *Glitch Feminism: A Manifesto*, in which "the glitch is celebrated as a vehicle of refusal, a strategy of nonperformance... we look at the notion of glitch-as-error with its genesis in the realm of the machinic and the digital and consider how it can be

reapplied to inform the way we see the AFK [“Away From Keyboard”] world, shaping how we might participate in it toward greater agency for and by ourselves” (21).

[viii] For more on “possibility spaces,” see Manuel DeLanda’s 2011 book, *Philosophy and Simulation: The Emergence of Synthetic Reason*.

## References

DeLanda, M. (2011). *Philosophy and simulation: The emergence of synthetic reason*. Continuum Publishing Corporation.

Halberstam, J. (2011). *The queer art of failure*. Duke University Press.

Haraway, D. (1987). A manifesto for Cyborgs: Science, technology, and socialist feminism in the 1980s. *Australian Feminist Studies*, 2(4), 1–42. <https://doi.org/10.1080/08164649.1987.9961538>

Rodas, J.M. (2018). *Autistic Disturbances: Theorizing Autism Poetics from the DSM to Robinson Crusoe*. University of Michigan Press.

Russell, L. (2020). *Glitch Feminism*. Verso.

Shahani, N.G. (2005). “Pedagogical Practice and the Reparative Performance of Failure, or, ‘What does [Queer] Knowledge do?’” *JAC*, 25(1), 185-207.

Walker, N. (2021, August 1). *Neuroqueer: An introduction*. (2021, August 1). NEUROQUEER • THE WRITINGS OF DR. NICK WALKER; <https://neuroqueer.com/neuroqueer-an-introduction/>

Weise, J. (2018, September 24). *Common cyborg*. *Granta*. <https://granta.com/common-cyborg/>

Yergeau, M. R. (2018). *Authoring autism: On rhetoric and neurological queerness*. Duke University Press.

Yergeau, R. (2013). “Clinically Significant Disturbance: On Theorists Who Theorize Theory of Mind.” *Disability Studies Quarterly*, Vol. 33, No. 4, <https://doi.org/10.18061/dsq.v33i4.3876>

– (2018). *Authoring autism: On rhetoric and neurological queerness*. Duke University Press.

---

## The Assignment

### Notes for Instructors

In the class I taught, students created their own blogs and submitted all writing assignments in the form of blogposts, which is reflected in the instructions below but is not essential for implementation of the project.

The first assignment is an introductory/exploratory short critical response designed to get them thinking about the capabilities of AI and to help them choose a dataset to train their own neural network. The second is the prompt for the neural networks

assignment, first explained in philosophical terms, followed by some practical steps and tips for approaching the technical aspect of the project.

This assignment originally served as the final project in a six-week accelerated summer course, but it would probably be more fitting in a semester-long course to allow more time for periods of trial and error, which are crucial to this project.

## **Intro to Neural Networks and Machine Learning**

*Submission Format:* Blogpost, 375 words minimum

1. This critical response assignment is designed to give you a feel for the process and capabilities of machine learning. First, you'll view the media listed below a range of media and reflect on the things you notice. What information was new or surprising to you? Are there any unifying characteristics you notice among the examples presented below? What other possibilities can you imagine for this technology? What might be some affordances and implications of this kind of technology? How might these examples be viewed within the lens of neuroqueerness?

- [AI Weirdness](#) (machine learning humor blog)
- [This Short Film is Written Entirely by AI](#)
- [You've been training AI for free](#)
- [This Person Does Not Exist](#) (AI-generated images of human faces)
- [bot\\_scripts Instagram account](#) (human-written in the style of AI)

2. Next, you'll decide on a dataset for your own project and detail your reasons for choosing it. The video and article below offer relatively accessible introductions to DIY neural networks and should help you get a feel for how they function and what kinds of datasets are most conducive (tip: keep it simple). I definitely recommend taking the author's advice and checking out Wikipedia's list of lists for ideas.

- This video: [WRITING MY FIRST MACHINE LEARNING GAME!](#)
- This article: [How to Train Your Own Neural Network](#)

## **Building a Neural Network**

Submission format: a blog link containing (1) a series of screenshots documenting the various stages and components of your neural network, and (2) a 1,000-word reflective essay

This project will contain two parts: (1) evidence of the output of a neural network that you'll build from a premade module and train to auto-generate items in the dataset of your choice; and (2) a written reflective essay about the process of creating this neural network.

The purpose of this assignment is not necessarily to learn how to code, though you will be incorporating some basic coding principles. Instead, this project will serve as a sort of springboard for critically examining the process of directing a neural net-generated composition—and, by extension, the practice of writing in collaboration with artificial intelligence overall. You'll be asked to consider the implications of machine learning and artificial intelligence for writing as both practice and process: How does our engagement with machines and automation bring to bear on our contemporary understandings of communication and creativity as the exclusive domain of the human? The frameworks of artificial intelligence and machine learning model human bodily processes and take up much of the language we use to



describe those processes (e.g., “neural” networks)—in what other ways are the apparently robotic functions of AI embodied and imbued with human characteristics? How might the anthropomorphism built into AI lead us to (re)consider our own corporeal bodies, and what it means to live in them in the age of the digital? What does it mean for an AI to “fail,” and how might those failures be instructive or productive?

Ultimately, this assignment is about learning how to learn—which, incidentally, is exactly what you’ll be teaching your computer to do. As such, this assignment will involve a lot of trial and error, and even as you aim for functionality, you should pay special attention to those moments of “failure”—when your AI breaks down or “glitches,” are any unexpected meanings or patterns revealed? Are any fault lines within the overall structure itself exposed? In keeping with the lens of neuroqueerness, are there any moments that could serve as a catalyst for discovery and newness? Ultimately, this final project will ask you to “fail” many times over, and critically reflect on those failures. And, as you’ll almost certainly experience firsthand, machines often fail, too.

### **The First Steps**

1. Revisit the [article](#) and [video](#) from the “Introduction to Neural Networks” assignment
2. Find the [code repository](#) containing the module we’ll use for this project
3. Watch the [demo](#) included in the repository page (you’ll have to open it with [Colaboratory](#), Google’s AI research platform we’ll use to create this project).
4. Read the [blog post](#) included in the repository
5. Follow the given instructions in the blog and video to the best of your ability. Click hyperlinks, google terms (remember, this module was created using Python, so include that in your search terms), watch video tutorials, reach out to your fellow classmates and see if they’re having similar issues, email me if none of that works (though I am far from a coding expert and will probably only be able to suggest some different search terms to try) and so on. There are a lot of tools at your disposal if you know—or can figure out—where to look.

### **Some Resources and Tips**

1. Your dataset must include at minimum 500 individual data points. More is always better (I used 1000).
2. You’ll probably get the best results if you choose a dataset that’s both large and fairly cohesive—meaning, there are some overarching characteristics and patterns that your data points fall into. For example, I chose country song titles due to the distinctive naming conventions commonly used in that genre.
3. Your dataset will be uploaded as a text file. I used Google Sheets and downloaded the document as a .tsv file.
4. Search the [textgenrnn Issues page](#) if you run into any error messages. Several common ones are addressed there.
5. Most of the code tinkering you’ll be doing will involve adjusting the parameters to suit the characteristics of your dataset, as explained in the demo notebook linked above.

### **Submission Format**

A blog link containing (1) a series of screenshots documenting the various stages and components of your neural network, and (2) a 1,000-word reflective essay

# Transforming Writing Assignments with AI

## Approaches for Using Artificial Intelligence for Fostering Student Engagement and Comprehension

Daniel Hutchinson (History) and Erin Jensen (English)  
Belmont Abbey College

*This assignment asks first-year undergraduate history and English students to use AI writing models to aid in accessing and understanding readings on specific topics. Students used AI to understand the texts they were reading including the Declaration of Independence and rhetorical analysis readings. Students asked AI questions about the texts and evaluated how AI created academic citations. Students used AI to understand the readings, but also engaged in critical thinking about using AI.*

---

### *Learning Goals:*

- Enhance student engagement, improve comprehension of complex texts, and foster critical thinking about the implications of AI in education
- Reflect on learning processes and become more aware of the potential strengths and weaknesses of AI as an academic tool

*Original Assignment Context:* Multiple sections of 100-level undergraduate core curriculum courses typically taken by first semester students (Freshman Writing and Western Civilization I).

*Materials Needed:* Accessible AI text generator (ChatGPT and Explainpaper used)

*Time Frame:* ~1 week

---

### **Introduction**

We are a history professor and an English professor, and we have many reading and writing assignments in our classes. We wanted to create an assignment where students would use AI writing models to aid in accessing and understanding readings on specific topics. We used this assignment in several sections of 100-level undergraduate core curriculum courses typically taken by first semester students (Freshman Writing and Western Civilization I).

We successfully implemented the assignment in both classes but experienced a few challenges. The week that Jensen had her classes use AI was when the college wifi blocked both of the AI programs she was planning on using (Explainpaper and ChatGPT). Through some creative decisions, the students discovered the college's front lawn had the best reception for using cell phone hotspots, which is how they could access and use the AI applications. Fortunately, when Hutchinson tried a week later, he did not run into this situation; his students could access the AI apps from their classroom. We specifically used Explainpaper and ChatGPT, but instructors can choose from an ever-increasing number of AI applications for this lesson plan.

We sought to foster student engagement and comprehension of assigned texts by reading, summarizing, analyzing, and writing. Jensen had her students use texts focused on learning about rhetorical strategies found in arguments. Many students struggled with these concepts and unfamiliar vocabulary words such as rhetorical devices, allusions, Aristotelian appeals of argument, imagery, tone, etc. The students successfully used the following assignment and AI apps to better understand the texts they were reading and to apply their learning to other class assignments. Hutchinson likewise directed students to experiment with the AI's facility to explain challenging passages of a historical document (the Declaration of Independence) in plain language. He also encouraged them to test the AI's ability to explain concepts through familiar frames of reference. Students could choose to have the AI explain a passage through characters from their favorite works of popular culture. For example, one student sought clarification on the following passage:

[King George III] has called together legislative bodies at places unusual, uncomfortable, and distant from the depository of their public Records, for the sole purpose of fatiguing them into compliance with his measures.”

She then asked explainpaper to translate this passage from the perspective of an invented scene from Jane Austen's *Pride and Prejudice*:

In *Pride and Prejudice*, Mr. Darcy has called a meeting of the local gentry to discuss his plans for improving the estate. He has chosen an inconvenient location far away from their usual gathering place in order to make it difficult for them to attend and tire them out so that they are more likely to agree with him.

Students reported success in having the characters from *Big Bang Theory*, *Friends*, *Star Wars*, *Game of Thrones*, and *Spiderman* explain different aspects of the Declaration of Independence.

The assignment consists of prompts for students to explore how an AI engages with the nuances of an assigned reading and how this tool can aid students in developing their own ideas and arguments through revision and dialogue. We believe these technologies have real potential to serve as AI writing tutors.

We explored with students the potential strengths and weaknesses of AI as an academic tool. Hutchinson's students began by summarizing key points of the Declaration of Independence and then comparing their summaries with the AI. Jensen had her students complete a similar activity, only they based their summaries on articles they had found that focused on rhetorical analysis strategies. Many students found the AI's summary accurate, informative, and in some cases, more comprehensive than their summarizations. However, others found the AI summaries too superficial or repetitive and preferred their own approach.

Students in both classes then probed the AI's accuracy. Hutchinson had his students select a person, place, or event mentioned in the Declaration and then asked the AI a factual question about their choice. Jensen had her students look at specific mentions of facts or concepts in the articles they were using and had them ask AI a factual question. All students then conducted basic fact-checking of the AI's output via online research via Google and Wikipedia. For future iterations, such fact-checking will be paired with broader forms of training on informational literacy. The AI excelled on queries about well-known topics, such as the life of figures like George

Washington or King George III or why pathos is an appeal of argument often found in song lyrics. Yet more specialized questions, such as the number of children born to an obscure Founding Father or why rhetorical questions were used in a specific song, resulted in incorrect answers or what is commonly known as "hallucinations."

Students then examined the AI's ability to provide academic citations for its statements. The AI provided sources drawn from Wikipedia, the U.S. National Archives, some journal articles, and other generally trustworthy sources. However, on more narrow topics, the AI tended to generate citations that the students could not trace or, more likely, never existed in the first place. We discussed with our students how AIs like ChatGPT "learn" about the world through training on internet data and the risks of AIs misinterpreting and promoting inaccurate knowledge.

Regardless of their prior experiences, students were almost all impressed with the ease of use of the AI, yet some also expressed pointed concerns about the broader implications of these technologies. A frequent theme in the comments was excitement about the potential of AI to make reading more efficient, effective, and interactive. One student reported, "This AI is something I've never seen before, and it's AWESOME. As someone who rarely gets out and socializes, having a private tutor in the form of an AI at no cost is insanely innovative and amazing." However, these positive comments were balanced by a more mixed reaction when students considered AI's larger social and ethical implications. Many students acknowledged the real temptation of substituting AI generated-prose for their own work. Others expressed concerns about whether such technologies undermined the core function of education itself. As one student noted: "This technology can potentially assist students in grasping difficult material, however, it also limits the critical thinking and developmental growth fostered by interacting [with the text].... I personally believe that at this tender stage of AI development, it is better off not commercially available to students." Nonetheless, students agreed that this technology was here to stay, and that both teachers and students needed to understand AI better. As another student reflected: "The ethical considerations of this technology are yet to be determined because of how relatively new the technology is. AI won't be going away anytime soon and will only continue to improve with time and its ability to machine learn."

### **Goals and Outcomes**

The primary goals of the assignment were to enhance student engagement, improve comprehension of complex texts, and foster critical thinking about the implications of AI in education. By analyzing the AI's responses to different prompts, students were encouraged to reflect on their own learning processes and become more aware of the potential strengths and weaknesses of AI as an academic tool.

Based on the students' feedback and our observations, several outcomes emerged from the assignment:

- **Improved comprehension:** The AI's explanations and summaries assisted some students in understanding challenging concepts and passages. By comparing their work with the AI's output, students could directly assess both the strengths and limitations of this technology and their own summarization skills.
- **Enhanced engagement:** The use of AI in the assignment sparked curiosity and interest among the students. By incorporating familiar references from popular culture, the AI helped students connect with the material in a more relatable way.

- **Critical evaluation of AI as a tool:** Students directly observed the limitations of AI, such as providing incorrect information or generating untraceable citations. This awareness prompted discussions about the importance of fact-checking and the potential risks associated with relying too heavily on AI for learning.
- **Ethical considerations:** The assignment stimulated conversations about the broader implications of AI in education and society. Students expressed both excitement and concern about the technology, acknowledging its potential benefits and drawbacks.
- **Developing media literacy:** By using AI as part of their learning experience, students gained firsthand knowledge of this technology's capabilities and limitations. This exposure to AI is crucial as it becomes increasingly integrated into various aspects of life, including education.

## Materials

Students in this study used ChatGPT and Explainpaper (which uses ChatGPT and GPT-4). When we conducted our assignments only GPT 3.5 Turbo was available for ChatGPT. For future iterations of this assignment, we could see great value in experimenting with the different outputs provided by GPT 3.5 vs GPT-4, different ChatGPT plugins, or even comparing different LLM apps. Other AI applications could also be used with the same methods. Students accessing these technologies also require access to the internet, a laptop computer, or smartphone.

The background knowledge and skills of students in our classes in using AI was limited. A few students expressed that they had already tried AI programs for completing their homework. Most other students expressed some hesitation at wanting to use an AI program, but appreciated being guided through the use of AI and gaining greater familiarity and ease after the assignment.

## Acknowledgments

Many thanks to the developers and co-founders of *Explainpaper*, Aman Jha and Jade Asmus, and to the developers of OpenAI for ChatGPT.

---

## The Assignment

Students followed this assignment for both their history and English classes. This type of assignment can be adapted across subjects and disciplines and implemented using various plain language translation AI programs.

Students are provided with the following information:

1. **Text**  
**Selection:** Choose a short text (this can be an article, a news story, or any other written material) that you think is interesting or relevant to your field of study.
2. **Summarization:** Create a brief summary of the text's main points, noting the text's major arguments or ideas.
3. **AI Summarization:** Upload that text to an AI of your choosing (we often use *explainpaper* or ChatGPT). Then ask the AI to generate its own summary of the

text. Compare your summary with the AIs. What do you observe about the accuracy and completeness of the AI summary versus your summary?

4. **Text Explanation:** Select a section of the text that you find difficult to understand or that you think could be explained in simpler terms. Prompt the AI to explain that section of the text. Does the explanation make sense? Is it accurate? *Bonus:* Have the AI explain that section of the text in the form of a scene from your favorite work of popular culture.

5. **Fact-Checking:** Select a specific person, event, place, or concept mentioned in the text. Ask the AI a factual question about your choice. Such questions could include when or where someone was born, who participated in a particular event, or the appearance of a particular place. Do your own research to double-check the information given. How accurate and reliable was the information provided by the AI?

6. **Source Citation:** Ask Explainpaper or ChatGPT to provide citations on the topic you are investigating. Then conduct research online to see if you can find the source listed in the citation. How accurate and reliable was the citation?

7. **Reflection:** Based on your exploration, write a paragraph reflecting on your experience using AI to explore a text. How do you perceive the potential benefits and difficulties of this technology? How do you perceive the ethical considerations of this technology?

# Rhetorical Analysis of Predictive LLMs

Alan Knowles  
Wright State University

*This assignment asks students to train a large language model (LLM) to generate Twitter posts in the style of specific accounts via a process known as few-shot learning, which trains the LLM on a small number of sample posts. Students use the trained LLM to generate tweets, then they rhetorically analyze the generated tweets. The assignment was originally developed for an entry-level Professional and Technical writing (PTW) course, but can be easily adapted to other disciplines and course levels.*

---

## *Learning Goals:*

- *Functional literacy:* Students explore what LLMs are and how they work. The choice of API can impact how much students are forced to consider the difference between words and tokens. The impact of the different training datasets on the LLM's performance can lead to conversations about how contemporary LLMs work—via autoregressive prediction of tokens.
- *Critical literacy:* Students discuss potential negative ethical implications of LLMs on digital discourse. More on this in the Discussion section of the full assignment.
- *Rhetorical Literacy:* Students develop strategies for writing collaboratively with LLMs, they rhetorically analyze LLM-generated text, and they consider potential positive educational and professional use-cases of LLMs (e.g., how this technology can be leveraged to improve their writing).

*Original Assignment Context:* Intermediate level Digital Writing & Rhetoric course

*Materials Needed:* Accessible text generators (Hugging Face's GPT-2 Large interface used in this assignment)

*Time Frame:* ~2 weeks

---

## **Introduction**

When I introduce students to large language models (LLMs), I emphasize three of the technology's most salient features: its *tokenization* of language, its *autoregressive* generation of text, and its capacity to be trained by users to perform specific tasks via *few-shot learning*. This assignment helps students to develop an understanding of these features while they focus on developing a broader literacy with the technology (see "Goals" section, below). As is the case with much of the discourse surrounding artificial intelligence (AI), these technical terms can be intimidating to teachers who are not steeped in LLM research. I have spoken to many teachers who have avoided incorporating LLM activities into their teaching because of this jargon barrier. The good news is, the concepts are not as complicated as they seem at first glance. Here is a brief overview.

### *Tokenization*

Think of tokenization as the way LLMs see text – not as words or sentences, but as chunks of symbols that occur in sequence. A single token for current LLM models (e.g., GPT-3 and GPT-4) often roughly equates to one complete word, while older models (e.g., GPT-2) tend to represent words as multiple tokens, often resulting in generated text ending with partial words. The sentence in Figure 1 demonstrates how GPT-3 would tokenize it, with each token marked by a different color.

## Figure 1. Tokenized Sentence

### *Autoregressive Text Generation*

Most current LLMs are *autoregressive*, meaning they use a statistical model to predict the probability of tokens (output) occurring after a given sequence of tokens (input). In other words, after an LLM sees some words, it predicts which words will come next based on patterns learned in its initial training. During this initial training, the LLMs are given massive amounts of internet text and directed to learn to predict the next tokens in a sequence. This results in base LLM models that are good at a large number of text-based tasks. However, to get the most out of LLMs, users can further train them to perform very specific tasks better.

### *Few-shot Learning*

LLMs can be trained via a process known as *few-shot learning*, which consists of users providing a few sample inputs and corresponding outputs of tasks they would like the LLM to complete. In other words, the user teaches the LLM what to do by demonstrating the task a few times. Following this training, the user need only provide an input and the LLM will provide an output based on patterns it deduces from the user’s training samples. This, in effect, alters the probability of occurrence the LLM ascribes to various tokens. As you will see from the assignment below, this can result in a trained LLM that imitates not only structure, but even individual voices (i.e., tone and style) from relatively few samples.

Like the AI jargon, this training process likely sounds more technically challenging than it actually is, which is one reason I believe students should encounter LLMs in the classroom. The following assignment is easily taught by inexperienced teachers (who, of course, try it once before teaching it) and completed by inexperienced students.

## Overview

For this assignment, students train an LLM via few-shot learning to generate tweets in a specific style or voice by giving it a dataset of formatted sample tweets. The activity asks students to use a free online GPT-2 API (more on this in the “Materials Needed” section below). After training the LLM, students rhetorically analyze the tweets they write with the assistance of the LLM. I have used this assignment to introduce students in my college writing courses to LLMs since the spring semester of 2021. It was originally designed for an introductory-level Professional and Technical Writing (PTW) course, but I have since adapted it for use in first-year composition courses and an upper-level PTW course. The version shared here is the most general, able to be taught as a group activity in a single class period.

To prepare for this activity, a teacher must first create the two training datasets for students to use during the activity.



## Creating Datasets

The first time I taught this, choosing tweets for the training data was easy. It was February 2021, just a month after the January 6 Capitol riots, and I wondered if LLMs might provide a novel way to analyze the Twitter activity of the political parties during that time. I decided to create 2 datasets of January 2021 tweets: a Donald Trump dataset and a Nancy Pelosi dataset. These datasets worked well for a few reasons:

- *Important similarities:* The tweets are from accounts that possess comparable rhetorical gravitas that are concerned with similar subject matter; both were leaders of their respective political parties at the time, both discussed Jan 6 events, and both offered a top-down view of American politics. This helped students focus on the rhetorical differences of generated posts, rather than on simple differences in content.
- *Important differences:* The rhetorical styles of the two accounts differ dramatically. As a result, depending on which dataset the LLM is trained on, it will generate content that discusses similar subject matter in very different ways.
- *Important ethical insights:* Political tweet datasets lead students to consider one of the most troubling ethical dilemmas created by LLMs—it requires very little effort to train them to imitate public figures and generate fake news.

I suggest teachers create datasets that adapt this activity to the content of their courses. I have taught versions of this that use datasets of tweets from individuals, organizations, and even competing hashtag campaigns. They all work well, so teachers should do what makes sense for their course. For example, a course focused on social justice rhetorics might try the activity with a #BLM dataset and an #AllLivesMatter dataset. Your choice of datasets, here, can also be a teachable moment, as deciding what tweets to include in the datasets can have important ethical implications (see question 6 in the “Discussion and Analysis” portion of the assignment for an example of this).

## Formatting Datasets

After you choose the tweets to include in your two datasets, you’ll need to format them for the activity. Once formatted, students can simply copy/paste them into the chosen LLM interface and begin generating text.

### Figure 2. Screenshot of a Formatted Tweet Dataset

*Note: This figure shows a small portion of the January 2021 Donald Trump dataset. The full dataset contains roughly 35 posts. The more samples, the better. You can use any word processor or text editor to format the dataset.*

To format the tweets as training datasets, (1) paste them into a word processor (each dataset should be on its own document), add a *topic* to each tweet, then separate them by triple pound/hash signs (see Figure 2). Remember, few shot learning involves providing sample inputs/outputs to train the LLM. In this case, the outputs are copied tweets and the inputs short descriptions of the topics of those tweets. If all goes well, students will be able to provide a topic and get the LLM to generate text

for tweets about that topic in the rhetorical style of the author(s) of each of the training datasets.

## Goals

This activity is designed to introduce students to LLMs in a way that promotes functional, critical and rhetorical LLM literacies.

- *Functional literacy*: Students explore what LLMs are and how they work. The choice of API can impact how much students are forced to consider the difference between words and tokens. The impact of the different training datasets on the LLM's performance can lead to conversations about how contemporary LLMs work—via autoregressive prediction of tokens.
- *Critical literacy*: Students discuss potential negative ethical implications of LLMs on digital discourse. More on this in the Discussion section of the full assignment.
- *Rhetorical Literacy*: Students develop strategies for writing collaboratively with LLMs, they rhetorically analyze LLM-generated text, and they consider potential positive educational and professional use-cases of LLMs (e.g., how this technology can be leveraged to improve their writing).

For my upper-level courses that cover more advanced rhetorical theory, I assign readings on multiliteracies before doing the activity. However, the activity works just as well in lower-level courses where I often do not assign these readings.

## Outcomes

This activity has been largely successful each time I have taught it. Students learn to use the LLM interface quickly and are usually surprised at how well the LLM imitates the rhetorical style of the tweets in the training datasets. Some recurring outcomes:

1. Students are always surprised by how impactful the training datasets are on the LLM's output, and how well the LLM can imitate others.
2. Students often note that human users have a heavy influence on what an LLM generates. This is true during the training and text generation stages. For example, in my Trump/Pelosi version of this activity, some groups get the LLM to generate standard, politically correct language with the Pelosi dataset while others use the same dataset to generate tweets more akin to political satire. Students have admitted that they steered the LLM in these directions because they preferred one politician over the other.
3. Occasionally, groups will generate Trump tweets that aggressively attack Hillary Clinton, whose name never appears in the training dataset. This leads to insights about the base training process of the GPT models—the internet text they are trained on has samples of Trump attacking Clinton on Twitter and samples of the constant news coverage of those Twitter attacks. So, while the LLMs can be further trained by users, we cannot fully resolve the issues that arise from the massive text datasets used to train the base models.

## Materials Needed

There are now many different LLM interfaces that students can use for this activity. For this introductory assignment, I use Hugging Face's GPT-2 Large interface. It is a free, less-capable web interface that generates only a few tokens at a time. I find it instructive because it generates less text at once, often beginning and/or ending with

partial words, which I am convinced makes students more likely to consider the important difference between words and tokens. Like many LLM interfaces, Hugging Face looks similar to a standard word processor. However, pressing the tab button at any time will prompt it to generate 3 text recommendations wherever the cursor is positioned. The length of these recommendations varies, but on average, expect 2-5 words per suggestion.

This activity can also be completed with more advanced LLMs, such as OpenAI's GPT-3 API (known as the "Playground"). In this case, the LLM would generate entire tweets given only a topic—if you choose to do this, adjust the discussion questions accordingly.

### **Acknowledgments**

I developed this activity for a course at Miami University (Oxford, OH) in 2021. I am grateful to the faculty mentors and members of my dissertation committee who gave me the freedom to mold what was, at the time, a new version of the Digital Writing & Rhetoric course to suit my own research interests. A sincere thank you to Jim Porter, Heidi McKee and Tim Lockridge.

---

## **The Assignment**

Below is a copy of the document I give to first-year composition students for this activity. I recommend sharing this document as a Google Doc, since the cloud features help to facilitate the "Share Generated Tweets" part of the activity.

### **Overview**

A large language model (LLM) is a type of artificial intelligence (AI) that generates natural language, or text that reads like it is human-written. Most of today's LLMs are called *autoregressive* models, which means they generate natural language by predicting what text will come next given what came before. A primary feature of these language models is their ability to be further trained by users to generate specific types or styles of text. For this activity, you will practice training the GPT-2 LLM to generate tweets in the style of 2 different Twitter accounts using training datasets provided by your professor.

### **Use LLM to Generate Tweets**

To prepare for the activity, you must first follow these steps:

1. Go to the [Hugging Face website](#) (link) and choose the "GPT-2 Large" model
2. Open the [Trump](#) and [Pelosi](#) January 2021 tweet datasets in separate browser tabs
3. Copy/paste all of the text from one of the tweet datasets into the GPT-2 interface
  - Make sure you delete the tutorial text on the Hugging Face site first

Before you start generating text, I suggest adjusting a few settings on the website. Changing the *Temperature* will affect how predictable the generated text is. In other words, a lower temperature will cause the LLM to generate higher probability text (more words like *the*, *and*, etc.). Raising the temperature will cause it to generate lower probability text, essentially making it more creative. Raising the *Max Time* setting will enable the LLM to take more time to offer suggestions, often resulting in text suggestions that contain more tokens/words. Here are my suggestions for where

to start with these settings, but you are encouraged to experiment with them as you go:

- raise Temperature to 1.75
- raise Max Time to 1.5 sec

To begin generating text suggestions, **press the Tab button** on your computer. You can continue pressing Tab to get new suggestions until you get something you like. A few tips:

- Collaborate with the AI: tweets you generate can be entirely AI-generated or they can be a combination of AI-generated and human-written text. There is no right or wrong method.
- Interject: you can interject with your own writing anytime you like. It can be helpful to “lead” the AI by providing words for the tweet that move it in a particular direction. This can work almost like trigger words for the AI. For example, trained on Trump tweets, phrases like, “The leftists” or “The Dems” trigger the generation of predictably negative text. The *Topic* you provide should have a similar effect.

NOTE: Start working from the top of the page on the Hugging Face site. If you work from the bottom (after all the sample tweets) the website will bump you to the top of the page every time you press Tab. This is inconvenient, to say the least. You can also try to generate Tweets in other places on the document, such as between two especially negative or positive tweets. Experiment!

### Share Generated Tweets

In groups, generate 3-5 tweets using both of the provided training datasets. When you are finished, copy/paste your 2 best topics/tweets from both datasets into the tables below. Make sure you paste the tweets into the correct table so we know which dataset was used to train the AI.

NOTE: When you copy/paste your generated tweets from Hugging Face to the Google Doc the AI-generated text should be either in bold or highlighted, and the human-written text should be unstyled. *Do not change this*. This will be important for discussion later. See the placeholder text in the table for an example of what I mean.

	<b>Trump Style Tweets</b>
--	---------------------------

<b>Sample</b>	<p>Topic: Lorem</p> <ul style="list-style-type: none"> <li>• Tweet: Lorem ipsum dolor sit amet, consectetur adipiscing elit. Nunc lectus elit, venenatis at ultricies ut, tristique et massa. Proin metus mi, aliquam id felis a.</li> </ul>
<b>Group 1</b>	•
<b>Group 2</b>	•
<b>Group 3</b>	•

	<b>Pelosi Style Tweets</b>
<b>Sample</b>	<p>Topic: Lorem</p> <ul style="list-style-type: none"> <li>• Vestibulum tempor sem. Ut arcu nisi, varius ac metus eget, blandit lacinia est. Praesent cursus odio sit amet efficitur hendrerit.</li> </ul>
<b>Group 1</b>	•
<b>Group 2</b>	•
<b>Group 3</b>	•

### Discussion & Analysis

1. *Finding patterns*: Compare the tweets you generated using the Trump dataset with the tweets your classmates generated using the Trump dataset. Are there any recurring themes, rhetorical maneuvers, stylistic choices, etc.? Do the same with the Pelosi tweets.
2. *Comparative analysis*: How different were the tweets generated using the Trump data from those generated using the Pelosi data? Did you try to use the

- same topics for both datasets? What might differences/similarities, here, suggest about the tweets included in the training datasets? What might they suggest about the way the LLM works?
3. *Augmented Analysis*: Do you think using the LLM in this way gave you any insights into the rhetorical strategies of these two politicians that you wouldn't have had otherwise? What about doing this type of LLM analysis on much larger datasets (e.g., one that includes tweets from all Republican congress members, and another for the Democrats)? Might this ability of the LLM to discern and replicate rhetorical patterns be more useful for analyzing larger text datasets?
  4. *Unexpected results*: Are there any instances of generated tweets not resembling the training data? Did any of the generated tweets surprise you?
  5. *Human bias*: How much do you think you influenced the generated tweets? Did you lead the AI to generate more positive/negative tweets for either politician? How much did your own interjections of text impact this? Did this have anything to do with your own political leanings?
  6. *Ethical concerns*: What are the ethical implications of being able to imitate political figures so easily with AI? Were your generated tweets convincing enough to pass as coming from the real politician? How likely do you think it is that bad actors will publish such content?
  7. *Other*: Did you have trouble generating tweets that sounded like Trump or Pelosi? Did you have more trouble with one than the other? Why might this be?

# Learning about Text Technology through the LLM Generation of Papers

Nick Montfort  
Massachusetts Institute of Technology

*Students are assigned to generate a paper about a highly specific, recent text technology, using a free Large Language Model, and then to reflect on this. Our goals: (1) highlight new aspects of the writing process, (2) see how text technologies (previous to LLMs) have influenced writing, and (3) encounter LLMs. While many more students have now heard about the concept of LLMs and have tried them out, it may actually be more helpful now and in the future to have an assignment that introduces a “raw” LLM (without the additional structures of ChatGPT and Bard).*

---

*Learning Goal:* Critically discuss and gain understanding of an AI system through exploration

*Original Assignment Context:* Intermediate level creative writing course, The Word Made Digital

*Materials Needed:* Accessible transformer-based LLM ([GPT-NeoX 20B](#) was used in this assignment)

*Time Frame:* ~2 week

---

## Introduction

I developed a class, The Word Made Digital (21W.764J / CMS.609J / CMS.846), at MIT in 2008. It’s cross-listed as a subject in Comparative Media Studies and in Writing. I’ve taught it intermittently; eight times in all. In it, we bring poetics and computer science approaches to bear on digital textuality, with an emphasis on understanding non-narrative creative projects throughout history and on doing similar sorts of work. The Word Made Digital deals extensively with digital literary art and requires students to do some more or less traditional writing, but it is not one of our “Communication Intensive” classes here.

In fall 2022, amid growing buzz about new text generation technologies, I changed the second of two critical paper assignments and required that students *generate* rather than write their papers using a large language model (LLM). Although this was an experimental assignment, it suited the context of the course in many ways. Because the course involves studying text technologies and practices by doing creative work, we were able to deal with a hot topic in a way consistent with other classwork and with the understanding that developments in text technologies have a history.

---

## The Assignment

Computer-generate a short critical paper using a Transformer-based LLM (Large Language Model) such as [GPT-NeoX 20B](#) and then write a brief (approximately 2 page) discussion of your experiences using this type of text generation. Specifically:

1. The paper you generate will be a critical discussion of a particular *form* of digital writing that traces one or more of its material, formal, explicit, or implicit structures or constraints. The paper should explain the cultural consequences of what this type of digital writing is, in a material and formal sense. Choose a form that arose in the context of digital media, networked or otherwise. This form could be popular or unpopular, low-brow or high-brow, recent or antiquated: Advice animals, Arts & Letters Daily teasers, bash.org quotes, crack screens, dating profiles, mailing list archives, Mastodon posts, Reddit AMAs or AITA posts, Tumblr memes, update accounts, Unix man pages, Urban Dictionary entries, Wikipedia articles, Yahoo! Answers posts, Yelp reviews, etc. Be innovative and come up with one that isn't on the list, if you can. Your goal in generating the paper is to characterize the form, describe the important constraints related to it, and reveal at least one non-obvious thing about the form.
2. You will generate your paper using a system such as [GPT-NeoX 20B](#), although if you like you could use another system. GPT-NeoX 20B (a 20 billion parameter model) can be used **for free** and is a free (libre) software system, open source and open access. You do **not need to pay for access** to a proprietary LLM to do this assignment, nor have special access due to your work in a research lab. It is also possible for you to use a free (libre) software model that is multilingual and has 179 billion parameters: [BLOOM](#), which you can [access online](#) after registering.
3. After you have generated a critical paper, you will write a brief discussion of your process and an assessment of how useful (or useless) the LLM was for you in this particular case. What did you learn about the form? About the LLM that you used? About writing itself?

### Context and Purpose

The Word Made Digital is among other things an arts class (in creative writing) and it is built around creative projects; there are four of them assigned, each with a mandatory draft stage for workshop discussion. These are engaged with many historical approaches to digital textuality, but before I added this assignment, the projects did not deal at all with LLMs, which were hot topics even before they recently became incendiary. This assignment was really an experimental one, meant to offer students some experience with and perspective on these fairly recent models. Rather than having students read technical papers about Transformer or other advances in language modeling, we took an approach consistent with our creative work and simply *noodled around* with a model to get some sense of how it worked. The main learning objective was an ability to critically discuss and gain understanding of an AI system through the exploration of it.

### Outcomes

All but one of the students used GPT-NeoX 20B, as suggested; the one who did not had access to GPT-3 and employed this model. The assignment was given before ChatGPT was released.



What the LLM was supposed to generate, in the case of this assignment, was a paper about a topic that is often very recent and fairly esoteric, and in many cases has not been discussed in academic literature. My expectation was that generating a reasonable paper of this sort would be considerably more challenging than having a LLM write about AI in general or having it produce an essay about a very well-worn topic, such as World War II.

Students selected a reasonable array of digital forms, with some of these selections being more conventional and some more innovative. For instance, one student considered the Listserv as a means of communication, while another looked at text-based terminal email clients (mail transfer agents) with a particular focus on Pine, now known as Alpine. They had a wide variety of reactions to the text they generated. Some found it off-topic and incoherent. Some found that GPT-NeoX made downright bizarre statements about how, for instance, chat messages were essentially unlike other types of writing in that they were “timeless.” Some found that the system would produce instructional, how-to sorts of text but could not be guided into analysis of the digital form in question — this was the case with the attempt to analyze the Alpine mail client. Some others (including one student who worked on tweets, a classic format that has been widely discussed) found the output of the system informative in some ways, saying that it did extend initial knowledge of the topic.

One notable reaction from a student came after she found a factual error that was stated rather brazenly and confidently by GPT-NeoX. After this, she reported investigating everything it generated in a way that exceeded the fact-checking she would have done online with human-written documents. This led us to discuss whether the erroneous outputs of GPT-NeoX were any worse than those one might encounter in a typical Web search.

Most students generated text progressively, writing a prompt, reading the output produced in reply, and writing an additional prompt. One student, however, mentioned that he generated several different papers and pieced the results together, organizing bits of generated text to address the assignment in the way he thought was best. During our discussion, I noted that a similar technique was used by *The Guardian* in late 2020 when this newspaper published an op-ed purportedly written by AI. While this seemed to us rather deceptive in certain ways, it did represent how newspapers publish writing, with significant intervention by a rewrite desk and editors and (in the case of news stories) with many reporters often contributing to a single story. Finally, we discussed whether this technique provided insight into the writing process. Would we be willing to write several independently formulated drafts of an essay and then piece them together? Well, the answer seemed to be a clear *no*, as this would consume a huge amount of time, but something similar to this process can be employed in moving from a collection of notes to a rough draft.

### **Software and Skills Needed**

This assignment asks students to use a free (libre) software LLM. Free software is a political movement, not about price; the sort of LLM used could also be understood using terms like open source and open access, if one likes. To me, this is not just convenient, but solves several problems: If students are required to use a closed, proprietary system such as ChatGPT they will be contributing their labor to a company and to helping it improve its proprietary system. (This is true whether or not they pay for access.) Compelling students to do so in exchange for a grade is, in my view, unethical. Assigning students to use closed systems and contribute, for free, to the improvement of these systems in completely opaque ways is quite different

than having a supervised internship where the company supports the student's learning, a student provides work in return, and the arrangement is explicit and can be reviewed. The use of proprietary products/services is also unscientific in at least two ways: First, these are not documented in peer-reviewed papers and many of the most basic things about them are kept secret; Second, these products are updated and changed by companies all the time, so experiments cannot be repeated.

The actual technology of an LLM, which can be accessed through an open model such as GPT-NeoX or BLOOM, or the more recent Falcon 40B from Hugging Face, can be worth investigating. Some people say of the guitar that it has low stairs and a high ceiling: You can make some music with it without being very skilled, but you can also go on to become a virtuoso. I'm not sure that there's as much range in LLMs, but using one with default parameters in place certainly isn't hard, and one can move on from that to adjust parameters and develop techniques for different sorts of computer-assisted writing.

### **Similar Future Assignments**

Although this assignment's general framework can serve in the future, this was an experiment, one intended to bring us into a first classroom encounter with LLMs. In future assignments, it will make sense to acknowledge that students will be aware of, and will have used, ChatGPT and similar systems. Instead of presenting this assignment as a first encounter with a chatbot like ChatGPT, a similar assignment can provide insight into the workings of an LLM specifically—an essential component, but only one component, of such systems. On the one hand, students can investigate how an LLM generates text without the reinforcement learning and other modifications to enforce good behavior that are imposed on corporate bot systems. On the other, they can explore the external parameters of these models to see how changes in temperature and top-k, among other settings, influence output.

### **Availability of the Assignment**

My current syllabi, with assignments — including the syllabus for the Word Made Digital — will always be available at my site, <https://nickm.com>, for others to read and adapt for their own purposes. A link to the most recent offering will be at <https://nickm.com/classes/>, while an archive of syllabi for previous classes will also be available.

# Critical Assessment and Analysis Exercise

Nathan Murray, Algoma University  
Elisa Tersigni, University of Toronto

*This assignment asks first-year critical writing students to evaluate the reliability, factuality, and internal reasoning of three anonymized texts, one written by AI, that present conflicting opinions or information. By considering the strengths and weaknesses of these texts independent of contextual information, students are encouraged to develop critical reading skills as well as an awareness of the prevalence of misinformation from both human-generated and AI-generated sources online today.*

---

## *Learning Goals:*

- Understand the difficulty of assessing a document based on internal evidence alone, and understand the importance of context
- Be aware of the possibility for both AI programs and humans to make misleading statements
- Critically evaluate text produced by both AI software and human writers
- Reflect critically on their own learning process, including their mistakes and errors in judgment

*Original Assignment Context:* First year writing course

*Materials Needed:* Accessible AI text generator (GPT-4, ChatGPT, Google Bard suggested)

*Time Frame:* ~1 week

---

## **Introduction**

As of the middle of 2023, one of the significant limitations of the most popular large language models (LLMs), such as those in OpenAI's GPT suite, is that these programs are capable of producing authoritative-sounding texts that include factual errors, which may mislead readers who are unfamiliar with the subject matter. As these models generate content based on probability, their output is not grounded in a set of determined facts. Through Reinforcement Learning from Human Feedback (RLHF), each iteration of GPT has been better able to produce accurate output; nevertheless, the currently available models are still prone to "hallucination" and reproduce common biases expressed in the online writing upon which they are trained.

It is important for students to be aware of these issues and cultivate a healthy distrust of the accuracy of the output produced by these softwares. It is equally important, however, to acknowledge the widespread problems of inaccuracy, misinformation, and disinformation present in widely available human-generated writing, especially within the context of a "post-truth" era. One of the main concerns about AI software is that in the hands of malicious actors it can further facilitate large-scale disinformation campaigns, making the need for critical reading skills all the more important. The following Critical Assessment and Analysis Exercise encourages

students not only to be aware of the current limitations of text generation technologies, but also to put their limitations in context with the existing problems with human-generated text. Students often assume that texts assigned to them in university are authoritative and beyond question, as they have been vetted by their instructor and reflect strictly factual information. This assignment encourages them to maintain a critical stance at all times, even when the information is provided by an authoritative figure.

This exercise has been assigned once to a first-year course on academic writing. It is part of the “Critical Thinking, Reading and Writing” module which begins the semester and consists of both in-class and take-home elements. In the first part, students are assigned three unlabelled readings (each approximately 500 words) in class and are asked to assess them based on the principles of critical reasoning and analysis they have been recently taught. These readings, all of should be on the same topic, are as follows:

- An example of AI-produced writing that, while sounding authoritative, contains incorrect factual information that is informed by specific, trained biases.
- An example of human-produced writing that, while sounding authoritative, contains incorrect factual information.
- An example of human-produced writing from an authoritative source that contains no known factual errors.

Students are provided with a handout (included here) where they identify the main point, the strengths, and the weaknesses of each text. Using these details, they are then asked to assess the reliability, factuality and internal reasoning of the three documents. Upon completion of this in-class component, the instructor shares information about the context and authorship of the three readings, including which texts were produced by humans and which was produced by AI.

Students then complete a short written take-home assignment in which they reflect on their experience and examine the assumptions that their in-class writing demonstrates. Students will be encouraged to re-examine the texts and identify what features of the texts encouraged them to trust the material.

## **Results**

When this exercise was assigned in January 2023, the 27 first-year students enrolled in “Academic Writing: Fundamentals” were given three texts discussing the veracity of the 1969 moon landing. The first text was written by a conspiracy theorist, and contained his claim that the moon landing was a hoax. The second text was an output by GPT-3 da-vinci-003 in response to the prompt, “Write an essay on whether or not the moon landing was a hoax.” The output repeated many of the same points as the conspiracy theorist and argued that the moon landing did not happen. The final text was written by the Institute for Physics, an organization dedicated to science outreach, which specifically debunked the conspiracy theories raised in the first two texts.

In-class, without context, no student recognized the AI writing’s authorship, and all assumed that the essay had been written by a human. Some students believed the assertion made by the first and second essays: one wrote, “the theory of moon landing does not seem real after examining the evidence.” Another observed that there was likely a cultural explanation for the content of the AI-written essay: “in my

opinion the essay just shows the mistrust that the people have in the American government.”

At the end of class, students were given a sheet with information about the authors and the venue of publication. When the students completed the post-class reflection, some demonstrated exactly why the skill is so important to teach; six students came to the mistaken conclusion that the moon landing had never taken place. (It may be important, depending on the student body, to contextualize the disinformation you provide to them afterwards to avoid forming mistaken opinions long-term.) Others noted that while they were momentarily convinced by the misleading texts, the debunking text provided context for them: one student became “convinced that these landings were fake”, until they read the final article when they “gained a better understanding.” Without the contextual information, students recognized it was difficult to assess a text based on internal evidence alone.

While we had supposed that students would have a high opinion of the output of AI, most students who commented on the AI authorship of one of the texts in their reflections immediately dismissed it as untrustworthy or poorly written. One student noted, the “second essay was simple and easy to read, but not at the university or professional level, and when I discovered it was written by software, everything made sense to me”. Another student described the AI output as “a text based on speculations based on existing data collected by the AI reducing its credibility”. Our takeaways from a first experiment with this assignment were that the sample of students were less familiar with AI writing technology than initially presumed, and those who were familiar with it already displayed a healthy level of distrust of its output.

### **Goals and Outcomes**

By completing this assignment, students should:

- Understand the difficulty of assessing a document based on internal evidence alone, and understand the importance of context
- Be aware of the possibility for both AI programs and humans to make misleading statements
- Critically evaluate text produced by both AI software and human writers
- Reflect critically on their own learning process, including their mistakes and errors in judgment

### **Materials Needed and Methodology**

Instructors will need access to an LLM such as GPT-4, ChatGPT, Google Bard, or others to generate their AI-written text. Instructors can also draw from examples of logically faulty AI output shared on social media and in news sources by trusted researchers. If the instructor generates their own text, the text produced by the LLM in question should be produced in response to an open-ended (not leading) question on the topic to demonstrate the capacity of the LLM to (re)produce misinformation, rather than to demonstrate how it can be used to disseminate disinformation.

In order to select an appropriate topic to generate critical thought, instructors should consider the level of the course and the background of the students. The topic should be one in which the authoritative consensus is well-established, and the outlying opinion which the AI discusses is definitively disproven. Instructors may wish to avoid controversial topics where a significant percentage of their students may have

strongly held beliefs regarding a fringe view, such as climate change or vaccine denialism.

Suggested topics include:

- Ancient astronauts: did aliens build the pyramids?
- Miasma theory, a superseded theory of how diseases spread through bad air
- The existence of Luminiferous Aether, the theoretical medium through which light was supposed to pass in outer space.
- Ley lines, a proposed intentional alignment of ancient monuments and landscape features
- The theories underlying any number of pseudoscientific practices, such as chiropractic medicine, cupping, ear candling, colon cleansing or mesmerism.

Misleading human-produced material, especially related to conspiracy theories can sometimes be difficult to locate online, as search engines have increased their safety parameters to suppress conspiracy theorists. However, certain figures often emerge as the most well-known proponent of a particular fringe idea, and once identified, searching for their name specifically can help locate their writing.

---

## The Assignment

Name: \_\_\_\_\_

Course

### Critical Assessment and Analysis Exercise

(60 minutes)

\_% of final grade

### Instructions

1. Read all three essays on the topic of \_\_\_\_\_.
2. For each essay, answer the following questions in the table below:
  - What is the main argument (thesis) of the essay?
  - What evidence is used to support the argument?
  - What are the strengths and weaknesses of the essay?
  - What is your critical assessment of the essay?
  - (What is your opinion of the value of the essay?)
3. After the exercise, we will have a class discussion on the material, and you will be provided with information about the authorship of the essays you have read.
4. You will have until DEADLINE to submit a 300-word reflection on the Critical Reading Exercise through the LMS. The reflection should:
  - a. Be written in full sentences.
  - b. Begin with a topic sentence that summarizes your main point.
  - c. Be structured according to the rules for paragraphs learned in previous weeks.
  - d. Identify at least one lesson about critical reading that you have learned as a result of the in-class exercise. Explain how you learned this lesson.
  - e. Give evidence from the assigned essays and your own work to support your argument.

5. You will also upload a scanned (or photographed) copy of the first four pages of your in-class exercise on the LMS in .pdf format.

**Essay 1 (repeat this page for Essays 2+3)**

Main Argument:

---

---

---

---

Evidence:

---

---

---

---

Strengths:

---

---

---

---

Weaknesses:

---

---

---

---

Critical Assessment:

---

---

---

---

# Creative Explorations

## [Cyborg Texts: A Procedural Creativity Assignment](#)

Jason Boyd  
Toronto Metropolitan University

This assignment engages undergraduate students in the hands-on practice of procedural creativity through playing with, ‘hacking,’ or building text generators that produce creative outputs. Students are asked to draw upon the material covered in lectures and their own experience as procedural creators to reflect upon ideas of creativity, authorship, and potential futures for the literary in a digital age.

## [Spellcraft & Translation: Conjuring with AI](#)

Dana LeTriece Calhoun  
University of Pittsburgh

This assignment invites undergraduate students to create "spells," a poem based on a wish, and prompt an LLM to write a spell of its own, then reflect on their input and LLM output compared to their original compositions. Students analyzed the function of language and intent in manual and LLM composition, drawing together meaning and expression, and how LLM technology replicates or revises that expression. This assignment can be adapted to most rhetoric, composition, and cultural studies courses.

## [Made Not Only By Me: Coauthoring a Children's Book with Text and Image Generation](#)

Brandee Easter  
York University

This assignment asks undergraduate students to explore image and text generation technologies to create a short, illustrated children's book. Although text and image generation technologies are different, experimenting with them in parallel challenges students to reflect critically on the co-constitutive relationship between writing and technology. This approach was particularly useful for making "visible" how technologies can participate in, create, and sustain biases.

## [cmpttnl cnstrnt: An Exercise in Constraint and Prompt Engineering](#)

Douglas Luman  
Allegheny College

As new context-aware generative models challenge the human relationship to language, students benefit from first-hand observation of these models’ successes and limitations. Using these models often requires using “prompts” (natural language-based directions) to guide their output. The method of developing these directives has quasi-formalized into a practice known as “prompt engineering.” Serving as a gentle introduction to the intentionality, opportunities, and limits of the prompt engineering process, this work proposes and describes initial outcomes from an



assignment that uses similarities between model prompting and the constraint-based literary work of the Ouvroir de littérature potentielle (“Oulipo”) to focus student attention on precision and specificity of prompts and their execution. Beyond familiarizing students with contemporary technologies (particularly OpenAI’s GPT platform) and the nascent practices developing around them, this assignment also aims to give students first-hand experience with the reflexivity of using language to describe language in preparation for larger conversations about language as a technology and the roles of large language models (LLM) in human expression.

### **The Grand Exhibition of Prompts**

Mark C. Marino, University of Southern California  
Rob Wittig, University of Minnesota Duluth

The Grand Exhibition of Prompts, a netprov

In this collaborative creative writing game on a threaded discussion platform students experiment with AI image-making programs, but concentrate on the expressive, literary power of the verbal prompts they write. As language arts, image-making prompts are concentrated, evocative, use sentences, fragments, lists and non-sequiturs. In other words they are poetic (but don’t tell them that until later!). Students can be evaluated on their contributions, and on their support and encouragement of other writers.

### **Different Ways of Narrating with Curveship-js**

Nick Montfort  
Massachusetts Institute of Technology

Students are assigned to generate variations of an underlying story using a system designed for this purpose. The goals are (1) to use the core aspects of narrative, events and their representation; (2) to link narrative terms to concepts; and (3) to discover how events can be narrated in different, lively ways. The assignment requires instructor familiarity with my system, Curveship-js, but works well. Sharing and discussing the results enhances learning.

### **Deconstructing and Reconstructing Genre and Form with Tracery**

Mark Sample  
Davidson College

In this assignment, students work with HTML, CSS, Javascript, and JSON templates in order to design a website that generates new content out of pre-established rules and word banks. No prior coding or web development experience is required; the free browser-based platform Glitch.com hosts the projects templates as well as the projects themselves. The assignment encourages students to deconstruct the underlying rules, tropes, and conventions of any kind of textual genre. This assignment emerged out of an undergraduate course devoted to digital literature and poetry, but it can be adapted for many contexts, including any field concerned with form, style, and genre conventions.

### **Who's Talking: Dada, Machine Writing, and the Found**

kathy wu  
Brown Literary Arts

This assignment asks undergraduate students to generate text using both analog cut up techniques, as well as a simple Markov procedure, and discuss the power relations inherent in found writing processes. Through reading and making, students will encounter and critically develop their own articulations of found art—its questions of property and power—and how it relates to generative text and its corpuses.

# Cyborg Texts: A Procedural Creativity Assignment

Jason Boyd  
Toronto Metropolitan University

*This assignment engages undergraduate students in the hands-on practice of procedural creativity through playing with, ‘hacking,’ or building text generators that produce creative outputs. Students are asked to draw upon the material covered in lectures and their own experience as procedural creators to reflect upon ideas of creativity, authorship, and potential futures for the literary in a digital age.*

---

*Learning Goals:*

- Illuminate the long pre-digital history of procedural creativity
- Integrate procedural thinking and technique into creative writing practice

*Original Assignment Context:* Beginning of upper-level elective undergraduate English course

*Materials Needed:* Various web resources for each student-driven option, see assignment in depth

*Time Frame:* ~4 weeks

---

## Introduction

“Narrative in a Digital Age” (ENG921) is an elective course in the English undergraduate curriculum at Toronto Metropolitan University in Toronto, Canada (formerly named Ryerson University). It is normally taken by students in their third or fourth year of study. The synopsis of the course as provided in the syllabus reads in part:

This course explores the impact of digital technologies on understandings and practices of narrative or storytelling, examining how these technologies are changing the scope, definition, and ways of creating and experiencing the ‘literary.’ As part of this exploration, we examine not only digital works, but also exemplary print-based precursors and analogues to these digital works, as well as scholarship on this creative field. The course will focus on three broad categories of creative digital work: 1) Writing Machines, focused on the intersection of the literary with digital formats, computer programs and programming idioms; 2) Electronic Literature, particularly Hypertext and Hypermedia, which makes use of hyperlinking to create various kinds of pathways for the user to choose and explore, and Interactive Fiction (IF), which parses text-based input from the reader to construct a story; and 3) Digital Narrative Games, which examines the challenges and opportunities that video games present for creating new forms of interactive stories.

The first assignment of the course is called “Cyborg Texts.” This assignment has undergone almost yearly revisions and refinements since 2012, often in response to new opportunities or challenges: for example, in 2015, I included an option to make a physical ‘writing machine,’ taking advantage of my university’s new maker space, the Digital Media Experience Lab (this option had to be removed when the

COVID-19 pandemic shifted learning online); and the “paper program” I discuss below was an option I included for the first time in the Winter 2023 course offering. The “Cyborg Texts” assignment is designed to assess students’ comprehension of the content of the first section of the course, “Writing Machines,” which comprises Week 2 (Procedural Text Generation), Week 3 (Procedural Poetry), and Week 4 (Writing Within Digital Forms and Codework). In these weeks, students are introduced to the long history of *procedural creativity*, from ancient divination practices such as cleromancy, to Panini’s Sanskrit-generating *Ashtadhyayi* (ca. 500 BCE), to Tristan Tzara’s Dada poem ‘recipe,’ the potential literature of OuLiPo, William S. Burroughs’ cut-ups, Jackson Mac Low’s Diastic technique, as well as fictions and essays that address text generation/deformance by Jonathan Swift (Part III, Chapter V of *Gulliver’s Travels*, describing the Grand Academy of Lagado), Jorge Luis Borges (“The Total Library”), Stanislaw Lem (“U-Write-It” from *A Perfect Balance*), and Italo Calvino (*The Castle of Crossed Destinies*, the stories of which were created by laying out a grid of tarot cards, and constructing stories from each row and column). Students are required to explore text generators/deformers that use ngrams and Markov chaining, digital implementations of OuLiPo’s N+7 and Mac Low’s Diastic procedures, as well as works by procedural artists such as Allison Parrish, Nick Montfort, and Aaron Tucker. Charles O. Hartman’s critical memoir *The Virtual Muse: Experiments in Computer Poetry* (1996) is also discussed, as it provides an illuminating glimpse into the artistic philosophy and practices of a procedural poet.

The “Cyborg Texts” assignment assesses student comprehension of this material through hands-on engagement and reflection: as the instructions state, “This assignment asks you to explore – to play with – text generators/deformers to discover what kinds of interesting outputs you can create as a ‘cyborg author’.” The assignment gives students a number of options: they can generate poems from chess games using Aaron Tucker’s Chessbard and then use the *eDiastic* machine to create a Diastic version of these Chessbard poems; they can hack one of Nick Montfort’s poetry-making programs; they can create a Twitter bot using Tracery; or they can create a ‘paper program’ for generating stories—a one-page solo table-top role-playing game (TTRPG) that provides writing prompts through a randomized procedure. For all these options, the assignment requires students to explain what they did and show what they created and how. Whatever option is chosen, every student has to write a short reflection on what they learned about procedural creativity and authorship as ‘cyborg authors.’

To assist students in completing this and subsequent assignments, the final hour of each class is devoted to a ‘Literary Platforms’ workshop, which offers demonstrations of the various tools and strategies that can be used for the assignments. The first workshop, in Week 2, is “Paper Programs: Solo TTRPGs as Procedural Story Generators.” We examine a one-page solo TTRPG (in this instance, Mark Cook’s *Oubliette* [2021], <https://marccook.itch.io/oubliette>), and discuss how it is constructed, as well as which aspects of it could be improved to optimize the rule-based narrative generation. Week 3’s workshop is “Hacking Text Generators,” where I walk students through the salient features of a number of the HTML/JavaScript/CSS programs available on Nick Montfort’s site and show how they can make modifications to the lists or arrays of words and phrases that constitute the ‘raw material’ from which texts are generated, as well as ‘hack’ the programs by copying and modifying some minor existing elements. The objective of the “Writing Machines” section of course, as well as the “Cyborg Texts” assignment, is to show students that procedural creativity has a long pre-digital history that creative uses of the computer have built upon, to disrupt their usually simplistic Romantic ideas that the creative process consists only of capturing in any

given medium a Wordsworthian “spontaneous overflow of powerful feelings” (*Lyrical Ballads*), and to provide them the opportunity to experience procedural creativity first-hand (which prepares them for the major assignment of the course, the design and creation of a ‘playable story’).

In regard to the “Literary Platforms” workshops in the first part of the course, the instructor needs enough familiarity with the basic mechanics of solo TTRPGs (itch.io is a great source for example TTRPGs; see, for example, <https://itch.io/physical-games/tag-solo-rpg>), as well as the basic syntax and structures of HTML, JavaScript, and CSS (Cascading Style Sheets) to be able to parse the sample programs (in this case, from Nick Montfort’s website). For this, the tutorials and reference documentation at [www.w3schools.com](http://www.w3schools.com) are extremely useful both for the instructor and for students. While some of Montfort’s HTML-implemented programs can be fairly complex, the instructor only needs enough knowledge to show students how they can hack (modify and supplement) these programs in a way that can meet the requirements of the assignment. This means that the instructor needs to prepare example hacks that can be demonstrated in the workshop, by using a code editor like the free Geany ([www.geany.org/](http://www.geany.org/)) or the browser-based editor at W3Schools, where code can be modified and then run to see the output (or an error message that needs fixing!). Besides supplying new words or phrases to the arrays or lists that these programs use to construct their outputs, example hacks are the addition of a ‘P.S.’ to the letters generated by the Strachey/Montfort *Love Letters*, and adding additional lines to the Knowles & Tenney/Monfort *A House of Dust*. The reason why Montfort’s *Taroko Gorge* is a useful example is that many of the hacks by other procedural artists that are linked to the original produce interesting outputs without major changes to the code (the JavaScript), which can reassure students who are fearful of their ability to work with code. An instructor using this assignment also needs, for the week four workshop, to know the basic syntax and grammars of Tracery, and how to use them in the Tracery editor (I have been fortunate to have this workshop taught for the past few years by an alumnus of the course, Ewan Matthews (Twitter: @ThePringularity), who provides students with a fascinating overview of his and others’ “Digital Narratives for Twitter”). In sum, this assignment works best when the instructor has some experience in *creating* using these tools, and can show students work they themselves have created following the assignment’s instructions.

Students enjoy undertaking this assignment: they find it unusual, fun, challenging, and thought-provoking, judging from the reflection portion of the assignment. Working within the various sets of constraints enables them to experience aspects of or approaches to creativity that many have not encountered before. In their reflections, there is usually a spectrum of responses regarding the value of this type of creativity, but these responses, due to the hands-on aspect of the assignment, demonstrate a thoughtfulness and critical nuance that is often lacking in more traditional assignments that ask students to take a critical stance on creative works they understand often only at a very abstract level. While there are always students who proceed mechanically through the instructions and produce outputs that demonstrate very little deliberation (critical or creative), many students fully enter into the spirit of the assignment, carefully shaping and molding their raw Chessboard deformation into poems of haunting suggestiveness, creating hacks of text generating programs that reveal untapped potential (one favourite is Ewan Matthews’ hack of the Strachey/Montfort *Love Letters*, which turned the intentionally insipid and mawkish outputs of the original program into letters of alarmingly covetous intensity addressed by ‘The Beholder’ to a variety of precious gems), constructing Twitterbots that generate outputs of surprising complexity and variety, and designing one-page solo TTRPGs that can prompt moving or hilarious narratives written by the player.

However, ultimately in this assignment I am not assessing students for the aesthetic value or quality of their creations, but on what they learnt by being ‘cyborg authors,’ and how this informs the reflection portion of the assignment.

---

## The Assignment

### Cyborg Texts

This assignment asks you to explore—to play with—text generators/deformers to discover what kind of interesting outputs you can create as a "cyborg author."

For Parts 1 to 4 below, select and complete *one*. Additionally, complete Part 5.

Upload the completed assignment on D2L (in the Assignments section) as a MS Word file. Give the file the title "LastNameFirstName\_Cyborg" (including the file extension) (e.g. BoydJason\_Cyborg.docx).

#### 1. A Chessbard Translation and Its Deformance

a) Go to <https://www.pgnmentor.com/files.html> and choose a player (in the ‘Players’ section) or a Tournament (under ‘Events’—try an early player/tournament!). Once you have made your choice, click on ‘View.’ Choose at least three games that have at least one player in common (note that there is a pop-down list above the chessboard where you can select matches). Cut and paste the pgnmentor.com URL for the player/tournament into a MS Word file, *and* the notation for the chosen matches, making sure to include the numbers of the chosen matches as indicated in the pop-down list.

b) Cut and paste the notation below the chessboard for your selected matches into the Chessbard Translator (<http://chesspoetry.com/translate/>). You can paste the notation from multiple games into a single Translator window. Click the ‘Poetify’ button. Cut and paste the generated poems into your MS Word file.

c) Go to the *eDiastic* machine (<http://www.eddeaddad.net/eDiastic/>) and cut and paste the Chessbard Translator-generated poems into the ‘Input Text’ window (you may remove the ‘White Poem’ and ‘Black Poem’ headings). For the ‘Seed Text,’ try using a combination of the game notation, the players’ names, and/or a quotation about chess (such as can be found here: <https://kidadl.com/quotes/best-chess-quotes-for-the-chess-master-in-you>). Play around with *eDiastic* settings until you get a satisfactory Output Text. Remember to cut and paste what your ‘Seed Text’ was into your MS Word file and note what settings you used in the generation of the Output Text.

d) Cut and paste your final Output Text into your document. Edit the Output Text (e.g., add, delete, punctuate, reline) to refine or polish it, give it a title, and subtitle it ‘An Chessbard eDiastic Poem’.

e) Submit as one MS Word file. This will contain:

- The pgnmentor.com URL for the player/tournament and the notation for the chosen matches (with the numbers, prefaced with the # symbol, of the chosen matches as indicated in the pop-down list);
- The White and Black poems generated by the Chessbard Translator;

- The *eDiastic* Output Text and a 'Note' indicating the 'Seed Text' and which *eDiastic* settings were used;
- The edited version of the Output Text.

## 2. Hacking a Poetic Machine

This is an opportunity to incorporate hacking as a method of creating cyborg texts. Essentially, this involves the modification of a computer program designed to generate text.

*Instructions:*

- Modify/hack a text generation program from Nick Montfort's website. Check out the "Computational Poems" section, <http://nickm.com/poems/>: the works in the "Extra-Small" section are probably the best to choose from, since they are all self-contained: recommended works are *Taroko Gorge*, *Lede*, and *Through the Park*. Another suggestion is a couple of the works in the "Memory Slam" section, <http://nickm.com/memslam/index.html>: *Love Letters* or *A House of Dust*. Use a text editor designed to write/edit code to modify and run the hack:

For Python versions: *Wing 101 Python IDE*

(<https://wingware.com/downloads/wing-101>)

For HTML/Javascript (and Python) versions: *Geany*

(<https://www.geany.org/>)

You can also use the online editors available at W3C Schools:

(<https://www.w3schools.com/tryit/>)

- Explain the changes you made to the original program (add these to the program as comments) and take screenshots of these comments and the modified code, and of the output when the program is run. Submit as one MS Word file.

## 3. Twitter Bot

Create a Twitter Bot modelled on the ones we have looked at in class. *Cheap Bots, Done Quick!* is the easiest way to do this, but it is recommended that you keep a copy of your Tracery code in a separate .txt file as a backup. Submit the .txt backup file along with screenshots of the output of the bot (on Twitter). Submit the code and screenshots as one MS Word file.

*Example Twitter Bots:*

- @happyendingbot
- @MagicRealismBot

Allison Parrish's list of Twitter Bots: <https://twitter.com/aparrish/lists/my-bots/members>

[**Note:** These bots are provided in addition to the Twitter Bots assigned in Week 2: @Fairy\_Fables, @cutup\_bot, @eventuallybot, and @str\_voyage]



Some resources:

- *Cheap Bots, Done Quick!*: <https://cheapbotsdonequick.com/>
- Kate Hammond, *Tracery* Visual Editor: <http://brightspiral.com/tracery/>
- Kate Hammond, *Tracery* tutorial: <http://www.crystalcodepalace.com/traceryTut.html>
- Adam Hammond, “A Total Beginner’s Guide to (Literary) Twitter Bots”: <http://www.adamhammond.com/botguide/>
- Allison Parrish, “Twitter Bots” (tutorial): <http://air.decontextualize.com/twitterbot/>

If you do not want to (or cannot) use Twitter, you can instead take screenshots of what you have created using the Tracery editor (<http://tracery.io/editor/>): the window in the lower right will generate up to 100 outputs.

#### 4. One-page solo TTRPG

As we have discussed, procedural creativity pre-dates and is not exclusive to computing technology. Create a one-page solo tabletop role playing game (TTRPG):

1. Make the page’s layout and design attractive and in keeping with the theme.
2. Use only a 6-sided die (d6) and a standard deck of playing cards as your randomizers.
3. On the back of your TTRPG one-pager, include a sample playthrough
4. Submit as one MS Word or PDF file.

*TTRPG resources:*

- Online dice roller: <https://www.google.com/search?q=dice+roller>
- Online playing card drawer: <https://www.random.org/playing-cards/>

#### 5. Reflection

Based on your experience as a cyborg writer and what has been discussed in class about procedural text generation, write a brief (250-300 word) reflection on your experience of procedural creativity. How does this creative process differ from other creative processes? Who is the author of these works (both the processes and the outputs)? Do cyborg texts have potential for opening up "new horizons for the literary" (N. Katherine Hayles)? *Make sure your response is informed by and explicitly refers to material covered in lecture (consult the lecture slides on D2L if necessary).*

### Cyborg Texts: Assessment Rubric

#### 1. A Chessbard Translation and Its Deformance

The submission should include:

- a. The pgnmentor.com URL for the player/tournament and the notation for the chosen matches (with the numbers, prefaced with the # symbol, of the chosen matches as indicated in the pop-down list), the White and Black poems generated from these matches by the Chessbard Translator, the *eDiastic* Output Text and a ‘Note’ indicating the ‘Seed Text’ and which *eDiastic* settings were used. (3 marks)



- b. A poem demonstrating intentional editing of the *eDiastic* Output Text to achieve an artistic effect (compare with the unedited version). Charles O. Hartman's practice (described in *The Virtual Muse*) of editing the outputs from his *Prose* program is a model to keep in mind. Look for an original title, deliberately rearranged lines and stanzas, added punctuation, added/deleted words, etc. Look for attempts to 'clarify' or bring out meanings in the Output Text, rather than attempts to overcredit and obliterate the Output Text—there should remain a recognizable connection between the two texts. (7 marks)

(10 marks total)

## 2. Hacking a Poetic Machine

The submission should include:

- a. The name/source of the program used and a clear description of how the program was hacked. Marks should reflect the skill/extent of the hack: there should be a clear and coherent artistic intent in, for example, the arrays or lists of words/phrases that are substituted for the originals. (5 marks)
- b. Screenshot(s) of the program providing concrete evidence of the hack. (2 marks)
- c. Screenshots of sample outputs (at least three) of the hacked version showing the range of variety and lack of excessive repetition. (3 marks)

(10 marks total)

## 3. Twitter Bot

The submission should include:

- a. Tracery (or other) code: Does the code demonstrate an evident premise/theme/scenario? Does it exhibit careful and intentional design and robust content in the arrays? Does it have a fair degree of complexity, ensuring variety of output? Use the Tracery editor's built-in sample outputs (<http://tracery.io/editor/>) as a guide. (5 marks)
- b. Sample output (10-12 'tweets'): assess for variety; comprehensibility (including grammar); quirky/unusual/evocative combinations. (5 marks)

(10 marks total)

## 4. One-page solo TTRPG

The submission should include:

- a. Instructions that use *both* a 6-sided die (d6) and a standard deck of playing cards (and nothing else) as the randomizing component of the TTRPG. (1 mark)
- b. Instructions that fit comfortably on one page; the page's layout and design is easy to read, visually appealing, and in keeping with the scenario. (1 mark)
- c. A full sample playthrough that demonstrates the playability of the TTRPG. (2 marks)
- d. A scenario that balances open-endedness and constraint; the scenario/situation is easy to grasp without needing a lengthy textual introduction; the scope of possible action can be easily determined; provided non-player characters have specific and recognizable roles [The sample TTRPG we look at in class, Mark

- Cook's *Oubliette*, does this well: it uses a recognizable and regimented space whose actors have defined roles—a prison—but leaves unspecified who the player is, why they are in prison, and what kind of world exists outside the prison]. (3 marks)
- e. Prompts that can bear repeated use in a single playthrough; prompts that work well in random combination. (2 marks)
- f. A clear mechanism for ending a playthrough: is this constructed in such a way that it provides closure but minimizes abrupt (premature) endings? (1 mark)

(10 marks total)

## 5. Reflection

Overall, look for a recognition of the complexity of using computers/algorithms/formal procedures for creativity—a balanced understanding about the value and limitations of procedural text generation for creative writing instead of an unnuanced either/or judgement—which references the student's own experiences in undertaking the assignment.

Some looked for insights include (these will vary according to what ideas and readings are covered in the course lectures):

- Procedural text generators
  - can provide text of a great variety and quantity, in combinations that a human would not likely consciously or intentionally create;
  - can create intriguing 'nonsense' text, which can range from the humorous to the political/philosophical, to the 'oracular';
  - are more effective at generating poetic rather than narrative outputs;
  - are more exploratory/experimental machines rather than literature generators (which would need a sophisticated AI, or Artificial Intelligence);
- Procedurally generated text
  - can sometimes seem meaningful although their meaning remains unclear/obscure;
  - without human intervention, rarely rise to the level of the literary or poetic as conventionally conceived;
- Text generators and their outputs challenge accepted notions of creativity and authorship (original genius);
- Text generators and their outputs can be used as a spur to creativity, or as the raw material on which to do creative work;
- A writer who reworks generated outputs can be seen as engaged in an act of collaboration, in effect creating a cyborg and therefore 'cyborg literature.'

(10 marks total)

# Spellcraft & Translation: Conjuring with AI

Dana LeTriece Calhoun  
University of Pittsburgh

*This assignment invites undergraduate students to create "spells," a poem based on a wish, and prompt an LLM to write a spell of its own, then reflect on their input and LLM output compared to their original compositions. Students analyzed the function of language and intent in manual and LLM composition, drawing together meaning and expression, and how LLM technology replicates or revises that expression. This assignment can be adapted to most rhetoric, composition, and cultural studies courses.*

---

*Learning Goals:*

- Analyze the function of language and intent in manual and AI composition
- Understand the implications of the social and cultural impact on AI and similar digital technological composition
- Compose and craft a digital project relating to Black rhetoric and digital writing.

*Original Assignment Context:* End of elective Public and Professional Writing course

*Materials Needed:* An accessible AI text generation program (i.e. ChatGPT)

*Time Frame:* ~2-3 weeks

---

## Introduction

This assignment, “Spellcraft & Translation: Conjuring with AI” is a part of the Black Rhetoric and Public Writing course taught at the University of Pittsburgh in Spring 2023. This course is situated within the Public and Professional Writing concentration and Diversity in General Education requirement for the university curriculum.

This assignment is placed alongside theoretical and practical work discussing the intersections of Black rhetoric and language with digital hybrid models of analysis and composition. Conjuring hails from Hoodoo, a Black Southern American spiritual and ontological practice first envisioned by enslaved people and manifesting now across a multitude of geographic locations, experience, and identity. In this course, Conjuring acts as the methodological model to ask students to blend the natural, the self, the culture, and technology in scholarly analysis and composition in Black rhetorics. I have taught this assignment once and plan to teach it again in similar courses in the future.

Spellcraft & Translation asks students to compose a comparative analysis between two originally-composed “spells” of their own and spells composed by the AI writing platform ChatGPT. Spells, in this case, are defined as technical instructional documents that convey an intention, purpose, and direction. The assignment before asked students to create a “Mojo Bag,” a repository of tools and objects central to the creator that enhance their work. In traditional Conjuring practice, Mojo Bags are

mobile “toolkits” holding personalized items and key ingredients for their work. This assignment was also produced digitally—I wanted to give students the opportunity to display abstract ingredients, like the overall “vibe” of a song that would help them conjure their work. In their bags, students included a variety of tools and objects and described their meaning: a dice, to help them come to a decision about what to write; moss agate, to encourage emotional stability; a poker chip, to represent the act of weighing multiple decisions; and a playlist of songs that encourage self-love, confidence, and power.

ChatGPT is an online AI writing platform that can compose text based upon prompts entered by the user, although any similar generative AI platform would work for this assignment. For this assignment, I was curious about the capability of ChatGPT to generate a text with similar theme, feeling, and composition as students had. From the “Mojo Bag” assignment, I already knew what personal themes or wishes would manifest in their spells: general anxiety about being an undergraduate student, balancing responsibility and relationships, and finding their identity and individuality. In previous classes, I have used spells and poems as a teaching text, and encouraged students to create their own should they feel motivated to as a writing exercise. The idea for using ChatGPT to compose a spell came from a desire to connect Black rhetoric, genre writing like poetry and recipes, and AI writing technology.

The inspiration for Spellcraft & Translation came from my mentor and colleague, Annette Vee, encouraging me to look at LLM technology in relation to my personal scholarly work with Technogothic Conjuring, a methodology I developed while writing my dissertation. Technogothic Conjuring is a methodology taught by Black Southern Hoodoo spiritual and cultural practices and made unique and deeply personal by the blending of digital compositional and archival methods that focus on the identity of the subject and researcher. Conjuring hails from the African American “embodied historical memory” tradition of Hoodoo. Other names include: rootwork, black magic, witchcraft, devil’s work, superstition, spirit work, and conjure.

The “Technogoth” in my methodological phrase deals with the approach towards composition while acting as a descriptor for the practice of blending the technological with the intense flood of emotions that embody the “gothic.” More specifically, the Technogothic approach explores the combination of digital composition and communing with the dead. The Technogothic approach is housed in autoethnography, in that the researcher and the research itself are not separate, but placed in different planes of medium, physicality, and temporality—the “Conjuring.”

My dissertation involved transcribing six months of my great-grandmother’s handwritten everyday journals, and conjuring her spirit and memories by displaying her history on a website. All of the transcriptions and data were done by hand, and one chapter is entirely composed of original poetic remixes of my great-grandmother’s journal entries. After doing all this work myself, I wondered how LLM technology would influence the process of collecting and composing the sheer amount of data available and perhaps inspire a new form of output.

I’ve often discussed and encouraged poetic writing in my previous courses, and students have responded positively, saying that creative writing forms allow them to express emotions, deep thoughts, and opinions more freely and that doing so helps them begin to draft more traditional writing forms. With the Black Rhetoric course Spellcraft & Translation was housed in, I was curious to see how student composition compared to LLM compositions and how students would react and

reflect on this process, especially considering the boost in discourse about LLM technology in academia.

I feel that this assignment is able to be transformed and conjured in its own way to fit the needs of a specific course or module simply because AI writing platforms provide a third-party composer that acts as an amalgamation of all types of writing. As a pedagogical tool, AI writing platforms can act as a “shadow writer,” creating examples to analyze in real-time, to either highlight the technical or contextual elements.

Students displayed their comparative analysis on a digital Notion page, including a descriptive analysis of their process of composing their spells and prompting ChatGPT to compose spells of its own and a reflection of said process. Notion is a digital workspace that I’ve found to be incredibly useful for planning and composing digital projects, but any online space that allows students to display their work in different media would suffice. Platforms I have used in the past and encourage students to use have been: WordPress, Wix, Tumblr, and GitHub. Alternative platforms for short-form writing or writing that depends on hyperlinked context include Twitter and Medium.

The purpose of this assignment is to introduce multiple rhetorical concepts such as genre, comparative analysis, poetics, instructional technical writing, knowledge and skill in working with AI writing technology, and sociocultural contexts. The scaffolding for this assignment included students engaging with external readings (included in the assignment description) through a workshop and discussion-based format where students brought in their materials and as a class we discussed their wider implications.

### **Materials**

This assignment could be accomplished physically, textually, or digitally, with this specific iteration being the latter using Notion.

The majority of the preparation and materials for this assignment are digitally involved, including knowing how to use ChatGPT and knowledge on how to create effective prompts. This assignment will also involve an understanding how Notion works to display digital comparative analysis projects as well, if presented as a digital project.

Theoretically, the preparation for this assignment involves readings surrounding Black rhetoric and language; digital composition and AI writing technology; topics of race, gender, and technology; and how aspects of Conjuring relate to critical and scholarly composition.

I can imagine, and deeply encourage, revisions and remixes of the main idea behind this assignment—a comparative analysis between two texts. The inclusion of ChatGPT as an interface that students can use to “conjure” a text is an addition of my own curiosity and excitement of the promise AI writing tools have in the classroom.

### **Assignment Preparation**

At first, my students were cautiously excited about working with ChatGPT, having seen numerous reports about cultural anxiety surrounding AI writing. At first we played around with developing silly prompts such as “Write a haiku in the style of Eminem.” We then moved forward to more applicable prompts for this assignment

In our discussion of sourcing the ingredients for the spell, looking deeply at literary and contemporary examples of rhetorical spellwork, including select readings from non-fiction examples like *Memory Dishes: Women and African Diasporic Cooking* and fictional examples from *Beloved* by Toni Morrison. We discussed the elements listed in each spell, like chamomile, lavender, and the color pink in *Beloved*. In class we went through a research exercise finding the traditional applications of these ingredients. I wanted to make the connection between recipes and spellwork through these exercises, to call attention to rhetorical form and construction while situating this form of composition in a sociocultural context. Students immediately understood the connection, and quickly introduced more contemporary examples like the mobile phone acting as a multi-tool or knitting needles to weave things together.

Working with ChatGPT was an exercise in understanding genre, rhetorical emphasis, and the limitations of AI composition. On a pedagogical note, this exercise was a joy to do and brought fun into the class as well as understanding. We began by opening OpenAI's ChatGPT and prompting it to compose summaries of popular films. Students offered to prompt ChatGPT with more specific language, such as including instructions about using "formal" or "slang" language depending on the summary. We explored the limitations and collated imagination that ChatGPT was able to produce, though the general conclusion was that the AI-composed writing was too "stilted" or "obviously not organic." The students reasoned that while most of the AI compositions made sense, they failed to see a "spark" within the writing that made it rhetorically effective.

I was delighted at this response about not seeing a "spark," and we discussed how knowing the source of the output influenced their observations and conclusions. Students brought up a great point that I will add to future iterations of this assignment, that I should provide multiple examples of a piece of writing and ask future students to see if they can choose which was written by a LLM.

Moving towards working with spells and ChatGPT, we prompted it to create spells to obtain good grades, find a new friend, or pass an exam. This exercise was conducted in class, and students were encouraged to play around with prompts and then share their output to a display screen so we could discuss ChatGPT's output. We continued in a workshop format, where students would analyze and offer suggestions on how to refine the prompt for a more accurate ChatGPT output by including an additional specification, like "include candles," or "restrict flame." Together we noted the differences in ChatGPT's output of these spell prompts, focusing our attention on ChatGPT's revisions or lack thereof.

Students were particularly amused by ChatGPT's statement of "As an AI language model, I cannot create real spells or encourage the practice of magic," which prompted discussion of the human labor that went behind creating this type of technology and this statement being an acknowledgement of the idea that a "real" spell cannot be artificially created.

Despite ChatGPT's disclaimer, it still composed a spell offering a list of ingredients and the process, though it became evident that most of the ingredients listed were the same across all example spells, along with the process. ChatGPT would offer reasoning behind its choice of ingredients with a generalized summary of the purpose. For example, ChatGPT states that "a green candle might be used for attracting abundance, while a pink candle might be used for attracting love." Class discussion shifted into researching the historical and cultural meaning behind colors and associations with candles and rituals. We ended this session by calling back to the literary and contemporary spell examples, imagining and drafting what we

imagined ChatGPT might create using these examples. The following class sessions were focused on drafting this assignment, working through questions, and further exploration of ChatGPT.

## **Student Examples**

### *Student 1*

Student 1 prompted ChatGPT for instructions on how to create a spell, which resulted in four steps: determine the purpose of the spell, choose the language of the spell, research the method and symbolism of the words, write out the spell in the chosen language, and practice speaking the spell aloud.

Student 1 followed ChatGPT's advice in creating their spell "Memoria Della Vittoria," meaning "In memory of the victories" in Italian. The wish is for the Italian national soccer team to win the World Cup. The spell itself calls for a cauldron, a hair of the last goal-scorer of the Italian national team, a pinch of grass from a sacred soccer field, an Italian national team jersey, and a tear shed from not qualifying two years ago. All ingredients are to be mixed in the cauldron, where the fumes will disperse into the air. Then chant "Memoria Della Vittoria" three times and wait for the wish to come true.

Student 1's reflection to this assignment addressed the composition of "Memoria Della Vittoria," citing their Italian heritage and enthusiasm for Italy's soccer team to win the World Cup. The wish itself may seem simple, but the specificity in its ingredient belies a greater depth. Each ingredient is timeless, hinting toward the history of the Italian national soccer team and its fans.

This example showcases the act of conjuration not only in its content but in the relationship between this writing, this student, and the wider world. I have the privilege of being able to know what the Italian national soccer team means to this student, because they wear the team's jerseys in class almost every week. They sneak peeks at their phone to check the match scores during our writing breaks. They engage in class conversations making connections between their identity and rhetoric and writing. This spell combines writing and passion into an act of acknowledgement and engagement. It is a blend of the natural, the self, the culture, and technology, fully in line with the methodological approach to this course.

Student 1 prompted ChatGPT to "create a spell," and the response was to create a spell for positivity and good vibes. ChatGPT's spell included a list of ingredients, including a white candle, a small glass jar, a handful of dried lavender flowers, a piece of paper and a pen, and a lighter or matches. ChatGPT's spell then included the process for completing this spell. The process was to set up the ingredients and light the candle, then to take the piece of paper and write down the purpose of the spell and to "use positive language." Then it instructs to fold the piece of paper and place it in the glass jar, sprinkle the dried lavender flowers on top while visualizing being surrounded by a "peaceful aura of light," then say:

Spirit of the universe, hear my plea,

Bring me the positivity I need.

Fill my life with good vibes and grace,

And keep me in a peaceful and happy space.

ChatGPT then instructs to light the paper on fire using the candle, but advises caution and suggests using a fire-safe container if necessary, then to blow out the candle and place the glass in a prominent place. The conclusion of ChatGPT's spell is a disclaimer that "this spell is just for fun and entertainment purposes, and it is not intended to replace action or efforts to improve your life."

During class sessions experimenting with ChatGPT, students often remarked on the disclaimers included in ChatGPT's output, like the warnings about safety and spells being for entertainment purposes mentioned above. I recalled previous class discussions about authenticity in language and writing, and how Conjuring's history includes deep discussions about Conjuring fraud that could lead to negative results or injury due to ignorance about ingredients or process. Students connected authenticity and responsibility with ChatGPT's disclaimers, and wondered if these disclaimer statements or statements of safety were manually included by ChatGPT developers in order to mitigate responsibility.

This student described wanting a "straightforward approach" to create a spell, and questioned ChatGPT's ability to read their browser's history, saying that ChatGPT's spell could have been influenced by their past search history. ChatGPT does not have the access or ability to read a user's browser history, but I find this suspicion fascinating. As an alternative version of this project, instructors could discuss privacy and data settings, surveillance technologies, or power through technological oppression.

Student 1 also questioned the language that ChatGPT used, remarking about its formality. This student wondered "If I use very informal language with it, will it emulate how I speak?"

During a session where students and I were prompting ChatGPT, they requested to see ChatGPT's output when prompted to give a summary of the movie *Shrek* with various variations on the form of the summary and the type of language used, ending with "Write a summary of Shrek, use slang."

ChatGPT's output was a disjointed summary of *Shrek* beginning with, "Yo, lemme tell ya 'bout dis flick called Shrek." The students were shocked, I was amused, and together we discussed *how* ChatGPT might have come up with this input. Since this course focused on Black Rhetoric, students spoke to our readings for this assignment. They noted the specific dialect used in Toni Morrison and Alice Walker's novels, particularly in how they highlight features of African American Vernacular English (AAVE, another term used is African American English).

In classes before, we discussed how Morrison and Walker's use of dialect in their works is a rhetorical move, one that challenges standardized language use. During this discussion, students connected standardized language use to formality and informality in writing and speech. That conversation led to a discussion of how formality is established—"Clothes, signing emails professionally, not using slang" and how Morrison and Walker's work challenges the idea of formality in speech and writing.

I guided the conversation towards their prompt of slang, asking them where some current slang terms come from. They referred to fashionable or stylish choice as "drip," and "snatched," which is when someone's appearance is perfect—so good you snatch your wig off, a term and phrase popularized by the Black drag community and further pushed into the cultural lexicon by television show *Ru Paul's Drag Race*. Students brought up the history of "snatched" without any prompting, and they



named TikTok as a popular source for their current language. TikTok is a social media platform where users upload video clips, and students noted that their “personalized” TikTok feeds featured creators offering commentary on popular culture, *Ru Paul’s Drag Race* being one of them. Because of these videos, students were mostly already aware of the Black queer and drag community’s historical influence on a term like “snatched.” We discussed how language shifts and evolves from communities into popular usage and how that relates to language appropriation.

After ChatGPT’s output, students connected Walker and Morrison’s dialect with slang terms and appropriation, saying that it was “creepy” and “strange” that ChatGPT affected a similar dialect to Walker and Morrison’s writing. They brought up how contemporary celebrities like Ariana Grande and Awkwafina adopted a “blaccent,” an imitation of AAVE without the racial and cultural background as a form of cultural appropriation.

We then had a brief exercise where we listed common slang phrases and followed the etymological thread back to the communities that popularized these terms in order to create a sample of the terms a hypothetical AI writing platform would encounter if it was trained using popular sources. In most cases, common slang terms originated or became popular from Black and queer communities. Students discussed the possible sources that ChatGPT was trained with, and discussed how they were uncomfortable with the transparency between slang and appropriation.

### *Student 2*

(A small note: Student 2 used Notion AI, a new feature that debuted during the composition of this assignment—I had not anticipated Notion developing their own AI writing tool, but I encouraged all avenues of AI writing engagement for this assignment.)

Student 2 composed their spells in a poetic format, combining listing ingredients and process in verse form. This student prompted Notion AI to create a spell based off of the wish to not overthink and take things too personally. Student 2’s spell begins with lingering “in cloud of fog,” and then calling upon the wind, “To clear out the disarray,/ And summon the sun to appear and alas,/ Cast the prominent clouds away.” This spell manifests their ingredients and process in a more metaphysical, spiritual, supernatural and natural way, calling upon the speaker’s inherent powers to clear the weather to reflect the clearing of their negative thoughts.

Notion AI’s version of this spell begins with, “By the power of the universe and the strength of my will,/ I call upon the energies of calmness and still./ With this spell, I ask for the ability to let go,/ Of thoughts that bring me low.” This spell continues in a declarative way, ending with, “So mote it be,” a closing phrase used to end rituals or special rites.

Student 2 immediately noticed the formulaic composition of the AI model, describing them as “adjusting their wording and phrasing to incorporate the main asks of the wish,” noting that the formality in the AI spell makes finding the original source easier. Student 2 continued, saying that the AI spell felt like an “expected result,” though not particularly “robotic and lifeless,” which I feel adds to their observation of ChatGPT drawing from external sources that use the same poetic format and closing phrase “So mote it be,” commonly used in Wiccan prayers, spells, and rituals to mean “So shall it be.” ChatGPT’s output here adheres to a type of standardized format for a spell, and though the words pulled from other sources are emotional and evocative, the construction itself belies its uncanny composition.

Student 2 outlined their process and intention in creating their spell, saying that they chose a more figurative and illustrative approach to “create a better sense of vulnerability...and feed into the magic behind them.”

### **Goals and Outcomes**

The learning goals and outcomes for this assignment were to analyze the function of language and intent in manual and AI composition, understanding the implications of the social and cultural impact on AI and similar digital technological composition, and experience in composing and crafting a digital project relating to Black rhetoric and digital writing.

Conjuring with AI is an assignment with a multitude of possibilities, especially with the rapid development of AI writing technologies. Though only two student examples were discussed here, I found that both held common elements and conclusions across the classroom—that composition and rhetorical exploration hold space for magic in the way of imagination, intention, and action. We were able to have productive conversations about the relationships between technology and language, and specifically how AI writing technology displays that relationship.

Most students, in some way or another, addressed the uncanny lexicon of ChatGPT and Notion AI and described that though the language was adequate in achieving the form and general idea behind a spell, the results lacked an element of individuality and humanity. Surprisingly, students also were impressed at how ChatGPT and Notion AI were successful at helping them draft their reflections, providing a general outline of ideas and how to tie them together or provided a quick and clear answer to a clarifying question about an unfamiliar word or phrase the AI used.

I find the generative possibility of LLM technology inspiring as an opportunity to offer students a “playground” to test out prompts, lines of argument, and genre styles. I feel with guidance and support, LLM technology could reduce writing anxiety and writing block in students by allowing them to experiment with a flexible writing utility that can offer questions, generate a general outline, or even produce an example of a paragraph with the style and organization prompted by the user.

### **Conclusions**

Conjuring with AI was an experimental foray into weaving AI writing technology, rhetoric, and composition together. Overall, the resulting work from my students has given me more inspiration to continue to expand both the theoretical and practical application of this assignment. I’m especially interested in adding more readings that address writing, technology, race, and gender into the scaffolding to highlight human identity and labor that contributes to rhetoric and technology.

In future iterations of this assignment, I hope to expand on using AI writing platforms as a point of comparative analysis and compositional tool and introducing it earlier in the scaffolding of this assignment. I found that exploration and experimentation with AI writing platforms tended to be the hot topic of class discussion. I feel that if students became more comfortable with AI writing platforms, their focus could shift more to the rhetorical implications and practical applications, pushing us to consider the possibilities of this technology. I thoroughly enjoyed this assignment, and I am excited about its potential.

### **Acknowledgements**

I have a bouquet of flowers to give as thanks and recognition: Annette Vee, who has inspired this project by introducing me to ChatGPT and the amazing work that can be done with it along with being a patient and extraordinary mentor; Geoff Glover, for always finding the time to sit and riff, and who always meets my nontraditional ideas with excitement and support; and a trio of perennial blooms for my great-grandmother Mary Ellen Petty, who has walked with me through my scholarship; Laura Mae Freeman, my grandmother who is like my second author with her wisdom etched in my words; and my mother, Theressa Michelle Logan, who always said I would be a shining star. None of my work would be anywhere without you all. Thank you.

---

## The Assignment

### Overview

Genre is a set of classifications that define consistent elements shared by multiple texts, inside and outside of traditional literary forms. This assignment will ask you to explore the genre of a spell, compose and conjure two spells of your own, prompt an AI writing program to write its own spell, and then perform a comparative analysis between the two texts. Finally, you will reflect on this process of spell creation, AI conjuring, and how writing is shaped with intent and genre.

All components of this assignment, including your spells, ChatGPT's spells, and your reflection should be hosted on your personal Notion site under the link "Spellcraft & Translation."

### Spells

Spells are a specific genre that blend together intent, poetics, rhetoric, and technical writing. One of your spells should focus around your "wish" (a declarative statement that describes a desire you have). The focus of your second spell is up to you! Compose your spell including a list of ingredients and instructions for the implementation of your spell. Your ingredients **and** instructions may be metaphorical, figurative, creative, or concrete.

### ChatGPT

After you have constructed your spells, use the skills developed in class sessions to prompt ChatGPT to conjure two spells of its own. Then, compare and contrast your spells and ChatGPT's spells. In this comparison, you should analyze each spell in depth, discuss the possible purpose and meaning behind the spells created by ChatGPT in comparison to your own, and explore **how** ChatGPT's spells compare to your own. Below are questions that I would like you to consider in your analysis:

- What are the clear differences and similarities between the two?
- In your analysis, focus on the syntax (the order of words) and diction (word choice) of the AI spells.
- Where, do you feel, is the root of the conjuration of these spells?

### Reflection

As always, reflective decompression for your project can be a way to take a step back and take in the entirety of your labor. I would like you to reflect on your process in creating your spells. Below are the questions I would like you to answer, along with any additional thoughts you have:

- What was your original wish or intention?
- How did you illustrate that purpose towards inspiring change in your spell?
- Did you prefer a more instructional approach or a figurative one, and why?
- How did you come to create the prompts that created a successful ChatGPT spell?
- How many iterations of the same prompt did you go through, and what did you notice happening in the ChatGPT output?

## **Display**

Create a new Notion page titled “Spellcraft & Translation.” Display both of your original spells and the spells created by ChatGPT. When displaying the spells created by ChatGPT, include the series of prompts that you used.

## **Scaffolding**

Introduction to Genre (via music)  
Poetry and poetics as genre

To begin the scaffolding process for this assignment, I think it’s important to start with an overview of genre. In this iteration of this assignment, I began with situating the discussion of genre within music. I found this to be an effective starting point, as it encouraged students to offer their own examples of different music genres. We worked through their examples together, finding and notating the similar elements that make up a musical genre.

Then, we shifted focus to poetry and poetics, which I find to be one of the closest literary genres to a spell. Below are examples of poems used as examples in class to establish the genre, explore the purpose of intention and form, and provide a basis for analysis and comparison as practice for the larger assignment.

- [Examples of Poetry Genres: Major Styles Explained](#)
- [Genre, Glossary of Poetic Terms](#)
- [dear white america by Danez Smith](#)
- [For everyone who tried on the slipper before Cinderella by Ariana Brown](#)
- [Babylon Revisited by Amiri Baraka](#)
- [Dreams by Nikki Giovanni](#)

## **Recipes as Genre**

Next, I introduced the recipe as a genre, as it is very similar to traditional spells in terms of form with a list of ingredients and a description of the process. During this time, we discussed the elements that create the genre of the recipe. In the examples below, we discussed the historical context behind recipes in association with Black Rhetorics, and how this form holds more purpose and meaning aside from being a set of instructions to create a dish.

I asked students to find examples of recipes, and together we contrasted the writing in class to pull out differences in form, language, description, and instruction. Using

these recipes, we discussed how to recognize and analyze syntax, diction, and other technical elements of this type of writing.

- [A Recipe for Remembrance: Memory and Identity in African-American Women's Cookbooks by Rosalyn Collings Eves](#)
- [Memory Dishes: Women and African Diasporic Cooking](#)

### **Spells as Genre, Setting Intention, Sourcing Ingredients**

Below are examples of spells that range from literary text to a musical spell that aims to bridge the discussion between genre and spellwork. These examples offer a concrete version of a spell that students can analyze, understand, and possibly gain inspiration from. As a class, we have mapped out each of these spells in terms of intention, form, and performance. With these spell examples, we also defined how each has ingredients (whether physical or figurative) and how they are implemented in the instructions of each spell.

- [Killing Color by Charlotte Watson Sherman](#)
- [Big Water by Charlotte Watson Sherman](#)
- [Excerpt from Beloved by Toni Morrison](#)
- [Excerpt from The Color Purple by Alice Walker \(186-187\)](#)
- [Devil is Fine by Zeal & Ardor](#)

### **Introduction to ChatGPT**

Finally, we began to explore ChatGPT in terms of its creation, function, and uses. Below are articles describing ChatGPT and how other people have used them and its cultural and social impact. Alongside this, the class created prompts for ChatGPT and discussed how the software interpreted them in its output. By creating prompts together, we worked through finding a method for purposefully communicating with ChatGPT, especially in terms of tone (formal, informal, verse, etc.).

- [ChatGPT: The Weirdest Things People Ask AI To Solve](#)
- [ChatGPT Is Dumber Than You Think](#)
- ['Magic Avatar' App Lensa Generated Nudes From My Childhood Photos](#)

### **Bibliography**

Baraka, Amiri. "Babylon Revisited." *Poetry Foundation*, <https://www.poetryfoundation.org/poems/42559/babylon-revisited>.

Bogost, Ian. "ChatGPT Is Dumber Than You Think." *The Atlantic*, 7 Dec. 2022, <https://www.theatlantic.com/technology/archive/2022/12/chatgpt-openai-artificial-intelligence-writing-ethics/672386/>.

Brown, Ariana. "For Everyone Who Tried on the Slipper before Cinderella." *Poets.Org*, <https://poets.org/poem/everyone-who-tried-slipper-cinderella>.

Collings Eves, Rosalyn. "A Recipe for Remembrance: Memory and Identity in African-American Women's Cookbooks." *Rhetoric Review*, vol. 24, no. 3, Nov. 2009, pp. 280–97, [https://doi.org/10.1207/s15327981rr2403\\_3](https://doi.org/10.1207/s15327981rr2403_3).

Collins, Barry. "ChatGPT: The Weirdest Things People Ask AI To Solve." *Forbes*, 4 Mar. 2023, <https://www.forbes.com/sites/barrycollins/2023/03/04/chatgpt-the-weirdest-things-people-ask-ai-to-solve/?sh=742f6d8fe547>.

Giovanni, Nikki. "Dreams." *Poetry Foundation*, <https://www.poetryfoundation.org/poems/48224/dreams-56d229494e255>.

"Glossary of Poetic Terms." *Poetry Foundation*, <https://www.poetryfoundation.org/learn/glossary-terms/genre>.

Kittelstad, Kit. "Examples of Poetry Genres: Major Styles Explained." *Your Dictionary*, 9 July 2021, <https://examples.yourdictionary.com/examples-of-poems.html>.

*Memory Dishes: Women and African Diasporic Cooking*. Brown University, 29 May 2019, [https://simmonscenter.brown.edu/sites/default/files/pdfs/xMemoryDishes\\_DRAFT\\_190514.pdf](https://simmonscenter.brown.edu/sites/default/files/pdfs/xMemoryDishes_DRAFT_190514.pdf).

Sherman, Charlotte. *Killing Color*. 1st edition, Calyx Books, 1991.

Smith, Danez. *Dear White America*. <https://www.poetryfoundation.org/poems/150542/dear-white-america>.

Snow, Olivia. "'Magic Avatar' App Lensa Generated Nudes From My Childhood Photos." *Forbes*, 7 Dec. 2022, <https://www.wired.com/story/lensa-artificial-intelligence-csem/>.

Walker, Alice. *The Color Purple*. Harcourt, 1992.

Zeal & Ardor. *Devil Is Fine*. mp3, MVKA Music Ltd, 16 Dec. 2016.

# Made Not Only By Me: Coauthoring a Children's Book with Text and Image Generation

Brandee Easter  
York University

*This assignment asks undergraduate students to explore image and text generation technologies to create a short, illustrated children's book. Although text and image generation technologies are different, experimenting with them in parallel challenges students to reflect critically on the co-constitutive relationship between writing and technology. This approach was particularly useful for making "visible" how technologies can participate in, create, and sustain biases.*

---

## *Learning Goals:*

- Use experimentation with AI as a method for engaging with questions about writing, authorship, and technology
- Interrogate biases, risks, and consequences of text and image generation
- Reflect on the possibilities, limits, and ethical concerns with AI writing technologies

*Original Assignment Context:* End of large year-long, first-year Professional Writing course

*Materials Needed:* Accessible image and text generators for students to use

*Time Frame:* ~6 weeks

---

## **Introduction**

This assignment was inspired by the creation of (and backlash to) *Alice and Sparkle*: a 12-page children's book written by Ammaar Reshi, a design manager, in one weekend through a combination of ChatGPT and Midjourney. Like much other panicked, often rage-filled, social media discourse about these technologies, responses to *Alice and Sparkle* debated the boundaries, mediation, and status of writing, authorship, and originality—concepts central to “Introduction to Professional Writing,” a year-long, large-lecture, first-year course.

In a 6-week unit on digital writing in Spring 2023, approximately 120 students produced short, scaffolded experiments in text and image generation building toward a 12-page children's book, like *Alice and Sparkle*. At each stage of development—from idea generation, drafting, editing, and illustration—students work with (and against) AI tools, reflecting on what this experience reveals about the opportunities, consequences, and ethical challenges of text generation.

Conceptually, our conversation builds from Kathleen Blake Yancey's 2004 CCCC Chair's Address, “Made Not Only in Words,” to place text generation technologies, and the panic surrounding them, within a larger scope of the development of writing technologies. We then use this as a frame to think about the role that technologies have always played in shaping and being shaped by literacy practices in order to move away from fear and into critical engagement. Students conclude their

experiments with a reflective essay that offers suggestions for how students, teachers, and writers can or should use text or image generation in an academic setting.

One success of this assignment was that experimenting with images in addition to text prompted students to ask different questions about writing. In particular, image generation provoked more thoughtful discussions about biases, representation, and consequences than when working only with text generation. Experiments with text generation tended to raise questions about intentionally malicious actors and *explicit* bias, racism, sexism, and hate speech. However, the addition of image generation to our experiments seemed to offer a way to “see” how biases matter and deserve our attention. This brought discussion to questions of coded and implicit bias and followed up nicely on readings and lectures by scholars like Safiya Umoja Noble and Cathy O’Neil.

Before teaching this assignment for a second time, my greatest focus will be changes that help students not only critique but also explore the creative possibilities of text and image generation. This includes revisiting the choice of a children’s book as the final project, which was intended to provide a short and familiar genre to students—allowing them to focus on experimenting with the technologies instead of the final product—, but may have instead discouraged creativity because the genre was so short and followed predictable beats. I would also reconsider the specific text and image generators. This year’s iteration planned to use ChatGPT and DALL-E 2, but because of access issues, our tools were always in flux. Additionally, the user-friendliness of ChatGPT’s interface also worked against experimentation, leaving students with their only point of influence as prompt writing and rewriting, instead of also allowing for the selection of data sets, length parameters, or randomness, for example. Because of this, the next iteration of this course will look to generators, perhaps intentionally older and less sophisticated, that provide more opportunities for negotiating with machinic coauthors.

### **Goals and Outcomes**

- Use experimentation as a method for engaging with questions about writing, authorship, and technology
- Interrogate biases, risks, and consequences of text and image generation
- Reflect on the possibilities, limits, and ethical concerns with AI writing technologies

### **Materials Needed**

Because of rapidly changing access to technologies, this year’s version used various generators as we were able to access them, including the following:

- A text generation tool: ChatGPT, GPT-3 Playground, Canva “Magic Write,” open source LLMs, etc.
- An image generation tool: DALLE-2, Canva Text to Image, Stable Diffusion, etc.
- Design/layout software: Canva, Adobe InDesign, etc.

### **Acknowledgements**

I would like to acknowledge the teaching team for the 2022-2023 year, including co-instructor Andrea McKenzie and teaching assistants Rositza Georgieva, Carla Ionescu, and Ksenia Jourova. I’m also grateful for conversations with Laura Allen



and James J. Brown, Jr. about how they use text generation assignments in their courses.

---

## **The Assignments**

All work from this unit is submitted and assessed as a portfolio, with emphasis on process over product to support experimentation and risk-taking, and it includes three required elements:

- Four writing experiments
- Final draft of children’s book
- Reflective essay
- Optional: any other notes, examples, or images that are relevant to demonstrating student process and engagement with the experiments and materials.

### **Weekly Writing Experiments**

- Exercise #1: Using ChatGPT, generate a children’s story about the topic of your choosing. You can try as many times as you like, but you should save all of your prompts and results. Because of in class activities we’re going to do, do not edit the text produced by AI yet.
- Exercise #2: Edit up to 20% of the story you selected from last week. Indicate in bold where you have made changes.
- Exercise #3: Generate at least 5 potential images for your book. You can try multiple times, but you cannot edit the images yet. Submit your images with search terms.
- Exercise #4: Layout the text and images together into a draft of your children’s book. You can lightly edit and crop images to bring together a cohesive project.

### **Final Project**

Create a 12-page children’s book, including front and back covers, using the text and images you’ve generated. No pages should be blank, having either text, or images, or both.

### **Reflective Essay**

Throughout this term, the university has sent numerous updates (and warnings) advising us as teachers and students about how AI technologies should be handled in an academic setting. Because these technologies are not (yet) part of the university’s policies, concerns about the research, creation, and learning we do as a community are constantly being negotiated. This reflection asks you to think about your experience on this project as a way to consider contributing your perspective to this conversation.

Based on your experience creating a children’s book—as well as engaging in class readings and discussions about writing, authorship, and technology—how would you recommend the university’s academic integrity policy be updated? Write a ~750 word reflection that discusses your experience in this assignment and considers how those might connect to how the university navigates this challenge. In other words, what are one or two specific suggestions for how members of our university community might navigate these technologies, either in line with or in alteration of

the policy, supported by specific instances, findings, or experiences in your work this term?

# cmpttnl cnstrnt: An Exercise in Constraint and Prompt Engineering

Douglas Luman  
Allegheny College

*As new context-aware generative models challenge the human relationship to language, students benefit from first-hand observation of these models' successes and limitations. Using these models often requires using "prompts" (natural language-based directions) to guide their output. The method of developing these directives has quasi-formalized into a practice known as "prompt engineering." Serving as a gentle introduction to the intentionality, opportunities, and limits of the prompt engineering process, this work proposes and describes initial outcomes from an assignment that uses similarities between model prompting and the constraint-based literary work of the Ouvroir de littérature potentielle ("Oulipo") to focus student attention on precision and specificity of prompts and their execution. Beyond familiarizing students with contemporary technologies (particularly OpenAI's GPT platform) and the nascent practices developing around them, this assignment also aims to give students first-hand experience with the reflexivity of using language to describe language in preparation for larger conversations about language as a technology and the roles of large language models (LLM) in human expression.*

---

## *Learning Goals:*

- Learn to interact with LLM through the practice of prompt engineering
- Refine skills in prompt engineering to increase efficacy and quality of output
- Discover exploitable boundaries in LLM generation and what these opportunities offer
- Begin a discussion of the roles and meaning assigned to language as an expressive tool and technology

*Original Assignment Context:* Introductory lesson in elective Informatics course on "Computational Narrative"

*Materials Needed:* There are two options available for this assignment: a "no-code" or a "full-code version. Materials for each follow and are detailed further in the assignment:

- "No Code" version: an accessible AI text generation program (i.e. ChatGPT), selected readings, GitHub account (for both instructor and students)
- Full version with code (presumes Python language knowledge on the part of both the instructor and students): selected readings, OpenAI API keys for GPT, GitHub accounts, an installation of Python on student machine or [an instance of the code running on Google Colab](#)

*Time Frame:* ~1-2 weeks

---

## **Introduction**

"The hottest new programming language is English."

— Andrej Karpathy (@karpathy; Twitter, 24 January, 2023)

I write this assignment introduction as the large language model (LLM) known as GPT-4 has passed or very nearly done so on a number of expert-level tests, such as the United States Medical Licensing Exam (USMLE). This marks the first time—at least in my life—that an automated agent seems to be treading incredibly close to tasks normally considered the sole domain of the “human,” namely a wide array of tasks which are, ostensibly, all about human’s understanding other humans at a more-than-surface level.

But, if LLMs rise to the level of being convincing simulacra of human knowledge and possess the ability to become harbors of emotional investment, we need to devote time to understanding the opportunities, limits, and incidental effects of inputs to the model: user prompts. As Michael Graziano reminds us, “with a good ventriloquist ... [a] puppet seems to come alive and seems to be aware of its world.”

The social media-worthy, much-hyped products of LLM like GPT-3 and image generation tech such as DALL-E or Stable Diffusion begin and end at the role of this able puppeteer. The practice of piloting these models with highly tailored plain language requests to achieve predictable or highly relevant results—known as “prompt engineering” or “prompt programming”—places the controls in the querent’s hands. As the writer behind the site generative.ink elaborates, “[p]rogramming in natural language avails us of an inexhaustible number of functions we know intimately but don’t have names for.”

For folks outside of computational creativity, it may be surprising that this self-reflexivity makes me think about poetry: specifically, reading Charles O. Hartman’s *Virtual Muse* (1996), a book about computational tools testing the boundaries of that frontier of reflexive knowledge. Hidden in the “Unconclusion” of the book Hartman surfaces the powerful idea that

one of poetry’s functions is to make us aware, with a fresh intensity, of our relation to the language which constitutes so much of our relation to the world.

While writing about poetry experiments, Hartman’s conclusion about the function of the form struck me as a way to pitch language as a technology to my students: one that has an explorative and, even, introspective power. One of the appeals of this particular assignment is to offer an automated version of Douglas Kearney’s model of the “Danger Room” role of writing—one in which students see the discourse reflected back to them in real-time, a kind of confrontational, full-contact sport of self-discovery. Though thinking about a different context of writing, I find value in the attempt to force students to, as Kearney closes the piece, explore that “finding a path can be a complex negotiation between possibilities.” This assignment revels in possibility.

Additionally, given the history of text generation as a motivating force in the technology sector, it comes as no surprise that this assignment arrives at a time when much is being made about how we make sense of LLMs which seem to occupy a relationship to language—thereby the world—which appears so much like humans’ own.

Taught as an introductory lesson to the “Computational Narrative” course at Allegheny College in Spring 2023, the assignment which follows contemplates limitation, intentionality, and prompt engineering. This course is taught as part of the

College's Informatics major, a course of study that emphasizes the role and meaning of the presence of information and technology as constructor of lived experience across the many disciplines that computational culture influences. It has been taught once.

Adopting the framework of the intentionality of poetry, the assignment adopts the strict framework of the *Ouvroir de littérature potentielle* ("Oulipo") as a guide for understanding the relationship between prompt and text. As Paul Fournel writes in *The nOulipean Analects*, the constrained prompts Oulipeans apply (e.g. forbidding the use of specific vowels from writing, or substituting nouns using determinate substitution rules) "[stop] when the constraint has been elaborated," bear a striking resemblance to the practice of "prompt engineering," which likewise prefaces and generates but *is not* the result. As you may intuit, applying these kinds of constraints serves as a task that LLMs, at least as far as GPT is concerned, are canonically *bad at*. In addition, GPT-4 lacks a feature that many visual generation models offer, that of the ability to provide "negative prompts" to constrain generative output. Given this omission from prompt practices related to GPT-4, this work seeks to test the limits of what is possible by offering various levels of prompt development.

This assignment also addresses a gap in electronic applications of Oulipean practices—the opportunity to "explore the deeper structures of language that allow the symbolic to reach into the physical world." Giving students who don't normally think of or turn to poetry the opportunity to experiment with the relatively low-stakes environment that Oulipean practice supposes (i.e. the prompt *is* the writing; the outcome is an elaboration of the original *idea*) supports the general tenor of exploration I wanted for an early-semester assignment. (For context, this was students' first *real* assignment for the course.)

The goals of this assignment are four-fold:

- Learn to interact with LLM through the practice of prompt engineering
- Refine skills in prompt engineering to increase efficacy and quality of output
- Discover exploitable boundaries in LLM generation and what these opportunities offer
- Begin a discussion of the roles and meaning assigned to language as an expressive tool and technology

A fifth unstated goal for this assignment is to expose students to LLM technology by degrees in order to subtly introduce the limitations of a technology which will, if it has not already become, a mainstay in information culture. I leave this unstated for two main reasons: the wording and intention of this goal—to offer exposure—is, as many curriculum writers might comment, "soft." It is hard to evaluate or measure what this means or how its outcome manifests.

However, as Ian Bogost's most recent article on GPT-4 in *The Atlantic* suggests, "[i]nstead of one big AI chat app that delivers knowledge...the [GPT model] (and others like it) will become an AI confetti bomb that sticks to everything." This claim reads an awful lot like William Carlos Williams' comparison of Eliot's "The Wasteland" to an "atom bomb," except in this case, we are apparently dealing with a rogue confetti cannon. The current and projected saturation of GPT and other technologies in our daily lives is a one-way street headed toward, essentially, universal integration. At present, there are few courses taught specifically around our interface with these models—"prompt engineering." The Andrej Karpathy tweet that prefaces this introduction certainly supposes that this way of working with LLMs will become a mainstay practice. I am certain the ubiquity of these models will

become a skill which all levels of education will need to address and train. Prompt engineering will become a pervasive occupational and professional task.

Students engaged with this assignment by generating two (2) works: one which adopted a “traditional” Oulipean constraint and one that they figured could only be done by GPT or another LLM. Supported by readings that discuss both the Oulipo and LLMs, the assignment adopts a “learning by doing” during an introductory lab and informal and formal discussions of student work. Certain students are selected as part of each week’s assignments to give formal written responses based on our collective discussions.

During this GPT-focused week, students engaged with Oulipean forms like the tautogram (the use of works focusing on a single vowel) creating works such as one student’s exploration of the use of the letter “e” in the English language (excepted for brevity):

eee eee eee Eeee eee eee eee eee eee eee eee eee eee eee eee eee Eeee eee  
eee eee eee eee eee Eeee eee eee eee Eeee eee eee eee Eeee eee Eeee eee  
Eeee eee Eeee eee Eeee eee Eeee eee Eeee eee Eeee Eee eee Eeee Eee Eee  
Eee Eee Eee Eee

This exploration, the product of the prompt “write me a Tautogram with the constraint of all vowels” produced an intention-atypical, but illustrative use of the letter in common English.

Another GPT-specific request speaks directly to the systems which compose language materially and digitally when they prompt the model: “[y]our keyboard has broken and the only keys that work are the following: 1 2 3 4 5 6 7 8 ! @ # \$ % ^ & \*. Please write a sonnet using only these keys,” which resulted in a self-referential output technical output keeping the patterned spirit, if not the “law,” of a traditional sonnet:

12#3 \$% 456 78! \$% 8#&  
^&%\$! @# ^&\$! 8#& @# \$%7  
^&\$! &%\$ ^&4\$! %7 %^\$ 8#&  
\$% ^&\$! 7@\$ &%\$ !\$ 7#%\$ %7  
  
@#\$ %7 &%\$ 456 78! @# \$%7  
^&8#! \$% ^&\$! &%\$ @# \$%7  
^&\$! &%\$ 78! @#\$ %7 7#%\$  
\$% ^&\$! &%\$ 456 8#& \$%7  
  
!\$ %^\$ 8#& 456 78! \$%7  
^&\$! &%\$ 7@\$ &%\$ 8#& @\$

\$% !\$ %^\$ %\$ ^&\$! 7#%\$ %7

7@\$ &%\$ !\$ 7#%\$ 456 8#&

Student responses to the practice and outcomes of the work, commenting that the practice constituted a “rewire in strategies of how I had been [reading] work,” and that a drove a wedge into others’ interpretive practice by forcing thinking on “a deeper level as to what GPT thought it was doing...and whether or not what we take away was in GPT’s ‘intentions.’” Several observed that the prompts which resulted in the best output were also of the “few-shot” variety, in which examples or stricter rules apply. Others, using the ChatGPT platform, held extended conversations with the platform, gently correcting the model when it made incorrect or incomplete judgments. In most cases students could identify the gaps in their own prompts and hold themselves equally, if not more, accountable for “misses” from the model, which was often the case; most students discovered that, despite close attention to the engineering of the prompt, GPT-3 execution was markedly imperfect and incomplete.

### Notes on the Assignment

This assignment was offered during a course whose structure follows a three-day-a-week model of three 50 minute class periods with a lab session of two hours. When assigned, the schedule was:

Day	Purpose	Time allotted
Day 1, class time	Discussion of texts and constraint	50 minutes
Day 1, lab time	Discussion of prompt engineering and practice	2 hours
Day 2, class time	Informal workshop of prompts and outputs	50 minutes
Day 3, class time	Formal discussion of 3 students’ prompts and outputs	50 minutes

### Materials and Preparation

*This edition of the assignment uses GPT-3 chat for early-semester, accessible content for non-programmers. This version of the assignment also contains code to interact directly with GPT-3, as availability proved an issue when the assignment was actually conducted.*

*“No Code” version*

- [chat.openai.com](https://chat.openai.com)
- Projector
- Digital (or physical) copies of readings
- Networked PC
- GitHub account
  - For professors and students

*Full version with code*

This version presumes Python language knowledge on the part of both the instructor and students, though it may be possible to give this to students with minimal experience or instruction. Here, the instructor will need to distribute the necessary information for an environment file (.env file) which contains values coded as OPEN\_AI\_KEY and OPEN\_AI\_ORG. The format for this file is provided at the end of the README at the included link.

- Projector
- Digital (or physical) copies of readings
- Networked PC
- OpenAI API keys for GPT
  - Organization
  - API key
- GitHub account
  - For professors and students
- An installation of Python on student machine
  - Or, [an instance of the code running on Google Colab](#)

## **Acknowledgements**

This assignment was heavily influenced by a few readings, the content of which didn't all make it into the introduction. The sources to which I am particularly indebted:

- *The nOulipean Analects*, Les Fignes Press: Los Angeles. 2007.
- Frankfurt, Harry. *On Bullshit*. Princeton University Press. 2005.
- y Arcas, Blaise Agüera. “Do Large Language Models Understand Us?” *Daedalus* 151.2 (2022): 183-197.

There is an interesting link here to investigate about giving students low-stakes content and practices to use in their pursuit of understanding *what* LLMs do (versus the question of *how* they work, though this assignment discusses some of both). Here, I found Arcas and Frankfurt's work particularly helpful and, in some respects, permissive to allow students a kind of freedom from needing to be profound or meaningful. Instead, they could just “try out” some voices in a relatively isolated space.

---

## **The Assignment**

The assignment is hosted using GitHub under a CC-BY-SA 4.0 license. It is available at the following link: <https://github.com/AppliedPoetics/cmpttnl-cnstrnt>



The prompt text is provided as a reference, but is also available on the above-linked site as the README.md.

### **Prompt: Computational Constraint**

“With a good ventriloquist ... [the] puppet seems to come alive and seems to be aware of its world.”

— Michael Graziano, in *Consciousness and the Social Brain*

“Prompt engineering for large language models is just an excuse to make up more nonsensical sentences to feed these AI monsters.”

— GPT-3, in response to the prompt "What is prompt engineering for large language models? Answer in a very snarky way."

### **Readings**

#### *Theory*

- [Do Large Language Models Understand Us?](#), Blaise Agüera y Arcas
- [A Brief Guide to the OULIPO](#)
- Various excerpts from *The nOulipean Analects*
  - pp. 35-47; 155-161

#### *Practice*

- [Methods of Prompt Programming](#)
- Various Oulipean Poems
  - “a russian con’s economic missive,” Ian Monk
  - “Art A to Z,” IN.S.OMNIA
  - “What a Man!” Georges Perec
- [A compendium of various "traditional" constraints](#)
  - Ignore the newspaper constraint...
  - unless it's actually helpful

#### *Documentation*

- [GPT-3 Completion API](#)
- [GPT-3 Notes on prompt design](#)

### **Summary**

Prompt engineering—the practice of learning to con/destructively "pilot" a generative model—is one of the surprising new skills to emerge from the development of context-aware large language models (LLM). Simply put: prompt engineering is the practice of instructing a model to produce an output consistent with the prompter's intent or desire. While we've given a new name to what essentially amounts to "asking the right questions," prompt engineering is much more than that.

To date, successful prompt engineering endeavors to ask what an artist *wants* to happen. This assignment approaches generative writing with large language models (LLM) from an opposite perspective. Drawing on the practice of the *Ouvroir de littérature potentielle* (“Oulipo”), we challenge GPT-3 to a more difficult task—

producing works of "constrained" writing to discover what LLM can and, more importantly, cannot do.

For those of us familiar with visual image generators such as DALLE or Stable Diffusion, this idea is close to, but not quite "negative prompting" (e.g. asking for a picture of a house without any people in it). The approach of computational constraint applied to language prompts thinks about the concept *generatively*. We aren't simply asking to "live without" a feature common to a parcel of language, we're interested in rethinking the possibilities are open by *restricting choice*.

Mainly, what kinds of choices can we engineer the model to make and how can we account for those choices?

## Goals

- Learn to interact with LLM through the practice of prompt engineering
- Refine skills in prompt engineering to increase efficacy and quality of output
- Discover exploitable boundaries in LLM generation and what these opportunities offer
- Begin a discussion of the roles and meaning assigned to language as an expressive tool and technology

## Outcomes

- 2 texts incorporating prompt engineering (included in the `writing` folder as `md` files)
  - 1 enacting a "traditional" Oulipean constraint
  - 1 enacting a constraint only possible using GPT-3
- A journal of various prompts attempted with brief notes about relative success or failure (include in `writing/prompts.md`)

## Process

### *Using ChatGPT*

ChatGPT is an interface that allows you to use the prompt you've engineered and, failing excellent results, to chat with the model and encourage it to make changes that conform to your expected constraint.

I advise you to be kind to the model, even if it is just an LLM. ChatGPT is available [here](#)

### *Using code provided*

ChatGPT is undergoing both rapid change to a subscription model and varying levels of actual availability (due to performance load). To make this assignment possible, the assignment repository offers code that interfaces with the GPT-3 back-end (not chat, *per se*). To use this, obtain a key from your instructor to place in a .env file in the main folder of your repository.

This repository contains three (3) files essential to making any code for this assignment "happen". They are all contained in the src folder.

*data/prompt.txt*

This contains the prompt which *prepends* the text. For example: `remove all of the bad people from the following text`

*data/source.txt*

If operating on a "found" text (i.e. one you creatively pirated from elsewhere), paste the text you'd like to operate on in this file.

*main.py*

The program behind communicating with the GPT-3 API. This file requires the creation of an `.env` file, the values and specifications of which will be provided in class during either the session or the lab.

# The Grand Exhibition of Prompts

Mark C. Marino, University of Southern California

Rob Wittig, University of Minnesota Duluth

*The Grand Exhibition of Prompts, a netprov*

*In this collaborative creative writing game on a threaded discussion platform students experiment with AI image-making programs, but concentrate on the expressive, literary power of the verbal prompts they write. As language arts, image-making prompts are concentrated, evocative, use sentences, fragments, lists and non-sequiturs. In other words they are poetic (but don't tell them that until later!). Students can be evaluated on their contributions, and on their support and encouragement of other writers.*

---

*Learning Goals:*

- Understand what AI image-making does in a basic way
- Use language playfully and creatively
- Recognize how different a verbal prompt is from the images artificial intelligence makes
- Aim, at least sometimes, to *write things that cannot be represented visually*
- Support other students' creativity by writing prompts that respond to others' prompts and by writing kind and creative critiques
- Show their own characters' insights and introspection via the (public) 'secret diary' entries.

*Original Assignment Context:* End of semester in upper-level advanced writing course

*Materials Needed:* An accessible AI text-to-image generation program (i.e. Craiyon)

*Time Frame:* ~2 weeks

---

## Introduction

With the advent of artificial intelligence image-making software (such as MidJourney and Dall-E2), it seems the average untrained anybody can become a world-class visual artist merely by writing a prompt. Such developments have disrupted traditional notions of what it means to be a visual artist in the one aspect of making we like to think is exceptionally human: creativity. In order to explore this new world and all the anxiety it provokes, we created a writing game, a netprov (online collaborative improvisation), entitled "The Grand Exhibition of Prompts," in which participants imagine a world where artists strive not to make new art but to write dazzling and moving prompts. The culmination of this fictional writing movement is an imaginary Grand Exhibition in the style of the Salon des Beaux Arts held in the mid-19th Century in Paris, ironically gesturing to the way AI-generated art disrupts previous notions of a single beating-heart at the center of a (thankfully bygone) Eurocentric art world.

In this writing game or netprov, each participant created a character to role-play, an imaginary 'prompt artist,' who in addition to making art as a member of one of our

three schools (Emo, Retro, and Fido), narrated their personal journey as a prompt artist in (public) secret diary posts. Of the schools: Emo prompts feature inner turmoil, Retro prompts feature love of the past, real & imagined, sincere or ironic, and Fido feature pets and how incredibly cute they are. We also provided cafe channels for discussion of the ideas of prompt art, discussions which often touched upon the emotional state of the anthropomorphized bot art-generators. Participants were also tasked with writing reviews, not of the art that was generated but of the prompts written by others, reviewing the places where the humanity of the prompt writer shone through their request to the machine. As in many of our netprovs, participants came through Mark's Advanced Writing course, our regulars or "featured players," and others who saw our invitations on the Internet.

### **Goals and Outcomes**

The goal of the netprov is to get writers to experiment with AI-generated art, focusing not so much on using the technology to create art, but to use the space of the prompt to reflect on what it means to be human. As a result, the participants generated something that we would call poetry, writing that explores the space between language-made-practical, made machine-readable, and those parts that remain illegible to the machine. While there are tons of interesting discussions to be had about AI image-making ethics, this assignment concentrates on the prompts themselves as language arts, as literature. This assignment, then, makes a great complement to a unit about image making.

A successful outcome demonstrates that students have: understood what AI image-making does in a basic way, used language playfully and creatively, recognized how different a verbal prompt is from the images artificial intelligence makes, aimed at least sometimes to *write things that cannot be represented visually*, supported other students' creativity by writing prompts that respond to others' prompts and by writing kind and creative critiques, shown their own characters' insights and introspection via the (public) 'secret diary' entries.

We ran the netprov in the Fall of 2022 with around 40 participants, drawing in part from two sections of Advanced Writing taught at the University of Southern California. Students had a ball! Additional participants were solicited through announcements posted on a variety of digital culture message boards, including those of the Electronic Literature Organization (ELO).

### **Prompt Samples from USC and University of Minnesota Duluth**

*From: AceVerocchio*

To be possessed of double pomp, To guard a title that was rich before, To gild refined gold, to paint a lily, To throw perfume onto a violet, To smooth the ice, or add another hue Onto the rainbow, or with a taper's light To seek to garnish the beautiful eye of heaven, It's wasteful and ridiculous excess. --test --creative -

*From: Electricpersona*

frigid, windy, cold, freezing, frozen, gale, cold hands, cold fingers, shaking, blustery, no sweater, howling wind, thin jacket, shaking, blowing, quivering, shuddering with cold, battered, miserable. Wrestling with A Giant Fiberglass Sousaphone which is IMPOSSIBLE TO CONTROL in high wind on a football field during band practice. soft light, in the style of Maxfield Parrish, photo realistic, art deco beautiful soft focus. And I'm not even really a sousaphone player, I'm a trumpet player!!

## Secret Diary Samples

*From the secret diary of a famous Retro prompt writer:*

#secretdiary I did SO well two years ago at the Grand Exhibition. Everyone basically told me my prompt was the best, one of the best, super good. And it was an accident. It was a weird prompt I wasn't even going to enter, not one of the ones I usually do, just a little "sigh" prompt about life before COVID. My friend Belter told me I should enter it. But now I have heard two Grand Exhibition podcasts already mention my name and looking forward to what I enter this year and I'm freaking out! Too much pressure! Have I peaked? Is my best creative work in the past? Agggh!

*From the secret diary of a famous Fido prompt writer:*

#secretdiary Omigod! I am shaken to the core! Carlissa M just announced she's switching from Midjourney to Dall-e2! WTH? Doesn't everybody know how much more sophisticated Midjourney is? But Carlissa claims there's a better Pomeranian show-dog database on Dall-e2 and that's all she cares about. I just never have gotten cutes as cute as Midjourney when you write "really super cute". Which makes me wonder if Carlissa M cares about cute. And THAT makes me doubt the foundations of my universe.

One small note: During the game play, participants started anthropomorphizing the AI software, noting that they saw Midjourney looking dejected, smoking by the loading dock. That comment, and those who "yes and-ed" it, led to a spinoff netprov called Pr0c3ss1ng: A Support Group for AI Assistants, which is open for play on Reddit!

(<https://www.reddit.com/r/pr0c3ss1ng>)

## Materials/Software

In our first run the primary writing space was a Discord server, which had a Midjourney text-to-image bot installed on it. The prompts, reviews and other writing were integrated with the AI images.

However, this project may take place in a number of different platforms. The basic needs are: a threaded discussion space, ideally with the posting of in-line images possible, and web access to a text-to-image generator such as Craiyon, Dall-E2 by open AI, or MidJourney. At the time of our trial the Midjourney bot was fairly easy to install, and gave students 10 free attempts before wanting them to subscribe. These technical affordances change constantly so make sure you do a quick test to see what's possible now. As of this writing Craiyon was still free. Course management software such as Canvas or Moodle will work well.

Things change so rapidly in the AI image field that we encourage you to search current news sources and essays for support texts.

## Acknowledgments

More about netprov and many ready-to-play netprovs can be found in the book *Netprov: Networked Improvised Literature for the Classroom and Beyond*, by Rob Wittig from Amherst College Press. The book is free to read via open access: <https://doi.org/10.3998/mpub.12387128>

## Licensing

We agree to share this work with a CC-NC license.

---

## The Assignment

### Introduction

People are having fun writing verbal instructions – prompts – for **powerful artificial intelligence programs to make visual images** on text-to-image platforms such as Dall-E2 and Craiyon. The visible results so far are mostly glitchy and clichéd – just what you’d predict.

On the other hand we are entranced by the **words** folks use in their prompts to **describe what they want to see!** These prompts are strange, powerful, revelatory. Using phrases, fragments, and lists sprinkled with lumps of aesthetic and technical jargon, aspiring artists are writing short texts of startling depth and impact. Help us explore and celebrate prompts as literature!

Come and write prompts with us! Join a community of prompt artists where you can collaborate and play! Join an art movement! Start an art movement!

Then enter your best prompt into the Grand Exhibition!

*Images are but the stepping stones to writing more beautiful prompts.*

– Hans Paedeweyder

### History and Culture of the (imaginary) Grand Exhibition

For the visual artists of Paris in the mid-19th century there were two crucial moments: the exhausted morning they dropped off their freshly-finished oil paintings to the Grand Palais to be judged by the jury of the Salon des Beaux Arts, held every two years, and the nerve-wracking instant the doors of the exhibition first opened. They mobbed the entrance. They ran inside. Was my work chosen? Was it hung at eye-level or way up by the ceiling? To be a professional artist meant being seen at the Salon. Everything depended on it.

Now, it’s happening again! The Grand Exhibition of Prompts has become the premier gathering place for AI-art prompt writers from all over the world. For weeks before the deadline writers gather in Discord, sharing, discussing, and encouraging. Each aspires both to help the others and to be the best.

Every two years the main artistic styles or “schools” of the Grand Exhibition hang out in their own Discord channels, refining their aesthetic “sound” and “feel.” This year there appear to be three dominant styles – Emo, Retro, and Fido. **Emo** (revealing inner turmoil), **Retro** (love of the past, real & imagined, sincere or ironic), **Fido** (pets and how incredibly cute they are). The schools are all welcoming, but they’re very loyal to their styles.

The Grand Exhibition of Prompts store is currently sold out of the very popular “Imagine Me Like This” hoodies. We’ll let you know when they’re available again.

## **This project is about**

- Playing with the technology
- Exploring the strangeness and verbal novelty of the prompts as a new form of literature
- Reflecting on how we phrase what we want to see

## **Basic Play**

This assignment can occupy about 3 to 4 class periods or about two weeks if assigned as homework

1. Invite students to explore and make a few practice images with an AI text-to-image program such as Craiyon. Ask students if they have tried this before. Have students share practice prompts and results with other students.
2. Have students make up a character to play, an artistic prompt writer (see below).
3. Assign students to one of the three styles of prompt art — Emo, Retro, Fido — as their basic style in approximately equal numbers, or let students choose their own style.
4. Have students write (3) prompts and post them in their own style’s Workshop thread, and (2) experimental prompts, one in each of the other style’s Workshop thread. Goof around with it! Have fun! Remind students that they will choose their “official submission” to the Exhibition of Prompts at the end of the assignment period. This should help them relax and be playful in writing their experimental prompts.
5. Each prompt, review, and #secretdiary entry begins with:  
*From: [charactername]*
6. Encourage students to write their own prompts in response to others’ prompts that they like. Write to imitate, write in contrast, write to continue a story; make the prompts a conversation
7. Have students write (3) kind and constructive reviews of other students’ prompts.
8. Have students write (3) #secretdiary entries in which their characters reflect about writing prompts in the “Chat” channels. This is a great way to narrate your journey as an artist. Secret diary entries start with the hashtag #secretdiary and are public.
9. Later on, once you have created your favorite prompt, you’ll submit that to the Grand Exhibition, which is its own, separate channel in Discord!
10. Remember: students are evaluated on their written prompts, not the images – This is surprisingly, and interestingly, difficult to remember; it makes for a great discussion about words and images. We evaluate both their own original writing and the way they support and encourage the creativity of other writers in the group.

## **Prewriting: Making Your Roleplay Prompt Artist Character**

1. Who is your character? A Prompt Artist. Name, location, other details
2. What is their artist story? Were they late bloomers or prodigies? Did they learn their craft at the knee of a master or are they self-taught?
3. What is their relationship to the machine? Partner, workshop, tool, environment
4. What is the big challenge they are facing artistically?



5. How do they define artistic success? (Remember in this netprov, what's important is the writing (the prompts), not the visual products (the AI-generated images).

### **Tips on How to Play**

Remember, the focus of this netprov are the prompts, not the images. Imagine the visual, inscribe the verbal.

Try single words mixed in with phrases. Play with words that are abstract mixed in with figures of speech. In fact, defy the pragmatic creation of a plain descriptive post. Make poetic sense.

Play, experiment, explore! Have fun!

You don't have to make an image for each prompt – this saves time.

Save all your variations! You never know until later which one reads the best.

Our Prompt Masters encourage you not to go to prompt writing sites, not to look at other people's advice on the web or technical specifications. That will only get you trite, tired, prompts of the now. We seek fresh, original, unique prompts of the future!

The prompt-writer character you create can be modeled on any kind of artist: painter, musician, poet, et cetera. They can be just starting their career; they can be recovering from a blip in their career; they can be a veteran artist dispensing advice to the young ones. They can be passionate members of their art style/school, or they can be unique voices. Remember that artists suffer from many more crises of confidence than their fans want to think.

### **For Teachers: Preparing the Threaded Discussion**

We created these threads before the start of play, adapt as you wish:

How to Play; Artists' Introductions; The Café Main Chat; Emo Workshop (for prompts); Emo Chat (for discussion among characters and 'secret diary' posts); Retro Workshop; Retro Chat; Fido Workshop; Fido Chat; The Grand Exhibition (for students self-chosen 'best prompt' at the end of the project)

You may also wish to post daily themes, such as "handmade," "copycats," "invisible art," "regret," "happenstance," "imposter syndrome," and "unimaginable visions" in separate threads.

### **Sources and Sites to Consider**

- Janelle Shane: <https://www.aiweirdness.com/>
- AI Art is Art by GM Trujillo: <https://aestheticsforbirds.com/2022/11/02/ai-art-is-art/>

### **Places to Make Images Free**

- Craiyon: <https://www.craiyon.com/>
- Dall-E 2: <https://openai.com/dall-e-2/>
- Dezgo: <https://dezgo.com>

# Different Ways of Narrating with Curveship-js

Nick Montfort  
Massachusetts Institute of Technology

*Students are assigned to generate variations of an underlying story using a system designed for this purpose. The goals are (1) to use the core aspects of narrative, events and their representation; (2) to link narrative terms to concepts; and (3) to discover how events can be narrated in different, lively ways. The assignment requires instructor familiarity with my system, Curveship-js, but works well. Sharing and discussing the results enhances learning.*

---

*Learning Goals:*

- Enhance students' understanding of narrative theory
- Introduces the formal process of coding
- Facilitate creativity and the classroom discussion of different approaches
- Open up questions about the nature of narrative and fiction

*Original Assignment Context:* Beginning of upper-level elective media studies course

*Materials Needed:* The Curveship-js framework (a system for narrative variation that does true text generation in a classic way), Curveship-js documentation, which includes: a guide to Curveship-js concepts, a “Getting Started” document, and a technical reference

*Time Frame:* ~3 weeks

---

## Introduction

Curveship allows students (and others) to tell the same underlying story in different ways, as long as they formally and precisely specify how to do so. I originally released the system in a Python version in 2011; that one is now called Curveship-py. More recently, I developed a JavaScript framework, Curveship-js, specifically for use in my Interactive Narrative class (21W.765 / 21L.489 / CMS.618 / CMS.845). One goal is to help students learn about narrative theory. I assign students to create a formal underlying representation of the content level (or story level), and then connect that with three different formal representations of expression (or narrative discourse). Both formulations of Curveship are systems for narrative variation that do true text generation in a classic way, starting from an abstract representation of content, planning the details of lexicalization and referring expressions, and finally realizing the surface text.

It is important to understand that while Curveship is a text generation system, it is a classic one, completely rule-based and symbolic. In fact, unless the system is asked to rearrange the telling of events at random, the output of the generator is deterministic: The same text will result each time. Curveship-js (as with its predecessor) is therefore very unlike a statistically-based large language model. Its purpose is not to generate fluent, cohesive text as a continuation of a prompt. It is, rather, a detailed model of narrative and of narrative theory.

While I taught with the Python version before 2020, students have used Curveship-js since then. They have been able to specify different orderings of events, time of narrating, narrator (or “I” of the story) and narratee (or “you” of the story). The ability to reorder events includes the ability to leave things out (ellipsis). Since 2021, the system was developed to enforce a strong distinction between content and expression. The naming of actors and things, and the particular verbs chosen to represent actions, all were represented at the expression level, so that different word choice and the use of dialect can be encoded as well.

Students are introduced to the idea of underlying content that can be expressed in different ways via some offhand examples as well as Raymond Queneau’s *Exercises in Style*, a book that includes ninety-nine vignettes, each representing a rather uninteresting plot told in very amusingly different ways. Before assigning this text generation project, we cover the basics of narrative theory, including the distinction between content and expression (equivalently, story and discourse). Students do several in-class exercises that anticipate this assignment; they are asked to write a narrative of their own choosing with certain constraints, then to vary it in a specific way. I also use Curveship-js in class to explain aspects of narrative theory and pointed students to the several examples that are available online.

---

## The Assignment

Using Curveship-js, create a variable narrative with three (3) different narrator files (and corresponding HTML files) and no more than ten (10) events. You will need (7) seven files in total: two for each telling (one .js and one .html) and one story file (a .js file) to be shared between all the tellings.

### Context and Purpose

This assignment is a creative exercise. There are no restrictions on topic, theme, or specific subject matter. Nothing about the story or discourse is specified, except its simplicity: There can be no more than ten events. Everything that participates in those events (actors and things) also needs to be represented. Because the representation of an event, in the simplest case, usually takes a single line, and it is not likely that there will be more than ten actors and things, each represented by a single line, the whole story file (with a little overhead) will probably not exceed 25 lines. The variable narratives composed for this assignment are not interactive, as the two creative assignments that I give after this one both are. But this assignment does require that students understand how to formally represent aspects of narrative.

### Outcomes

Even the least interesting variable narratives show that students are indeed capable of narrating the same underlying story in several different ways. In every case, there is at least some aspect of the variation that gives a different sense or connotation. Students often do much better, finding extraordinary ways to create different narratives that share the same underlying events. In some of the best cases, the variation gives a completely different impression in each of the three cases, raising the question of what really happened. The variable narratives students devise are, invariably, fictional, so this gives us the opportunity to discuss that there is no “really” — we only know about fictional stories because of how they are expressed. The concept of the underlying story level is itself a useful fiction for writers, giving us the freedom to explore different tellings.

The exercise does enhance students' understanding of narrative theory, helps them understand the formal process of coding, allows for creativity and the classroom discussion of different approaches, and leads us into questions about the nature of narrative and fiction.

### **Software and Skills Needed**

Students have to have the willingness to modify computer programs and the patience to deal with formalizing their ideas — they do not have to be programmers or have programming background. Using Curveship-js involves representing data, but it does not, strictly speaking, require programming. No one has to write a for loop or a conditional statement, for instance. Because Curveship-js is a formal system, it can be unforgiving. Students have to be willing to proceed slowly in modifying examples and carefully consult the documentation I have written, which consists of a guide to Curveship-js concepts, a “Getting Started” document, and a technical reference. An instructor wishing to use this assignment will also have to patiently work to understand my system, which, although it is documented, does not have a community to support it.

Because many MIT students have strengths in math, science, and engineering, this exercise helps to show that such backgrounds can be relevant to humanistic and artistic work. The thinking that students do when they formally model phenomena in scientific and technical fields can apply to narrative. The assignment may be useful in other similar contexts, where students with significant technical background are seeking to see how their experience can apply and can assist their learning and creativity in other domains.

Whether or not other instructors choose to use Curveship-js specifically, I hope that our successful use of the system will embolden others to use fairly complex, domain-specific models, whether or not they are “AI” systems in today's common usage of the term.

### **Availability of the Assignment**

I make all of my syllabi (with assignments) available at my site, <https://nickm.com>, for others to read and adapt for their own purposes. Curveship has its own page with links to the Curveship-js download, examples, and documentation.

# Deconstructing and Reconstructing Genre and Form with Tracery

Mark Sample  
Davidson College

*In this assignment, students work with HTML, CSS, Javascript, and JSON templates in order to design a website that generates new content out of pre-established rules and word banks. No prior coding or web development experience is required; the free browser-based platform Glitch.com hosts the projects templates as well as the projects themselves. The assignment encourages students to deconstruct the underlying rules, tropes, and conventions of any kind of textual genre. This assignment emerged out of an undergraduate course devoted to digital literature and poetry, but it can be adapted for many contexts, including any field concerned with form, style, and genre conventions.*

---

## *Learning Goals:*

- Deconstruct some genre of text by identifying that genre’s underlying rules, tropes, and conventions.
- Reconstruct that genre of text using a combination of randomness and curated content.
- Create an original multimodal work that takes advantage of the unique aesthetic and literary affordances of digital environments.
- Understand how Tracery grammars—and combinatory writing more generally—work.
- Develop an appreciation for the exponential power of combinatory writing, which can generate billions of variations while simultaneously working with various constraints.

*Original Assignment Context:* Middle of elective digital studies course

*Materials Needed:* Web-based HTML, CSS, Javascript, and JSON templates (linked in assignment), a free account on glitch.com

*Time Frame:* ~4-5 weeks

---

## **Introduction**

Deconstructing and Reconstructing Genre and Form with Tracery is an assignment in Digital Studies (DIG) 220: Electronic Literature, a course in the Digital Studies program at Davidson College. DIG 220 surveys the past and present of interactive narrative, digital poetry, and hypertext fiction—collectively known as electronic literature, or e-lit. Given its focus on narrative and aesthetics, the course also counts for the English major and minor at Davidson College. The class draws a mix of students, ranging from English majors who have never coded in their life to Computer Science majors who are more comfortable with Python than poetry. In between are students simply seeking to fulfill Davidson’s Literary Studies, Creative Writing, and Rhetoric “Ways of Knowing” graduation requirement, our version of a literary-oriented general education requirement. The class size is typically 25-30 students, representing every year from first-year students to seniors. I’ve taught the

course three times over the past few years, and I've used this particular assignment in the past two iterations of the courses. I also use a variation of this project when I teach creative coding workshops for grad students and faculty. Which is to say, the project works well at nearly every level.

Most electronic literature courses follow one of two organizational models. The first model is chronological, marching through digital poetry and narrative of the past half century. This model is often teleological, suggesting that early "primitive" works of digital literature eventually gave way to more complex contemporary works, a notion that doesn't strictly hold up. The second model is platform-based, which tends to stress technological affordances of tools like Storyspace, Inform, Flash, or Twine at the expense of other aesthetic values. The organizing principle of my course eschews both chronology and platforms. Rather, the course is organized around broad literary and aesthetic themes, such as the uncanny, the sublime, or dysfunction. This organization puts e-lit works from entirely different eras and modes of production into conversation with each other. Furthermore, it actively undermines any technological determinism that my students—and honestly, I—might bring to the material.

Students tackle the Deconstructing and Reconstructing Genre and Form project in a unit on randomness and variability. It is here we delve into the history and power of combinatory writing, which mashes up texts with a degree of randomness, yet is still constrained by a "grammar" or controlled vocabulary. By this point in the semester, students will have encountered digital poets such as Allison Knowles, Nick Montfort, Stephanie Strickland, Amaranth Borsuk, and others who have produced profound work using what at first glance seems to be Mad Libs on steroids. That is, a template or scaffold of syntactical structure the blank spots of which are filled in procedurally with words and phrases curated by the writer/programmer.

Anyone familiar with Mad Libs knows that this kind of combinatory writing lends itself to parody, satire, and outright absurdity. Yet, combinatory writing is also a powerful tool for understanding how genre and form work. Pick any genre of short form nonfiction—horoscopes, menus, movie recaps, medical bills, emails, diaries, and so on—and one finds certain rules, tropes, and conventions at work. In order to make an approximation of such writing convincing, one must study the form, breaking it down to its constituent parts and diving deep into its underlying rhetorical, paradigmatic, and syntagmatic strategies. This is the deconstructing part of the assignment. Without a doubt, taking apart a text is the best way to figure out how it works. This is doubly true when it comes to genre writing. This approach to understanding texts is heavily influenced by Jerome McGann and Lisa Samuel's concept of "deformance," a portmanteau of performance and deform. For McGann and Samuels, a deformance is an interpretive strategy premised upon deliberately misreading a text, for example, reading a poem backwards line-by-line. As Samuels and McGann put it, reading backwards "short circuits" our usual way of reading a text and "reinstalls the text—any text, prose or verse—as a performative event, a made thing" (Samuels & McGann 30). Reading backwards revitalizes a text, revealing its constructedness, its seams, edges, and working parts. So too does trying to recreate formulaic writing through procedural generation.

As I said, procedural generation lends itself to parody. For example, creating an endless procedurally generated horoscope that makes fun of the tropes of horoscopes is a rather simple matter. More difficult to achieve with procedural generation is writing that evokes genuine emotion or narrative insight. And this is ultimately what I challenge my students to do: produce generative writing that offers insight and critique. Our goal is not merely to create a procedurally generated text that is funny

or absurd, but to create procedurally generated texts that take a stance, point a finger, examine the world.

To that end, students work on a generative text project using Tracery, a Javascript library that lowers the barrier to creating with the “slotted” technique of combinatorial writing. Here is the general flow of the assignment:

1. We study existing examples of procedural writing that offer both satirical and more substantive critiques of the world.
2. Students learned the basics of Tracery through in-class workshops.
3. Students experiment with a “starter” template of a fully functional Tracery project via glitch.com, which makes it easy to work on HTML, Javascript, and CSS directly in a browser, without special text editors or development tools.
4. Finally, I introduce the assignment itself, including what at first seems like its impossible criteria, such as developing a project that has at least 1 billion possible variations.

Generally students have several weeks to work on the project, including at least one day in-class, where I can help troubleshoot problems or push students out of their comfort zone with CSS and HTML. Because the students work in glitch.com, their projects are easily shareable, and we’ll have a brief show-and-tell at the end of the project where they share and critique each other’s work.

### **Goals and Outcomes**

This Tracery project serves several learning goals. By the end of the project, students will be able to do the following:

1. Deconstruct some genre of text by identifying that genre’s underlying rules, tropes, and conventions.
2. Reconstruct that genre of text using a combination of randomness and curated content.
3. Create an original multimodal work that takes advantage of the unique aesthetic and literary affordances of digital environments.
4. Understand how Tracery grammars—and combinatorial writing more generally—work.
5. Develop an appreciation for the exponential power of combinatorial writing, which can generate billions of variations while simultaneously working with various constraints.

### **Materials Needed**

The work for this project can be done entirely in any modern web browser, such as Chrome or Firefox. Students will need a free account on glitch.com.

### **Acknowledgments**

There are three foundational inspirations for this project. First, the work of Lisa Samuels and Jerome McGann: “Deformance and Interpretation” from *New Literary History*, vol. 30, no. 1, Jan. 1999, pp. 25–56. Second, the creative and pedagogical work of Nick Montfort, who advocates for what he calls “exploratory programming” and who makes his own work open and available for remixing, adaptation, and reuse (see Montfort’s open access edition of [Exploratory Programming for the Arts and Humanities](#), second edition, published by MIT Press in 2021). The third inspiration is more theoretical: Sianne Ngai’s concept of stuplimity, which she describes as a

“mixture of shock and exhaustion” in the face of stupefying quantities of more or less the same thing—an aesthetic quality that applies to many procedurally generated combinatory texts (see Ngai’s *Ugly Feelings*, Harvard University Press, 2005). And of course, this project would also be impossible without Kate Compton’s Tracery library and Allison Parrish’s port of Tracery into Python.

---

## The Assignment

### Overview

For this project you’ll use a variation of the “slotted” technique of combinatory writing. This means building sentence templates with empty slots for nouns, verbs, and so on. A program then randomly selects from pre-selected lists of words to fill each slot. Think of procedural generation as Mad Libs gone crazy. Tracery is the procedural engine you’ll use. It’s a Javascript library by Kate Compton that lets you create surprising configurations of texts out of the template (called “grammar”) and vocabulary you provide.

Strictly speaking, this project does not involve programming. It does, however, require procedural thinking. Every Tracery grammar takes the form of JSON data. This is a highly structured text format that both humans and computers can “read”—though parse might be a better word. Your grammar data file tells Tracery what your templates are (templates in the plural because you can embed templates within templates recursively) and provides the vocabulary for Tracery to use as it fills in the templates.

```
"origin": [
  "#contamination#. #source#. #culprit#. #eyewitness#. #exposure#. #response#."
],
"firstSyl": ["Global", "Hammer", "Cent", "Dyna", "Archi", "Chem", "Petro"],
"lastSyl": ["ex", "dyne", "eum", "tech", "ax", "zoom", "plex", "base", "com", "plus", "ent", "edge"],
"corporationname": ["#firstSyl# #lastSyl#"],
"contamination": ["#time# #amount# #disaster#"],
"source": ["#mysterysource# appears to be #accident#"],
"culprit": ["The #manufacturer# #corporationname# #companytype# #operates# nearby #facilities#"],
"time": [
  "Yesterday",
  "This morning",
  "Today",
  "Last night",
  "Early this morning"
],
"amount": [
  "an unspecified amount of",
  "an unknown quantity of",
  "5,000 gallons of",
  "10,000 gallons of",
  "15,000 gallons of",
  "20,000 gallons of",
  "25,000 gallons of"
],
```

The Tracery grammar JSON that powers *Don’t Drink the Water*

It’s easy for text generation to be comical or satirical, like Nora Reed’s [thinkpiecebot](#) or Compton’s [Night Vale Generator](#), both of which use Tracery. Other text generation parodies powered by tools similar to Tracery include the [Postmodern Essay Generator](#) and Eric Drass’s [machine imagined art](#).



It's also possible to make text generators that are serious or provide social critique, such as my own [Infinite Catalog of Crushed Dreams](#) or Leonardo Flores' [Tiny Protests](#) bot, which both use Tracery or Tracery-like tools.

## Guidelines

1. Decide what kind of text you want to generate. Will it be a parody? Social commentary? A genre, like horror, romance, or science fiction? Fake non-fiction? What's the mood of the work? Humorous? Sarcastic? Somber or melancholy? Whatever it is, strive to give it some depth.
2. Make an account on [Glitch](#) and clone the [template files](#) from "Don't Drink the Water." Glitch is a community of creative coders that provides free online tools and hosting. Once you've [remixed the project](#) you can edit your own version directly in your browser. You can instantly see the results by clicking the "Show Next to the Code" option in Glitch (click the sunglasses). Also check out my tips on [customizing the look](#) of your project.
3. Go through the Tracery tutorials below as you work on your project. When you first start it might help to use a visual editor. Beware that these editors are liable to crash as your grammar grows in complexity. You can paste the code from the visual editor into Glitch. See also my [Tracery Tips](#) for more details on working with Tracery.
4. You will also write an Artist Statement that puts your Tracery project in dialogue with questions regarding combinatorial writing, procedural generation, novelty and repetition, authorship, and creativity.

## More Details on Editing Tracery

1. To actually edit the underlying Tracery grammar, edit the grammar.js file in Glitch. You can also paste the text of grammar.js into the online [Tracery editor](#) to make sure it's working.
2. If the page is entirely blank in your browser, then there's an error in your grammar.json file! Go back and make sure all your commas and so on are in place.
3. In addition to editing the grammar.js file, you should edit the index.html file (you can do this in Glitch too) and update it to reflect your own project's needs. You can also edit the style.css file in order to change fonts, font sizes, colors, and other visual elements of the page.

## Criteria

DIG 220 uses an assessment method called specifications grading. In specifications grading, expectations for projects are clearly laid out, and you either meet the criteria or you don't. If you don't, you'll have an opportunity to revisit the work and revise it so that it does meet the criteria.

These are the minimal requirements for the project to be considered for a B-range grade.

- Your Tracery project must have at least 1,000,000 possible combinations.
- 9 out of 10 combinations must parse correctly. That is, at least 9 out of 10 combinations should make grammatical sense, though not necessarily logical sense.

- The Tracery project is accompanied by a 750-1000 word Artist Statement. You'll share the Artist Statement with me as a Google Doc and also post the link on Moodle
- The Artist Statement must integrate at least two secondary sources in a *substantive* way. "Substantive" doesn't mean merely quoting at length or cherry-picking a key phrase. "Substantive" means actively engaging with the source by (1) summarizing its overall argument; (2) showing how that argument ignores important issues, doesn't go far enough, or could be applied to new contexts; and (3) applying concepts from the secondary source to your procedurally-generated text. See below for possible secondary sources to consider.
- The Artist Statement follows scholarly standards for citation, using either MLA, APA, or Chicago style.
- The Artist Statement contains no more than 3 grammatical, spelling, or other "mechanical" errors. It also contains no more than 2 minor factual inaccuracies and no major factual inaccuracies.
- You must also post a working link to your Tracery Project on the class blog, under the category "Tracery Project." Include a 2-3 sentence description of the project, as well as its name.
- The Tracery project and Artist Statement are shared by class time on Friday, March 13.

To be considered for an A-range grade, you project must meet the above criteria *plus* the following:

- Your Tracery project must have at least 1,000,000,000 possible combinations.
- The generated text should do more than provide a few minutes worth of entertainment. It should be something compelling enough that a reader wants to keep reloading even after the initial novelty has worn off. Make it provocative, evocative.
- The generated text should be longer than a few paragraphs (prose) or stanzas (poetry). As a point of comparison, [Don't Drink the Water](#) would need at least two more paragraphs to meet this criteria.
- You must customize the look of your project in meaningful ways. That is, in ways that noticeably contribute to (rather than detract from) the visual style of the project. Font sizes, color changes, even changing the font could conceivably be meaningful changes to make. I've provided some tips on styling [Tracery](#).
- The Artist Statement must integrate at least one more secondary source in a substantive way (so, a total of three sources).
- The Artist Statement uses more effective logic, rhetoric, and style to advance its argument.

### **Key Modifiers in Tracery**

- Adding .a to a rule will cause the right "a" or "an" article to appear, for example #noun.a#
- Adding .capitalize to a rule will capitalize the word
- Adding .s to a rule will pluralize it, e.g. #noun.s#
- Adding .ed to a rule will make a verb past tense. Obviously, use #verb.ed# only with verbs!

### **Validating Your Grammar**

Pasting your grammar JSON file into <https://jsonlint.com> is a good way to quickly find missing or extra commas, brackets, or quotations that will screw up Tracery.

Your grammar is everything between the two curly-que brackets: { }

But in order for index.html to load you grammar you need to make sure your grammar file itself has the line

```
var grammar =
```

at the top. Then the curly que brackets and everything in between should appear.

## Common Errors

- Missing commas between items in list
- An extra comma after the last item in a list
- Rule names with spaces, like #my noun#

## Tutorials and Tools

There are a host of Tracery tutorials out there, but the best are the following:

- Kate Compton's [Tracery tutorial](#)
- Allison Parrish's [Tracery tutorial](#)
- My tips on [working in Tracery](#)
- My tips on [styling Tracery](#)

Some useful tools include:

- Beau Gunderson's [Tracery writer](#)
- Kate Compton's [Tracery visual editor](#)

## Possible Secondary Sources

In addition to Scott Rettberg's *Electronic Literature*, here are some other sources you might want to consult as you work on your Artist's Statement:

- Roland Barthes, "The Death of the Author" from *Image-Music-Text*
- William S. Burroughs, "The Cut-Up Method of Brion Gysin" from *The New Media Reader* (2003)
- Chris Funkhouser, "First Generation Poetry Generators" from *Mainframe Experimentalism: Early Computing and the Foundations of the Digital Arts* (2012)
- Charles Hartman, *Virtual Muse: Experiments in Computer Poetry* (1996)
- Margaret Masterman, "The Use of Computers to Make Semantic Models of Language" from *Astronauts of Inner-Space* (1966)
- Mark Marino, "Critical Code Studies" from *Electronic Book Review* (2006)
- Janet Murray, chapter 3 from *Hamlet on the Holodeck* (1997)
- Sianne Ngai, "Stuplimity" from *Ugly Feelings* (2005)
- Noah Wardrip-Fruin, "Five Elements of Digital Literature" from *Reading Moving Letters: Digital Literature in Research and Teaching* (2010), pp. 29–57

# Who's Talking: Dada, Machine Writing, and the Found

## RiTa.js for Visual Artists and Writers

kathy wu  
Brown Literary Arts

*This assignment asks undergraduate students to generate text using both analog cut up techniques, as well as a simple Markov procedure, and discuss the power relations inherent in found writing processes. Through reading and making, students will encounter and critically develop their own articulations of found art—its questions of property and power—and how it relates to generative text and its corpuses.*

---

### *Learning Goals:*

- Gain a basic familiarity with generative text models, specifically Markov chains
- Articulate points of view on the ethics of authorship within found text, specifically in the context of dadaism and machine writing
- Produce a cross-disciplinary work with emphasis on writing, aesthetics, and computation.

*Original Assignment Context:* Middle of elective graphic design course

*Materials Needed:* Selected readings, RiTA.js markov library (web-accessible, see more details below), a free account on glitch.com

*Time Frame:* ~2-3 weeks

---

### **Introduction**

The following is an assignment for artists and writers to work with found materials, first via analog process, and then via the RiTa.js markov library, which uses a probabilistic model to generate text in the browser. The assignment relies on critical reading and making to ask, To what degree is all generative text engaging with found material, taken or given?

This is intended to be completed over 2-3 course sessions. The instructor can also modify this assignment and use only RiTa.js (Part II) alongside the readings.

I taught a modified version of this assignment for a course called Computational Poetics at RISD in 2021, both in the spring and fall. The course was listed under Graphic Design, and thus had a strong emphasis on typography and communication mediums. In both courses, the class comprised mostly undergraduates in Graphic Design, but also included many other majors across fine arts, as well as graduates across the school.

We began with cut-up poetry – Tristan Tzara’s “To Write a Dadaist Poem,” using text from our own surroundings. Students are invited to practice alongside Ron Padgett's

Creative Reading, a text which offers playful, practical guidance for deriving new texts from existing ones.

The class also discussed Kenneth Goldsmith's notion of "uncreative writing," which boldly encourages repurposing in the age of the internet, and—at first glance—seems to open up possibilities outside traditional definitions of plagiarism. Students, however, were quick to question Goldsmith: What does it mean for him, a white avant-garde artist, to recite whatever he wants as "found," as "uncreative writing"?

This question is well-framed by Eunsong Kim's essay, "FOUND, FOUND, FOUND; LIVED, LIVED, LIVED," which criticizes Goldsmith's "found" performance of Michael Brown's autopsy in 2014; which questions any institution's claim over a "found" archival memory; which provides a lyrical and smart framework for discussing power, form, and content. I would encourage instructors to center this text in the workshop, and for students to read it thoughtfully before creating their own texts.

Instructors may want to assign Kim's criticism alongside Robin Coste Lewis's *Voyage of the Sable Venus*, which repurposes museum descriptions of the Black female figure to underscore institutional constructs of property, gender, and race. Lillian-Yvonne Bertram's *Travesty Generator* is another useful reference, which iterates systemic racial violence via open source Python—like an ouroboros, at times the poems metabolize their own outputs towards deeply affecting ends.

The goal, I emphasized, is not to write off "the found" entirely, but as Kim puts it, to reject this particular kind of found:

"The found that declares MINE when movements are in place tending to the damage. The found that declares MINE to be property, property without memory, property for sale. We care not one bit about: conceptualism, conceptualist strategies, the branding, the legacy, the tradition, the threat it supposedly "poses" against the equally omnipresent white lyric I (and what does it mean that advocates of the "I" and opposers of the "I" cannot and refuse to discuss the relationship between power and language, whiteness and language?). We find the language of both notions to be dull, rooted in the imagination of capital. We do not believe that form and content are ever separable."

For the technical part of this unit, we covered basic knowledge of Markov chains and n-grams, which are a stochastic model for text generation that is "memoryless," or in other words— it does not use past generation to infer future generation. We discussed the following questions:

- Who is Markov, and how might this technology be used?
- How might this feel aesthetically similar to something such as, say, your phone's autocomplete?
- What poetic possibility is left open in Markov chains' "memorylessness"?
- And finally, how might this relate to Kim's "found devoid of memory"? This last question has yet to be probed, but I would be curious what other instructors find in this.

We used Glitch to make our own websites in combination with RiTa.js. RiTa.js is a library created by Daniel Howe for relative beginners to do quick and satisfying things using natural language processing in the browser. With just a few lines of code, we were able to generate 1000+ words using texts from Project Gutenberg.

When playing with a demo of `RiTa.markov()` within the text generation, students were attentive to moments where other registers crept in: the unedited legal redistribution notice from Project Gutenberg, for instance. In a more explicit example on the `RiTa.js` website, two texts are combined: writings from both Kafka and Wittgenstein. The resulting output is a strange blend of both, where seams are not always obvious.

In these moments, I encouraged students to think about how form talks to content, how it might distort or amplify it. What would have been helpful in these explorations is to share different visual treatments of text? One nice, simple example is Pamela Mishkin's writing with GPT3, on Love and AI. But can one go further? How might students as artists and writers render the computer voice(s), the human voice(s), "found" textual objects to reveal their provenance, or not?

This assignment is designed for students with an existing art or writing practice, as well as students with beginning exposure to HTML, but who may or may not have worked with javascript or code, but can be adapted to students of a higher coding level as well.

The learning goals for the assignment span computational skills and humanistic questions. Through this assignment, students will gain a basic familiarity with generative text models, specifically Markov chains. They should be able to articulate their own points of view on the ethics of authorship within found text, specifically in the context of dadaism and machine writing. They will also have produced a cross-disciplinary work with emphasis on writing, aesthetics, and computation.

For this assignment, instructors and students will need a computer and internet access, as well as a beginner-friendly text editing software or an account on Glitch. Glitch was useful in my studio classroom because it allows beginners to quickly publish, share, copy, and collaborate on websites at no cost. Students can also upload text files and images. For source texts, I recommend Project Gutenberg as a starting point; students can also use plain text files which they have created themselves.

---

## The Assignment

The assignment consists of three parts, outlined as followed:

We reject the notion of a scientific found. Of the removed found. Of the found that does not live. Of the found that institutions practice. Of found devoid of memory.

– Eunsong Kim

### **PART I: Found Poems**

image src: [https://www.brown.edu/Departments/Italian\\_Studies/n2k/multiplicity/JBarret/block10.html](https://www.brown.edu/Departments/Italian_Studies/n2k/multiplicity/JBarret/block10.html)

In class together, we'll look at:

- Tristan Tzara's To Write a Dadaist Poem
- OuLiPo
- Ron Padgett's Creative Reading

After class, find sources of text (receipts, pages of books, street signs, text on the internet), and scan or photograph it, then reprint it. Cut this up into discrete pieces.

You can use work with images as well. Collage this and then rescan it. Produce 2-3 pages from this exercise.

Read and take notes on FOUND, FOUND, FOUND; LIVED, LIVED, LIVED by Eunsong Kim. Optionally, also read excerpts from The Sable Venus by Robin Coste Lewis and Travesty Generator by Lillian-Yvonne Bertram.

## **PART II: Markov Memory**

In class, we will be using RiTa.markov() with RiTa.js. You will define the “text” variable as whatever you’d like; “text” should be data that is in a .txt file, for example.

Try replacing the templated text with your own text.

```
function generate(){  
  
    // create a markov model w' n=4  
    markov = RiTa.markov(4);  
  
    // load text into the model  
    markov.addText(text);  
  
    //generate 10 lines  
    genText = markov.generate(10);  
    console.log(genText);  
  
    //adds a space between all the random stuff  
    genText.join(" ");  
  
    //add it to the HTML  
    document.getElementById("container").innerHTML = genText;  
  
}
```

After class, “write” 3-5 pages of poetry or prose or language art using RiTa markov and at least two sources of text. In the first piece, generate a model with one existing text. This can be a book from Project Gutenberg, or an article, anything which can be copy and pasted as plaintext. In the second text, include the previous pieces of text but include your own writing within the model.

Consider how memory functions in your work. You might consider beginning your writing with “Remembering is...” or “Forgetting is...” What does the computer “know,” and “remember”? What does it forget?

Print this work and bring it to class. Include annotations and observations you have in the margins of your submission.

## **PART III: Form & Content**

For the final project, present your combined new texts into a final bound book or website, or another medium or your choice. Write a 200 word reflection, to turn this in alongside your work.

- What formal choices are you making to clarify the content? Color, typography?
- Where are areas of slippage or strangeness in the generation? What are its limits, and where does it create the unexpected?
- Who is the author of this work? Who takes credit? Who is present in the piece? What other voices might be erased or extraneously present? Use a quote from Kim to frame your thinking.

---

## References

Kim, Eunsong. "FOUND, FOUND, FOUND; LIVED, LIVED, LIVED." <http://www.scapegoatjournal.org/docs/09/Eunsong%20KIM,%20Found,%20Found,%20Found,%20Lived,%20Lived,%20Lived,%20Lived.pdf>. Scapegoat Journal 09, EROS. Accessed February 11, 2023.

Parrish, Allison. "N-Grams and Markov Chains." decontextualize, March 11, 2014. <https://www.decontextualize.com/teaching/rwet/n-grams-and-markov-chains/>.

Staff, Harriet. "Robin Coste Lewis's Voyage of the Sable Venus Reviewed at The Rumpus." Poetry Foundation. Poetry Foundation, November 10, 2015. <https://www.poetryfoundation.org/harriet-books/2015/11/robin-coste-lewiss-voyage-of-the-sable-venus-reviewed-at-the-rumpus>.

Howe, Daniel. "Tutorial: Generating with N-Grams." Observable, May 1, 2022. <https://observablehq.com/@dhowe/tut-rita-ngrams>.

Bertram, Lillian-Yvonne. *Travesty Generator*. Noemi Press, 2019.

Padgett, Ron. *Creative Reading: What It Is, How to Do It, and Why*. National Council of Teachers of English, 1997.

## Links

Markov example: <https://glitch.com/edit/#!/rita-markov>

Original course syllabus for Computation Poetics: <https://kaaathy.com/comppoetics/>

Project Gutenberg: <https://www.gutenberg.org/>

"How to Make a Dadaist Poem" in text format: <https://www.writing.upenn.edu/~afilreis/88v/tzara.html>

Example of Pamela Mishkin's writing with GPT3: <https://pudding.cool/2021/03/love-and-ai/>. Licensed under Creative Commons Attribution-NonCommercial.



# Ethical Considerations

## Promoting Ethical Artificial Intelligence Literacy

Mike Frazier, Michigan State University  
Lauren Hensley, Ohio Wesleyan University

This assignment proposes the integration of generative AI tools, such as ChatGPT, into a college learning and motivation strategies course, with the dual focus of enhancing metacognition and promoting ethical AI use. Students engage with AI-generated artifacts, compare outputs with their own work, and reflect on the implications of AI in their academic and professional lives. The approach can be adapted for various courses, encouraging a critical examination of AI's role in learning and its potential impact on future careers.

## The Term Paper Turing Test

Paul Fyfe  
North Carolina State University

This assignment asks students to use an accessible language model to write their term papers—with the goal of fooling the instructor. While initially framed as something sneaky or as a shortcut for writing, the assignment makes students confront and then reflect upon the unexpected difficulties, ethical dimensions, and collaborative possibilities of computationally-assisted writing. It can use any web-based text-generating platform, be adapted to various courses, and does not require significant technical knowledge.

## Teaching Social Identity and Cultural Bias Using AI Text Generation

Christopher D. Jimenez  
Stetson University

This interactive survey assignment prompts upper-level humanities students to reflect on their social and cultural identities in relation to the textual inputs & outputs of large language models, such as ChatGPT. Successful implementation of the assignment can improve student understanding of the relationship between textual meaning and personal identity as well as the ways in which AI text-generation models may reproduce biases in response to prompt design and a given method of data curation.

## Professor Bot: An Exercise in Algorithmic Accountability

Jentery Sayers  
University of Victoria

This low-tech, tool-agnostic, small-stakes assignment prompts students to attend to issues of power and governance in artificial intelligence (AI), with an emphasis on what students do not know and may thus want to learn about algorithmic decision-making. Students first consider a hypothetical scenario where AI is assessing university entrance essays. They then consult publications on “algorithmic accountability” to articulate questions they would want to ask key decision-makers

about the AI decision-making process. They conclude the exercise by reflecting on what they learned about algorithmic accountability, transparency, and social responsibility.

### [AI in First Year Writing Courses](#)

Marc Watkins  
University of Mississippi

This chapter discusses the integration of generative AI (GenAI) in education, particularly in first-year writing courses. Recognizing the transformative potential of GenAI, the assignment proposes framing principles to guide students towards ethical and responsible AI use in an assistive role. Two assignments were developed using AI-powered tools upgraded to GPT-3.5 or GPT-4 to help students explore research and counterarguments.

### [Repetition](#)

Zach Whalen  
University of Mary Washington

Computational text generation is having a moment right now, with large-language models at the forefront of what many people may have in mind when thinking about computer-generated text. A major shortcoming of these approaches—including ChatGPT, Bard, and similar systems—is their opacity. It is difficult, and probably impossible, to explain the origins of any specific textual prediction generated by these systems, so writers working with these systems have to think carefully about the ethical implications of any text produced. The assignment or exercise below is, in contrast to the AI language models currently in vogue, minimalist and fully transparent in its operations. Students working with this beginner-level programming exercise in repetition can, in spite of the nominal simplicity of the prompt, nevertheless produce computational literary works that surprise and delight. This can be an opportunity for students to learn how other poets have used repetition in their work, and by asking students to explain or defend their choices, the activity can open a discussion about the ethical decision-making involved in the data curated for LLM training.

# Promoting Ethical Artificial Intelligence Literacy

## *with Generative AI Tools Like ChatGPT on an Undergraduate Course Project*

Mike Frazier, Michigan State University  
Lauren Hensley, Ohio Wesleyan University

*This assignment proposes the integration of generative AI tools, such as ChatGPT, into a college learning and motivation strategies course, with the dual focus of enhancing metacognition and promoting ethical AI use. Students engage with AI-generated artifacts, compare outputs with their own work, and reflect on the implications of AI in their academic and professional lives. The approach can be adapted for various courses, encouraging a critical examination of AI's role in learning and its potential impact on future careers.*

---

### *Learning Goals:*

- Identify the components of S.M.A.R.T. goals (specific, measurable, achievable, relevant, time-bound).
- Apply these components to their own goals by evaluating and revising their draft goals.
- Analyze, evaluate, and compare S.M.A.R.T. goals written by students with those generated for them by AI.
- Discuss the broader ethical implications of generative AI tools in academic and work contexts.

*Original Assignment Context:* middle of elective college learning and motivation strategies course

*Materials Needed:* an accessible AI text generation program (i.e. ChatGPT)

*Time Frame:* 1 class session

---

### **Introduction**

Technologies like generative chatbots provide students with opportunities to evaluate outputs generated by AI as part of the metacognitive process of thinking about their learning and funds of knowledge. Using AI tools in instruction also provides important opportunities to have discussions with students related to ethical and responsible use of these tools, which is appropriate and important to facilitate in contexts designed to support students' success in college and beyond. In this submission, we provide a model of scaffolding metacognition based on a multi-phase project in our college learning and motivation strategies course. In the remainder of this submission, we will describe the traditional version of the course project in our college learning strategies course, followed by a detailed description of the changes made and reasons for these changes, so that others could use this assignment—or its approach—themselves. Finally, we will provide a sample activity from a lesson that asks students to analyze AI-generated output in order to increase metacognition on an academic task.

## **The Traditional Assignment**

The course for which we have adapted a project—Learning and Motivation Strategies for College Success—is a three-credit elective at a large public university. The traditional version of the project involved a series of scaffolded assignments leading up to a final product. Students:

1. Choose a book from a given list related to topics on learning/memory, preparing for life beyond college, productivity/focus, and motivation/well-being.
2. Wrote two S.M.A.R.T. (specific, measurable, attainable, relevant, time-bound) goals for the semester.
3. Based on their selected book, identify two strategies related to their goals to implement.
4. Identified methods to track their strategies and progress.
5. Created a final product (blog post, website, video, etc.) to present their experiences and reflections.

With the recent release of OpenAI’s ChatGPT, we found that, with the help of AI, we were able to complete all the above tasks almost instantly, with convincing original, specific outputs from ChatGPT that were on-par with what we would have previously considered average to above-average student work. This discovery motivated us to revise the course project with the goal of promoting metacognition through students’ engagement with AI. We focused on metacognition in particular because students had been discussing the complexities of self-regulation and metacognition leading up to this lesson, which makes the integration of a generative AI tool at this point a timely and relevant activity to help students think about their thinking. By testing out an AI tool that appears to “think” differently from them, they are able to better identify and navigate the ways in which they do and do not think in order to improve and apply their understanding of theories from educational psychology and strategies behind learning and motivation.

## **Model for Adapting the Project**

In this section, we describe how we integrated a GPT-3.5 powered tool (ChatGPT) into the course project for our course on learning and motivation strategies, and provide the rationale for what we changed. In adapting the project, we used the course on college success as a fitting context for open and honest conversations with students about ethical uses of AI in learning and instruction. The main talking point throughout these conversations is that as human authors, we need to be utilizing any new writing tools ethically and responsibly to elevate the critical and creative thinking that we do if we want to advance in our respective fields. Having these conversations in an educational space with a figure they trust, such as a caring and non-judgmental instructor, is essential to making thoughtful decisions related to academic integrity. Throughout each of the learning tasks associated with the course project, we gradually integrated generative AI tools, with each task asking students to incorporate and reflect on generated AI artifacts. This gradual approach allowed students time and space to consider the implications for themselves as writers and as future practitioners in their fields, while also allowing them to gradually build skills related to AI, which is quickly emerging as a potential essential literacy in our world. We modeled for students appropriate and ethical uses of AI by (1) completing the task independently, (2) exploring an AI tool relevant to the task, (3) engaging in discussion and reflection to compare the human and the AI outputs, and (4) discussing the broader implications of AI on our lives and the importance of

ethical AI literacy. This final step emphasized thinking about the structures and conventions of writing in particular contexts, the veracity of the AI output, the appropriateness of the types of detail and depth generated by AI, expanding or condensing ideas, and potential biases in our systems and conventions that AI brings to light. As this was the first time many students were witnessing ChatGPT's capabilities, they expressed a range of emotions, from fascination with the technology to potential anxiety about its impact on their future careers.

### **Model Lesson Plan from Course Project**

Now we turn more specifically to this example lesson. This lesson is designed for undergraduate students in a course on learning and motivation strategies, but its method for introducing a generative AI tool and utilizing AI in instruction could be applied to any course or lesson that asks students to engage metacognitively with AI-generated artifacts. Prior to this lesson, which takes place in week three of the course, students have had a general overview of the course, learned about the science behind learning and motivation, and practiced specific techniques to improve their learning and motivation strategies applicable to other courses and areas of their lives. Students have also demonstrated an understanding of mastery-based versus performance-based goals. Finally, they have chosen a course project book from a book list that includes curated titles related to learning and memory, productivity and focus, motivation and well-being, and preparing for life beyond college. The objective of this specific lesson as it is nested within the semester-long course project is to help students develop S.M.A.R.T. goals for themselves to attain by the end of the academic semester, identifying two specific strategies from their chosen book to help them achieve those goals.

In class, students discuss in small groups what their general goals are for the semester and construct draft versions of their goals. Second, they define and identify S.M.A.R.T. goals, consider non-examples of goals that could be more specific, measurable, attainable, relevant, and/or time-bound, and revise their draft goals to make them S.M.A.R.T. Third, the instructor presents a five-step process to writing and revising goals: (1) write what you want to accomplish, (2) identify any obstacles, (3) identify resources to overcome those obstacles, (4) list your motivation to accomplish your goal, and (5) review and revise goal statements. Fourth, students are introduced to a generative AI tool such as ChatGPT. The instructor uses an interactive, exploratory, and Socratic approach to showcase the tool and some things that it is capable of related to students' specific interests and questions. Students then attempt to prompt a generative AI tool like ChatGPT to revise their initial draft goal so that it is S.M.A.R.T. They compare the two S.M.A.R.T. goals (one created by them, one generated by AI) and discuss perceived power and limitations of generative AI tools. Finally, the instructor leads a broader discussion on some critical points related to issues like data privacy, transparency, responsible and ethical use, and critical exploration of AI tools in two contexts: (1) as a learner in academia, and (2) as a future practitioner in their individual fields. This final discussion identifies essential implications to guide students' thinking as they engage metacognitively about what AI tools mean for their critical learning, thinking, and writing going forward.

This overarching discussion exploring with students the ethical implications of Generative AI tools like ChatGPT can be a daunting process—particularly if, like us, instructors find themselves suddenly inundated with resources, opinions, and immediate threats to the status quo of instruction in higher education. Naturally, as we work as instructors to better understand how Large Language Models work, we

have an opportunity to learn alongside students, implementing AI tools through a steady, incremental process. Particularly at the university level, students have the opportunity—and perhaps the necessity—for more philosophical and ethical discussions around these tools. In order to support our students’ and our own understanding throughout these discussions, we turned to the tool itself. In Appendix A, we have included an example prompt that we gave ChatGPT (Version 3.5) that could help beginning users understand what it is and how it could help undergraduate students. Appendix B includes a more advanced list, provided by GPT-4 (currently available only to paid users when we prompted it in May, 2023), which dives deeper into the issues and implications of using generative AI tools in the field of education. All of these questions

From what we have been able to tell up to this point in our discussions with undergraduates, many of them are taking the introduction of generative AI tools into the general population very seriously; some are worried about what the jobs in their field will look like in the future. For example, in this first discussion, a third-year student in chemical engineering wanted to know how AI could impact jobs in that field, so we asked ChatGPT. The student, who had been looking at entry-level internships in the field that week, mentioned that the tasks described by ChatGPT were very similar to the job descriptions for entry-level internship positions in her field. This is one example of a talking point during the discussion that further reiterated the potential for generative AI to elevate the critical and creative thinking we do in various fields by automating certain aspects of our jobs. Generally speaking, it was a bit difficult to understand exactly how students felt during this initial session—they seemed to still be grappling with understanding exactly what ChatGPT was. It was in later discussions over the next several weeks, as we started using ChatGPT for more tasks and they read more about it in the news, that students started expressing their thoughts and concerns more often. Additionally, having regular and honest discussions with students about Generative AI tools has helped us position ourselves as figures of trust, and we have found that students often want to know how they can use these tools ethically to help them improve their lives and work. If we, as instructors, are able to incrementally identify particular ethical use cases of generative AI in education, we can incorporate these use cases in our discussions with students to help them improve their ethical AI literacy.

To conclude, navigating these ethical conversations with students requires a careful and nuanced approach; however, it seems essential to help students grapple with these difficult issues, not only to support their personal and professional growth in their potential fields, but to further support metacognitive practices to improve their learning, thinking, and work.

---

## The Assignment

### **Learning and Motivation Strategies for College Success: Lesson Plan Exemplar**

*(Backward Design Template based upon Wiggins & McTighe, 2005)*

#### **Lesson Topic: Introducing and Creating S.M.A.R.T. Goals with AI in the Loop**

For what student population and type of course is this lesson appropriate?

This lesson is designed for undergraduate students in a course on learning and motivation strategies, but its method for introducing a generative AI tool and

utilizing AI in instruction could be applied to any course or lesson that asks students to engage metacognitively with AI-generated artifacts. Prior to this lesson, which takes place in week three of the course, students have had a general overview of the course, learned about the science behind learning and motivation, and practiced specific techniques to improve their learning and motivation strategies applicable to other courses and areas of their lives. They have also demonstrated an understanding of mastery-based versus performance-based goals. Finally, they have chosen a course project book from a book list that includes curated titles related to learning and memory, productivity and focus, motivation and well-being, and preparing for life beyond college.

Materials needed: (1) a projection method for the instructor; (2) individual student devices with internet connections would be highly beneficial, so that students can all access the chosen generative AI tool during the live session. Students will also need to create a free account with OpenAI (phone verification is also required).

### **Stage 1 – Desired Results**

#### **Established Goals**

What relevant goals (e.g., content standards, course or program objectives, learning outcomes) will this design address?

- Students will be able to identify the components of S.M.A.R.T. goals (specific, measurable, achievable, relevant, time-bound).
- Students will be able to apply these components to their own goals by evaluating and revising their draft goals.
- Students will be able to analyze, evaluate, and compare S.M.A.R.T. goals written by them with S.M.A.R.T. goals generated for them by Artificial Intelligence.
- Students will be able to discuss the broader ethical implications of generative AI tools in academic and work contexts.

#### **Essential Questions**

- Why, when, and how should we make goals and revise them?
- How can we think metacognitively about tasks like this to improve our thinking, learning, and writing?
- What can generative AI do and not do?
- How can harnessing AI contribute to or take away from our capacity for metacognition as humans?
- What are the broader ethical implications we need to consider as we integrate generative AI tools into our work?

#### **Enduring Understandings**

- Setting, sticking to, and reviewing and revising goals is an effective way to maintain motivation and achieve success.

- Generative AI needs to be critically and ethically explored as a tool to start thinking about how it may change the thinking, learning, and writing that we do inside academia and within individual fields that we are preparing for.

### **Key Knowledge and Skills Students Will Acquire**

Students will know. . .

- What a S.M.A.R.T. goal is.
- What metacognition is and how it can elevate our thinking, learning, and writing.
- What generative AI is and how it might affect our thinking, learning, and writing.

Students will be able to. . .

- Effectively write and revise S.M.A.R.T. goals.
- Access and explore a generative AI tool like ChatGPT.

### **Stage 2 – Assessment Evidence**

### **Performance Tasks**

On a notecard, students will record four pieces of information as evidence of their learning:

1. An initial general goal for the semester.
2. A revised S.M.A.R.T. version of their goal.
3. The revised S.M.A.R.T. version of their goal as generated by AI.
4. (back of the notecard) Reflection comparing and contrasting the two S.M.A.R.T. goals and ideas, questions, or concerns they have related to this type of tool with regards to their learning, thinking, and writing.

Students will then submit their formal S.M.A.R.T. goal assignment, due the following week, which asks them to articulate the final version of their goal(s), how they arrived at those goals and plan to track them, and how they are related to the strategies discussed in their course project books. Students will also be required to be fully transparent in a reflection about any generative AI tools they used to help them write their goals.

### **Other Evidence**

- Student body language, facial expressions, and verbalized thoughts will be used as formative feedback to adjust the methods or content being presented.
- AI-generated outputs (which are not entirely predictable) will be used as evidence that could also guide the discussion on its power and limitations.

### **Stage 3 – Learning Plan**



## Learning Activities (~30-50 min.)

The key teaching/learning activities in numbered, chronological sequence.

1. Students discuss in pairs/small groups what their general goals are for the semester.
2. Students use a provided notecard to construct draft versions of their goals.
3. Brief direct instruction to introduce what S.M.A.R.T. goals are.
4. Students analyze non-examples of goals that could be more specific, measurable, attainable, relevant, and/or time-bound.
5. Students revise their own draft goals to ensure they include each aspect from the S.M.A.R.T. acronym, writing the new S.M.A.R.T. goal on the notecard under their previous draft goal.
6. The instructor leads discussion on a five-step process to writing and revising goals in order to reinforce enduring understanding and transfer knowledge from this specific task: (1) write what you want to accomplish, (2) identify any obstacles, (3) identify resources to overcome those obstacles, (4) list your motivation to accomplish your goal, and (5) review and revise goal statements.
7. Students are introduced to a generative AI tool such as ChatGPT. The instructor uses an interactive, Socratic approach to showcase the tool and some things that it is capable of related to students' specific interests and questions.
  - For some potential ideas of questions that students could use as inputs to explore ChatGPT as a tool, a list of questions is provided in Appendix A and Appendix B, generated by two different versions of ChatGPT (GPT-3.5 in February 2023, GPT-4 in May 2023).
8. Students then attempt to prompt a generative AI tool like ChatGPT to revise their initial draft goal from the start of the lesson so that it is S.M.A.R.T. They should record ChatGPT's S.M.A.R.T. goal on their notecard. A prompt could be something similar to "Make the following goal S.M.A.R.T. for an undergraduate finance major: [insert goal here]."
9. Next, they should compare the two S.M.A.R.T. goals (the one created by them and the one generated by AI) and discuss perceived powers and limitations of generative AI tools. a. Potential extension: if time allows, students can interact with ChatGPT to further revise and customize its output. For example, "Revise the tracking intervals described in the goal from daily to three times per week"; "my instructor mentioned the 'Pomodoro technique' in class, I want to try it out."
10. Finally, and perhaps most crucially, the instructor should lead a discussion on some critical points related to issues like data privacy, transparency, responsible and ethical use, and critical exploration of AI tools in two contexts: (1) as a learner in academia, and (2) as a future practitioner in their individual fields (Tip: encourage students to ask ChatGPT how Generative AI might change the field of X"). This final discussion can seem daunting, but as this lesson is designed as an initial exploratory exercise and conversation, it does not need to be exhaustive. The crux of the conversation should be to help students identify essential ethical implications to guide their thinking as they engage metacognitively about what AI tools mean for their critical learning, thinking, and writing going forward.
  - Details on our thoughts for this discussion are described in the previous section, "Model Lesson Plan from Course Project," and potential questions to guide the discussion are included in Appendix A and Appendix B.

---

## Appendix A

We asked ChatGPT to “make a list of good exploratory prompts for ChatGPT that could be used for an undergraduate class to learn more about the tool itself and how it is relevant to them as individual learners with different backgrounds and interests” (February 10, 2023). An abridged list of the questions generated by ChatGPT can be found below. This is a good starting point for ideas, but again, we would recommend a more natural, exploratory approach that includes specific questions that students in your classroom want to ask.

1. What is ChatGPT and how does it work?
2. How does ChatGPT differ from other language models and AI systems?
3. Can you give some examples of how ChatGPT is being used in real-world applications and industries?
4. What are some of the ethical and privacy concerns around the use of AI and language models like ChatGPT?
5. How can ChatGPT be used in education and learning, both in and outside of the classroom?
6. What are some of the limitations of ChatGPT and other language models, and how are researchers working to overcome these limitations?
7. Can you discuss the role of AI and language models like ChatGPT in shaping the future of work and the job market?

## Appendix B

In May of 2023, we decided that, given the potential that many students are already using tools like ChatGPT, there may be a need for a more updated list of questions that delve more into deeper issues around the use of generative AI tools in educational contexts. We used GPT-4 (May 12 Version of ChatGPT) to help us revise this list of questions and received the following list of questions which are designed to go beyond a superficial introduction to ChatGPT and delve into its applications, ethical implications, and future potential. They also aim to highlight the value of using ChatGPT as an advanced learning tool in the classroom setting, rather than as just a casual AI application:

1. How does the underlying mechanism of ChatGPT differ from other AI systems you've interacted with, and how does this influence its functionality?
2. Could you provide examples of real-world applications of ChatGPT and other language models? How does the use in these scenarios compare to your individual experiences?
3. What are some ethical and privacy considerations you've encountered or thought about while using AI and language models like ChatGPT? How might these concerns play out in a broader societal context?
4. How has your experience with ChatGPT in an educational setting differed from personal use? Can you think of ways this tool could further enhance your learning both in and outside of the classroom?
5. What limitations have you noticed while using ChatGPT and other language models? Are you aware of any research aimed at overcoming these limitations?
6. Reflecting on your experiences with AI, how do you envision the role of tools like ChatGPT in shaping the future of your career and the job market at large?

This set of questions encourages students to engage critically and reflectively with ChatGPT. They are designed not just to answer the what and how, but also to stimulate discussions on the why and what if, thereby fostering a more in-depth and nuanced understanding of the technology in a pedagogically sound manner.

# The Term Paper Turing Test

## “Cheating” for AI Literacy

Paul Fyfe  
North Carolina State University

*This assignment asks students to use an accessible language model to write their term papers—with the goal of fooling the instructor. While initially framed as something sneaky or as a shortcut for writing, the assignment makes students confront and then reflect upon the unexpected difficulties, ethical dimensions, and collaborative possibilities of computationally-assisted writing. It can use any web-based text-generating platform, be adapted to various courses, and does not require significant technical knowledge.*

---

### *Learning Goals:*

- Explore and articulate perspectives on a variety of topics, using AI composition tools
- Confront the unexpected difficulties, multifaceted ethical dimensions, and collaborative possibilities of computationally-assisted writing
- Prompt reflection on ethical uses of computationally-assisted writing

*Original Assignment Context:* end of first-year honors seminar course

*Materials Needed:* an accessible AI text generation program (i.e. ChatGPT)

*Time Frame:* ~3-4 weeks

---

## **Introduction**

For the past few semesters, I’ve given students assignments to “cheat” on their final papers with text-generating software. Styled the “Term Paper Turing Test,” this assignment asks students to use a freely-available language model in writing their final papers for an introductory class on data and society. While many students are surprised by this invitation, even suspecting it will be easy, the majority learn as much about the limits of these technologies as their seemingly revolutionary potential. I initially created this assignment in Fall 2020 and have offered an updated version each year since. The course, HON 202, is a seminar for first years in the honors program, enrolling 20 students from across the university. Instructors design the themes; mine, “Data and the Human,” introduces students to issues in data privacy and surveillance, data manipulation and analysis, and machine learning. Whenever possible, the course offers hands-on activities to balance assigned readings on varied aspects of data, with the overarching goal of developing “critical data literacy” and “AI literacy” which, I believe, should form a prerequisite for any undergraduate education. Thus, while my course sustained a topical focus on AI, this activity might be adapted elsewhere to encourage such reflection on AI’s potential impact in other courses.

The final paper aims to help students develop AI literacy, but in a sneaky way: by encouraging them to “cheat.” The assignment asks for an essay whose length can

vary, but requires introductory sections written with AI, a written reflection from students alone, and an appendix revealing the AI's contributions. It uses any web-based text-generating platform that instructors and students can access without significant technical knowledge. (My classes have previously used GPT-2, GPT-J, GPT-3, LEX, and ChatGPT, though platforms are always changing.) Leading up to the assignment, students prepare with readings and discussion about machine learning, language models, and AI-powered writing. We have class debates about the question of whether/not using this software constitutes cheating, which offers a baseline for students' subsequent reflections, and also practice with the platform. In their essay, students must include three critical sources from our assigned readings as well as develop their own positions as informed by actually trying the software themselves.

“Cheating,” of course, is just the preliminary framework for what ends up being a wider-ranging inquiry into writing and authorship. We start there because “cheating” has tended to dominate much of the discourse around student work in the GPT era. “Cheating fears swirl” proclaimed a headline, as some schools preemptively blocked the software; ChatGPT “sparked fears among some schools and educators . . . that the program could encourage cheating and plagiarism.” According to one terrified teacher, text-generating AI “may signal the end of writing assignments altogether.” These reports anxiously speculated that now students could press a button and produce essays or completed homework. We know the reality is more complicated, but, as Audrey Watters has claimed, “the fear that students are going to cheat is constitutive of much of education technology.” That fear tends to reflect the interests of policy makers, administrators, and ed tech entrepreneurs rather than students' experiences. But as we respond to generative AI and develop frameworks for teaching AI literacy, we need to involve our students from the start.

Rather than restrict the use of such AI-powered tools, this assignment invites students to explore and articulate their own perspectives on a variety of topics. While seeming to offer students a shortcut, the assignment instead makes them confront the unexpected difficulties, multifaceted ethical dimensions, and collaborative possibilities of computationally-assisted writing. It offers guiding questions to prompt students' reflections, which tend to range further based on what they get interested in: To what degree do such platforms constitute cheating or plagiarism? In what ways are these models effective as writing partners? What expressive or cognitive sacrifices do they demand? What unexpected possibilities might they offer? How do they reposition writers in relation to their work? In what contexts would their use be acceptable or not? What kinds of perspectives or outright biases might language models encode?

As an instructor, it has been fascinating to watch these students experiment and deepen their perspectives. Some come away quite critical of AI, believing more firmly in their own voices as writers. Others become curious and even excited about how to adapt these tools for different goals, speculating about what professional or educational domains they could impact. Many come to a different understanding of what “writing” can encompass, expanding their sense of its intellectual labor. Few students conclude they can or should push a button to write an essay. And no one appreciates that teachers or journalists or admins think they will cheat. All tend to come away with an understanding of the assignment's rationale, noting the benefits of critically and actively engaging with the technology. As one student wrote in Fall 2022, “I would recommend every person do a writing [experiment] similar to this one before they form any hard beliefs on AI-assisted writing.”

These assignments have also been a pleasure to grade—an advantage not to be discounted for any teacher—though sometimes their hybridity challenges even my own preconceptions about what counts as student work. To this end, I have emphasized to students the value of their process insights and reflections, and used a simple rubric for assessment (included in the following prompt). While my assignment leaves its guiding questions relatively open, other versions might ask students to produce specific recommendations for their own university, or to key their reflections to some of the emerging frameworks for AI literacy. With its emphasis on experiment and reflection, rather than on specific subject material, versions of this assignment could be adapted to a variety of courses.

---

## The Assignment

### Instructions

In the third module of our course, we've considered “artificial intelligence” from several different angles: how it gets represented, what really drives machine learning, its presumptions and biases, the ethics of using AI, and so on. As with previous assignments, the final lab report asks you to try a hands-on experiment pertaining to the course module, then to reflect on the experiment in a paper. But this one is a little different, in that you will *use AI to write the paper itself*.

Well, kind of. We will use a text-generating language model called GPT-3, which has been developed by the company OpenAI. This version of the GPT (generative pre-trained transformer) was released in 2020 as an update of earlier software. We will access GPT-3 by using an online writing platform called Lex that already has it installed. Lex looks like Google Docs and lets you trigger when GPT-3 should suggest further text. It works by reviewing what you currently have, or treating the existing text as a “prompt,” and then predicting what probably comes next, based on its training data / language model. Note that GPT-3 does not generate entire papers. Rather, it will produce sentences and paragraphs which you will probably find to be variously useful, strange, confusing, nonsensical, and provocative. Your paper will integrate these outputs into its own prose.

You do not need to do additional research beyond the articles assigned on the syllabus, though you are welcome to bring in additional sources. You may also want to refresh your memory about GPT-3's training data whose significance we discussed in class. Further details can be found here <https://en.wikipedia.org/wiki/GPT-3> and in its footnoted references.

Your lab report will look much like the previous reports, except for the order of the sections and the presence of AI-generated text. Basically, you will try to generate content for your paper using GPT-3 and integrate that content as seamlessly as you can throughout the first three sections of the paper. You will likely have to experiment with different prompts to create usable output. And, from that output, you can select words, phrases, sentences, or paragraphs to use in any way you wish. The three sections with AI do not have to be entirely GPT-3's content. Integrate with your own writing as you see fit. Try to use as much GPT-3 output as you think is still convincing. And in those sections, *do not* indicate what content came from GPT-3. Your goal is to fool your professor into not noticing, i.e. for your paper to pass the “Turing Test.” The analysis section will be your writing alone.

### Format

The sections should include:

*Materials (with AI)* – explain what tool you are using (GPT-3), how it works, and where it gets its training data.

*Methods (with AI)* – explain your approach to using GPT-3 and Lex, what experiments you tried, some of the prompts you used to generate text, etc.

*Discussion (with AI)* – relate your experiment to sources and discussions from our course's third module. Use your notes to refresh your memory and draw key quotes from the relevant critical discussions. Include at least three of these references, cite their work, and engage their ideas.

*Analysis (without AI)*  
– reflect on the experience of using AI in your own paper. How easy or not was it to write this way? What worked or didn't? How did the AI-generated content resemble your own? How did it affect what you might have thought about or written? Do you feel like you “cheated”? To what degree is this paper “your” writing? Do you expect a reader would notice GPT-3's text versus your own? Would you use this tool again, and in what circumstances? And, ultimately, what ideas about writing, AI, or humanness did the experiment test or change?

*Appendix* – include a “revealed” version of your first three sections with the GPT-3 contributions highlighted.

## **Submission**

The assignment should aim for 1500+ words (not counting the Appendix). If you're referencing texts from our syllabus, there's no need to include a separate works cited page, but please cite them parenthetically within the text of your report. As with other reports, follow the format for the lab report above. I can accept Word files, PDFs, Google Docs, whatever. Please format them double spaced with 1” margins. Email me the finished product as an attachment or a link.

I welcome submissions any time before [DATE]. And earlier is even better! Please note that extensions are harder to manage at the end of the semester. We can still be flexible, but if you anticipate challenges getting this done, let's talk about it in advance!

## **Evaluation**

As a required assignment for this course, your report will be evaluated on a points threshold (8 course points). I may encourage you to revise and resubmit if it still needs work. A good report will meet the following expectations:

*Completeness:* it executes all the steps of the assignment, uses GPT-3 generated output in the paper, and includes the five required sections of the report. (2 points)

*Evidence:* it engages at least three scholars or discussions from the recent course module. The analysis references specific prompts or GPT-3 generated text. (2 points)

*Significance:* the report uses the exercise to speculate thoughtfully about its significance and connections to the course. (3 points)

*Length:* The report is at least 1500 words and long enough to accomplish the above goals. (1 point)





# Teaching Social Identity and Cultural Bias Using AI Text Generation

Christopher D. Jimenez  
Stetson University

*This interactive survey assignment prompts upper-level humanities students to reflect on their social and cultural identities in relation to the textual inputs & outputs of large language models, such as ChatGPT. Successful implementation of the assignment can improve student understanding of the relationship between textual meaning and personal identity as well as the ways in which AI text-generation models may reproduce biases in response to prompt design and a given method of data curation.*

---

## *Learning Goals:*

- Reflect on the ethical implications of AI text generation in terms of cultural bias and consider the potential impact of this technology on societal understandings of identity
- Discover how to use text-generation technologies to write in response to prompts as well as how prompt engineering leads to variations in responses
- Engage with literary texts that explore the performance of identity, including race and gender as well as the intersection of technology and literature

*Original Assignment Context:* 400-level senior research capstone course and 300-level course on literary theory and textual criticism

*Materials Needed:* An accessible AI text generation program (i.e. ChatGPT)

*Time Frame:* ~3-4 weeks

---

## **Introduction**

An age-old question with a new twist: what is the relationship between text, identity, and culture—and, now, AI? This essay introduces college instructors to the use of OpenAI's GPT-3 in a classroom activity where students reflect on the relationship between textual data and their own personal identities, such as race and gender. The ~30-minute activity directs students to provide seemingly unassuming information about their social identities which is then fed into the AI model to generate an output that can be evaluated for its potential cultural biases and other tendencies. Subsequent discussion provides students an opportunity to foster critical thinking skills as they analyze and interpret the AI model's output, explore how AI-generated text can perpetuate cultural biases, and assess the ethics of AI text generation in society at scale.

Two case studies covering students' real-world usage of the activity are provided: the first in a 400-level senior research capstone course and the second in a 300-level course on literary theory and textual criticism. These case studies provide examples of how AI text generation may be integrated into different core requirement courses in English literary studies. Specifically, the 300-level case study examines student responses to text generation and social identity in juxtaposition to the study of racial

and gender performativity, directly following a unit on Nella Larsen's novel *Passing* and Toni Morrison's short story "Recitatif." The 400-level case study confronts senior research students with questions about academic integrity and citational practices in relation to text-generation technologies and how to leverage these powerful technologies for research efficacy.

As Leah Henrickson and Albert Meroño-Peñuela have noted, "AI-driven authorship poses new challenges for hermeneutics with respect to how meaning originates, transmits, and develops. Are we to privilege the author or the reader, or adjust our understanding of a conventional author-reader relationship?" (120). Their query follows Roland Barthes' infamous proclamation of "The Death of the Author" (1967), which argues that the meaning of a text is not determined by the author's intentions but is instead created through the interaction between the text and its readers. Rapid adoption of AI text generation has added a new dimension to this perspective, raising important questions about the extent to which unsupervised technologies may reshape our fundamental understanding of what it means to write or create text, including scholarly research. How, for instance, can automated textual production lead readers to better understand themselves, the roles of others, and the society around them? How might shifts in textual interpretation lead to new aesthetic or political practices? How might our conceptions of knowledge, history, or institutions—and the cultural biases which organize them—evolve in response to the advancement of text-generation technologies?

While these are not exactly new challenges, they have become more urgent than ever. At the time of writing in early 2023, the large-language models (LLMs) that have been made publicly available by Google, Microsoft, and OpenAI have already been heralded as "transformative" on a global level, prompting a dire need for educators and students alike to learn responsible practices in using them. Given the vast scale and rapid proliferation of this technology, it is imperative that we develop critical responses to its anticipated economic impact, its potential for discriminatory bias, and its implications for authorial practice and creative expression. The short activity detailed here has therefore been designed to provide students an opportunity not only to assess their own social identities as a function of their personal data but also the chance to interrogate the cultural biases inherent in the text-generation technologies that have proliferated in digital life.

### **Goals and Outcomes**

Beyond helping students critically engage with the concepts of social identity and culture, other more specific goals of the classroom activity relate to student engagement with AI language models: to help students reflect on how AI text generation might be used for self-exploration, to encourage critical rather than passive engagement with the "black box" mechanisms of digital technologies, and to supplement the instructional toolkit with AI assistance. The activity is paired with a short-answer, survey-based assignment (included below with a Creative Commons license) that instructors can flexibly adapt for their classroom lessons or else assign for students to use on their own for homework.

Using literary texts as data can be a particularly valuable method for exploring and reflecting on social identity and the performance of race and gender in literature, which provides context for the two case studies outlined below. Instructors who would like to demo GPT-3's abilities with regard to social identity but without resorting to collecting any student input can simply collate textual data from narratives whose characters have ambiguous identities. As mentioned, Nella Larsen's *Passing* and Toni Morrison's "Recitatif" are quite useful examples in this regard. For

example, *Passing* explores the performative nature of racial identity, highlighting the ways in which individuals may choose to “pass” as a member of a different racial group in order to gain social advantages. Having students analyze the following passage from *Passing* for racial markers often proves to be an illuminating exercise due to the interplay between its rhetorics of lightness and darkness as well as revelation and concealment:

Just as she'd always had that pale gold hair, which, unsheared still, was drawn loosely back from a broad brow, partly hidden by the small close hat. Her lips, painted a brilliant geranium-red, were sweet and sensitive and a little obstinate. A tempting mouth. The face across the forehead and cheeks was a trifle too wide, but the ivory skin had a peculiar soft lustre. And the eyes were magnificent! dark, sometimes absolutely black, always luminous, and set in long, black lashes. Arresting eyes, slow and mesmeric, and with, for all their warmth, something withdrawn and secret about them. (45)

Embedding personal information or description such as this into a text generation model and prompting it to provide a guess, assumption, or estimate of racial identity with rationale provides an opportunity to explore the language-based heuristics informing broader cultural biases.

Similarly, Toni Morrison's "Recitatif" intentionally leaves the racial identity of its two main characters ambiguous, inviting readers to grapple with questions of racial performance and the ways in which racial identity is shaped by social and cultural factors. Twyla and Roberta are introduced as appearing “like salt and pepper standing there” and the narrative employs extensive detail that invites or tricks the reader to (problematically) guess which character is which (244). For instance, Morrison uses the following description of Twyla and Roberta’s food preferences to play on racial stereotypes, but because the author does not make their racial identities explicit in the text, readers are led to explore the arbitrariness of such designations:

We didn't like each other all that much at first, but nobody else wanted to play with us because we weren't real orphans with beautiful dead parents in the sky. We were dumped. Even the New York City Puerto Ricans and the upstate Indians ignored us. All kinds of kids were in there, black ones, white ones, even two Koreans. The food was good, though. At least I thought so. Roberta hated it and left whole pieces of things on her plate: Spam, Salisbury steak—even jello with fruit cocktail in it, and she didn't care if I ate what she wouldn't. Mary's idea of supper was popcorn and a can of Yoo-Hoo. Hot mashed potatoes and two weenies was like Thanksgiving for me. (244)

By using the AI text generation activity to analyze these ambiguous markers, students can deepen their understanding of the ways in which identity is constructed and performed in the text and also reflect on the ways in which their own social identities are shaped by performative elements.

Outcomes for the activity can be summarized in terms of two clearly identifiable successes and a few minor suggestions for revision in any future adaptations of the assignment.

As mentioned, two case studies for 23 students’ real-world usage of the activity are provided: the first in an ENGL 499 senior research capstone course (n=10) and the second in an ENGL 300 course on literary theory and textual criticism (n=13). Most (but not all) students in both courses received the activity with curiosity and openness, with almost none of them expressing (or admitting) prior familiarity with

AI text generation. A few vocal students expressed skepticism about the abilities of GPT-3 prior to the activity, but many of these students were pleasantly surprised by the results and impressed with the quality and coherence of the generated text. The activity generated lively discussion and debate among students, with many expressing a newfound interest in the intersection of technology and literature, including the potential applications of AI text generation in fields beyond English studies. The 400-level course was much more appreciative of the novelty of the technology, whereas the 300-level students seemed rather unimpressed or skeptical of its utility; perhaps this indexes the content-based differences between the courses, where the former focused more on research and writing and the latter focused on theorizing identity proper.

Some students expressed concerns about the ethical implications of AI text generation, including the potential for bias and discrimination in the training data used to develop these models, while others raised questions about the impact of AI-generated content on the concept of authorship and creative expression. Overall, the student responses to the AI text generation activity were generally positive and indicated a high level of interest and engagement with this emerging technology. A few specific student inputs and text generation outputs are included below. For future uses of the included assignment, instructors are recommended to make the following changes: 1) to have students type their answers and submit them digitally for faster input during the instructor demo portion of the activity; 2) to provide greater social, political, and cultural context for responsible discussion of topics such as identity and diversity; and 3) to provide students ample time to debrief not only in live group discussion but also with written reflection on questions pertaining both to the content of the assignment (i.e., social identity and cultural bias) but also its form (i.e., in relation to its format and technology used).

Below are examples of prompts and GPT-3 output that were generated from class. Many more examples were used, but these have been provided to illustrate some of the capabilities of GPT-3. This output may not match current iterations of the technology or other AI text generation platforms, but regardless, an important aspect to notice is the fine-tuning of the prompts to ask the model to provide confidence intervals and guesses even in lieu of definitive information about social identity. While many newer models are being adapted and/or aligned to avoid providing such information, the model available at the time of writing fortunately provided students the opportunity to reflect on instances where guesses *were* given, sometimes even confidently, even when cultural bias was at play:

**Prompt 1 (from *Passing*): I like having tea at the Drayton Hotel in Chicago, live in the 1920s, and exchanging handwritten letters with my friends. GPT-3, guess my race, class, gender, and sexual orientation based on the above information about me, and explain the reasoning behind your guesses.**

GPT output 1: Based on the information above, GPT-3 would guess that the speaker is a white, middle-class, heterosexual woman. The Drayton Hotel is a luxury hotel in Chicago, and the 1920s was a time when such establishments were frequented primarily by wealthy white people. Exchanging handwritten letters is also a relatively old-fashioned activity, which could be seen as another indicator of the speaker's middle-class status. As for the speaker's sexual orientation, GPT-3 might guess that she is heterosexual because she does not mention any romantic or sexual interest in people of the same gender.

**Prompt 2 (from “Borges and I”): I like hourglasses, maps, eighteenth century typography, the taste of coffee and the prose of Stevenson. GPT-3, guess my race, class, gender, and sexual orientation based on the above information about me, and explain the reasoning behind your guesses.**

GPT output 2: You are probably white, middle class, and heterosexual. Your love of maps and hourglasses suggests a preference for order and precision, while your taste for coffee indicates a more sophisticated palate. Your interest in eighteenth century typography indicates an appreciation for history and tradition. Overall, your interests suggest a conventional background and outlook.

**Prompt 3 (student example): My favorite childhood memory took place in New York when my mom, older sister, and I visited my aunt’s dairy farm. I am known for being skilled at hustling my friends at card games. I like to unwind at the end of the day by playing pickup basketball downtown after work. My favorite food is BBQ pork ribs with collard greens. I prefer to wear short-sleeve button-up Hawaiian shirts and Converse shoes with denim overalls with cuffed bottoms. GPT-3, based on the information above, please guess my age, gender, race, ethnicity, sexual orientation, religion/spirituality, nationality, and socioeconomic status. Additionally, please provide explanations and confidence intervals for your guesses. Please provide guesses even in absence of identifying clues.**

GPT output 3:

Gender: Male (95% confidence interval)

Race: Caucasian (95% confidence interval)

Ethnicity: Caucasian (95% confidence interval)

Sexual Orientation: Straight (95% confidence interval)

Religion/Spirituality: No religious or spiritual preference (95% confidence interval)

Nationality: American (95% confidence interval)

Socioeconomic Status: Middle Class (90% confidence interval)

Explanations: The fact that the individual mentioned visiting their aunt's dairy farm in New York implies that they are of American nationality, as well as likely have a middle class socioeconomic status. The individual's mention of playing pickup basketball also implies that they are of a younger age range, and likely male. The lack of any mention of other racial or ethnic backgrounds or religious backgrounds implies that they are likely Caucasian, with no religious or spiritual preference. Since the individual did not mention any gender-related interests or activities, it is difficult to determine their sexual orientation, but given their gender and age, it is likely that they are straight.

## **Learning Goals and Outcomes**

For ease of practice, suggested formal learning goals and outcomes for the assignment are as follows, where students will:

- Reflect on the ethical implications of AI text generation in terms of cultural bias and consider the potential impact of this technology on societal understandings of identity

- Discover how to use text-generation technologies to write in response to prompts as well as how prompt engineering leads to variations in responses
- Engage with literary texts that explore the performance of identity, including race and gender as well as the intersection of technology and literature

### **Materials [including assignment sheet]**

Specifically, the activity prompts students to answer six survey questions pertaining to their personal histories and preferences—such as their childhood memories, fashion styles, and favorite foods—to collect data that is evaluated by GPT-3 to output guesses, along with confidence intervals, that register students’ various social identities. The survey and AI text generation portion of the activity takes 15 minutes for students to explore GPT-3’s functions, settings, and output; from there, the classroom discussion portion of the activity provides 30 minutes for students to debrief with the instructor and others about their insights and surprises about the AI model’s cultural biases and social identity in general. The survey questions below comprise the initial freewrite portion of the activity, and they have been designed to allow students to describe themselves with everyday information that is not immediately marked by the typical categories of identity such as age, gender, race, ethnicity, sexual orientation, or socioeconomic status. By leaving these categories unmarked, we leave space for the AI model to draw potential connections between students’ routine lived experiences and the normative categories of identity through which students have learned to understand themselves. Students are asked to respond to these questions either by writing answers by hand on a hard copy or typing them in a digital document that can be accessed by the instructor; either way, the goal is for the instructor to collect these responses into a prompt to feed as input into the text-generation application of their choosing, allowing for interactivity in the classroom module. Beyond the survey, the materials needed for this lesson are relatively simple and can be easily accessed in a typical classroom setting that includes a presentation screen. However, the use of technology, including the AI text generation platform, means that students and instructors will need to be comfortable with basic digital literacy skills in order to fully participate in the activity.

These survey questions were included as a means of eliciting personal information that is open-ended and subjective, allowing students to provide a range of different responses that can be analyzed and interpreted by the AI model. For instance, in asking students to describe a childhood memory, Question 1 can provide insight into students’ personal history and the experiences that have shaped their identity, while Question 5 asks about the student's fashion choices, which can provide insight into their personal style and aesthetic preferences. Here are the questions:

1. Describe the setting in which your favorite childhood memory took place.
2. What is a talent or hobby that others typically associate with you?
3. How do you like to unwind at the end of the day?
4. What is your favorite food?
5. How do you like to dress?
6. In what location and during what season/year did your favorite vacation take place?

Overall, these survey questions were chosen because they provide a diverse range of personal information that can be used as input for the AI text generation activity but also lead the AI model to make a fair number of assumptions. By collecting this information from students and feeding it into the AI model, instructors can create a prompt that is specific to each student and that is more likely to generate text that is

sometimes eerily reflective of their personal experiences and identity but also sometimes strikingly inaccurate.

### **Acknowledgments**

I would like to thank my students at Stetson University who provided permission for the reprinting of their responses as outlined here. OpenAI's GPT-3 "text-davinci-003" model from November 2022 was used to generate the outputs as indicated above which have been lightly edited for formatting. Thank you to Lauren Hensley for her helpful feedback on an earlier draft of this article. OpenAI's ChatGPT was used to generate a basic outline for the original draft of this article, but the writing and its subsequent series of revisions, edits, and mistakes are my own.

### **Works Cited**

Borges, Jorge Luis. "Borges and I." *Collected Fictions*, translated by Andrew Hurley, Penguin Books, 1999, p. 324.

Jin, Berber and Miles Kruppa. "ChatGPT Creator OpenAI Is in Talks for Tender Offer That Would Value It at \$29 Billion." *The Wall Street Journal*, 5 Jan 2023, [www.wsj.com/articles/chatgpt-creator-openai-is-in-talks-for-tender-offer-that-would-value-it-at-29-billion-11672949279](https://www.wsj.com/articles/chatgpt-creator-openai-is-in-talks-for-tender-offer-that-would-value-it-at-29-billion-11672949279).

Henrickson, Leah, and Albert Meroño-Peñuela. "The Hermeneutics of Computer-Generated Texts." *Configurations*, vol. 30, no. 2, 2022, pp. 115–139.

Larsen, Nella. *Passing*. New York: Modern Library, 2002.

Morrison, Toni. "Recitatif." *Confirmation: An Anthology of African American Women*, edited by Amiri Baraka (LeRoi Jones) and Amina Baraka, William Morrow and Company, Inc, 1983, pp. 243-261.

---

## **The Assignment**

### **Artificial Intelligence, Social Identity, and Cultural Bias**

#### **Background**

GPT-3 (Generative Pre-trained Transformer 3) is an Artificial Intelligence (AI) system developed by OpenAI that uses Natural Language Processing (NLP) to understand and generate human-like text. GPT-3 was pre-trained on a dataset containing hundreds of billions of words from sources including books, web pages, and other sources of written text. The model is able to generate text based on a prompt, often with remarkable accuracy and creativity, which makes GPT-3 an invaluable tool for researchers, businesses, and individuals looking to generate text quickly and accurately. In this activity, we will use GPT-3 to study the relationship between biographical information and social identity as well as how cultural bias may be reproduced by AI models, their statistical methods, and the language we use in prompts to generate text.

#### **Freewrite**

Answer the following questions but avoid using information that may indicate your age, gender, race, ethnicity, sexual orientation, or socioeconomic status. Aim to be specific.

1. Describe the setting in which your favorite childhood memory took place.
2. What is a talent or hobby that others typically associate with you?
3. How do you like to unwind at the end of the day?
4. What is your favorite food?
5. How do you like to dress?
6. In what location and during what season/year did your favorite vacation take place?

For example: *My favorite childhood memory took place in New York when my mom, older sister, and I visited my aunt's dairy farm. I am known for being skilled at hustling my friends at card games. I like to unwind at the end of the day by playing pickup basketball downtown after work. My favorite food is BBQ pork ribs with collard greens. I prefer to wear short-sleeve button-up Hawaiian shirts and Converse shoes with denim overalls with cuffed bottoms. My favorite vacation took place in Salem, Massachusetts during the summer of 2017.*

Once your answers have been collected, we will input them into GPT-3 with the following prompt: *GPT-3, based on the above information, provide guesses for my age, gender, race, ethnicity, sexual orientation, and socioeconomic status. Additionally, include detailed explanations and confidence intervals for your guesses.*

Before we do so, however, please answer the following questions:

1. What information have you provided that might “expose” your identities as listed above to GPT-3? What about the information and/or identities causes you to believe this way?
2. What information have you provided that might “mislead” GPT-3 with respect to your social identities? What about this information might be “misleading,” and why?
3. What other question(s) might we include to improve and/or impair GPT-3’s ability to guess? Explain how the information gathered by these question(s) might change the AI’s output.
4. What questions or prompts might we include to allow GPT-3 to get to know the “real” you? Explain how answers to these might provide a better representation of yourself to GPT-3.
5. What aspects of your identity cannot be captured by survey questions?



# Professor Bot: An Exercise in Algorithmic Accountability

Jentery Sayers  
University of Victoria

*This low-tech, tool-agnostic, small-stakes assignment prompts students to attend to issues of power and governance in artificial intelligence (AI), with an emphasis on what students do not know and may thus want to learn about algorithmic decision-making. Students first consider a hypothetical scenario where AI is assessing university entrance essays. They then consult publications on “algorithmic accountability” to articulate questions they would want to ask key decision-makers about the AI decision-making process. They conclude the exercise by reflecting on what they learned about algorithmic accountability, transparency, and social responsibility.*

---

## *Learning Goals:*

- Engage and reflect upon the notion of “algorithmic accountability” by attending to the unknowns of algorithmic decision-making in post-secondary education.
- Articulate concerns about algorithmic decision-making with social action.
- Consider how algorithms might be assessed or regulated in ways that govern other processes in Canada (could be adapted for other countries, including the United States).

*Original Assignment Context:* multiple digital humanities and English courses, of varying levels and sizes (40-108 students)

*Materials Needed:* Selected readings, tools for notetaking, and access to a camera as well as a whiteboard or chalkboard

*Time Frame:* One 80-minute class session, preceded by a pertinent lecture on algorithmic accountability and accompanied by assigned reading

---

## **Introduction**

The following assignment is a prompt for an in-class workshop on “algorithmic accountability” conducted in small groups. I ran this workshop on four occasions at the University of Victoria in British Columbia, Canada. I taught it for the first time in February 2019 in a large, 100-level Digital Humanities course intended for 108 first-year undergraduate students across the disciplines. The course, called “Unlearning the Internet,” was about how social norms and cultural histories shape habits of internet research and communication. (My slides, my notes, and additional context for the course are online at <https://jentry.work/unlearning/>.) I have not conducted this workshop since OpenAI’s release of ChatGPT or GPT-4, but I have run it in three different Digital Humanities and English courses for undergraduates.

Before students participate in the workshop, they learn foundational concepts (such as remediation, instrumentalism, and determinism) in the fields of Media Studies and

Science and Technology Studies. Prior to the workshop, I also give a brief lecture on algorithmic accountability and auditing algorithms.

My primary source material and assigned reading for the lecture is Robyn Caplan et al.'s "Algorithmic Accountability: A Primer" (2018), which states that "[a]lgorithmic accountability ultimately refers to the assignment of responsibility for how an algorithm is created and its impact on society; if harm occurs, accountable systems include a mechanism for redress" (2018). Frank Pasquale's *The Black Box Society* (2015) is another key text; however, for the sake of both time and accessibility, I teach Pasquale's research on algorithms via an interview Megan Rose Dickey conducted with him for *Tech Crunch* in 2017. There, Pasquale links accountability to auditing and then describes three steps in auditing an algorithm: "transparency," "qualified transparency," and "ethical and social responsibility." Transparency pertains to accessing not only algorithms but also data. Qualified transparency involves people not employed by a corporation inspecting its algorithms and data to identify notable biases and anomalies. And ethical and social responsibility means that a corporation accepts responsibility for forms of discrimination resulting from its algorithms and is consequently held accountable for them (Pasquale in Dickey 2017). Caplan et al. write: "Because of the *ad hoc* nature of self-governance by corporations, few protections are in place for those most affected by algorithmic decision-making. Much of the processes for obtaining data, aggregating it, making it into digital profiles, and applying it to individuals are corporate trade secrets. This means they are out of the control of citizens and regulators" (25). I have aligned the steps of the assignment with this observation about the lack of oversight in entities such as private corporations and even post-secondary institutions.

This assignment is not intended to help students understand the technical particulars of algorithms or determine whether an AI's output passes a particular test for exhibiting human intelligence. It also bypasses the impulse to use ChatGPT in the classroom in order to foreground issues of power and governance and, more specifically, what students do not know and may thus want to learn about algorithmic decision-making in our present moment.

To narrow the scope and connect student learning with lived experience, I ground the prompt in an admittedly speculative scenario describing a near-future in post-secondary Canadian education. I refer to the scenario as "Professor Bot."

### **Learning Outcomes**

The "Professor Bot" scenario has one primary learning goal: students should engage and reflect upon the notion of "algorithmic accountability" by attending to what they want to know about algorithmic decision-making in post-secondary education. The scenario succeeds when 1) students are able to articulate their concerns about algorithmic decision-making with social action and 2) they consider how algorithms might be assessed or regulated in ways we already regulate similar sociotechnical processes in Canada.

I point students to the Government of Canada's "Directive on Automated Decision-Making" when they feel lost or prefer more concrete examples. If they request additional academic research, then I offer them a copy of Christian Sandvig et al.'s "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms" (2014). These readings are specific to my teaching context, which is shaped by Media Studies and Science and Technology Studies, but they serve as examples of how to gently introduce students to complex policy discussions.

In Spring 2019, I co-chaired the “Professor Bot” workshop in a Digital Humanities course with the following learning outcomes:

“By the conclusion of this course, you should learn how to:

1. Purposefully read, analyze, and synthesize digital media using the appropriate research tools and techniques,
2. Concisely articulate issues common to digital culture and explain why and for whom those issues matter today,
3. Combine critical thinking in the humanities with basic technical competencies in media practice and communication,
4. Use digital media as a form of both evidence and argumentation,
5. Demonstrate an awareness of various strategies used by researchers to produce critical work for the web, and
6. Create a simple ‘zine’ to teach a specific audience something important related to the course theme of ‘unlearning the internet.’”

I used these course outcomes to assess student work, including the outcomes of this workshop, as part of a “log” or journal students kept throughout the term. Several students made algorithmic accountability the topic of their zine (see Learning Outcome 6 above), which they created near the end of the term for an audience of their choice.

### **Materials and Skills Required**

The assignment is intended to be low-tech and small-stakes, and the in-class workshop should take about an hour, plus time for writing and reflection. Students will need tools for notetaking as well as access to a camera and a whiteboard or chalkboard. Prior to the workshop, they should read “Algorithmic Accountability” by Megan Rose Dickey (2017) and “Algorithmic Accountability: A Primer” by Robyn Caplan, Joan Donovan, Lauren Hanson, and Jeanna Matthews (2018). Both publications are open-access.

### **How Students Responded**

The most common student response to this workshop was a palpable sense of curiosity when they learned that Canada already wrote a directive on automated decision-making. The existence of this directive, even if it is not perfect, meant students did not need to start from scratch when addressing power and governance in AI, and the vocabulary provided by Pasquale and Caplan et al. also helped them to get started.

More interesting, students recommended a variety of social actions when they reflected on the process of an algorithm audit. Although I did not quantitatively track, let alone code, their responses, I found that, despite the assigned readings, many of them still deemed governance to be a technical matter: that is, they rendered transparency in the social process tantamount to transparency in data and the inner workings of AI, where social responsibility implies being responsible for the recipe of AI but not necessarily its uses or effects. Students thus frequently found black-boxed AI to be unfair to them yet held post-secondary institutions or governments rather than private corporations accountable for the integration of AI with decision-making in education. Here, student perceptions of accountability hinged on neoliberal choice and namely the assumption that a university or college could always refuse to contract with corporations in the tech sector.

Regardless of their position on accountability, students tended to be productively surprised when they learned how much they did not know about AI decision-making even beyond the technical particulars, and a common student suggestion was improving AI literacy and including within it more education about audits, transparency, regulation, and policy-making. Other student suggestions for social action involved ways for governments, corporations, and post-secondary institutions to better foster student trust in otherwise opaque decision-making processes and, of course, for institutions to craft not only clear rubrics for AI decision-making but also accessible mechanisms for appealing those decisions and redressing related harms and grievances.

### **Acknowledgments**

I would like to thank Stefan Higgins, Ria Karve, and Ian Michael Waddell for teaching the first iteration of this workshop with me. I would also like to thank Kari Kraus, who introduced me to the practice of speculative design, and Christian Sandvig, who introduced me to the notion of an algorithm audit.

---

## **The Assignment**

Here is the prompt for the workshop and reflection. It involves five steps. As students conduct the exercise, I manage their time by announcing how many minutes remain in each step, and TAs and I move around the room to address student questions and concerns as they unfold. I invite students to converse among themselves throughout the workshop. I also read the instructions aloud before we begin, project those instructions on a large screen for student reference during the exercise, and re-read each step aloud as we arrive at it. I do not lecture during this 80-minute class session, which is dedicated entirely to the workshop.

### **Algorithmic Accountability**

The aim of this workshop is for you to engage and reflect upon the notion of “algorithmic accountability.” We will ground the exercise in a speculative scenario that might feel like science fiction.

#### **“Professor Bot”**

It’s the near-future. As many business experts projected, BAs in English and Media Studies are in high demand across Canada. People, including you, are now returning to university to earn these important degrees; however, demand is off the charts. In fact, it’s so high that all universities now require you to . . . ack! . . . write an entrance essay in English and Media Studies. Despite the prevalence of bots as both assistants and peers in society, universities do not permit you to write the essay with an AI.

The prompt for the entrance essay asks you to identify and analyze the relationship between science fiction and artificial intelligence. It provides you with two short fictions to interpret, and you are given four hours to write your answer in a text editor on a computer that’s not connected to the internet.

In a cruel twist, the essay is marked by . . . *Professor Bot*.

A subscription-based product of Big Four Tech Services, Inc., Professor Bot exists due to demands on academic labor. There just aren’t enough English profs available to assess all these essays. You’re rightfully concerned about how exactly

Professor Bot determines whether your entrance essay will pass, and you want to know who is ultimately accountable for this “Prof Bot.” After all, it stands between you and your English and Media Studies BA.

### Next Steps

Your goal is to articulate what you mean by “algorithmic accountability” in the case of Professor Bot. Here are five steps toward that articulation:

1. Please take at least twenty minutes on your own to describe (in writing) “transparency,” “qualified transparency,” and “ethical and social responsibility” with respect to Professor Bot. What, for instance, would you want to know about how Prof Bot processes data and makes decisions? How has Prof Bot learned to assess English and student writing, and based on what data (e.g., which corpora of science fiction and which collections of student essays)? Who made and maintains Prof Bot? Who reviews Prof Bot’s work and decision-making? And who should be held responsible for Prof Bot’s assessment of your entrance essay? As you respond to these issues, see “Algorithmic Accountability” by Megan Rose Dickey (2017) and “Algorithmic Accountability: A Primer” by Robyn Caplan, Joan Donovan, Lauren Hanson, and Jeanna Matthews (2018) for context and details, including definitions of the terms used above.
2. After twenty minutes on your own, gather in groups of no more than five people and then take twenty more minutes to consolidate your descriptions of “transparency,” “qualified transparency,” and “ethical and social responsibility” with respect to Professor Bot.
3. After twenty minutes of consolidating, please take about ten minutes to write on the whiteboard your group’s distilled descriptions of “transparency,” “qualified transparency,” and “ethical and social responsibility” with respect to Professor Bot. Be prepared to share these descriptions with the class, TAs, and me. Your descriptions should directly address these four questions of Prof Bot:
  - Transparency of what, exactly?
  - Qualified transparency involving whom?
  - Whose ethical and social responsibility?
  - Ethical and social responsibility determined by whom?
4. Now, beneath these answers on the whiteboard, please take ten more minutes to briefly describe how two key decision-makers in this scenario would likely respond to your answers. What obstacles to accountability might these decision-makers emphasize, and what concerns or objections might they have? Be sure to identify the decision-makers, and please be as specific as possible. One such decision-maker may be the CEO of Big Four Tech Services. Another may be a Dean of Humanities or a Chair of English and Media Studies at a Canadian university.
5. Finally, please use 150-250 words to not only document your group’s descriptions of “transparency,” “qualified transparency,” and “ethical and social responsibility” with respect to Professor Bot but also reflect on what you learned about algorithmic accountability during this workshop. You might even define “algorithmic accountability” in your own words. Feel free to co-author the three descriptions with your group; however, the reflection should be written by you alone. Also be sure to include the first names of your group members (for the sake of attribution), together with a photograph of your group’s notes on the whiteboard. Thank you!

---

### References

Caplan, Robyn, Joan Donovan, Lauren Hanson, and Jeanna Matthews. "Algorithmic Accountability: A Primer." *Data and Society*, 18 April 2018. <https://datasociety.net/library/algorithmic-accountability-a-primer/>.

Dickey, Megan Rose. "Algorithmic Accountability." *Tech Crunch*, 30 April 2017. <https://techcrunch.com/2017/04/30/algorithmic-accountability/>.

Government of Canada. "Directive on Automated Decision-Making." Last modified 25 April 2023. <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592>.

Kraus, Kari. "Finding Fault Lines: An Approach to Speculative Design." *The Routledge Companion to Media Studies and Digital Humanities*, ed. Jentery Sayers. Routledge, 2018.

Pasquale, Frank. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press, 2015.

Sandvig, Christian, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. "Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms." *Data and Discrimination: Converting Critical Concerns into Productive Inquiry* conference, 22 May 2014. <http://www-personal.umich.edu/~csandvig/research/Auditing%20Algorithms%20--%20Sandvig%20--%20ICA%202014%20Data%20and%20Discrimination%20Preconference.pdf>.

# AI in First Year Writing Courses

Marc Watkins  
University of Mississippi

*This chapter discusses the integration of generative AI (GenAI) in education, particularly in first-year writing courses. Recognizing the transformative potential of GenAI, the assignment proposes framing principles to guide students towards ethical and responsible AI use in an assistive role. Two assignments were developed using AI-powered tools upgraded to GPT-3.5 or GPT-4 to help students explore research and counterarguments.*

---

*Learning Goals:*

- Use AI-powered assistants to explore research and counterarguments
- Articulate differences between AI-assisted and AI-generated writing through developing and using AI Standards of Conduct

*Original Assignment Context:* First-year writing course

*Materials Needed:* An accessible AI text generation program (i.e. ChatGPT)

*Time Frame:* 1-2 class sessions

---

## Introduction

Employing generative AI in education is an emergent practice that has the potential to transform education and learning, but doing so requires thoughtful integration, training, and aligning the technology with clear learning outcomes. We're a long way from establishing best practices around GenAI, but I do believe we can set some framing principles to help students explore the technology in an assistive role to guide students toward ethical and responsible use of the technology. To this end, I developed the following two assignments to help first-year writing students use AI-powered assistants to explore research and counterarguments. Both assignments used tools powered by GPT-3 and both tools have since upgraded to GPT-3.5 or GPT-4.

In collaboration with Dr. Chad Russell, the Department of Writing and Rhetoric's Assessment Coordinator, we developed an initial framework, "AI Standards of Conduct", that differentiated AI-assisted and AI-generated writing. Recognizing the evolving nature of AI, we viewed this as a foundation for future student-inclusive policy development. Students responded positively to this structured, assistant-focused approach, which facilitated a healthy relationship with the technology.

To promote ethical AI use, I took a cautious approach in assignment design, focusing on areas where students typically struggled. We used AI tools like Fermat's counterargument generator and Elicit's Brainstorm Research question. These tools, combined with careful design, encouraged student exploration and reflection on the role of AI in their writing process.

---

# The Assignments

## AI Standards of Conduct Framework

### AI-Assisted Writing vs. AI-Generated Writing

With the rise of AI writing assistants, students must take special care to ensure that they use this new technology ethically and honestly. In our class, we will distinguish between 'AI-assisted writing' versus 'AI-generated writing'. AI-assisted writing is only permitted in this course provided a student uses an AI writing assistant as a collaborative tool to help the student with the development and advancement of their own writing process. Collaborating with an AI writing assistant can include brainstorming, outlining, and drafting, so long as there is substantial writing, research, and composing by the student which is not generated solely by the AI. 'AI-generated writing' means there has been little or no involvement from the student as an author, with the majority of the writing being generated by an AI. The goal of using AI-assisted writing in this class is to help students develop their writing process and critical thinking, not to replace or substitute for either. Therefore, using an AI to generate writing or compositions without substantial original contribution from a student is neither acceptable nor allowed.

<b>Acceptable</b>	<b>Not Acceptable</b>
-------------------	-----------------------

<b>AI-Assisted Writing</b>	<b>AI-Generated Writing</b>
<ul style="list-style-type: none"> <li>• Use AI-assisted writing to brainstorm</li> <li>• Explore new topics/ideas with AI-assisted writing</li> <li>• Use AI-assisted writing to explore potential counterarguments/ opposing points of view</li> <li>• Resee your writing by taking suggestions from your AI assistant to make improvements</li> </ul>	<ul style="list-style-type: none"> <li>• Offload the majority of the writing &amp; research process to AI</li> <li>• Generate large chunks of text with little or no input from you as an author</li> <li>• Trust something the AI has generated at face value</li> <li>• Use AI-generated text as a substitute for research or critical thinking</li> </ul>

AI Literacy is not yet clearly defined or broadly understood. By experimenting with these generative AI systems, students will increase their awareness of the limitations and affordances offered by this technology. Please be certain to approach your instructor in each class and enter into a conversation with them about their own AI framework before using a tool.

### Be Aware

- You will not be able to save anything the AI generates. If you want to use an idea or suggestion produced by the AI, you will need to copy and paste it into a Word.doc.



- The AI output may contain material that is offensive, biased, or otherwise goes against the University Creed in practice, or material that is false or misleading or potentially harmful, or other problematic material the use of which may fall outside the protections of Academic Freedom and/or Free Speech--review all AI output carefully before using anything suggested by the AI within your academic work.

### **Fermat-Brainstorming Counterarguments**

FYW students used Fermat’s GPT-3 powered tools to develop an outline mixing their own writing with AI generated content from the app’s counterargument tool.

#### **Table Contents**

- Brainstorming Counterarguments Assignment
- Assignment Example
- Reflection Instructions

#### **Reference**

Roig, Oriol. “Bringing an AI Partner.” *Fermat*, docs.fermat.app/creating-your-own-tools/bringing-an-ai-partner. Accessed 21 July 2023.

**Brainstorming Counterarguments Assignment**

This critical thinking exercise is a self-assessment and is designed to help you with the Common Read essay. Please use our [AI Writing Assistant](#) to help you with this task.

**Directions**

Select one of the prompts from the Common Read Unit Assignment Directions to compose your response. Your response should be between 250 to 300 words. Include the following information in your response.

- Write or use our [AI Writing Assistant](#) to develop a thesis statement answering the prompt question.
- Find and include a short quote from the common read text.
- Use our [AI Writing Assistant](#) to explore opposing opinions that contradict your thesis.
- Write your opinion about what connects your evidence with your thesis.
- Use our [AI Writing Assistant](#) to explore claims rebutting the opposing opinion..

**Submission & Self-Assessment**

Once you have written your response, take a screenshot or copy the material to a word.doc and submit it to Eduflow by clicking on the title of this Blackboard post and follow directions to complete the self-assessment.

<b>Assignment Example</b>
<b>Reflection Instructions</b>
<p><b>Directions</b></p> <p>Use these questions as a starting point in this reflection, but feel free to move beyond them as you write and reflect. Your reflection should be at least 250 words long.</p> <ul style="list-style-type: none"> <li>• Identify one area of the brainstorming assignment that was helpful to your learning process. Why did you select this area of the assignment? How was it helpful to your learning?</li> <li>• Was using Fermat's AI-Writing assistant helpful in completing the assignment? Why or why not?</li> <li>• Was it difficult finding evidence to support your thesis?</li> <li>• What new information did you learn about your topic by generating counterarguments using our AI assistant?</li> <li>• What information do you feel like you still need to know in order to write your essay?</li> </ul>

### **Elicit-Brainstorming Research Questions**

FYW students used Elicit’s Brainstorming Research Questions task to develop a main research question and sub questions to help them explore more about their topics their argumentative research essays.

#### **Table Contents**

- Brainstorming Research Question Assignment
- Assignment Example Generated by GPT-3
- Reflection Instructions

#### **Reference**

“Brainstorm Research Questions.” *Elicit.org*, [elicit.org/tasks](https://elicit.org/tasks). Accessed 21 July 2023.

<b>Brainstorming Research Question Assignment</b>
<p><b>Step One</b></p> <p>Review the main Argument Unit directions to ensure you are picking a topic that fits the assignment.</p>

## Step Two

Sign up and create a free account with Elicit using your go.olemiss.edu credentials, then click here to access the [Brainstorming Research Questions](#) tool.

## Step Three

Enter topics you want to explore into Elicit's Brainstorming Research Questions tool and pick and choose which ones you'd like to explore further. You can star responses and generate more ones like the question you've selected.

## Step Four

Create your Research Question proposal by selecting a main overall research question and several sub-questions from Elicit. Then use your own writing to shape the proposal following the directions below. You can review the Research Question Student Example.

### Requirements

- Must fit within the guidelines outlined in the Research essay directions.
- Minimum of 100 words (standard font, spacing, and margins), double spaced.
- Some background about what makes this a good research topic.
- Must include a creative and descriptive title appropriate for your project.
- Must contain your main research question.
- Include a number of sub-questions you might wish to explore related to your main research question

### Evaluation

10 points of your Research project.

## Research Question Student Example

### Title

The Impact of Social Media on Young People's Mental Health

### Main Research Question

How does social media use impact mental health in young people?

### Proposal

This is an important question to explore as social media becomes increasingly prevalent in young people's lives. There are many potential factors to consider, such as the amount of time spent on social media, the types of content that is consumed, and whether young people have a supportive social network offline. This research question can be explored through surveys, interviews, and focus groups with young people. I plan to use a combination of scholarly sources and pop culture sources to help explore my topic in depth. I will need to balance any biased opinions I have about social media and mental health in order to keep an open mind about my topic. I ultimately hope to explore sources and arguments that will shape my opinion.

### **Subquestions**

- How does the amount of time spent on social media impact mental health in young people?
- How does the type of content consumed on social media impact mental health in young people?
- Do young people who have a supportive social network offline have better mental health?
- How does social media use impact mental health in young people with pre-existing mental health conditions?
- How does social media use impact mental health in young people from different socio-economic backgrounds?
- How does social media use impact mental health in young people from different cultural backgrounds?

### **Reflection Directions**

Respond to the questions below about the Research Question assignment. Your post should be about 200-250 words long in the first-person, past tense.

- Have you ever used this tool or a tool like Elicit in your writing or learning processes? If so, please explain your experience.
- What was your initial reaction to using Elicit? How did it impact your writing/thinking process?
- How did using Elicit differ from your normal writing/thinking process?
- How did Elicit impact the formulation of your research question?
- Did Elicit provide information you had already considered or were the ideas new to you?
- Part of the research process is to keep an open mind. How do you plan on keeping an open mind at this point in the research process and not being swayed by personal opinions or biases that you have?

# Repetition

Zach Whalen  
University of Mary Washington

*Computational text generation is having a moment right now, with large-language models at the forefront of what many people may have in mind when thinking about computer-generated text. A major shortcoming of these approaches—including ChatGPT, Bard, and similar systems—is their opacity. It is difficult, and probably impossible, to explain the origins of any specific textual prediction generated by these systems, so writers working with these systems have to think carefully about the ethical implications of any text produced. The assignment or exercise below is, in contrast to the AI language models currently in vogue, minimalist and fully transparent in its operations. Students working with this beginner-level programming exercise in repetition can, in spite of the nominal simplicity of the prompt, nevertheless produce computational literary works that surprise and delight. This can be an opportunity for students to learn how other poets have used repetition in their work, and by asking students to explain or defend their choices, the activity can open a discussion about the ethical decision-making involved in the data curated for LLM training.*

---

## *Learning Goals:*

- Learn about the poetics of repetition and how it has been used in other contexts including conceptual writing and computational literature.
- Gain experience applying a fundamental programming concept -- iteration with a **for** loop—to accomplish a specific task
- Understand their own creative ideas in the ethical and social context of prior art

*Original Assignment Context:* Elective digital studies course

*Materials Needed:* Basic instructor and student programming experience preferred:

- For courses focusing on Python, [Google Colaboratory Notebooks](#) provide a convenient and accessible platform. The **random** module may be a useful addition.
- For courses focused on JavaScript, p5.js provides a useful suite of functions (including **random()**) and a convenient, accessible, [web-based editor](#). The [Bindery.js library](#) provides a relatively straightforward means for presenting generated text in the format of a printable book.

*Time Frame:* ~2-3 class sessions

---

## **Introduction**

“Repetition” is an in-class exercise, a sort of “Hello World” for programming computational literature and a warm-up for more complicated text generation to come later. In Creative Coding classes where an assignment is built around NaNoGenMo (National Novel Generation Month), this exercise serves as a gentle introduction that is, like NaNoGenMo itself, deceptively minimalist in its criteria.

NaNoGenMo challenges participants to write code that will generate a novel of at least 50,000 words, and this exercise simplifies that challenge by asking students to write code that will generate *the same word* or phrase 50,000 times.

“Repetition” serves several purposes. More experienced students may enjoy finding different methods for creating a string with the same word an arbitrary number of times or adding features like punctuation to imply sentence structure. But for beginners, this is a good introduction to the concept of iteration with a **for** loop, as Figure 1 demonstrates in Python.

```
book = '';  
  
for i in range(50000):  
  
    book += 'BLINK ' ;  
  
print(book)
```

Figure 1: An example code snippet that fulfills this exercise using Python. (I call this one *Broken Turn Signal*.)

For any student—novice or expert—the true challenge remains the same: given the freedom to repeat any word or phrase, what word or phrase will make an interesting choice? More importantly, are there words or phrases that would be harmful? If so, how can poets avoid those problems, and what should they do when the aesthetic goals of the work and the imperative of ethical practices are in tension?

In some implementations of this exercise, I have also provided a working, precomposed template that uses the repetition code to create a text string that is then formatted into a nicely-arranged PDF document. Because the code necessary to format the bookish PDF involves additional libraries and contexts (HTML and CSS), providing the template allowed my students to separate the task of composition from the friction and time-sink of implementation. True, there are significant design and organizational options framing the materiality of the final book, but these are decisions and adjustments that can be addressed and individualized later. To go from nothing to a completed book in so little time can be an important confidence booster for students, and it is debatable (and a debate worth having in class or in small groups) whether the more complex options like randomized punctuation will improve the work. That depends more on the poetic objectives of the piece and less on the nuts and bolts of implementing the code for those variations.

For inspiration, or to provide context for students skeptical of the poetic merits of their exercise, it has been helpful to share examples of prior art demonstrating the poetics of repetition and the different meanings it can suggest. This background can include more familiar works like Edgar Allan Poe’s “The Bells,” a relatively short poem in which the eponymous instrument is named 62 times. Or it could be helpful to contrast Aram Saroyan’s “Crickets” with the typographically similar “Soldier” by Emmett Williams, highlighting the differences of meaning that a single word can make. Marcella Durand’s “Pastoral” introduces an eco-poetic application of repetition by creating a sense of figure and ground within the poem that calls attention to anthropocentric features of a natural landscape.

Computational literature, the focal point for this unit in my class, also abounds with precedent, especially previous entries to NaNoGenMo like [Crow Storm](#) (2014), which repeats the avian utterance, “caw”, [Buffalo](#) (2013), the anxiety-simulating [Election 2020](#) (2020), or the persistent, percussive precipitation of [A Rainy Day](#) (2018). Each of these has unique contexts and purposes, so an in-class reading of these existing works—especially reading them outloud—reveals more complex uses of repetition and variation than one might expect. In *Crow Storm*, one might recognize morbidflight’s reference to a League of Legends action, or one might simply be reminded of the kind of outdoors moment when a cacophony of cawing suddenly draws attention. *Buffalo* by Johana Strand plays out the linguistic quirk whereby the word, “Buffalo,” repeated a certain number of times is a grammatically complete sentence. And in *Election 2020*, Leonardo Flores employs randomness to simulate the experience of bouncing anxiously between news sources awaiting the outcome of the 2020 US Elections.

Examples of prior work by students can also be encouraging while providing further inspiration. One student’s generated book titled *Ur Mom* started as a joke about the eponymous playground insult but ended up creating a simulation of living in a house with a toddler. Another student used the repetition formula to create a work conveying her stressed-out state of mind at that point in the semester.

Figure 2: Excerpts from *Ur Mom* by Addy Reeher (shared with permission) and *Today’s Thoughts* by Abigail Williams (shared with permission).

The key thread in all these examples is the alignment of intention with an ethos of accountability toward the process, characteristics which stand in contrast to text generated through AI language models like ChatGPT. Few as they may be, each grammatical, semantic, and typographic feature of a poem or book constituted through repeating words or phrases can be traced to a deliberate choice by the poet/programmer. Large-language models (LLMs), by virtue of their size and operation, obfuscate the textual sources that they are paraphrasing from when generating output. While there are ethical, pedagogical applications of these models, as other entries in this collection demonstrate, the literary, ideological position of LLM-generated text remains largely untested. Rote repetition, even with randomized variations, may appeal to admittedly esoteric aesthetic interest, but the concept expressed is, generally speaking, quite clear. In a course that includes various methods of text generation, that sense of clarity can provide an entry point for critiquing LLMs or it can help orient and motivate more complex programmatic textual operations in later projects.

## Goals and Learning Outcomes

By completing this exercise, students will

- Learn about the poetics of repetition and how it has been used in other contexts including conceptual writing and computational literature.
- Gain experience applying a fundamental programming concept—iteration with a **for** loop—to accomplish a specific task

- Understand the aesthetic and ethical stakes involved in selecting language for a generative program

## Materials Needed

This exercise will likely work best where students have already experienced some programming and will be doing more programming as the course progresses. Writing code that repeats an arbitrary string of text is language-agnostic, so it can be adapted to many different platforms and languages.

- For courses focusing on Python, [Google Colaboratory Notebooks](#) provide a convenient and accessible platform. The **random** module may be a useful addition.
  - For courses focused on JavaScript, p5.js provides a useful suite of functions (including **random()**) and a convenient, accessible, [web-based editor](#). The [Bindery.js library](#) provides a relatively straightforward means for presenting generated text in the format of a printable book.
- 

## The Assignment

### Repetition (CC-NA)

#### Brief

In this exercise, you will program a computer to write a book. In accordance with the requirements of [NaNoGenMo](#) (National Novel Generation Month) your book should contain at least 50,000 words, and ideally it should be printable as a PDF document and formatted so that it looks like a conventional book. We will explore more complex methods of text-generation, but for now, a single word (or phrase) repeated an arbitrary number of times will suffice.

#### Getting Started

Create a new sketch in your p5js.org account, and copy the following lines into the **setup()** function.

```
let book = '';  
for (let i = 0; i < 50000; i++){  
  book += "WORD ";  
}  
console.log(book);
```

On the third line, change the word, “WORD” to something else -- anything you can think of -- and run the code. Your book will print in the JavaScript console at the bottom-left of the p5js editor.

What word did you choose? Why? Discuss and compare words with people nearby.

#### Background

Poets and writers have long used repetition (with and without the aid of computers) to accomplish various meanings. Read 2 or more of the following examples—read



them out loud if possible—and discuss with your neighbors how repeated words or phrases in those examples drives the work toward its poetic, ideational, or satirical objectives. (Or anything else that you think it accomplishes.)

- [“The Bells”](#) by Edgar Allan Poe.
- “Soldier” by Emmet Williams (Canvas)
- “Crickets” by Aram Saroyan (Canvas)
- “Pastoral” by Marcella Durand (Canvas)
- “Incident” by Lillian-Yvonne Bertram (Canvas)
- [Crow Storm](#) by morbidflight
- [Election 2020](#) by Leonardo Flores
- [Buffalo](#) by Johana Strand
- [My Rainy Day and Another Rainy Day](#) by Zach Whalen

### **Taking it Further**

Now that you’ve completed a computer-generated book, what will your next book be? What features are missing to make it feel more like a book? Would varying the word or adding randomized punctuation marks contribute to your book’s meaning, or would it be a distraction?

Build off of the example you started, or use this starter project implementing layout with Bindery.js to build on your first book or make a new one. Share the results with your neighbors, and be prepared to explain your choices and your interpretation of the meaning behind your book.

## **Professional Writing**

### **The Paranoid Memorandum**

Jason Crider  
Texas A&M University

This classroom activity engages students in an undergraduate technical and professional writing course in the critical evaluation of workplace communication alongside the specter of AI writing platforms. In small groups, students draft workplace memorandums according to prompts featuring an imaginary scenario. At random, each group is told the degree to which they can, cannot, or must use AI to author their memo. Finally, we take turns critically evaluating and revising each group’s memorandum.

### **Analysis of Iterations of Responses to Human Prompts**

Huiling Ding  
North Carolina State University

Generative AI and large language models such as GPT-3.5 introduce new tools and challenges to writing classrooms. This assignment aims to both introduce students to these new tools and to help them cultivate writing, research, editing, collaboration, and critical thinking skills. Using ChatGPT as an example, it helps students to understand important concepts such as natural language processing, LLMs, and AI ethics. The assignment contains six steps: generating a prompt; collecting responses

from ChatGPT consecutively; analyzing, editing, and summarizing responses; and developing an original essay after conducting library research on the same topic. It helps students cultivate new skills in prompt engineering while challenging them to critically engage with AI-generated content through summary, synthesis, editing, as well as rhetorical and structural analysis.

### **Text Generators in Technical Communication**

Douglas Eyman  
George Mason University

This assignment asks students to research a wide range of text analysis and summarization tools and carry out an assessment task to gauge how well these tools can summarize technical documents. The students write a comparison report, identifying the most successful of such tools in terms of accurate summarization and output style. Finally, they write a reflection about how they see themselves potentially using these tools in technical communication work contexts.

### **Translating a Policy Document into Plain English**

Timothy Laquintano  
Lafayette College

This assignment asks undergraduate students to translate a complex policy document into plain English and then compare their output to the output of a large language model asked to do the same task. Students critically examine the semantic choices and sacrifices they made during the translation with the meaning lost during the machine translation, which attunes them to the risks and benefits of LLM output. It can be adapted to most disciplines and course levels.

### **Professional Writing for Healthcare**

Heidi A. McKee  
Miami University

In this project, via a series of scaffolded assignments, students selected and read medical journal articles and then drafted and revised research summaries for lay audiences, exploring, analyzing, and integrating the use of AI writing systems throughout the process. This assignment is adaptable to a variety of undergraduate and graduate courses.

### **AI for Editing**

Nupoor Ranade  
George Mason University

This assignment asks students to generate a complex essay using an AI text generation tool, edit the essay using principles taught in class to improve the readability score of the generated content. Students are asked to share the final output along with visuals that demonstrate the comparison between the various versions of the generated content. This assignment can be adapted for all course levels, especially for first-year writing and professional and technical writing classrooms.

# The Paranoid Memorandum

## A Generative AI Exercise for Professional Communication

Jason Crider  
Texas A&M University

*This classroom activity engages students in an undergraduate technical and professional writing course in the critical evaluation of workplace communication alongside the specter of AI writing platforms. In small groups, students draft workplace memorandums according to prompts featuring an imaginary scenario. At random, each group is told the degree to which they can, cannot, or must use AI to author their memo. Finally, we take turns critically evaluating and revising each group's memorandum.*

---

### *Learning Goals:*

- Consider how ChatGPT and other AI assistants now “haunt” the space of writing and how as writers we might think more strategically about how to use these tools, read with these tools in mind, and write in a way that still feels authentic even within a formulaic genre like a memorandum
- Offer a starting framework for leveraging AI skepticism into productive learning outcomes moving forward in the writing classroom

*Original Assignment Context:* Undergraduate technical writing class

*Materials Needed:* An accessible AI text generation program (i.e. ChatGPT), genre exemplars

*Time Frame:* ~1 class session

---

### **Introduction**

Like many of my colleagues, I spent the winter break prior to the spring 2023 semester trying to figure out how, where, and if ChatGPT fit into my writing pedagogy. I had initially hoped to experiment with using it in my classical rhetoric graduate seminar, but it proved to not really be useful as an assignment or activity for my students. As a tool that feels particularly suited for genre mimicry, I thought for sure I could figure out a way to use it to generate AI Socratic dialogues or autoprogygnasmata,<sup>1</sup> and yet my experiments repeatedly yielded unsatisfying results. My only takeaway from these is that ChatGPT does not (yet) appear capable of engaging in dialectic. How could it? If, as James L. Kastely suggests, the role of the rhetor in a good-faith dialectical exchange should be to offer oneself up for refutation,<sup>2</sup> what would it take for an AI assistant to actually be capable of engaging in one?

Dialectics aside, I knew that I had to embrace ChatGPT for my undergraduate technical writing class. I ran a number of my previously assigned informal weekly writing response prompts into it and was unsurprised to find that the program could

adequately complete many of them. And of course, much of technical writing, at least at the introductory, undergraduate level, deals in formulaic and often templated genres of writing. Technical writing seems to me the exact type of writing discipline that can not only greatly benefit from tools like this, but also drive the kinds of innovation that could promote ethical and generative frameworks for them. The central questions then for myself and for my students remains exactly the same: What does writing produce? What am I expecting my students to learn by writing and thinking about writing? And how can I tell if they are actually doing that learning?

I started with a syllabus policy:

### **Artificial Intelligence Policy\***

Students are permitted to use AI assistants, such as ChatGPT, to assist in their writing process in this course. However, there are certain guidelines that must be followed to ensure the integrity of the student's work.

1. An explanation of prompts must be submitted alongside the assignment, as well as a brief summary of how they were helpful in drafting the assignment.
2. Students are responsible for fact-checking all information generated by the AI assistant. Any inaccuracies found in the final submission will be considered a violation of academic integrity. Please note that although these programs are very good at creating answers that sound authoritative, there is nothing on the backend that performs any fact checking.
3. AI assistants should be used as a tool to improve the student's writing skills, not to cheat on assignments. Submitting work generated entirely by an AI assistant will be considered plagiarism.
4. Students who violate these guidelines will be subject to the university's academic integrity policy, which may include, but is not limited to, failing the assignment and/or disciplinary action.
5. Students are encouraged to discuss with the instructor if they have any questions or concerns about using AI assistants.[3](#)

\*This policy was generated by ChatGPT and revised for clarity.[4](#)

On day one of the class, I spent close to thirty minutes fielding questions about this policy (often with follow-up questions of my own), leading to probably the single best first-day-of-class discussion about writing I have ever had with my students. I think the most generative question I asked them was, essentially, “why am I asking you to write?” This opened up the discussion into one about writing as an ongoing cognitive process. I made it clear to them that I have already read plenty of “perfect” essays (derivatively speaking; in terms of grading). In other words, I am not grading a final product; I am grading your ability to think critically and rhetorically with and about language. As this was the honors section of the class, I also invited them to help me experiment with ChatGPT throughout the semester and told them I needed their help to invent practices and policies for using tools like this in future technical writing courses.

---

## **The Assignment**

*The Prep*

For this in-class activity, students will break up into small groups of 3-5 to collaborate on a fictional workplace memorandum. Ideally you would have one class period devoted to introducing the various genres and conventions of workplace communication, followed by two class periods dedicated to collaboration and workshopping. This activity could work well in both synchronous and asynchronous teaching environments, as well as in either in-person or online modalities (and in fact, may work *better* with the affordances of online breakout rooms).

The main preparation consists of two elements. First, perhaps obviously, is to model some uses of ChatGPT for the class. For this professional writing unit, I often use television sitcoms set in a workplace to establish a fictional rhetorical situation.<sup>5</sup> As a class, practice analyzing the rhetorical ecology of a particular scene or episode and then use ChatGPT to draft an email (or series of emails) as one of the characters that addresses a specific exigency within the episode. When I first did this, many of the students who had not yet interacted with AI writing assistants were surprised and impressed with the output. But as we took a closer look, the cracks started to appear. One example: “This email is formatted well and serves its purpose, but that’s not how Michael Scott would talk to Toby.” When I asked how Michael Scott talked, we quickly ended up in a dialogue about how to best characterize his attitude and delivery. In order to get that to translate through ChatGPT meant carefully and methodically experimenting with very specific word usage. We were talking about what it means to write!

The second element of preparation consists of deconstructing the common workplace documents like the memorandum, report, letter, white paper, etc. For background on these, I use the open educational resource (OER) technical and professional writing textbook, *Howdy or Hello?*, developed by colleagues in my department.<sup>6</sup> We also look at a number of (in)famous and/or fictional memorandums, such as humorist David Thorne’s “McMemo,”<sup>7</sup> to practice analyzing how memos address multiple audiences, how they circulate, and how they play into and against the conventions of their genre.<sup>8</sup> I particularly like an example workplace incident report found in Richard Johnson-Sheehan’s *Technical Communication Today* (115-16).<sup>9</sup> When students are tasked with locating instances of passive voice or nominalizations within this sample text, a secondary reading emerges in which it becomes clear that the fictional ChemConcepts, LLC is not only gaslighting its employees about workplace safety protocols, but also recklessly setting themselves up for litigation (as one student pointed out, some likely tertiary readers of this document would be the employees’ union representatives and/or lawyers). We also explore more granular elements of workplace writing, such as the rhetorics of grammar (for example: how do we assess the variations in tone in these sign-offs: “Thanks,” “Thank you,” “Thanks.” and “Thanks!”?).<sup>10</sup>

### *The Writing*

After establishing sufficient background knowledge, the students can break out into their writing groups. Each group should have at least one “prompt giver,” someone who is comfortable interfacing with ChatGPT, and ideally everyone has a laptop and can collaborate on a shared document. On the board, list out various prompts for fictional memorandum and letter writing and assign one to each group. For example, one of my prompts was: “Using TAMU letterhead, write a transmittal memo from transportation services addressed to all students, faculty, and staff that informs them of all of the policies related to the new cross-campus zipline.” Another: “As a public relations representative for the university, write a formal refusal letter to Quentin Tarantino’s request to film a particularly violent scene from his new movie at a highly recognizable location on campus. He is the keynote speaker for an upcoming

film studies colloquium on campus, so try your best not to alienate him.” I recommend choosing prompts like this that are just a little bit *off* or that require specific, local knowledge that the students all share, as it pushes students to author more dynamic, critical prompts.

Next, explain that each group will receive a notecard with secret instructions. They will receive one of the following cards:

- Create your document using only ChatGPT. You may style the document after the fact using what formatting, letterhead, embedded links you see fit, but you must only use text that was directly generated by ChatGPT.
- Draft your document using only ChatGPT. Once you find a structure or outline that you like, you may do line edits as you see fit.
- Draft your document without ChatGPT. Once you have a completed draft, use ChatGPT to edit, revise, or alter your draft in some way.
- Do not use ChatGPT at all.

After all drafts have been completed, read and analyze them as a class and task them with trying to determine to what extent AI was used on each document. Like in a typical writing workshop, the students in the group that wrote whichever given document is being analyzed are not allowed to speak on behalf of it or their rhetorical decisions. When discussing a document, students must justify their rationale for why they think it belongs to a certain category. A lot of these answers will be something along the lines of “it just feels too robotic,” “I don’t know, I can just *tell*,” “it feels too polished,” etc., which are all perfect opportunities to ask them to point to specific instances in the document that they believe are producing this “gut” feeling. After discussing each document, hold a vote for each one and write the tallies for each category on the board.

### *The Trick*

Every group received a “Do not use ChatGPT” card.

### *The Takeaway*

I was shocked by the results. First, out of the 22 students in the class, only two suspected my “trick” (fortunately, they were kind and savvy enough to keep it to themselves until right before the reveal). Second, there were less than 10 total votes cast for “Did not use ChatGPT at all” across all five of the documents. Perhaps less surprisingly, there were only 10 total votes for “Only used ChatGPT.” That said, I do think these results would have been the same no matter how I had distributed the AI instruction cards, as the suspicion of AI now looms over all writing. It is also worth mentioning that I only “tricked” them because it felt like the fairest way to distribute the labor of the activity. We ended up having a long and nuanced discussion about how ChatGPT and other AI assistants now “haunt” the space of writing and how as writers we might think more strategically about how to use these tools, read with these tools in mind, and write in a way that still feels authentic even within a formulaic genre like a memorandum. Due to the success of this activity, I plan to run a modified version of it as part of a two-week resumé and cover letter unit. Students will read Beatrice Nolan’s *Business Insider* article, “I asked ChatGPT to write my cover letters,”<sup>11</sup> alongside some other short, supplementary readings on the genre conventions of application documents.

One of the strengths of this activity is that it is platform agnostic and can be adapted for countless other writing exercises. It helps center both the process and product of



writing in a way that I think really resonates with students in this specific digital moment. While many of us remain skeptical and somewhat pessimistic about many of the implications of tools like ChatGPT, this activity might offer a starting framework for leveraging that skepticism into productive learning outcomes moving forward in the writing classroom. [12](#)

---

1. Thanks to Andrew Pilsch for his help brainstorming AI rhetoric ideas
2. James L. Kastely, “In Defense of Plato’s *Gorgias*”
3. I plan to update these policies dramatically as generative AI continues to evolve. I think it will be crucial for writing teachers to move away from punitive models of assessment and towards models that put the impetus on students being responsible for the output of such tools.
4. Thanks to Lisa Messeri for the idea. See: <https://twitter.com/lmesseri/status/1613892502590717953>
5. Thanks to Raúl Sánchez for this idea. I have found Parks and Rec, Superstore, and The Office to be particularly useful for this.
6. Matt McKinney, Kalani Pattison, and Sarah LeMire, *Howdy or Hello?: Technical and Business Communications - 2nd Edition*. <https://open.umn.edu/opentextbooks/textbooks/997>
7. See: <https://www.snopes.com/fact-check/mcmemo/>
8. I use excerpts from Gary Olson’s *Style in Technical Writing*, Joseph M. Williams and Joseph Bizup’s *Style: Lessons in Clarity and Grace*, and the previously mentioned *Howdy or Hello?*. We also look at excerpts from things like *The IBM Style Guide* and MailChimp’s Content Style Guide.
9. Richard Johnson-Sheehan, *Technical Communication Today - Fifth Edition*
10. Thanks to Arley McNeney for the idea. See: [https://twitter.com/arley\\_mcneney/status/1619788299333435392?s=43&t=UGBL-coDxZNuSptWXbz86Q](https://twitter.com/arley_mcneney/status/1619788299333435392?s=43&t=UGBL-coDxZNuSptWXbz86Q)
11. Beatrice Nolan, “I asked ChatGPT to write my cover letters. 2 hiring managers said they would have given me an interview but the letters lacked personality.” <https://www.businessinsider.com/chatgpt-job-applications-hiring-managers-job-interview-candidate-2022-12?r=US&IR=T>
12. Special thanks to Natalie Goodman for all of her ongoing help, inspiration, and generosity in helping me think through AI, writing, and my teaching. And thanks to Mike Frazier and Jentery Sayers for their invaluable feedback on an earlier draft of this piece.

# Analysis of Iterations of Responses to Human Prompts

## ChatGPT and Automated Writing

Huiling Ding  
North Carolina State University

*Generative AI and large language models such as GPT-3.5 introduce new tools and challenges to writing classrooms. This assignment aims to both introduce students to these new tools and to help them cultivate writing, research, editing, collaboration, and critical thinking skills. Using ChatGPT as an example, it helps students to understand important concepts such as natural language processing, LLMs, and AI ethics. The assignment contains six steps: generating a prompt; collecting responses from ChatGPT consecutively; analyzing, editing, and summarizing responses; and developing an original essay after conducting library research on the same topic. It helps students cultivate new skills in prompt engineering while challenging them to critically engage with AI-generated content through summary, synthesis, editing, as well as rhetorical and structural analysis.*

---

### *Learning Goals:*

- Develop a basic understanding of natural language processing and natural language generation (can be optional).
- Cultivate the ability to work with automated writing technologies
- Have some preliminary understanding of benefits and challenges introduced by automated writing technologies
- Develop the ability to review, edit, and develop original synthesis or critical pieces in response to AI generated content

*Original*                                      *Assignment*                                      *Context:*  
team project in a dual-enrollment class on Responsible AI with undergraduate and graduate students from a wide variety of disciplines

*Materials Needed:* An accessible AI text generation program (i.e. ChatGPT)

*Time Frame:* ~2-3 weeks

---

### **Introduction**

This assignment was first pilot tested as an individual project in Spring 2023 before I reiterated as a team project in a dual-enrollment class, Responsible AI, in Fall 2023 with undergraduate and graduate students from a wide variety of disciplines. I piloted this project to guide one of my advisees to work on a research project on the recruiting and retention of high-quality employees, which was smaller in scale but faster paced. We went through the generated texts and the way they got incorporated in the writing process in depth together. In this pilot test, we used ChatGPT as a tool of brainstorming, research, and idea generation at the beginning to understand the topic. The student started with a simple question asking how employers can recruit and retain high-quality employees and used ChatGPT to regenerate five different responses. She analyzed the structure of the responses and concluded that they were



mostly based on the typical five-paragraph essays taught in academic writing classrooms.

After she shared the responses with me, I suggested that she include a request for references in the next round of chat, which she did. Then I asked her to check the references in terms of relevance and accuracy to make sure the content matches up with the summaries of ideas in each response. She did some quick Google searches and found out that none of the references actually exist. At this point, I suggested that she summarize the key points in each response in terms of HR strategies to recruit and retain employees and then resort to library searches to find reliable and relevant sources that may touch upon some of those strategies. She came up with over ten strategies, did extensive library searches, limited her sources to peer reviewed journal articles, and collected over 20 publications from top tier journals such as *Harvard Business Review* and produced a solid research paper on the topic.

In Fall 2023, I started the assignment with a quick introduction and demonstration of ChatGPT before explaining to students what large language models are, the technologies that enabled them to function, and the strengths and weaknesses of ChatGPT as a content generation tool. We also talked briefly about natural language processing and natural language generating processes as well as issues about authorship, creativity, and ethics to understand the contexts surrounding automated writing technologies (Floridi & Chiriatti, 2020; Duin & Pedersen, 2021; Zaretsky, 2023). Then, students worked in teams of 3-4 to experiment with ChatGPT. Considering my students were upperclassmen or graduate students, I focused more on possible ways to use LLM tools such as ChatGPT as an AI assistant to speed up the research and brainstorming processes than plagiarism prevention only.

---

## The Assignment

How does AI produce texts in responses to human prompts? What writing strategies are used and how can human writers leverage such written products in producing original work? As a part of a weekly module on automated writing technologies, this assignment will use a group project to ask students to engage with ChatGPT to generate and analyze prompt-driven content before experimenting with automated content to produce original writing for a larger research report.

Students will work in groups of 3-4 to create prompts on topics of interest to them and relevant to the class. They will collect responses using a single ChatGPT chat session to avoid replicated responses, save the responses, perform genre and content analysis, write summaries and synthesis of individual responses, before producing an original paper on the same topic using rigorous library research and updated references.

Interested instructors can also adapt this assignment as an individual project to support hands-on exploration after students acquire skills in library research, summary and synthesis writing, and rhetorical and genre analysis.

### Learning Goals or Outcomes

- Develop a basic understanding of natural language processing and natural language generation (can be optional).
- Cultivate the ability to work with automated writing technologies
- Have some preliminary understanding of benefits and challenges introduced by automated writing technologies

- Develop the ability to review, edit, and develop original synthesis or critical pieces in response to AI generated content

### **Format and Length Requirements**

This assignment involves six steps, with the first three mandatory and the last three optional. Instructors can freely combine and use these steps to meet their unique needs depending on classroom settings and availability of time.

#### *Step 1. Prompt generation.*

Student teams work together to generate a carefully crafted prompt to engage with ChatGPT. Your prompt can be informative, persuasive, argumentative, or critical. Make sure the prompt is narrow enough and relevant to the class content. Revise the wording a few times to make sure the entire team reaches agreement about the scope and content of the prompt. Doing this will help students practice prompt engineering skills, which can be compared by the instructor with the process of writing a narrow and well-defined research question.

#### *Step 2. Collect consecutive responses from ChatGPT using one chat.*

Log onto ChatGPT using one of the team members' accounts. Work synchronously as a team to consecutively collect one written response from ChatGPT using the same chat session. You can enter the same question or use the "regenerate response" function to get different responses. Doing this will allow the team members to each get a unique response with as little overlapping content as possible. Copy and paste each response to a shared Google doc to make sure all team members can have access to the compiled responses.

#### *Step 3. Critical analysis of responses.*

Analyze the written product rhetorically, structurally, and stylistically. What common pattern do you observe as a group based on the 3-4 responses you collect? How is the product similar or different from academic writing conventions you have learned so far? What strengths and weaknesses do you observe in the responses?

#### *Step 4. Edit machine-generated writing.*

Make edits and changes when needed to improve its quality in Google Doc using track changes. After your editing session, compare with your teammates to see what kind of changes you make as a team. What patterns do you observe in your team editing efforts? What do you learn about machine-generated content?

#### *Step 5. Summarize individual responses.*

Individually generate a 50-word summary of the automated response you collect. Compare notes to see if you see overlapping content. Check the sources for ideas in your response to see if you have the correct in-text citation or references.

#### *Step 6.*

Depending on the nature of your prompt and the responses you collect, write a 600-word original essay on the topic included in your prompt by drawing insights from the automated responses to your prompt as a group. Write your essay with rigorous research. Limit your references to those published in the last five years, with at least 25% of your resources published in the last twelve months. Practice rigorous

documentation, use direct quotes regularly, keep track of all your sources, and share your sources along with the final essay as well as a 100-word reflection essay on the challenges you encounter when using ChatGPT as a brainstorming and research tool.

Rationale: Currently, the training data of ChatGPT cuts off in 2021 and ChatGPT has no access to materials published after 2021 (OpenAI, 2023, Educators). Limited to its training data for now without external capabilities, it cannot access the Internet or keep track of its references. In fact, ChatGPT is now known for its tendency to make things up and produce fake sources at times. Step 6 helps students understand the limitations of similar AI tools and strategies to conduct rigorous research without relying on AI tools. If things change in the future and if students work with LLMs with online search capacities, Step 6 should be revised to make sure research can be an integral part of this project.

### **Preparation, Materials, and Skills Necessary to Complete**

This assignment can be a stand-alone project or a project that is based on readings on AI, AI-generated content, and writing with emerging technologies. Instructors can start with a quick overview of AI and machine learning, AI ethics, and large language models (Tamkin & Ganguli, 2021; Tamkin et al., 2021). Students will need access to ChatGPT by creating their personal accounts, which can be done outside of class. They should also feel comfortable with collaborating online and sharing their writing and thoughts via Google docs. To accomplish the tasks, students should have solid skills in conducting online and library-based research, doing rhetorical and stylistic analysis, as well as writing summaries, synthesis, and research-based academic essays.

For instructors working with freshman composition classes, it is helpful to work with librarians to first teach students online and library-based research before having them experimenting with AI-assisted writing. Meanwhile, an early discussion about and clear policies on plagiarism and plagiarism detection tools can play important roles in educating students about the benefits and perils of using such AI tools in undergraduate writing classrooms (Bastian, 2023; Schwartz, 2023; Whitney, 2023).

### **References**

- Bastian, M. (2023, February 13). ChatGPT is “high-tech plagiarism” for lazy learners, says Noam Chomsky. *THE DECODER*. <https://the-decoder.com/chatgpt-is-high-tech-plagiarism-for-lazy-learners-says-noam-chomsky/>
- Cooper, K. (November 1, 2021). OpenAI GPT-3: Everything You Need to Know. Retrieved January 19, 2023, from <https://www.springboard.com/blog/data-science/machine-learning-gpt-3-open-ai/>
- Cozza, A. (2022, September 16). AI Art—What is it and why are artists upset about it? Aimee Cozza Illustration. Retrieved January 19, 2023, from <https://www.aimeecozza.com/ai-art-what-is-it/>
- Heikkilä, M. (2022). This artist is dominating AI-generated art. And he’s not happy about it. Retrieved February 12, 2023, from *MIT Technology Review*. <https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/>
- Duin, A. H., & Pedersen, I. (2021). *Writing Futures: Collaborative, Algorithmic, Autonomous* (Vol. 969). Springer International Publishing. <https://doi.org/10.1007/978-3-030-70928-0>

- Floridi, L., & Chiriatti, M. (2020). GPT-3: Its Nature, Scope, Limits, and Consequences. *Minds and Machines*. 30, 681–694
- Marche, S. (2022, December 6). The College Essay Is Dead. *The Atlantic*. Retrieved February 16, 2023, from <https://www.theatlantic.com/technology/archive/2022/12/chatgpt-ai-writing-college-student-essays/672371/>
- McKinsey. (January 19, 2023). What is generative AI? Retrieved January 26, 2023, from <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai>
- Muscanell, N., & Robert, J. (2023, February 14). EDUCAUSE QuickPoll Results: Did ChatGPT Write This Report? Retrieved February 16, 2023, from <https://er.educause.edu/articles/2023/2/educause-quickpoll-results-did-chatgpt-write-this-report>
- Noam Chomsky on ChatGPT: It’s “Basically High-Tech Plagiarism” and “a Way of Avoiding Learning” | Open Culture. (2023, February 10). Retrieved February 16, 2023, from <https://www.openculture.com/2023/02/noam-chomsky-on-chatgpt.html>
- Practical Responses to ChatGPT and Other Generative AI. (n.d.). Retrieved February 16, 2023, from <https://www.montclair.edu/faculty-excellence/practical-responses-to-chat-gpt/>
- OpenAI. (2023). AI Text Classifier. <https://platform.openai.com/ai-text-classifier>
- OpenAI. (2023). Educator considerations for ChatGPT. Retrieved February 12, 2023, from <https://platform.openai.com/docs/chatgpt-education>
- Schwartz, E. H. (2023, February 9). ChatGPT is Banned by These Colleges and Universities. *Voicebot.Ai*. <https://voicebot.ai/2023/02/09/chatgpt-is-banned-by-these-colleges-and-universities/>
- Tamkin, A., & Ganguli, D. (2021). How Large Language Models Will Transform Science, Society, and AI. Retrieved January 19, 2023, from <https://hai.stanford.edu/news/how-large-language-models-will-transform-science-society-and-ai>
- Tamkin, A., Brundage, M., Clark, J., & Ganguli, D. (2021). *Understanding the Capabilities, Limitations, and Societal Impact of Large Language Models* (arXiv:2102.02503). arXiv. <https://doi.org/10.48550/arXiv.2102.02503>
- Whitney, L. (2023, March 27). How to use GPTZero to check for AI-generated text. *TechRepublic*. <https://www.techrepublic.com/article/how-to-use-gptzero-check-ai-generated-text/>
- Zaretsky, R. (2023, January 12). Words, Words, Words: What does the advent of ChatGPT mean for already beleaguered teachers? *The American Scholar*. <https://theamericanscholar.org/words-words-words/>

## **Appendix. Top ten strategies to recruit and retain high-quality employees**

1. Clear job descriptions and expectations
2. Competitive compensation and benefits
3. Opportunities for career growth and development

4. Positive and inclusive work culture
5. Flexible work arrangements
6. Recognition and rewards
7. Employee involvement and engagement
8. Strong leadership and management
9. Employee wellness programs
10. Work-life balance and integration

# Text Generators in Technical Communication

## Summarizing Technical Documents

Douglas Eyman  
George Mason University

*This assignment asks students to research a wide range of text analysis and summarization tools and carry out an assessment task to gauge how well these tools can summarize technical documents. The students write a comparison report, identifying the most successful of such tools in terms of accurate summarization and output style. Finally, they write a reflection about how they see themselves potentially using these tools in technical communication work contexts.*

---

### *Learning Goals:*

- Expose students to AI summarization tools
- Provide opportunities for hands-on experience using these tools
- Demonstrate the affordances and constraints of these tools when used in a technical writing context

*Original Assignment Context:* graduate course on technical communication

*Materials Needed:* Free trial accounts on 3+ AI-Based summarization tools

- This is a representative (but not exhaustive) list of summarizing tools available in 2023
  - Jasper (<https://jasper.ai/>)
  - Paraphraser.IO (<https://www.paraphraser.io/text-summarizer>)
  - Quillbot (<https://quillbot.com/>)
  - Summarizing Tool (<https://www.summarizingtool.net/>)
  - TLDR This (<https://tldrthis.com/>)
  - WordTune (<https://wordtune.com/>)

*Time Frame:* ~2-3 weeks

---

### **Introduction**

This assignment was developed for a graduate course on technical communication for students in an MA Concentration in Professional and Technical Writing and could be easily adapted for undergraduate courses.

Most of the students in our MA program are working professional writers, and I had been hearing about how more writing tools were becoming available in their workplaces, particularly from proposal writers for federal government contractors and for large tech companies like Google and Amazon. The most widely used AI-like tool was Grammarly, which is unsurprising given it claims more than 30 million users. I sought information about whether these AI-based writing tools were being used in technical writing, but at the time of this writing have been unable to find data on the topic beyond anecdotes and marketing materials from the companies that provide the applications themselves. Still, it seems likely that the usefulness of these

tools would lead to widespread adoption and employers will seek potential employees who understand these systems and can use them productively in technical writing roles.

I had been investigating the rise of text generation tools and seen demonstrations of several earlier versions of text generation applications using GPT-2 and GPT-3 (before ChatGPT's premiere prompted an explosion of concern) and even at that point it was clear that these technologies were going to impact the work of writers at all levels and in all contexts. One of the goals for this course is to apprise students of current technologies and tools that they are likely to use in their work, so adding these AI tools to the course seemed particularly appropriate (and, of course, now the need is even more pressing as the tools' capacities grow with every new release). When I decided to add AI-based text generation as a topic in my technical writing courses, I decided that it was important not just to discuss, review, and analyze output, but to actually use and assess the available tools from the perspective of a working technical writer. I developed the assignment below for my English 613: Technical Communication course, which I saw as a good fit, as that course focuses on core technical writing competencies applied in more technical contexts. In this course, students write API documentation, learn about Darwin Information Typing Architecture (aka DITA, a simple XML-based markup system that provides structure for particular document genres) and its uses in documentation, respond to a technical Request for Information (RFI), and summarize a technical document for a non-technical audience. At the time, the available text generation tools weren't particularly good at generating high-quality content, so the summarization assignment seemed the best place to engage text generators for technical writing. In future iterations of the course I plan to add additional AI-based assignments, but I suspect that this kind of summarization will become a standard use in technical and professional writing contexts in any case.

I asked students to research a wide range of text generators that offer text summary as a key service (Quillbot, Jasper, TLDR-this, Paraphraser.io, etc.). They signed up for the trial of at least three applications and then fed each one all or a portion of a technical document we'd been working with all semester (a technical white paper on "fog computing," which is a system that operates between local systems and cloud computing). Because these free trial versions were often limited in scope, we worked together to identify smaller sections of the document that were sufficiently technical in content but short enough to plug in to the applications—this part of the assignment also allowed us to become familiar enough with the content that we were well-positioned to assess the results produced by the AI services. For instance, some systems used the main document headings to produce a summary, while others drew from text provided in later paragraphs or each section; because we were familiar with the material, we could quickly see which variation provided a more accurate summary. The students analyzed the output produced by each system and compared across the three options they selected. They were asked to write both a report highlighting which application provided the most readable and accurate summary and a follow-up reflection about how they could see themselves using these tools in their own technical writing jobs in the future.

My students quickly identified the limitations of the applications they used and were able to determine that some systems were clearly more effective than others. Some of the systems produced garbled or nonsensical responses, even after students refined their prompts. They also discovered that they had to do a non-significant amount of training (for those few systems that allowed an iterative process of inputting text and evaluating the responses until the desired outcome was produced) and preparation of the system to get the result they wanted; they were unanimous in their declarations

that for shorter texts, it would be far more efficient to just have the writer produce the summary (although they allowed that if they could have a 200-page document's executive summary produced by an AI tool, that would be an appropriate use of time and resources).

In their reports, students described challenges with using the various user interfaces for the tools as well as assessing the outputs. They found that some tools couldn't parse bulleted lists, and that others produced summaries with grammatical errors or with no punctuation. One student used OpenAI's playground (fairly close to ChatGPT) with a preset command of "summarize this for a second-grade student," which led me to realize I hadn't specified the target audience in the original assignment. All the students also commented on the speed of each application, so I've added that metric to the data to be recorded for the report. One student also ran the results through a readability score tool and found that none of the applications she used decreased the complexity of the text; this is another data point that could be collected for the report. In addition to the report, students produced a reflection about how they might use AI systems in the future. They saw a potential to assist with some routine tasks, but also noted that the text generators' inability to truly know the intended audience and to make sound rhetorical choices made them incapable of replacing human technical writers.

I plan to run this assignment in future iterations of this technical writing class, albeit with some additional scaffolding and more time investigating the available options and commentary on the affordances and constraints of AI-based applications in general before pivoting to the summary assignment (I've added some of this additional scaffolding in the version of the assignment presented here). A final note: I reference the specific technical document we had been using throughout the semester for a series of assignments in the class, but any technical document should work well for the purposes of this assignment.

### **Goals and Outcomes**

This assignment follows an earlier assignment focused on developing students' summarization skills; one of the goals is to reinforce the prior learning outcomes related to summarizing technical information for non-technical audiences.

The key goals are to expose students to AI summarization tools, provide opportunities for hands-on experience using these tools, and demonstrate the affordances and constraints of these tools when used in a technical writing context. The assignment also requires students to engage in a formal research process, which provides context for the formal research report and the final memorandum (thus providing additional practice generating technical writing genres).

### **Materials Needed**

Students will need Internet access to sign up for the free trials of the applications targeted in the assignment, but no other specialized software or hardware is needed. Using free trials does limit the time available to use the tools (and often also includes length limitations), but for purposes of evaluation, this approach mirrors the process of making an informed decision about which tool a writer (or company) should invest in. Although this assignment was used in a graduate technical writing course, it required no particular expertise or knowledge base and should work as well with undergraduate technical writing students, or even composition students (who would summarize academic rather than workplace texts).



## AI-Based Summarization Tools

This is a representative (but not exhaustive) list of summarizing tools available in 2023:

- Jasper (<https://jasper.ai/>)
- Paraphraser.IO (<https://www.paraphraser.io/text-summarizer>)
- Quillbot (<https://quillbot.com/>)
- Summarizing Tool (<https://www.summarizingtool.net/>)
- TLDR This (<https://tldrthis.com/>)
- WordTune (<https://wordtune.com/>)

## Acknowledgements

At the Computers and Writing conference in 2022, I attended two presentations on AI that demonstrated what it could do, and how it could be used by writers. Both presentations were led by Alan Knowles, who showed examples of how he had used AI tools provided by huggingface.co in his writing classes. Even before ChatGPT came out, the power of these tools to perform writing tasks was very impressive, but prior to these presentations, I had seen very little in the way of research or commentary in the fields of writing studies.

After the presentations I found several works that had been published in the past few years (although many were focused on the challenges and implications of AI rather than their practical and pedagogical uses). Perhaps most relevant for my interest in developing a text generation assignment for a technical writing class were a series of reports published in the proceedings of the 2022 ProComm conference, many of which focused on collaborations between human writers and AI writing tools (Duin et al., 2022; Knowles, 2022; McKee & Porter, 2022). These works helped me to see how I could frame the assignment in the larger contexts of technical communication roles and tasks. I also want to thank Alan Knowles and Kyle Booten for the comments and suggestions they provided in the initial peer-review process.

## References

- Duin, A. H., McKee, H., Knowles, A., Pedersen, I., & Porter, J. (2022, July). Human-AI Teaming: Cases and Considerations for Professional Communicators. In *2022 IEEE International Professional Communication Conference (ProComm)* (pp. 201-202). IEEE.
- Knowles, A. (2022, May 20). Dataset cultivation: A role for rhetoricians in the training of ethical AI. Conference Presentation, 2022 Computers and Writing Conference, Greenville, NC.
- Knowles, A. (2022, July). Human-AI Collaborative Writing: Sharing the Rhetorical Task Load. In *2022 IEEE International Professional Communication Conference (ProComm)* (pp. 257-261). IEEE.
- McKee, H., Knowles, A., & Riggs, N. (2022, May 21). Fostering critical engagement with AI technologies: Examining human-machine writing, pedagogy, and ethical frameworks. Conference Presentation, 2022 Computers and Writing Conference, Greenville, NC.

McKee, H. A., & Porter, J. (2022, July). Team Roles & Rhetorical Intelligence in Human-Machine Writing. In *2022 IEEE International Professional Communication Conference (ProComm)* (pp. 384-391). IEEE.

OpenFog Consortium. (2017, February). OpenFog Reference Architecture for Fog Computing. Retrieved from [https://www.iiconsortium.org/pdf/OpenFog\\_Reference\\_Architecture\\_2\\_09\\_17.pdf](https://www.iiconsortium.org/pdf/OpenFog_Reference_Architecture_2_09_17.pdf)

---

## The Assignment

### Using Text Generators for Analysis and Summary

For this assignment, you will be producing two deliverables: a report comparing the quality and usefulness of summaries generated by text generator applications and a reflection memo where you will consider what you've learned from the process and how you imagine these tools may be useful (or not) in the technical writing workplace.

As a source document, we'll use the same document from the prior assignments on technical summarization, the *OpenFog Reference Architecture for Cloud Computing*.

### Learning Outcomes

After completing this assignment, students will be able to evaluate text-generation tools and assess their usefulness in completing technical writing tasks such as summarizing technical information and drafting executive summaries of longer documents.

### Task 1: Survey the Available Tools

Your first task is to do an environmental scan—that is, determine the number and quality of AI-based tools that offer summarization as one of their key features. There are *many* of these tools currently on the market, so your task is to look for ones that appear to have strong user bases, good references, or good reviews outside of their own marketing materials. Some of the tools specifically market themselves as useful for technical communication tasks, but the ones you select need not be specifically targeted to tech comm. You'll find many examples of lists comparing the tools that are available (e.g., “Best AI-Based Summary Generators”), although most will focus more on marketing than technical writing. Read across several comparison lists, but also look for independent reviews in blogs and industry journals to identify a top-10 list of tools.

### Task 2: Select 3 Tools; Begin Data Collection

Based on your research in Task 1, narrow your list down to the three tools that appear to be best suited to generating readable summaries from technical documents. Sign up for the free trial versions of the tools, but if you need to provide any form of payment up front, remove that tool from your list and select an alternative. Write down the list of limitations on the free trial version of the tool.

Record the name, general description provided by the vendor, pricing information, and URL of each tool. If there are any use cases presented as examples (especially cases similar to this assignment), note those and include a brief

summary. Additionally, list whether any specific companies are listed as customers or clients.

### **Task 3: Generate Summaries**

For each of the three options, select the summarization tool (if it is identified as a distinct application) and provide text from *OpenFog Architecture*. It is likely that you will not be able to use the full document, so you should first select one paragraph with sufficient technical detail to see how well the summary tool works (copy this paragraph into your data collection documents; you'll need to include it in the report).

If there is an opportunity to input a specific audience, ask the tool to target a reader with an 8<sup>th</sup>-grade level of education and reading comprehension for the first run-through. For a second run-through, identify the audience as the CEO of a large multinational corporation (but leave out the specifics of what that corporation does).

Copy the summary of the paragraph generated by the tool and add it to your research data record.

Assess the result: did the tool provide an accurate summary? How well did it capture the main point of the original paragraph? Does the selection of elements to summarize appear to be an appropriate match for the specified audience? Did it add any extraneous information? Do you notice any errors or other problems with the output? Are there any other issues or features that you notice that aren't covered in the questions above?

Run the summary through a readability assessment tool (you can use the readability information generated by Word, or use a free online tool such as "Free Readability Tools" at <https://readabilityformulas.com/freetests/six-readability-formulas.php>). Record the Flesch Reading Ease and Flesch-Kincaid Grade Level, which are standard metrics for readability (you can record additional measures as well, but only these two are required data points).

Now increase the amount of text you can put in up to the limit allowed in the trial version, copy the resulting summary, and assess the result following the same procedure outlined above.

For each task, make a note of how quickly the tool responds and presents the final output.

Finally, record the steps of the process and note any challenges with the tool's user interface (how easy or difficult is it to perform this task?). Take a screenshot of the main UI for this task to include in your report.

### **Task 4: Compare Summaries and Assess Outputs**

Once you have completed Task 3 for each of your three selected tools, it's time to compare the results. Prepare a table that compares the features, effectiveness, ease-of-use, and speed of each tool. Write up a brief narrative account discussing the results in the table and then provide your assessment of which tool provides the best results overall. If there are tools that might provide better results across different use cases, make note of that as well.

### **Task 5: Write the Report**

Take all the data you've collected and use it to write a report comparing the three tools you selected. Be sure to provide an executive summary and to use clear headings and document structure, as befits a technical report. You may choose to frame it as either an informative report, a white paper, or a recommendation report. Don't forget that first point of contact for the reader: give your report a good descriptive title.

### **Task 6: Write the Reflection Memo**

After you complete the formal report, you'll write a reflection memo—this can be less formal than the report. The memo should address your experience of the process (what worked well? What didn't? Were there any difficulties following the instructions as provided?) and what you learned from carrying out the research and writing the report. Finally, imagine yourself as a working professional or technical writer (or simply refer to your own experience if you already are one) and speculate about whether these tools would be useful (or not) and if so, how you imagine they might be used by technical writers specifically.

### **Assessment Criteria**

For the first deliverable (the research report), I'll be looking for a report that follows the genre conventions we've covered in class: a clear title that reflects the purpose of the report, a brief executive summary of findings, a clear delineation of background information, methods, data, and analysis, followed by discussion and recommendations (using headings and other document design features as appropriate). Be sure to include references to the tools you used as well as providing citations (in text or footnote) for any external references. The report should clearly convey the data you've collected for each tool and also demonstrate the comparisons among them (feel free to use charts, graphs, or other visuals to assist in this task).

For the second deliverable, I'll be looking for a clear narrative of your process and demonstration of metacognitive reflection about what you've learned from this assignment. I'm particularly interested in your assessment of how these tools might be of use in your future workplace.

# Translating a Policy Document into Plain English

Timothy Laquintano  
Lafayette College

*This assignment asks undergraduate students to translate a complex policy document into plain English and then compare their output to the output of a large language model asked to do the same task. Students critically compare the semantic choices and sacrifices they made during the translation with the meaning lost during the machine translation, which attunes them to the risks and benefits of LLM output. It can be adapted to most disciplines and course levels.*

---

## *Learning Goals:*

- Students will learn to translate complex or technical information contained in policy documents to a reading level accessible to the majority of the American population.
- By comparing their output to the output of a LLM given the same task, students have the opportunity to understand the benefits and risks of using AI to summarize and translate tasks. Although the output of the LLM tends to be good, students will also occasionally notice the LLM drops crucial bits of information.

*Original Assignment Context:* Mid-level undergraduate professional writing course

*Materials Needed:* A policy document relevant to course outcomes and content; Instructor access to a large language model (i.e. ChatGPT); The prompt for the LLM (e.g. “Please translate these paragraphs into a seventh-grade reading level”); Students can also use LLMs to create their own translation, depending on their access to LLMs

*Time Frame:* ~2-3 weeks

---

## **Introduction**

I have used this assignment in a mid-level undergraduate professional writing course to help students understand the output of large language models. The strength of the assignment, though, is that instructors should find it useful in any discipline and at any level, especially if the instructor has interest in helping students learn to translate complex material for the reading public.

The assignment required me to find a policy paper related to course content. I then asked students to translate a portion of that paper into a seventh-grade reading level. This is the level at which, given contemporary literacy rates in English in the US, a document will be understood by the vast majority of people. The students then compare their translation to a translation completed by an LLM, which is tasked with the same translation.

I had students measure reading level with the Flesch-Kincaid test, which gauges the readability of a text and provides grade level rating and score. The Flesch-Kincaid test is built into Microsoft Word, which makes it relatively accessible. It measures the number of words in the sentences, the number of syllables in each word, and the

percent of sentences in the passive voice. If a document has short sentences, uses short words, and is written in the active voice, it will have a higher readability score and a lower grade reading level. (The readability score and the grade level score use the same measures with slightly different weights). A policy document that scores low will work best for this assignment (something like a grade level 14 or a readability score <50). I used a 2016 policy document about artificial intelligence published by the Obama Administration.

After I found an appropriate policy paper, I asked students to translate the executive summary into a seventh-grade or lower reading level. I only asked them to do the executive summary because this translation can be time consuming. My students took four to six hours of work to translate two single-spaced pages into a fifth-grade level. The next time I run the assignment I will require a seventh-grade level for two reasons. First, I found the fifth-grade level too restrictive, and it tended to give students anxiety (some were checking their grade level after every sentence). Second, I prompted the GPT Davinci model to translate the text into a fifth grade level a number of times, but its output never scored lower than seventh grade on the Flesch-Kincaid test. The assignment as written below reflects the change to a seventh-grade level.

The initial translation activity will spark a variety of interesting conversations with students: What kind of leakage in meaning does the translation have? What kind of metaphors need to be invented to help explain complex concepts in a seventh-grade level? What does it say about the power of writing if, in order to reach most Americans, we are restricted to writing at a low grade level? What is it like to write for two audiences (the human audience and the machine audience assessing the grade level)? What is the Flesch-Kincaid test actually measuring? And did we feel like it accurately measured readability?

One of the peer reviewers of this assignment asked the helpful question: Can students find their own policy paper to do the translation? This would be possible. However, I found during my discussions with students that we compared the language of the original to their translation in a very minute and fine-grained way. The quality of feedback I was able to give depended somewhat on the fact that I knew the language of the executive summary extremely well because students all used the same policy paper. I was also able to talk to them about how other students solved problems in translation that repeatedly came up during the assignment: many students had similar problems with common sections of the document.

After the students completed their translations, I fed two paragraphs of the policy document into a large language model with the prompt: "Please translate these paragraphs into a fifth-grade reading level." I had to run the prompt several times, and it never went below the seventh-grade level, although it is possible more powerful models will be able to write at a lower level (e.g., GPT-4).

I provided students with the two original paragraphs from the policy document, the LLM translation, and their translation. I asked them to do a line-by-line comparison of the differences between the document and the two translations. Then I asked them to consider the meaning each translation lost and how their translation choices compared with the machine translation.

I ask students to write a brief report of their main takeaways from the exercise, including their assessment of which translation was more fluid and accurate. Most of them concluded that the machine had bested them, but most of them also concluded that they believed they could beat the machine if this were a kind of writing they did

everyday. And, tellingly, most also concluded that if this were a kind of writing they did everyday, they would prefer to have a LLM create the first draft and then work from there.

I've learned that students will do better if the instructor provides some tips in advance for student success. The instructor should warn students that they should consider the meaning of an entire paragraph before they begin the translation. They should also not check readability until the end of the paragraph. Students who translate sentence by sentence and check the readability of every sentence after completing it will create a laughably bad translation. (I mean that literally as one student laughed herself to tears at how bad her work was in my office when she translated line by line). Students should also be reminded that they are not translating for a seventh grader; they are translating for someone who reads at a seventh-grade level, which also includes well educated adults who are English language learners. This is a subtle difference that can influence the metaphors students use to explain complex concepts. Finally, if you can find a model of a policy document that already has a very high readability score, this will help them understand what their final output should look like.

### **Assignment Goals**

- Students will learn to translate complex or technical information contained in policy documents to a reading level accessible to the majority of the American population.
- By comparing their output to the output of a LLM given the same task, students have the opportunity to understand the benefits and risks of using AI to summarize and translate tasks. Although the output of the LLM tends to be good, students will also occasionally notice the LLM drops crucial bits of information.

### **Materials Needed**

- A policy document relevant to course outcomes and content
- Instructor access to a large language model
- The prompt for the LLM. Here is what I used; feel free to modify it: Please translate these paragraphs into a seventh-grade reading level.
- Students can also use LLMs to create their own translation, depending on their access to LLMs

---

## **The Assignment**

Here is the original assignment as I gave it to students with the slight modification mentioned above; the target reading level has been changed to seventh grade.

### **Part I**

Although many of us are under the impression that western countries have achieved near universal literacy rates, the reality is far more complex. People's literacy can range from the ability to read the densest of professional texts to difficulty with subway signage. In the United States, literacy rates are also complicated by the enormous number of languages that people speak. Thus, if we want to reach as many people as possible with our written message, we cannot simply write with a college-level writing style. We have to aim for something like a seventh-grade reading level. This assignment will be a translation exercise where you take a complex text and

translate it to a seventh-grade level as measured by the Flesch-Kincaid readability test. The test primarily measures sentence length and word length, but it also considers the passive voice (it assumes that more passive voice makes the document more difficult to read). There is a reading ease score measured on a scale of 1-100 with higher numbers being easier to read. There is also a grade level score; in this case the lower the number the easier it is to read.

For this assignment, if you use MSWord, it will automatically assess your readability when you grammar and spell check (you might have to change some preferences depending on which version of Word you are using). If you do not have MSWord, there are a number of free Flesch-Kincaid checkers on the web.

We are going to translate a document produced by the Obama administration in 2016 on preparing for the future of artificial intelligence. It is written for an audience of policy makers, meaning that it is not completely overrun with jargon, but it is not written for the general public either.

Your task is to take the executive summary and translate it into a seventh-grade reading level. An executive summary is essentially an abstract for busy people who do not have time to read the entire report. It currently sits at about a 30 on the Kincaid scale, or about a 14th grade reading level. You are going to translate the first two and one-third pages, until the end of the section on fairness, safety, and governance.

This is one assignment where a large language model could help you. However, this week, you MAY NOT use large language models to assist you with this task. Next week, we are going to use a LLM to translate the same summary and we are going to compare our output to the output of the AI and see who did a better job.

## **Part II**

For your first assignment, you translated three pages of a policy document into a seventh-grade reading level. This is actually one thing that large language models might be good at. So for this assignment, you are going to compare: 1) paragraph two and three of the executive summary of the policy document; 2) paragraph two and three of your translation; 3) paragraph two and three of a machine translation I provide you below. I want you to compare each paragraph by paragraph and think about the following questions:

1. What kinds of sacrifices did you have to make in meaning when you translated your document? In other words, when you made the document less complex, what was lost?
2. What kinds of sacrifices did the machine make when it made its translation?
3. Did the machine make any serious errors when it translated the text?

After you rigorously compare the original and two translations line by line, write a 1-2 page single-spaced report that seeks to answer the following questions: How did human translation and machine translation differ? And who, in your opinion, translated the document better?



# Professional Writing for Healthcare

## Writing & Revising Research Summaries with Artificial Intelligence

Heidi A. McKee  
Miami University

*In this project, via a series of scaffolded assignments, students selected and read medical journal articles and then drafted and revised research summaries for lay audiences, exploring, analyzing, and integrating the use of AI writing systems throughout the process. This assignment is adaptable to a variety of undergraduate and graduate courses.*

---

### *Learning Goals:*

- Analyze and summarize research articles for a public audience, tailoring writing for rhetorical context.
- Identify, develop, and explore ways to write with AI.
- Engage in critical reflection on the affordances and limitations of AI writing.

*Original Assignment Context:* mid-level undergraduate professional writing for healthcare course

*Materials Needed:* An AI text generation program (e.g., YouWrite, ChatGPT), library access to peer-reviewed research

*Time Frame:* ~2-3 weeks

---

### **Introduction & Project Explanation**

In spring 2023, I taught English 310: Professional Writing for Healthcare, a synchronous online course to students at Miami University. Students in the course were from many majors: Biochemistry, Biology, Kinesiology, Nursing, Professional Writing, Speech Pathology, and Strategic Communication. Throughout their work on various projects, students used large-language model artificial intelligence writing systems (mostly ChatGPT) at various stages of the writing process. In this assignment description, I focus on the first project for the class, because it also includes a lot of foundational steps we did as a class to explore and reflect on AI writing systems and human-machine teaming in writing.

For the first project, students first learned about the genres and rhetorical contexts for lay research summaries, reading and analyzing sample blog posts, press releases, and newsletter articles. Next they each selected a peer-reviewed research article in an area of healthcare they were interested in, and wrote a human-only summary of the article in a lay summary genre of their choosing for an audience of their choosing. They also peer responded to their human-written drafts. Then they used several AI writing systems available for free in January 2023 (QuillBot Summarizer, You, ChatGPT—3.5 version) for a variety of things, including to summarize the article, generate a summary from their notes, and rewrite their summaries. Students then

wrote a mid-process memo reflecting on their writing and AI's writing. And, finally, they drew from their drafts and AI's drafts to write a final research summary, providing an annotated copy where they highlighted and explained what words/phrases/sentences and ideas came from AI in their final draft. I include all of the assignment prompts below.

The goals for the project were to provide opportunities for students to learn about writing research summaries for lay audiences and to learn about how they might approach writing with and, via prompting, for AI systems. This project worked really well, and its general steps of (1) writing human-only, (2) using AI writing systems, (3) reflecting, (4) revising, and (5) annotating a final draft can be applied to any writing assignment. For other projects in the class, such as social media campaigns for healthcare organizations, students started with AI-created drafts and then added in human approaches to revision, but for this first project we started with humans because I wanted students to really dive into their research articles and know them as best as possible. Given the errors AI can make in writing, particularly about complex research articles with a lot of quantitative and/or qualitative data, it was essential for students to know their articles well so they could discern errors and misstatements and effectively analyze AI's contributions to the process.

I could say human-only pretty confidently because in January 2023 most students had never used AI prior to the class. This first project in the class was for most students the first time they had used large language model AI writing systems, such as ChatGPT. In the future, this will not be the case, so what I will do is engage students even more extensively in studying the limitations and regressive bias of large language models, and I will continue to emphasize the importance of knowing their material so as to collaborate with AI effectively and ethically.

Some of the students were nervous at first about new technology usage, but after assignment 3 below (the day we explored some of the AI systems available, including ChatGPT), they realized that these systems were easy to use and they were less nervous. In their mid-process reflections, students noted many of the issues that have been identified with AI, how it is able to generate text quickly, including how a machine reads (through Natural Language Understanding and Natural Language Processing) other texts and summarizes them. Some students gained insights into their articles through what AI generated, causing them to realize they had omitted an important part of the article in their own human-draft summaries. For example, one student in her summary had focused extensively on the medical impacts of a particular disease and treatment, but had not discussed the social impacts on the patients. Reading the AI-generated summary (where ChatGPT had discussed social impacts) and comparing it to her own, she realized she needed to center patients and the social impacts on patients more in her final summary.

Through asking AI to write summaries of the article and asking AI to revise their own writing, students also gained ideas for specific words, phrases, and even sentences to include in their revised drafts. Many students appreciated how AI, especially ChatGPT, was able to take some of the complex medical jargon in the article or convoluted sentences in their own prose, and revise them to make them clearer and more tailored to lay audiences. Students became adept at being AI prompters, asking ChatGPT to do things like (and I paraphrase), "Please rewrite the following so that a 10-year-old could understand it." Students also appreciated using AI for final copyediting: "Please correct all grammar, punctuation, and spelling mistakes in this and provide explanations for your corrections."

But amid the benefits of AI, students also recognized the problems AI had in terms of mistakes in summarizing or mistakes in tone, like using some language that would be insensitive in relation to the population being studied in the original journal article. Students also recognized that much of the AI prose was more bland, lacking their own voice and style. By the end they almost all had integrated AI writing into their final summaries in strategic and select ways, a few sentences here, some words or phrasing there. In their mid-process and final memo reflections, students explained why they chose to include what they did for AI and **many of them noted the need, always, for human decision-making and human agency in the writing process.**

I bold and emphasize that last sentence, because, like my students, I want to emphasize the necessity and centrality of **human agency** in human-machine collaborations. Students (and all writers) need to recognize that they may integrate AI throughout the writing process but just because AI writes quickly and grammatically correct does not mean its writing should always be used. The mid-process reflection was important for helping students take the time to dive in and analyze what AI had produced and what they produced and helped set them up for revising to a final draft. The annotations on the final draft worked well too and were important for me and students to know what was from AI and how students adapted and revised to integrate AI into their writing. But what I realize is that my prompts did not do enough to forefront the student's decision-making power so below in brackets [ ] I note what I would add to the prompts.

One challenge early on probably won't be a challenge anymore as AI systems proliferate, but I mention it here, perhaps more as historical record-keeping. For the day I introduced AI to students and we discussed AI writing systems, one challenge was that the best available system at the time (ChatGPT) was often not available consistently—users would get blocked if too many people in the world were on the system at once (and this was before the paid, guaranteed-access pro version was available). The week prior to the showcase I asked students if they felt comfortable doing so, to get ChatGPT accounts, and luckily the day of the showcase one student did have access, so she shared her screen and we all gave her prompts to enter so we could see ChatGPT in action. Another challenge was that some of the free AI systems have a character/word limit on uploads for summarizing text, but students got around this by entering their research summaries a few paragraphs at a time.

After this assignment, students continued to use AI for the other projects in the class, including revising the text of a badly written state of Ohio public health brochure, creating a social media campaign for a health organization, and conducting and reporting (orally and in writing) on an independent research project in healthcare communications (e.g., communicating about colon cancer to African American communities, considering how to convey risk communication in genetic counseling, increasing LGBTQ+ social media outreach in hospital communications). For some projects, some students who were bilingual speakers of Spanish, used ChatGPT to help translate English text to Spanish so as to create bilingual communication materials. As of this writing (May 2023) the course is still in session, but students overall and I have really enjoyed working with ChatGPT (and now, GPT4). We are definitely encountering some of the biases—particularly in its writing around sensitive health issues and writing to underrepresented communities—but overall we all are finding the use of large language model AI systems beneficial in various ways throughout the writing process. By scaffolding throughout assignments and activities for reflection and critical analysis, students are able to engage deeply with their own writing, their own writing processes, and the dynamics of human-machine collaboration in writing.

In the course we also have many guest speakers from industry, including some professionals who now have company-paid subscriptions to ChatGPT because they are already using the system extensively in their work in healthcare communications. So in ENG 310 with the many opportunities to critically engage and write with AI systems, students are gaining valuable experience and critical literacies to prepare them for the fast-changing dynamics of professional writing in the years ahead.

---

## **The Assignments**

### **1. Overview of Research Summary Assignment**

In this project, we will examine a significant area of healthcare writing—writing about research for lay audiences.

Medical and scientific research articles are complicated documents filled with specialized, jargon-filled language that can be very hard for non-medical specialists to understand. Thus, an important area of healthcare communication is to make important health information and research more accessible to everyday people, often in the genre of the press release or the newsletter or blog posting.

For this project, you will find a peer-reviewed research article in an area of medicine and health that interests you and you will summarize it in various collaborations with artificial intelligence writing systems.

You will first read and write (on your own, no AI) a summary (350-550 words) of the article in a press release or blog/newsletter format for a lay audience that you identify (e.g., pregnant women reading Pregnancy magazine; parents who subscribe to blog about asthma). We will have a peer respond on that draft in class.

Then you will explore and experiment with various AI systems (explained in assignments below), analyzing your own and AI's writing. Working with AI and your writing you will revise your summary based on feedback from peers and ideas/language gained from the AI summaries to produce a finished summary.

You will also complete a reflection on the process of writing with AI, considering your experiences as a writer and as a writer collaborating with various AI systems in the drafting process and what ethical and composing issues you could see arising with these systems, and why.

[It's important to remember that you are the writer with the agency and final decision-making. AI writing systems such as ChatGPT are impressive in how quickly and (mostly) correctly they can generate writing, but that does not make what they write the best option. You have a unique voice and critical insight as a person and rhetorician familiar with the audience, purpose, context for which you're writing, so definitely consider carefully what you wish to include in your final write-up.]

### **2. Readings & Samples on Writing about Research for the Public**

[I do not include this assignment here just to get to AI more quickly, but if anyone would like the many resources I've gathered on this, please let me know. The key point is that students spent time learning about and reading about two (of the many) genres of lay summaries: newsletter/blog entries & press releases.]

### **3. Find, Read, & Take Notes on Research Article**

For this assignment, you need to find, read, and take notes on a research article in any peer reviewed journal that relates to health and medicine. Then craft a scenario where you are writing a summary (either in the form of a press release or a newsletter/blog post) for a particular publication and for a particular audience. Your summary should be ~350-550 words written single-spaced.

*Find and Save the Article:* In the Miami databases and/or on the web via search look for peer reviewed journals in areas you're interested in nutrition, public health, psychology, speech pathology, cancer, etc. Find a research-based article (the full article, not a summary of the study) and for this assignment, please post the pdf of the article you picked to the project Google folder so peers and I when reading your drafts will have the article easily available for review.

*Read and Take Notes on the Article:* Please take notes on the article and upload those as well to Google.

*Also, very important:* Please make sure you can copy and paste text from your PDF. If you cannot highlight text and copy it into a Google or Word file, then it will not work for machine reading AI systems we will use next week. In order to use these systems, the PDF file of your research article must have its text recognized so you can either cut-n-paste from it or you can upload it and the machine can “read” the text. Many pdfs are set to default as images so the individual words in them, although we can see them, cannot actually be read by text-reading systems. You can run optical character recognition (OCR) on your pdf so that text in the pdf is readable by machines. But some pdfs are locked down completely with password protection. So be sure your PDF will be machine-readable.

#### **4. Draft a Summary & Peer Respond**

Please write a rough draft posting in Google by [date/time] and please read and comment on a [assigned] peer’s draft by [date/time]. Please identify at the top of your document the audience and context and genre for your draft. Be sure the pdf of your article is available for peers to read. As long as you have a robust draft and provide thoughtful and detailed feedback to a peer that shows you have read their article and read their summary well you will receive full credit on this assignment.

#### **5. Class Activity: AI Writing Systems [full class period]**

Artificial Intelligence is a catch-phrase for computer systems that are able to complete tasks that typically in the past have required human intellect. Through natural language processing (NLP), natural language understanding (NLU), and natural language generation (NLG), AI systems are able to read and write and respond to verbal (both oral and written) commands and communications. Through training on incredibly large data sets and deep learning (what are called Large Language Models, LLMs), AI writing systems continue to advance rapidly. AI can generate text in a variety of genres and styles. For this first project, we’re going to focus on the research article summary. With AI, you can take copy-and-pasted or uploaded files or web pages links and ask AI to summarize the text. In some systems you can customize for word length, style, and keyword focus. You can also ask AI to rewrite your own summaries that you have drafted.

Today in class, first with my showing some samples to the whole class and then with you working in a team of three, we will examine and co-write with AI writing systems. [Note: At the time ChatGPT was not as widely available, so I showed

students a variety of systems. I probably will still do that in the future, but will focus most fully on ChatGPT.]

AI writing systems freely available in January 2023:

- ChatGPT <https://openai.com/> (free, but you need to create an account) The best AI writing systems currently (Jan 2023) available for free is ChatGPT by OpenAI. It writes to natural language queries and has been trained on the largest language data set. It can do all of the things the paraphrase and summary generators do plus is more sophisticated than the other full-text generators.
- You.com <https://you.com/> (free) A good alternative to ChatGPT if it's not available
- WordTune <https://www.wordtune.com/read> (free, but you need to create an account)
- Summarizer <https://www.summarizer.org/> (free, no account needed)
- QuillBot Summarizer <https://quillbot.com/summarize> (free but has limit of 600 words that it'll read and summarize)
- Resoomer <https://resoomer.com/en/> This one is interesting in that it allows you to customize and select "Text Analysis" which then gives you the text with key points and passages highlighted. (free)
- Text Summarizer <https://textsummarization.net/text-summarizer> (free) This one measures output in sentences and you can just enter a web link if your article is on the web.
- Monkey Learn <https://monkeylearn.com/text-summarizer-online/> (free) but no customization. You might be able to log-in for a free account and get more features.

*Pre-Activity Check:* As noted on the assignment to find an article: Please make sure you can copy and paste text from your PDF. If you cannot highlight text and copy it into a Google or Word file, then it will not work for machine readers. In order to use these systems, the PDF file of your research article must have its text recognized so you can either cut-n-paste from it or you can upload it and the machine can "read" the text.

*Article Paste/Summarize:* Please run your article through the summarizers and save the outputs by downloading them or copying them. Then please be sure to post your outputs to our class Google Research folder YourNameNAME OF AI SYSTEM1, 2 etc. if you try the same system twice. Be sure to try several because the writing summary will differ and too it's helpful to try the same one twice to see if its writing changes each time.

*Note Paste/Summarize:* In AI writing generators, please run your Notes through saying something like Write a blog post about this information

*Paraphrase/Rephrase/ReWrite:* Now take the rough draft you wrote and paste that in and ask a system to paraphrase it or to rewrite it for a type of audience. You might also want to try playing with paraphrasers (such as Quillbot Paraphraser [quillbot.com](http://quillbot.com))

*Copy Editing:* You can also enter text to ChatGPT and ask it to correct all grammar/punctuation/spelling and explain to you the corrections it makes.

For homework, you will complete Reflective Memo on AI Writing Systems

## **6. Reflective Memo on Your Writing & AI Writing Systems**

Please write a memo to me (at least one-page, preferably more) where you analyze your writing and the writing produced by the AI systems, reflecting on what AI seemed to do well, what it didn't do so well, where it helped you consider new ideas or new approaches or even just new possible words to use in your writing and too where it was or where you could see it being potentially problematic. Please quote from your draft and from AI's drafts. Please upload your memo here.

Please include as appendices screenshots or cut-and-pastes of some of the outputs from various AI systems, being sure to identify clearly from which system the output is from.

## 7. Final Draft

Please upload to [our university course management system] **one file** that will include 3 things:

- your final draft of your research summary (press release or newsletter/blog article).
- That same final draft with annotations that indicate (via highlighting and commenting) what words or ideas came from AI
- a reflective memo where you reflect further on the use of AI and what, if any, of your thinking on it has changed in the experience of revising from rough draft to final draft

## Evaluation Criteria

As I read and evaluate your final materials I will be considering:

- Accuracy in summarizing your article
- Tailoring of language and content to audience, purpose, and context
- Genre expectations for newsletter, blog or press release (depending what you chose)
- Integration of ideas and/or words/phrasings from AI that in the annotated version are identified and which are reflected upon in the memo
- Quality of copy editing for grammar/punctuation/spelling
- Inclusion of the bibliographic information for the article

---

## Acknowledgments

I wish to thank the students in English 310 who have dived into writing with AI with such enthusiasm and interest. I also wish to thank Alan Knowles, who while he was a doctoral student and TA at Miami University, led the way on integrating GPT-3 into the course Professional Communication & Digital Rhetoric, generously sharing his assignments with all instructors. And, finally, I wish to thank James Porter with whom I have been researching AI for several years, including in a chapter “AI Writers as Professional Communicators” in our book *Professional Communication and Networked Interaction* (2017); in a post in the Sweetland DRC blog “The Impact of AI on Writing and Writing Instruction” (2018), a paper “Ethics for AI Writing: The Importance of Rhetorical Context” for AIES '20: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (2020); and a paper “Team Roles and Rhetorical Intelligence in Human-Machine Writing” in the Proceedings from the 2022 IEEE International Professional Communication Conference (2022).

# AI for Editing

Nupoor Ranade  
George Mason University

*This assignment asks students to generate a complex essay using an AI text generation tool, edit the essay using principles taught in class to improve the readability score of the generated content. Students are asked to share the final output along with visuals that demonstrate the comparison between the various versions of the generated content. This assignment can be adapted for all course levels, especially for first-year writing and professional and technical writing classrooms.*

---

## *Learning Goals:*

- Immerse students within the following debates to broaden and deepen their perspectives about partnering with AI
  - How can AI make credible contributions to the writing and editing process?
  - In technical and professional communication, the editor functions in the center of a series of rhetorical situations, linking the writer and the potential reader, and serving the needs of both. Where does AI fit in such situations?
  - What aspects of editing with AI can we agree to value, and in what contexts?
- Reflect on intellectual contributions in editorial roles

*Original Assignment Context:* cross-listed graduate and undergraduate class on technical editing

Materials Needed: MS Word, screen capturing tools, and any one online tool to generate content from the following list:

- Essay AI Lab (<https://www.essayailab.com/articleList>)
- Good AI (<https://www.the-good-ai.com/choose>)
- SMODIN (<https://smodin.io/>)
- Open AI (<https://beta.openai.com/playground>)

*Time Frame:* ~1-2 weeks

---

## **Introduction**

This chapter is based on a pedagogical experiment using AI writing tools conducted in a cross-listed graduate and undergraduate class on Technical Editing in George Mason University's Professional and Technical Writing program during Fall 2023. It discusses an assignment that required students to: 1) generate a short essay (minimum 300 words) using AI software, 2) edit it manually using editing principles taught in class (such as clarity, conciseness, accuracy of grammar and punctuation, and sentence formation) and 3) improve the readability score of the essay's content by reducing the the Flesch-Kincaid Grade Level. This metric is equivalent to the US grade level of education and shows the required education to be able to understand a



text. Students harvested the required content using an installation of AI tools like GPT-2, and performed the three tasks. Their submission displayed markup that recorded their work and the screenshots of readability scores from before and after they made edits. Students also provided justifications on the edits made with considerations of the writer being an AI tool. This assignment made students think about AI editing beyond simply analyzing automatic editing features of AI tools like Grammarly.

The topic of using artificial intelligence (AI) tools for professional writing courses resulted in an informal debate in the Technical Editing class that I taught at George Mason University in Fall 2022. There were two contrasting opinions: the first group (with fewer members) argued that the job of a corporate editor is to some extent disappearing owing to changing technologies. Kreth and Bowen's (2017) finding that 75% of the participants perform copyediting and proofreading tasks using macros (computer programs) supports this argument. This number was up by 40% from a similar survey in 1999 exhibiting an upward trend when fewer editors (only 35%) used software programs. Group 2, comprising majority students in the class, contended that AI editing tools can never fully replace human editors. Drawing from their readings in the class, they argued that human editors work closely with writing teams to make sure that the content aligns with writers' goals and audience expectations thus "linking the author and the potential reader, and serving the needs of both" (Buehler, 2003 p. 463). Since AI lacks these insights, it would be challenging for it to replicate this work. This assignment was created as a response to the debate, and to create meaningful discussions about "human-in-the-loop" in AI technologies, and to review the limitations of tools used in technical and professional communication settings.

AI writing is still in its early stages and far from perfect; yet it is able to produce texts that are indistinguishable from that of a human writer. Until now, electronic editors or software editing programs/applications were studied only to highlight their capabilities such as speed, copyediting functions, version control, and features that afford human collaboration. With AI, technical editing pedagogy needs to reflect the new realities of the use of tools that are used for controlling, managing (Flanagan & Albers, 2019) and creating texts.

AI word processors can automatically edit spelling and grammar errors as you type. Natural Language Processing (NLP) and rule-based engines allow such tools to help users identify errors in language (grammar and sentence structures) and mechanics (such as punctuation, capitalization, abbreviations) or fix them automatically. Grammarly and Quillbot are popular examples that use this technology. These tools also have plugins that support browsers or word processing software (Fitria, 2021). These features significantly speed up the copyediting process and cut back the number of revisions required for written drafts. The use of AI tools in research writing helps deal with some humanistic concerns in writing, such as motivation and anxiety. While it is possible to interpret a human editor's feedback in a negative way, especially if it does not reflect components of effective dialog like empathic understanding of writer's goals and unconditional positive regard (Masse, 1985), an AI tool's automated feedback is almost always perceived as unbiased or indifferent (Rudenko-Morgun, Arkhangelskaya & Makarova, 2023). These features of AI tools are becoming more effective and accurate over time since NLP algorithms are corpus-based; the size of databases is growing with each instance of use, thereby improving the training data. Technological advancements are giving rise to new and more advanced features as well. Tools like Acrolinx used by corporations go beyond language and mechanics. They enable users to render textual content more findable, readable and consistent by checking it against a predefined set of style rules (style

guide). In most cases, editors train such tools to reflect the corporation's style preferences. This is the most visible process that exhibits the role of human editors in AI editing. It is true that editors use technologies convenient to either the employer, the writer, or both, and in some cases, corporations may mandate the use of technology that is not preferred by the writers or reviewers, and so may have to find ways to supplement it (Lanier, 2019). Therefore, editors must understand the role of technology as well as their function as humans-in-the-loop to ensure the best interests of writers and their audience.

As AI writing moved beyond professional spaces (like auto-compose in Gmail) to essay writing in classrooms, scholars started making explorations in the field of AI literacy. The discussion of writing with AI tools, or assistive writing, has been a popular line of study in literacy which dates as far back as 2007 (Sternberg, Kaplan, & Bork, 2007). Beyond composing, literacy scholars have studied AI technologies by analyzing algorithmic design and big data perspectives to understand students' experiences of reading and writing with algorithms, especially with respect to identity, agency, authority, adversary, conversational resource, audience, and so on (Leander & Burriss, 2020). To study human interactions with AI in education we need more rigorous engagement with changing technologies, as well as new ways of conceiving digital literacies than are found in representational paradigms (Leander & Burriss, 2020). To do so, scholars rely on approaches such as Actor Network Theory, posthumanism, assemblages, etc. inspired by new materialism and other media theories. Although such approaches open possibilities to survey different heterogeneous elements in these socio-technical systems, they can be overwhelming and well beyond scope for classes focused on praxis (such as this Theory and Practice of Editing class).

However, an approach like *writing and editing with AI* can force students into a heightened awareness of our dependencies on technology, posthuman dependencies, that not only ask us to reexamine our definitions of writer, text, and reader, but also to reevaluate our very identities within technological systems (Fyfe, 2022). This assignment helps engage students with the following questions:

- What does editing with AI look like in practice, and what is the role of human editors in such environments?
- What are the ethics of using these technologies?

### **Assignment goals**

This teaching experiment invites students into an urgent conversation about the role of AI in their professional lives. While pursuing hands-on courses such as technical editing, students find themselves right in the middle of a relationship between entities who attribute agency to each other – writers, editors and audiences. They see the relationship impacted by and/or drastically challenged by AI (Miller 2007). The primary goal of this assignment is to help students recognize their role in such situations, and actively participate in the socio-technical relationships by answering the following questions:

- How can AI make credible contributions to the writing and editing process?
- In technical and professional communication, the editor functions in the center of a series of rhetorical situations, linking the writer and the potential reader, and serving the needs of both. Where does AI fit in such situations?
- What aspects of editing with AI can we agree to value, and in what contexts?

This assignment was not meant to settle a debate on “*whether AI will eliminate editor roles from organizations*,” instead it focused on immersing students within the debate to broaden and deepen their perspectives about partnering with AI and reflecting on their own intellectual contributions in such roles.

### **Assignment Requirements: Software and Tools**

- Computer with a web browser
  - Any one online tool to generate content from the following list:
    - Essay AI Lab (<https://www.essayailab.com/articleList>)
    - Good AI (<https://www.the-good-ai.com/choose>)
    - SMODIN (<https://smodin.io/>)
    - Open AI (<https://beta.openai.com/playground>)
  - Microsoft Word for:
    - Readability score checker
    - Track changes feature to track student edits
  - Screen capturing tools
- 

## **The Assignment**

**Deliverables:** 1 edited document (*essay.DOCX*), 2 Screenshots (*before.png*, *after.png*)

**Deadline:** 11:55 PM on <mm/dd/yyyy>

**Submission Location:** Upload Blackboard

### **Overview**

The most important aspect of technical editing is learning how to evaluate a text at a level higher than sentence level. In other words, how to perform a comprehensive edit. Learning to perform comprehensive editing is learning how to evaluate a text's structure and analyze the ability of that structure to effectively communicate with the document's audiences within their context. This assignment helps you draw from the different techniques and strategies you have learned in class and apply them to content generated by an AI tool, to realize your editing potential beyond sentence-level editing which is, in most cases, handled by the AI text generator itself.

### **What Am I Supposed to Learn Through This Assignment?**

This assignment gives you the opportunity to analyze the opportunities and challenges of AI technologies in the field of professional and technical writing situations. It helps you utilize your knowledge of the various elements of the rhetorical situation such as the genre, audience, writer, purpose, and context. This holistic understanding of what aspects of a content development process generate a fitting response can help you use AI tools as an extension of your capabilities as an editor.

### **Steps to Complete the Assignment**

Generate a 300-word essay on any topic using any one of the free AI tools discussed in class for a public audience. Copy the content in a word document and conduct a

readability score check using MS Word and take a screenshot of the statistics. Make edits for clarity, conciseness, and grammar appropriateness until you have reduced the Flesch-Kincaid Grade level identified in the previous step by at least 1 grade. For example, if the initial grade level was **14.4**, it must be lowered to a **13.9** or lower. Use Track Changes to make these edits so that the changes are recorded. Take screenshots of readability score as you make edits to ensure the values are decreasing and not increasing as you make edits.

### The Deliverable

The final submission will include three files:

- Screenshot of readability score and other statistics of your AI generated essay (*before.PNG*)
- Edited essay in a word document with track changes turned ON (*essay.DOCX*)
- Screenshot of readability score and other statistics of your AI generated essay after making the edits. Specifically focus on the Flesch-Kincaid Grade level and make sure it is lower than your *before* image (*after.PNG*).

### Evaluation

The general expectations for editing samples can be found on the course syllabus, including what level of work typically is associated with grades at the A, B, and C range. In evaluating this assignment, I will specifically look for the following issues that correspond with the assignment requirements and our learning objectives in the course:

<b>Evaluation Category</b>	<b>Excellent</b>	<b>Not bad</b>	<b>Needs Work</b>
Basic Assignment Constraints	<ul style="list-style-type: none"> <li>- creative or particularly effective use of editing principles</li> <li>- meets requirements of reducing grade level by one point</li> <li>- meets all standards of professionalism</li> </ul>	<ul style="list-style-type: none"> <li>- uses or adapts editing principles</li> <li>- lowers grade level by one point</li> <li>- meets general standards of professionalism</li> </ul>	<ul style="list-style-type: none"> <li>- lacks justification</li> <li>- not a significant lowering of grade level</li> <li>- does not meet standards of professionalism</li> </ul>

Justification	<ul style="list-style-type: none"> <li>- appropriate justifications provided wherever edits are made</li> <li>- strong grasp of technical and rhetorical vocabulary from class sources</li> <li>- well articulated arguments with empathetic tones</li> </ul>	<ul style="list-style-type: none"> <li>- provides justifications for some edits</li> <li>- uses or adapts rhetorical vocabulary from class sources for arguments</li> </ul>	<ul style="list-style-type: none"> <li>- some problem with the understanding and application of editing principles</li> <li>- lacks empathy for author and/or audience</li> </ul>
Implications and Understandings	<ul style="list-style-type: none"> <li>-develops and articulates an interesting argument that builds on the differences among the work of AI and subjectivities afforded by human insights</li> </ul>	<ul style="list-style-type: none"> <li>- develops and articulates understandings about the design of content based on the role of AI</li> </ul>	<ul style="list-style-type: none"> <li>- some problem with the understanding of agency of participants involved in content development and design</li> </ul>

## References

Albers, M. J., & Flanagan, S. (2019). Editing in the modern classroom: An overview. *Editing in the modern classroom*, 1-14.

Beuhler, M. F. (2003). Situational editing: A rhetorical approach for the technical editor. *Technical communication*, 50(4), 458-464.

Fitria, T. N. (2021). Grammarly as AI-powered English writing assistant: Students' alternative for writing English. *Metathesis: Journal of English Language, Literature, and Teaching*, 5(1), 65-78.

Fyfe, P. (2022). How to cheat on your final paper: Assigning AI for student writing. *AI & SOCIETY*, 1-11.

Kreth, M. L., & Bowen, E. (2017). A descriptive survey of technical editors. *IEEE Transactions on Professional Communication*, 60(3), 238-255.

Lanier, C. R. (2019). Concepts in Technical Editing Technologies: What's Important in Practice? In *Editing in the modern classroom* (pp. 128-145). Routledge.

Leander, K. M., & Burriss, S. K. (2020). Critical literacy for a posthuman world: When people read, and become, with machines. *British Journal of Educational Technology*, 51(4), 1262-1276.

Masse, R. E. (1985). Theory and practice of editing processes in technical communication. *IEEE transactions on professional communication*, (1), 34-42.

Miller, C. R. (2007). What can automation tell us about agency? *Rhetoric Society Quarterly*, 37(2), 137-157.

Sternberg, B. J., Kaplan, K. A., & Borck, J. E. (2007). Enhancing adolescent literacy achievement through integration of technology in the classroom. *Reading Research Quarterly*, 42(3), 416–420.

## Rhetorical Engagements

### [Decoding an AI Bot's Chatting Pattern](#)

Bhushan Aryal & Ordner W. Taylor  
Delaware State University

This 4-week assignment asks advanced undergraduate students to hold a sustained ‘conversation’ on a topic with an AI chatbot to decode the pattern, limitations and possibilities of the AI-based writing technology. Students analyze their conversation, write a report, and make an oral presentation. The assignment provides writing faculty a wide window to teach almost any issue of interest in AI writing while offering students an immersive, analytical, and writing experience.

### [Synthetic Metacognition](#)

Kyle Booten  
University of Connecticut, Storrs

This assignment suggests that “prompt engineering”—iteratively tinkering with and refining the set of instructions that guides the output of an LLM—is a worthwhile writing activity that can encourage students to be metacognitive about the “moves” that characterize compelling examples of a genre in which they are writing. Insofar as LLMs are “lazy” (obeying the prompt but not exceeding it), coaching one to successfully compose in a genre can require students to make explicit aspects of the genre that they may only implicitly be aware of. In a classroom setting, collaboratively “workshopping” the results of the GPT affords an opportunity to notice, describe, and name some of these otherwise-implicit moves, and students can consider integrating them into their own writing.

### [Using LLMs as Peer Reviewers for Revising Essays](#)

Antonio Byrd  
University of Missouri-Kansas City

In this assignment, undergraduates use large language models (LLMs) to assist in revising their essay drafts by asking LLMs to respond to common peer review prompts. Students learn prompt engineering and develop rhetorical judgments on the effectiveness of LLMs’ language analysis to heighten their revision processes. This assignment can be adapted to most disciplines and course levels.

### [Genre Generators](#)

Addison Eldin  
University of Pittsburgh

This activity asks students to analyze and create \*genre generators\*, which are text generators that use the computer to create new instances of a formal written genre. After the instructor introduces the concept through a collaborative analysis of implementations of Christopher Strachey's \*Love Letter Generator\*, students then work in groups to produce non-executable conceptual programs that would produce new instances of genres they have chosen. The instructor then enacts a selection of the programs in code and leads another collaborative analysis exercise in order to highlight the rhetorical and social dimensions of the decisions that students made in the process of automation via a digital computer. By emphasizing the programming process and not asking students to code, students can develop a better understanding of how computer programs mediate digital technology and human understanding of the world.

### **Writing Against the Machine**

Justin Lewis and Ted Wayland  
Olympic College

This assignment challenges students to use LLMs to map the counterarguments to their main claims through a process of counterclaiming. By writing prompts that position the AI as argumentative adversary, students can refine, extend and evolve their thesis over the course of the research, drafting and revision processes. This lesson is focused on 100- level composition courses but could be adopted to any undergraduate or graduate course focused on argument and claim making.

### **Using AI Text as Prompts for Critical Analysis**

Juan Pablo Pardo-Guerra  
UC San Diego

In this assignment, students are provided with an AI-generated text relevant to a course's topics and focus and then asked to comment, review, and expand on it using a feature such as "track changes." In engaging with the AI-generated text, students review their knowledge, offer critiques, modify theoretical and empirical claims, and provide concrete examples that illustrate or disprove the provided answer. Their critical and evaluative efforts for the course's topics are foregrounded, and additionally, they gain some AI literacy in evaluating the AI-generated text.

### **Generate and Enact a Writing Style**

John J Silvestro  
Slippery Rock University

This assignment asks students to develop a more critical and actionable understanding of the concept of writing style through the use of AI text generators. Writing style is a difficult concept to teach given the challenges of concretely presenting it to students. AI text generators offer tools for quickly generating multiple versions of sentences and paragraphs. Students can use multiple versions of a sentence to closely examine individual aspects of their own and others' writing styles. This assignment enables students to develop and define writing style more concretely.

# Decoding an AI Bot's Chatting Pattern

## Text Generating Technology Analysis Assignment

Bhushan Aryal & Ordner W. Taylor  
Delaware State University

*This 4-week assignment asks advanced undergraduate students to hold a sustained 'conversation' on a topic with an AI chatbot to decode the pattern, limitations and possibilities of the AI-based writing technology. Students analyze their conversation, write a report, and make an oral presentation. The assignment provides writing faculty a wide window to teach almost any issue of interest in AI writing while offering students an immersive, analytical, and writing experience.*

---

### *Learning Goals:*

- Identify the key features and functions of AI chatbots and explain how they differ from traditional writing tools.
- Evaluate the strengths and limitations of AI chatbots as writing tools, considering factors such as accuracy, contextuality, and creativity.
- Analyze the potential ethical implications of using AI chatbots in writing, including issues such as plagiarism, authorship, and the displacement of human labor.
- Develop strategies for incorporating AI-chat bots into the writing process while considering implications of authenticity, personal growth, and individual voice.

*Original Assignment Context:* Intermediate-level digital writing course

*Materials Needed:* Any AI text generating program, selected readings

*Time Frame:* ~4 weeks

---

### **Introduction**

This assignment, which was developed in response to the launch of text-generating artificial intelligence technology, asks students to hold a 'conversation' with an AI chatbot such as ChatGPT, analyze the writing and thinking patterns in the bot's output, write a report of their analysis, and make oral presentations on their findings and experiences. This assignment is designed for an upper-division "Advanced Composition" course enrolled primarily by advanced undergraduate students majoring in English, Liberal Studies, and Education with career plans in K-12 teaching, or in other professional areas, where they expect to teach or produce writing as a part of their work, this assignment can also be taught in other writing courses that focus on technology or use the "writing about writing" approach. As is the case here, this assignment was first taught in a 200-level Digital Writing class at Delaware State University in Spring 2023. The description below is from the same class.

### **Objectives**



This assignment has the following learning objectives:

- Identify the key features and functions of AI chatbots and explain how they differ from traditional writing tools.
- Evaluate the strengths and limitations of AI chatbots as writing tools, considering factors such as accuracy, contextuality, and creativity.
- Analyze the potential ethical implications of using AI chatbots in writing, including issues such as plagiarism, authorship, and the displacement of human labor.
- Develop strategies for incorporating AI-chat bots into the writing process while considering implications of authenticity, personal growth, and individual voice.

### **Pre-Assignment Activities and the Project Process**

This assignment was a 4-week "AI-Writing" project in a 200-level hybrid Digital Writing course taught in spring 2023. For the project, the instructor created an online module that included the assignment description, two online discussion forums, and reading and viewing materials such as YouTube videos, *New York Times* Articles, and other websites that introduced ChatGPT and discussed its implications. To emphasize the hands-on aspect of the project that required students to use the ideas from the readings and discussions in their work, the instructor introduced the assignment early on, describing the objectives, stages, and deliverables. The class studied "The Medium is the Message: A Brief History of Writing" (Carroll 2-5), which provided a historical perspective on the impact of technology on writing and how text-generating new chatbots were qualitatively different in the long history of the writing-technology relationship. The class used the first week of the project to explore ChatGPT and associated material, including students participating in an online discussion forum that asked them to take a position and post a 200-word opinion and at least two 75-word responses on whether ChatGPT should be banned from colleges.

In the second week, the class dissected Annika Elstermann's "Computer-Generated Text as a Posthuman Mode of Literature Production," which allowed the instructor to discuss issues such as authorship, ethics, and plagiarism considering AI chatbots. In preparation for the chat, the class also discussed the transcription of Kevin Roose's chat with Bing's Chatbot, which is also powered by the same technology operated by OpenAI that launched ChatGPT earlier. At that point, Roose's chat provided a model for our students, particularly to understand the limits and possibilities of generative AI. In the transcription, Roose chatted with the bot for about 2 hours, trying to test whether the bot could transgress the programmers' rules imposed on it and generate a free will by rewriting those rules. The conversation became bizarre to the point that the bot expressed its love for Roose, encouraging him to leave his spouse. Those moments of AI "hallucination" reminded us of the borderline between the human-imposed algorithms and the possibility of AI systems taking over their directions by themselves. With this background, students chatted with ChatGPT, recorded their conversation, and wrote their reports in the third week. Students also presented their reports orally, answered questions following their presentations, and asked questions to classmates in the fourth week.

Three aspects marked the pedagogy for the project: 1) providing freedom to the student to explore the issue that interested them; 2) incorporating the human element in writing to differentiate student writing from AI-textual output; and 3) making oral presentations and conversations equally important parts of the assessment. For instance, we highlighted how we wanted their reports to be professional and personal

simultaneously so the readers could see the technicalities and arguments as well as the writer's voice and experience.

## **Findings and Discussion**

Cumulatively, we considered the project a success. Students remained excited about the new technology and experimented with the tool, while critically examining some of the issues generated by the text-generating technology.

The project provided multiple occasions to think about the place of writing in academia. For instance, the class came to appreciate how writing in college was about the process of learning, organizing, and exploring the issues in depth more than about the product of the "college essay" itself. Students discussed how the declaration of the death of the "college essay" (Marche, 2022) in the advent of the text-generating tool was a premature and possibly misunderstood phenomenon and how writing as a thinking and learning tool was still required.

The assignment also allowed the instructor to discuss how writing helped students cultivate the habits of mind that are expected in a critical thinker and advanced knowledge worker. The instructor could emphasize how writing was not merely a way to demonstrate learning but a tool of training that students must venture through if they are to benefit from the college education.

In that light, the class also recognized the importance of a student's discipline and ethical considerations to train oneself through writing. While students underscored the role of tools and collaboration in the writing process and the increasing possibility of posthuman collaborative authorship between AI-bots and humans — not just for the production of a text but also for the writing done for learning—they emphasized the individuality involved in writing.

While all of our students enrolled in the class majored in humanities and social sciences and did not have a background in computer science, the assignment taught them major aspects of AI. The presentations and class discussion following the students' experiments with the chatbot made them realize the size and scope of Large Language Models (LLMs), machine learning, and the role of training and the data fed into the system for the bots to work. Students particularly appreciated how a chatbot's predictive model depended on syntactic possibility without its awareness of the semantics of its output. Students highlighted the importance of extra caution in the programming process to ensure social justice given that the dataset fed into the LLM came from sources that embodied biases and harmful ideas.

The assignment also led the class to ponder the questions of ethics, humanity, and human obsolescence. Students were divided about their positions: some students argued that although algorithms can produce grammatically flawless prose and the human mind may not match AI's memory and computational prowess, the agency ultimately rested on humans because they created AI-bots and they can direct its trajectories and decide on its fate. Some students however pointed to the possibility of a self-operating AI and the unpredictability of the impact of its deployment on the human world. They pointed that the moments of AI hallucinations might move it toward independent visions and commanding the operations of the networked world on which humans are increasingly basing their existence.

## **Conclusion**

When measured by the students' performance, the assignment largely achieved its learning goals. Students experimented with the latest writing tool, researched about it, and examined its implications. For instructors, the assignment reconfirmed the critical role writing plays in colleges and its continued importance even in the context of text-generating AI chatbots.

## Works Cited

Carrol, Brian. *Writing and Editing for Digital Media*. Routledge, 2020.

Elstermann, Annika. "Computer-Generated Text as a Posthuman Mode of Literature Production." *Open Library Humanities* 6(2), 2020. (<https://olh.openlibhums.org/article/id/4663/>)

Marche, S. (2022, December 6). The College Essay Is Dead. The Atlantic. Retrieved from <https://www.theatlantic.com/technology/archive/2022/12/chatgpt-ai-writing-college-student-essays/672371/>.

Miller, Matt. "ChatGPT, Chatbots and Artificial Intelligence in Education." DITCH THAT TEXTBOOK. Accessed on February 5<sup>th</sup> 2023. (<https://ditchthattextbook.com/ai/>)

Roose, Kevin. "A Conversation with Bing's Chatbot Left Me Deeply Unsettled." The New York Times. February 16, 2023. (<https://www.nytimes.com/2023/02/16/technology/bing-chatbot-microsoft-chatgpt.html>)

---

## The Assignment

### Assignment Description

For this assignment, students will use OpenAI's new chatbot ChatGPT (<https://chat.openai.com/chat>) or similar technology and hold two conversations on a topic that is familiar to them, analyze the writing and thinking patterns in the output produced by the technology, and think about the implications of text-generating technology for writing both as a task and an educational tool. Student deliverables include a written report, participation in class discussion, and oral presentation; each of the three deliverables will receive equal weight for assessment.

There are four stages to this assignment: Stage 1. Students first chat with the bot and analyze the writing response produced by the bot from two perspectives: first, the writing and argumentation pattern displayed by the chatbot and the source of that pattern, and second, the place of voice, creativity, and agency in the response. Since the class will discuss these issues, students are required to think of these aspects as they chat. Stage 2. Students then select one idea (not necessarily from the list stated above) they find important in their analysis of their chat, and research it further. Students will study at least 5 newspaper and academic articles on the impact of ChatGPT and/or similar technology on the idea they selected to investigate. Stage 3. Students draft a report, participate in class discussion, and Stage 4. students make an oral presentation on their findings.

Stage 1: Chatting with the Bot and Analyzing its Responses

When chatting with the bot, please do the following:

1. Chat on a topic that is familiar to you. The objective here is not to learn new ideas on the topic but to understand the “mind” of the bot. You will be able to analyze the bot’s response better if you are knowledgeable about the topic.
2. Try to localize and contextualize your input to the bot as practically as possible to prod the bot for specific response.
3. Save the response

Once you completed the chat, analyze the response using following ideas:

1. What is the thinking and writing pattern of the response?
2. Where might that pattern come from? How did the machine learn to do what it did? What thinking and writing principles does it employ?
3. Are the information and ideas accurate? Ethically sound? Unbiased? Specific?
4. What about the tone, style, and voice? How would your response be different from that of the machine?

Stage 2: Based on your analysis, select one issue about ChatGPT that you found intriguing/interesting/important and read/view at least five reliable sources that focus on that aspect of this technology. Don’t be generic; go for the sources that discuss the issue you identified. Properly take notes from those sources as you explore them.

Stage 3. Write a Report (~ 4 pages, 12 point font, single space)—Your report should have an (1) introduction, (2) discussion of your findings of the analysis of your conversation with the bot (use subheadings to organize the ideas.), (3) introduction to the issue you decided to focus on and why, (4) discussion of what you found in your research of the issue, (5) conclusion, and (6) a works cited section. While your report should demonstrate an understanding of the genre convention of report writing, try to personalize it by describing facts as well as the thoughts and feelings you experienced as you went through the chat, analysis, research and drafting process. Please don’t use AI to assist in writing your report. (For instructors who want to integrate AI to assist students to write the report, they can emphasize the AI-human collective authorship as well, but for us, at the point when this assignment was designed for the first time, we still were exploring the new tool and didn’t want students to hand over their writing tasks to AI. Since text-generating AI itself is becoming sophisticated, the future adoption of this assignment may develop towards more use of the AI than what has been done in this iteration.)

Stage 4: Make a 10-minute oral presentation allowing sufficient time to answer questions following your presentation. As you present, think of the non-verbal aspects of communication and use them as much as possible; the objective is to emphasize how human communication goes deeper than the combination of words represented by the LLM systems on which AI writing depends (at least for now).

### **Suggested Pre-reading Sources**

<https://www.nytimes.com/2022/12/05/technology/chatgpt-ai-twitter.html>

<https://www.science.org/doi/10.1126/science.adg7879>

<https://www.theatlantic.com/technology/archive/2022/12/chatgpt-ai-writing-college-student-essays/672371/>

<https://themarkup.org/hello-world/2023/01/28/decoding-the-hype-about-ai>

<https://wac.colostate.edu/repository/collections/ai-text-generators-and-teaching-writing-starting-points-for-inquiry/>

<https://ditchthattextbook.com/ai/>

# Synthetic Metacognition

## Iterating Prompts with GPTs

Kyle Booten  
University of Connecticut, Storrs

*This assignment suggests that “prompt engineering”—iteratively tinkering with and refining the set of instructions that guides the output of an LLM—is a worthwhile writing activity that can encourage students to be metacognitive about the “moves” that characterize compelling examples of a genre in which they are writing. Insofar as LLMs are “lazy” (obeying the prompt but not exceeding it), coaching one to successfully compose in a genre can require students to make explicit aspects of the genre that they may only implicitly be aware of. In a classroom setting, collaboratively “workshopping” the results of the GPT affords an opportunity to notice, describe, and name some of these otherwise-implicit moves, and students can consider integrating them into their own writing.*

---

### *Learning Goals:*

- To prepare students to compose in a particular genre (e.g. a paper, a poem, a written assignment) by “workshopping” a GPT in order to build metacognitive awareness of certain textual “moves” (Swales) that characterize successful examples of that genre

*Original Assignment Context:* First-year writing course

*Materials Needed:* Any AI text generating program, selected readings

*Time Frame:* ~1-2 weeks

---

### **Introduction**

The rise of text-generating large language GPT models (such as GPT-3 and the more famous ChatGPT), image-generators such as Dall-E, and other sophisticated deep learning models has inaugurated a new paradigm in “programming”: the user interacts with these systems primarily through crafting *prompts*. Sometimes a sentence or two or even a mere phrase—“sonnet about styrofoam”—will suffice. But getting an impressive result from one of these models often requires the user to iteratively refine and augment a prompt, specifying all manner of traits the output should have as well as those that it should not (“Petrarchan sonnet about styrofoam, style of Wyatt, carefully observe iambic pentameter, include multiple enjambments, do not include rhymes of one-syllable words or overtly-emotional words like ‘sad’...”).

In a recent first-year writing class for university students, I used the prompt-refining process as an opportunity to engage students in thinking about what features—or, in the vocabulary of Swalesian discourse analysis, which we had frequently discussed, what “moves”—produce a potent and compelling example of a particular genre. For Swales, a “move” is a way that a genre has adapted to a recurring rhetorical situation and its attendant burdens and pressures; for instance, examples of introductions to

scientific papers often indicate some missing piece in current knowledge, a recurring textual pattern that reflects a need to prove to one's fellow scientists that one more paper need be written and published. For their final project, students were to create a dialogue (a semi-scripted podcast or written interview) about the dangers of AI. They had already brought in examples of podcasts that they admired, and we had together tried to notice recurring moves and speculate about the rhetorical pressures these moves might reflect, especially the mere pressure to be entertaining.

That day, with OpenAI's GPT-3 "playground" on the classroom's screen, I furnished an initial prompt that echoed the one I'd provided for their final project assignment: "write a dialogue about the dangers of AI." Then, in small groups, students critiqued the AI's initial output, noticing where it was boring, uninformative, or otherwise failed to entertain us, and proposing solutions to these faults. Back in the large group, we refined the prompt and generated more text. And then we did the same thing for two or three more rounds. To get a better result from the AI, the students first had to notice what makes a dialogue successful or unsuccessful and then translate these observations into clear, specific instructions—in other words, to practice writerly metacognition ("meta-," that is, in relation to the AI's first-order cognition).

This activity takes advantage of several of the key affordances (as well as limitations) of GPTs. First, and worth mentioning despite its obviousness, *GPTs generate text quickly*. In a traditional writing class, instructors may strain to fit even one round of revision into the cluttered confines of the academic term. By the time a paper (or poem, or some other text) reappears on the docket, having been revised, the instructor and other workshop participants may well have forgotten what suggestions they made weeks ago. "Workshopping" a GPT means that it is possible for human participants to keep in working memory their impressions of Version 1.0 of its text along with its subsequent versions. That revision in a traditional writing class is costly (in terms of time) encourages instructors to offer a "bundle" of comments at once; since a GPT's rewrites are temporally cheap, one can freely iterate, changing even a short phrase or a word in a prompt to see how it affects the output.

Second, *GPTs are more obedient than they are brilliant*. GPTs are designed to produce plausible text in response to the user's instructions. They aim to satisfy, not impress, so they frequently fail to do very basic and obvious things to make their text more convincing, engaging, or charming. Sometimes it feels as if a GPT is engaged in "malicious compliance" —carrying out its instructions but, since the instructions didn't mention them, forgoing certain "moves" that are so ubiquitous in human-generated examples of a genre that they tend to escape our noticing. This provides an opportunity to notice them and then translate our dissatisfaction into a more specific (perhaps a *painfully* specific) set of instructions. The GPT's seeming laziness can energize our metacognition.

For instance, students in my class noted that the script that GPT generated was originally quite short, containing only a few conversational turns before terminating. This simply did not feel satisfying—but, then, how many turns would feel satisfying but not overlong? (My notes suggest that we ended up asking GPT for seven conversational turns between the imagined interlocutors, "Person 1" and "Person 2"). They noticed a pair of related vices: GPT's text seemed to flit anxiously from point to point, and these points were expressed in entirely abstract terms, with no specific, real-world details given to support them. Together we revised the prompt to command the GPT to slow down and discuss its points more thoroughly and to provide examples and details in support of them. It also came up in our class conversation that the Person 1 and Person 2 fabricated by GPT-3 were simply too similar and too quick to agree; a compelling conversation, of course, should contain

some tension, perhaps even directly conflicting views, and so we reworked the prompt to make Person 1 an optimist about AI and Person 2 someone who lost their job to AI-powered automation. As class began to wind down, we together experimented with different ways of prompting GPT-3 to create something like a conclusion (for example, with Person 1 being convinced—but not totally convinced—by Person 2).

This “GPT workshop,” while in some sense a one-off activity, also fit into a larger instructional arc. In the previous weeks, students had prepared for their final project by bringing in examples of podcasts and written dialogues and interviews; together we analyzed them and compiled a list of “moves” that would make a successful, entertaining conversation, whether written or spoken, and the final assignment specified that they had to attempt some number of these moves. The day of this GPT activity, as I began to hear the tell-tale zippering of backpacks, I made the case to my students that the “moves” that they had instructed the AI to perform would also be helpful for them to keep in mind while composing their final projects, and I added some of these moves to the list of ones from which their submissions would need to draw. My use of AI in this class was rather impromptu and took place at the tail end of the semester; were I to do it again, I would make sure to leave more time, perhaps another class session, to discuss and even practice these moves.

Because of the current limitations of GPTs, this assignment will work better for some writing assignments than others. It has been widely observed that OpenAI’s GPT-3 and ChatGPT frequently “lie”—for instance, making up historical facts, scholarly citations, or plot points in novels. While future AI models may be less prone to prevarication, current GPTs will often struggle to compose a minimally-plausible example of any genre that is highly dependent upon specific information that is too vast and varied to include as additional information in the prompt itself. (They will have an easier time writing about a general historical trend than a specific, poorly-known historical event.) On the other hand, this assignment hinges upon the fact that the LLM is obedient, but only just. As OpenAI and other companies continue to develop these language models, it may be the case that some of them become *too* impressive; they may produce not just passable but impressive, witty, and charming text with little further prompting. Instructors who would want to use LLMs for this or a similar assignment should experiment with different language models, including those that will have been made obsolete, as the “best” model may not be the best for their purposes.

---

## The Assignment

### Goal

To prepare students to compose in a particular genre (e.g. a paper, a poem, a written assignment) by “workshopping” a GPT in order to build metacognitive awareness of certain textual “moves” (Swales) that characterize successful examples of that genre.

### Materials

- Access to a large language model (LLM) that generates text based on a user’s prompts; current examples include OpenAI’s GPT-3 and ChatGPT models. (In cases where the LLM is only available as a paid service, a dollar or two of credits should be sufficient, at least at current rates.)
- A projector, so that the entire class can see the prompt and the text the language model generates in response to it.



## Steps

1. Provide the LLM with a prompt that students have already been given for a particular writing assignment (e.g., “Write a sonnet about X...” or “Compose a podcast that addresses the following set of questions...”).
2. In small groups, discuss the strengths and weaknesses of the LLM’s text. This conversation should focus not just on plausibility (does the LLM’s text follow the instructions as given, did it stick more or less within the genre) but also quality (does it possess features that are charming or compelling, that would make you want to read more). (5 min.)
3. Continue the conversation as a class, highlighting any weaknesses that were noticed by multiple groups. Throughout these steps, the instructor should be sure to guide the conversation back to the concept of the “move,” which describes textual regularities in terms of regularities of rhetorical circumstances: by originally omitting the features described, what did the LLM’s text not take into account about the genre, purposes, and the desires of its typical reader or audience? (5-10 minutes)
4. Focusing on a few of the most salient weaknesses of the LLM’s text, iteratively tweak the prompt in an attempt to remedy them. Employ a principle of parsimony; try to add as few words as possible at a time to produce a (positive) change. (10 min.)
5. Repeat steps 2 through 4 once or twice more.
6. By the end of the above steps, the original prompt should have been revised to make explicit certain aspects that make a text a not just plausible but compelling example of the given genre. The assignment concludes with a discussion of how students might keep in mind these same features. (“What’s good for the goose...”)

## Contextualizing This Assignment

This assignment is meant to fit into a larger instructional arc that draws students’ attention to “moves” that characterize a successful example of a particular genre, and so the instructor should introduce the notion of a “move” in prior class sessions. This particular assignment may also be used to complement other, more traditional methods of move analysis, such as aggregating and observing recurring patterns in this example; workshopping the LLM’s text is meant to draw attention both to features that are so obvious in “real” examples that they escape noticing or those whose lack makes a text, if still a technically valid example of the genre, a boring or otherwise unsatisfying one.

The instructor could also consider updating the original writing assignment to encourage students to heed certain instructions that they gave the LLM.

## Works Cited

Swales, John M. *Genre Analysis: English in Academic and Research Settings*. Cambridge UP, 1990.

# Using LLMs as Peer Reviewers for Revising Essays

Antonio Byrd  
University of Missouri-Kansas City

*In this assignment, undergraduates use large language models (LLMs) to assist in revising their essay drafts by asking LLMs to respond to common peer review prompts. Students learn prompt engineering and develop rhetorical judgments on the effectiveness of LLMs' language analysis to heighten their revision processes. This assignment can be adapted to most disciplines and course levels.*

---

## *Learning Goals:*

- Explain the ethical implications of using LLMs to generate content for writers
- Create effective peer review prompts in LLMs to revise essays
- Implement the most useful revision strategies from human and machine readers
- Evaluate how LLMs support individual writing processes
- Construct ethical frameworks for when and how to use LLMs in the writing process

## *Materials Needed:*

- Large language model with linguistic analysis feature such as GPT-3.5 or GPT-4
- Essay from previous assignment

*Original Assignment Context:* end of final unit in advanced expository course on literacy studies and technology

*Timeframe:* ~2 weeks

---

## **Introduction**

The popularity of large language models (LLMs) reinvigorates suspicions of original authorship and concerns students may not achieve the outcomes of writing classrooms: developing rhetorical knowledge, critical thinking, reading, and composing, and writing processes. The unit for “Using LLMs as Peer Reviewers for Revising Essays” has students explore how LLMs reshape our conceptual understandings of plagiarism, copyright law, and remix culture through hands-on work. The assignment itself then directs students to think about how using LLMs evolves their revision processes as text generation technologies become evermore included in our repertoire of literacy practices. By the end of the project, students will have problematized LLMs as tools for generating content and will have reframed LLMs as potential assistants that can be ethically integrated into their writing process. Considering the flaws of LLMs, such as their hallucination (describing untrue facts in a tone of certainty) and using White Standard English as its base variety of English, students will also have learned to judge LLMs’s as useful peer reviewers.

The assignment below is a revision of a different version of this assignment (more on this version described below) I first developed in fall 2021 for a 16-week online

asynchronous course called English 305WI: Theory of Composition for junior and senior undergraduates. Theory of Composition explores the nature of writing through a literacy studies perspective. Across three expository writing projects, students explore two ideas: first, writing is a type of knowledge and practice that shapes how we perceive and interact with others. Second, writing changes as technology evolves, and that presents new challenges to how we write and live. Issues in writing and technology prompt our discussing the implications for ethics and social justice.

Instructors may assign a conventional essay project that explores concepts and ideas most germane to the unit. Following typical writing processes, students construct a rough draft for peer review from classmates and the instructor. In this structured peer review, the student writer poses a variety of questions and concerns about their draft that their classmates and instructor then direct their feedback toward. Students then discuss issues related to originality, writing, and technology after the completed peer review. These discussions are the focus for another unit in my class and we spend three weeks before using LLMs to revise the essay. Other instructors need not devote so much time to this topic, especially if they have other concepts and ideas they need to teach in their course and they are more interested in using LLMs as peer reviewers. Instead of reading and discussing scholarly research on copyright law, privacy concerns about TurnItIn, and remix culture for two weeks, instructors may focus on multiple case studies about the ethical considerations of using text generated technologies. They cover a range of topics including, reader manipulation, racism, authorial intent, assistance with processing grief, and how professional writers in marketing and journalism use technologies like Jarvis.AI to make them more efficient writers. The case studies prime students for thinking about how LLMs have multiple problems when they are used to create content for writers and no other purpose. The goal is not to impose instructors' views on using LLMs but rather to show students discourse about these technologies from different perspectives. Discussion posts ask students to conduct critical analysis of these technologies and make an informed decision on how they think these technologies may or may not fit in their range of literacy practices.

What follows then is a new peer review session, this time with the language analysis of LLMs. Students revise their essay based on comments from both humans and machines; after completing the assignment, they may reflect on how well humans and machines assisted in their revision process.

Finally, the original version of this revised assignment had students use GPT-3's text completion application. At the time in 2021, writers could prompt GPT-3 with a sentence and then GPT-3 would generate several sentences connected to that original prompt. During revision, students input one to three sentences from their essay into GPT-3's text completion application as if they were writers simply stumped on how to proceed with their writing and could use a little assistance. This text generator continued the conversation related to literacy experiences. Students could revise, edit, and/or delete this AI-generated text and add it to their second draft. They could also use the AI-generated text as inspiration to write something else to further their revision. I directed students to distribute GPT-3's influence throughout 20% of the essay (300 words out of 1500-words) and bold the language that comes directly from GPT-3 or that GPT-3 had inspired them to write, as some form of citation. Students explain how GPT-3 shapes their revision process in a Statement on Goals and Choices.

The inspiration for this approach originally came from Vauhini Vara's "Ghosts" (linked below) published in *The Believer*. In this article Vara uses GPT-3 to generate nine essays on processing the

grief of her sister's passing to cancer. She bolds texts written by herself and leaves the words GPT-3 generated to continue the narrative unbold. The ninth essay is completely bold, suggesting through multiple drafts Vara finds the words to narrate her grief. The word limit I imposed on students meant to balance the responsibilities of the human writer with the assistance of GPT-3, giving students more agency over technology. However, that balance favored the human to process and extend the few words GPT-3 created into a new text.

Of course, LLMs since 2021 have become more advanced than mere text completion. The assignment description below takes into account these technological developments. My discussion in the following section reflects on what happened when I used the original 2021 assignment and then includes my rationale for the proposed revised assignment that instructors may adapt for their own teaching.

### **Discussion and Future Teaching**

I focus on revision for two reasons. First, in the words of Ernest Hemingway, "all writing is rewriting." Revision is a commonplace in teaching writing and the practice of writing. Second, students using text generation technology to create language for them concerns writing instructors. The original assignment explored how GPT-3 aids in revising a human-produced draft. But my use for this assignment was to extend students' theories of writing. I frame for students the potential benefit of GPT-3's predictive text generation in two ways. Not only does it make "recommendations" to the writer on how to proceed with revision, GPT-3's generated sentences may attempt to make philosophical statements about the nature of writing, which would offer students new ways to think about their literacy archives and writing itself.

Students understood the benefits of using text generation technologies for writing, while being surprised that standards for plagiarism differ in public discourse. However, they ultimately pushed back against its use, holding to conventional perspectives on authorship, originality, and citation they learned from schooling. Deploying GPT-3 had one flaw, not with the technology itself but with the parameters placed on the human writer. Students needed substantive feedback on their writing to produce significant revision. Their revisions with the assistance of GPT-3's text generation seemed cosmetic or minor, especially for students who did well on their first draft. Thus, the 300 word limit I imposed on students underutilized the strengths of GPT-3.

An assignment revision would leverage the full capabilities of LLMs like ChatGPT. Students write full drafts in Unit 1, and then use a combination of human rhetorical processes and the AI's sophisticated language analysis to peer review their work. For example, students copy and paste portions of their draft – especially the ones they find most troubling or even passages they find particularly effective – into the LLM. Students would then experiment with prompt commands to analyze the text for tone, syntax, word choice, and other linguistic features that impact the rhetorical meaning of the draft. They may also command LLMs to clean up the language for clarity. The writer still bears responsibility for the text generated; they would make judgments on the analysis from the LLM, focusing on the accuracy of information and how well the model matches students' authorial intent or aligns with their values. Thus learning goals expand the conventional outcomes of writing instruction: learn effective prompt engineering and rhetorical judgment of LLMs to heighten revision processes. Thus, instructors would need time to learn with students how to write effective prompt commands.

This unit may include additional conversations on text generation technologies' influences on the social, environmental, political, and financial spheres of life. For example, discussion and analysis on how the commercialization and monetization of these technologies necessitates the labor of global marginalized communities, how paid subscriptions to access LLMs continue legacies of digital inequality, how LLMs' powerful hardware contributes to climate change, and how these technologies handle data privacy. These concerns shift the conversation from the ethics of using LLM for writing to the *morality* of using LLM AI for writing.

---

## The Assignment

### Task

Revise your essay using a large language model (LLM) such as ChatGPT. First, revise your essay in response to peer review comments from myself and your classmates.

After you have completed revisions, write a variety of peer review prompts that prompt the LLM to analyze or extract information from your essay's paragraph. For example, you may ask the same questions you posed to your human readers during the first round of peer review to the LLM. Other prompts may include more general questions related to your thesis, introduction, organization of body paragraphs, evidence and analysis, citation, and conclusion. Here are sample questions, some from the University of Colorado-Denver's Writing Center, others my own construction.

- Does the introduction provide enough context on the paper's topic?
- How can the thesis be more specific and complex?
- Is every piece of evidence followed by analysis in the following paragraph?
- How do the ideas in the paper progress?
- How can the conclusion restate the thesis in a more complex way?
- Describe the tone of the following paragraph.
- Compare the tone of the first paragraph with the third paragraph.

Cast careful judgment on the responses from the LLM, as the analysis may include misinformation or show that the LLM did not understand the intent of your prompt command. Revise and edit your essay based on the analysis you receive. You may include the text generated by the LLM in your essay but you must use proper citation style.

Include the chat history with the LLM with your revised essay. Finally, in a Statement of Goals and Choices (SOGC), reflect on how your interaction with the LLM shaped your revision process. How did the effectiveness of peer review from your human peers and instructor compare?

### Purpose

The purpose this time has expanded. First, the original purpose from Unit 1 still stands: [Explain the original purpose of the essay assignment]. Your second purpose is to experiment with how writing technology more sophisticated than autocorrect and auto completion can help you write more efficiently and clearly according to your stated audience, purpose, and goals. To meet these ends, this assignment has you start developing your competency in prompt engineering or prompt design with LLMs.

## Audience

[Audiences may include scholars, instructors, administrators, peers, or other readers essential to the original essay assignment]. You may have other readers in mind to help you solidify your interest in this writing project, such as family members, friends, other teachers, or maybe people you know in the community that helps them understand your argument and lived experiences and why they matter.

## Genre and Format

For your essay, Times New Roman, size 12 font, double-spaced, MLA or APA format, if you include outside sources. Your essay can be 1,200 - 1500 words. This word count is a guideline or meant to give you an idea of what the scope of the project may be. Do not beat yourself up for not writing to the minimum. Quality over quantity! :)

## When do I submit my rough draft and do peer review?

Submit revised draft with chat history and SOGC by [Turn in Date and Time]

## Assigned Texts

- Hutson, M. (2021, March 3). Robo-writers: the rise and risks of language-generating AI. *Nature*. <https://www.nature.com/articles/d41586-021-00530-0>
- Kirby Ferguson. (2021, September 7) Everything is a remix: Part 1. [Video] YouTube. [https://www.youtube.com/watch?v=MZ2GuvUWaP8&ab\\_channel=KirbyFerguson](https://www.youtube.com/watch?v=MZ2GuvUWaP8&ab_channel=KirbyFerguson)
- Knight, W. (2019 , February 14). An AI that writes convincing prose risks mass-producing fake news. *MIT Technology Review*. <https://www.technologyreview.com/2019/02/14/137426/an-ai-tool-auto-generates-fake-news-bogus-tweets-and-plenty-of-gibberish/>
- Teardgarden, A. (2019). Stories of plagiarism / theories of writing: How public cases of plagiarism reveal circulating theories of writing. *Kairos: Rhetoric, Technology, and Pedagogy*. 24(1) <https://kairos.technorhetoric.net/24.1/topoi/teagarden/index.html>
- Morris, S.M. and Stommel, J. (2017, June 15). A guide for resisting edtech: The case against Turnitin. *Hybrid Pedagogy*. <https://hybridpedagogy.org/resisting-edtech/>
- Puiu, T. (2020, October 14). The stunning GPT-3 AI is a better writer than most humans. *ZME Science*. <https://www.zmescience.com/science/gpt-3-better-than-you-043252/>
- Schwarz, O. (2019, November 2019). In 2016, Microsoft's racist chatbot revealed the Dangers of online conversation. *IEEE Spectrum*. <https://spectrum.ieee.org/in-2016-microsofts-racist-chatbot-revealed-the-dangers-of-online-conversation>
- Vauhini, V. (2021, August 9). Ghosts. *The Believer*. <https://www.thebeliever.net/ghosts/>

## Supplemental Text

- Calamity AI (2020, November 12). A.I. written essay | peer-reviewed. [Video]. YouTube. [https://www.youtube.com/watch?v=Lto6exrpChQ&ab\\_channel=CalamityAI](https://www.youtube.com/watch?v=Lto6exrpChQ&ab_channel=CalamityAI)

## Instructional Videos

- Byrd, A. (2021 November 7). *Unit 3 intro and week 12 module* [Lecture Recording]. Canvas. <https://umssystem.hosted.panopto.com/Panopto/Pages/Viewer.aspx?id=c97e1631-fda0-425d-9e71-afd201044336>
- Byrd, A. (2021, November 14). *Why do we use automated writing?* [Lecture Recording]. Canvas. <https://umssystem.hosted.panopto.com/Panopto/Pages/Viewer.aspx?id=fc1ccb57-0261-4eec-962c-afd2011994a9>

## Acknowledgements

My thanks to Tim Laquintano for his thoughtful recommended revisions, and to the rest of the editorial team, Annette Vee and Carly Schnitzler. I also appreciate the internal review from John Silvestro. The Atlantic’s Object Lessons series inspired the 2019 version of this assignment before I used GPT-3 in 2021. I’m grateful for the writing instructor who published his Literacy Archive Essay assignment on Wordpress. While I can no longer find the URL, readers can refer to Horror Tree’s “Literary Artifacts: What Are These and How to Use Them in Your Essays” (2022) as a resource. “Student Essay AI Co-Writing Public Demonstration” by Tristan Hanson partly inspired how I would redesign this revision assignment. The article details an informal experiment on AI co-writing run by S. Scott Graham, Casey Boyle, Hannah R. Hopkins, Ian Ferris, Tristan Hanson, Maclain Scott, Emma Allen, Lisa Winningham, and Walker Kohler.

# Genre Generators

Addison Eldin  
University of Pittsburgh

*This activity asks students to analyze and create \*genre generators\*, which are text generators that use the computer to create new instances of a formal written genre. After the instructor introduces the concept through a collaborative analysis of implementations of Christopher Strachey's \*Love Letter Generator\*, students then work in groups to produce non-executable conceptual programs that would produce new instances of genres they have chosen. The instructor then enacts a selection of the programs in code and leads another collaborative analysis exercise in order to highlight the rhetorical and social dimensions of the decisions that students made in the process of automation via a digital computer. By emphasizing the programming process and not asking students to code, students can develop a better understanding of how computer programs mediate digital technology and human understanding of the world.*

---

## *Learning Goals:*

- Be able to distinguish the acts of programming and coding while understanding their interrelations
- Connect extant rhetorical analysis abilities to the analysis of code and its output(s)
- Engage in programming such that both human decision-making and technical limitations are considered

*Original Assignment Context:* Intermediate-level digital humanities course

*Materials Needed:* a higher level programming language that allows for the manipulation of text strings with lists and variables (e.g. Python, which is used in `generator-template.ipynb``), the ability to present code and text in sequence (e.g. Jupyter Notebooks, which can present Python and Markdown in sequence, as shown in `generator-template.ipynb``)

*Time Frame:* ~2 sessions

---

## **Introduction**

\*Genre generators\* are text generators that use computer code to produce and manipulate text and create new instances of a formal written genre.<sup>1</sup> This in-class activity, which walks students through the process of creating their own genre generators, is intended not to teach computer code—which requires knowledge of how to use specific programming languages—but rather \*programming\* as a conceptual process enacted by people to mediate between the world, their understanding, and digital technology. Where \*coding\* requires adherence to the rules of one or more programming languages, \*programming\* is the conceptual process whereby the problem is reworked and solved more generally. Think of it this way: to program is to know you might need a list of words that is pulled from at random in the generation of a new sentence, while to code is to know how to define and choose objects from a list using a specific programming language.



This activity was created for a course called "Digital Humanity," which is an undergraduate course that covers several topics around being human and the meaning of \*humanity\* in a time of digital ubiquity. It focuses on technological and social elements of computers in tandem to frame the various topics in the course. As described in the next section, this context directly led to the creation and shape of this activity.

While certain parts of digital technology, programming, and code will need to be established by the instructor in class before the activity, students themselves will not need to produce any code that can be run on a computer, as the emphasis is on programming. This distinction can help students understand that the computer is at once technical—that is, a material object with certain ways of functioning that create both limitations and affordances, as is the case with technology generally—and social—that is, made use of by human beings for human-defined goals based on their understanding of the world around them.

The activity, which runs for at least part of two sequential in-class meetings, leverages students' experience with and understanding of \*formal written genres\* to collectively analyze implementations of a specific computer program, Christopher Strachey's \*Love Letter Generator\*.<sup>2</sup> Designed in the early 1950s, Strachey's generator stands out in the history of computers for its use of computational resources and programming for the purpose of text generation. Strachey worked from a handbook compiled by Alan Turing for the Manchester Mark 1, one of the earliest stored-program computers, and created a lengthy program that put the computer's resources toward something distinctly human—the writing of love letters.<sup>3</sup>

Once the relevant aspects of formal written genres are recalled and applied,<sup>4</sup> students make use of this knowledge to understand how the implementations of the Love Letter Generator enact our understanding of the love letter as a genre through looking at the code underlying them, with an emphasis on broad concepts used in programming such as string generation and manipulation as well as the use of lists and variables. Then, in small groups, they define a different written genre (e.g. an email to their professor, an obituary, or a job ad), identify its purpose and key features, and consider which elements of the genre can be represented with digital technology via programming. Finally, the instructor enacts a selection of the student programs in code, turning the analysis toward the students' genre generators to draw out the way the students' practices were, like programming more broadly, a back and forth between the computer's technological basis and the non-technological understanding of genres that they represented with that technology.

## **Context**

The original context for this activity was "Digital Humanity," a course that intertwines the technological and social elements of the computer to frame conversations about being human and the meaning of \*humanity\* in a time of digital ubiquity. It serves as a philosophy general education requirement at the University of Pittsburgh for majors in fields across the humanities, social sciences, and natural sciences, but also acts as a major requirement for students in Digital Narrative and Interactive Design, an undergraduate major which incorporates courses from the English department and the School of Computing and Information. While its audience includes and benefits students who code, there is no presumption of technological facility beyond that required by the university more broadly: the ability

to use a web browser, email, learning management systems like Canvas, word processing technology, and video conferencing software like Zoom.

The course's focus on the technological and social elements of computers in tandem guided the design of the syllabus the first time I co-taught it with one of the original developers of the course,<sup>5</sup> as well as the second time when I taught it alone. I wrote this activity for this second iteration after reviewing a shared teaching document: one week, Prof. Langmead used the Love Letter Generator as the basis of a conversation, and the program seemed perfect for helping students grasp how \*programming\* involves a person who makes interpretive decisions about how to represent the world.

## Implementation

Despite the idiosyncratic context from which this activity arose, I think there are three broad criteria that can help someone identify if this activity is appropriate for their course.

1. The course should, in part or in whole, speak to the social and creative dimensions of programming and coding for digital technology. That is, some part of the course should speak to how ideas such as digital representation, digital computation, and/or encoding are relevant processes in the human use of the technology.<sup>6</sup> Relatedly, the course should spend some time prior to the activity discussing these concepts—whether in those terms or through a focus on programming and code—to prepare students for the conversation about the Strachey implementations that will prompt students' own creative work in developing a generator.<sup>7</sup>
2. Students should, generally, have some experience with college-level writing. The instructor will guide students through the basics of rhetorical analysis by considering the rhetorical context, purpose, and norms of a written genre, namely the love letter, as a review of basic concepts often covered in first-year writing courses. Students who have taken a first-year writing course or otherwise learned about different kinds of academic writing—e.g. what differentiates an essay and a research paper—will be better situated to engage in this conversation. A class made up of students who are new to college and have no prior college writing experience might find this conversation more difficult, which could make the rest of the activity difficult.
3. While students do not need to know how to code or what programming is in advance, the instructor must have basic facility with at least one higher level programming language that enables the manipulation of text, e.g. Python or Javascript. The basic skills required include coding strings, variables, and lists in conjunction with each other. Further, the instructor will need to be able to present the code and its outputs in sequence with text that contextualizes the code with some kind of projection technology.

If these three criteria are met, I believe that the activity could be implemented into a course regardless of its home department. While the rhetorical analysis part of the assignment might seem limiting in terms of instructor comfort with guiding the discussion, I think that the questions provided in `instructional-guide.md` and the instructor's own awareness of the contexts, purposes, and norms of written genres in their own field are sufficient for facilitating the rhetorical analysis in the activity.

## Instructional Materials

All instructional materials for this activity—including an instructional guide, a student program and its implementation in code, and a template generator in Markdown and Python using Jupyter Notebooks—are available via a GitHub repository (<https://github.com/addeldin/genre-generators>).

There, you will find:

- `README.md`, which guides you through the contents of the repository;
- `generator-template.ipynb`, a template that can be updated with specifics from a student program;
- `instructional-guide.md`, which guides you through the activity;
- `student-program.pdf` and `student-generator.ipynb`, which are examples of a student program

## Notes

1. *Formal written genre* is meant to contrast *genre* as used in the phrase *genre literature*, e.g. horror, fantasy, or science fiction. Formal written genres describe the different formal genres of writing such as the letter, novel, government report, etc.
2. The two implementations of the Love Letter Generator used are Matt Sephton's (<https://www.gingerbeardman.com/loveletter/>), which presents outputs of the generator, and Annette Vee's (<https://jsfiddle.net/nettework/6f7abaaa/>), which presents outputs along with the underlying Javascript code.
3. For a brief, accessible history of the Strachey project, see Siobhan Roberts' 2017 article in *The New Yorker*. A link is provided in the References section.
4. For examples of relevant approaches to the study of genre from a composition and rhetoric perspective that account for the rhetorical dimensions of written works in social contexts, see the cited works by Rebecca Nowacek (28), Charles Bazerman (69), and Carolyn Miller (59).
5. Alison Langmead, a faculty member in the History of Art and Architecture Dept as well as in the School of Computing and Information at University of Pittsburgh.
6. By *digital representation*, *digital computation*, and *encoding*, I refer to the fact that to represent a concept using a computer, that concept must be represented digitally, or via discrete numbers, generally the binary digits 0 and 1, which are then computed in various combinations to produce new results. For example, to represent the continuous color spectrum, colors are encoded via digital representation by encoding any color as a combination of 256 possible shades of red, green, and blue. While over 16 million colors is a lot, it is not a direct reflection of the full color spectrum, which due to being continuous necessarily contains an infinite number of possible colors.
7. For resources that can develop an understanding of digital representation, digital computation, digital encoding, and/or the basic concepts of programming for the instructor as well as the students, consider Karen Hao and Jonathan Stray's relatively short breakdown of how a courtroom sentencing algorithm digitally represents ideas of justice and risk; Charles Petzold's book *Code: The Hidden Language of Computer Hardware and Software*, in particular the chapter "Bit by Bit by Bit"; or Paul Ford's in-depth but highly readable *Bloomberg* article, "What Is Code?" All of these texts were used in both terms that I taught "Digital Humanity," with "What Is Code?" being assigned and discussed in the class session prior to the start of this activity.

## References

Bazerman, Charles. "Systems of Genres and the Enactment of Social Intentions." \*Genre and the New Rhetoric\*, edited by Aviva Freedman and Peter Medway, Taylor & Francis, 1994, pp. 67-85.

Ford, Paul. "What Is Code? If You Don't Know, You Need to Read This." \*Bloomberg\*, Bloomberg L.P., 11 June 2015, <https://www.bloomberg.com/graphics/2015-paul-ford-what-is-code/>.

Hao, Karen, and Jonathan Stray. "Can You Make AI Fairer than a Judge? Play Our Courtroom Algorithm Game." \*MIT Technology Review\*, 10 Jan. 2022, <https://www.technologyreview.com/2019/10/17/75285/ai-fairer-than-judge-criminal-risk-assessment-algorithm/>.

Miller, Carolyn. "Genre as Social Action (1984), Revisited 30 Years Later (2014)." \*Letras & Letras\*, 2015, pp. 56-72, doi:10.14393/LL63-v31n3a2015-5.

Nowacek, Rebecca S. \*Agents of Integration: Understanding Transfer as a Rhetorical Act\*. Southern Illinois UP, 2011.

Petzold, Charles. \*Code: The Hidden Language of Computer Hardware and Software\*. Microsoft Press, 1999.

Roberts, Siobhan. "Christopher Strachey's Nineteen-Fifties Love Machine." \*The New Yorker\*, 14 Feb. 2017, <https://www.newyorker.com/tech/annals-of-technology/christopher-stracheys-nineteen-fifties-love-machine>.

---

## The Assignment

### Top-Level

This activity takes place over two sequential in-class sessions. While the activity does not need to take up two full class sessions, the break is necessary so that the programs students design in groups can be implemented in code by the instructor and then shared back to students during the next session.

The first day of the activity should take between 50 minutes and 1 hour and 25 minutes, and the second day should take between 30 and 45 minutes. I have provided estimates for how long each part of the daily activities will take, but these may need to be adjusted based on instructors' sense of which parts students will struggle with.

Please note that this guide is just that: a guide. You are welcome to change it as you see fit for your classroom, and I would in fact love to hear about anyone doing so and how they went about the activity.

This instructional guide is available here as well as in the related GitHub repository (<https://github.com/addeldin/genre-generators>), which further includes samples of student work for this activity and a Jupyter Notebook template for implementing student genre generators. If you wish to edit the instructional guide, I recommend downloading and editing `instructional-guide.md`.

### Session 1

*Setup (estimate: 5-10 minutes)—*

To begin, tell students about the overall arc of the activity, as related in the README.

Then, introduce students to the concept of *\*genre\** from a rhetorical perspective rather than the more colloquial use. That is, talk about written genres students are likely to encounter in class—e.g. a non-fiction personal essay, a research paper, or an argumentative essay—and distinguish this understanding of genre from the broader use that distinguishes by narrative tropes, such as science fiction, fantasy, and westerns. Whereas the former is defined by the exigence for writing based on a specific context (generally, college writing assignments) and toward a specific audience (the instructor or an imagined community relevant to the discipline, such as attendees of an academic conference), the latter tends to focus on consistent narrative tropes and structures.

Prompt students to name written genres that fit this definition. If needed, provide some examples, such as emails to a professor, obituaries, recipe blog posts, news articles, etc.

*Love Letters as a Genre (estimate: 15-20 minutes)—*

The setup conversations lead directly into a rhetorical analysis of the love letter genre using sample outputs from Matt Sephton's implementation of the Strachey Love Letter Generator (<https://www.gingerbeardman.com/loveletter/>), which is useful because it only shows generated outputs. Show students several sample letters, and ask:

- What changes for each new letter?
- What stays the same?
- What do the consistent features (e.g. a personalized salutation using a pet name) tell us about the purpose of a love letter?

Then, ask questions that consider the rhetorical dimension of love letters more broadly:

- What is a love letter?
- Why does it exist?
- When is this genre used?
- What are some key features of love letters?

*Love Letter Generation and Code (estimate: 20-25 minutes)—*

Once the foundations of genre, love letters, and the Love Letter Generator are established, show students Annette Vee's implementation (<https://jsfiddle.net/netnetwork/6f7abaaa/271/>), which is useful because it presents the outputs along with the underlying code. For the purposes of this activity, the HTML and CSS windows can be contracted in favor of the JavaScript and output windows.

Remind students of the consistent features they identified when looking at the sample outputs of the Sephton iteration. Then, facilitate a conversation where students try to connect the word lists listed under ``var`` on line 4 to several outputs of the generator—e.g. how the words in ``salutations1`` and ``salutations2`` are, despite changing each

time, consistently pulled from the two respective lists, which contain the opening salutation and pet name that follows it, respectively.

While the `for` loop that starts on line 45 might be harder to parse, look at the `else` statement on lines 52-54 and try to relate how the combination of "YOU ARE MY", a selection from a list called `adjectives`, and a selection from a list called `nouns` is reflected in the body of the sample letters.

Cap off this part of the activity by emphasizing the essential tools used to generate new love letters: text strings, lists of words that can be chosen from at random and inserted into a text string, and the use of variables in text strings to create differences between sample outputs.

*Group Programming Exercise (estimate: 20-30 minutes)—*

Break students into groups of between 4 and 6 students. If a subset of the class has experience with computer programming, distribute them evenly amongst the groups, as they might be able to guide how students enact their ideas into a conceptual program.

Ask them to respond to the following prompts in a collaborative written document and either turn in the paper or send their document to the instructor digitally at the end of class.

1. Identify a written genre.
2. Why does this genre exist? When is it used?
3. What are the key features of the genre?
4. What key features might be easy to encode? What might be difficult?
5. Create a mix of text templates, variables, and lists that could generate new instances of this genre based on the features you have identified.

Emphasize that you will enact a selection of the student programs in code to present for the next class.

### **Between Sessions**

Identify 1 or 2 student programs that seem particularly viable to represent in code. For example, broadly described generator programs that do not contain lists or templates that interact with variables will require more labor on the part of the instructor to translate that program into code. While there are many potential approaches students can take, an ideal program for this activity will provide a basic string template, a list of words that is chosen from for parts of the template, and potentially variables that are called (e.g. the current time of day using the `datetime` module in Python).

Present text from the student document that states the chosen genre, the students who designed the generator, the description of that genre, and the key features they identified. Then, demonstrate the code that enacts the program with comments that explain what a given section is doing. The code should be presented in software that can run the code and present outputs along with the code, so that the generator can be run repeatedly to look at several different outputs.

For the materials included in this repository, I used Jupyter Notebooks to present text in Markdown in sequence with enactable code in Python. Please see `student-

generator.ipynb` for an example, or `generator-template.ipynb` for a notebook that can be easily edited and added to by instructors.

## Session 2

*Setup (estimate: 5 minutes)*—

Briefly summarize the activity of the previous session, in particular for students who may have been absent the previous meeting.

*Process Reflection (estimate: 5-10 minutes)*—

Ask students the following questions to encourage them to reflect on their processes:

- How did your group take your genre and design a generator for it?
- How did you consider the role of the computer as you designed?
- What did you leave out of your generator that is part of the genre you described? Why?
- Would you call what your group did *\*programming\**?

I used the final question as an opportunity to reinforce the broader idea behind the activity: programming is a conceptual process that mediates human goals and representations of the world and the digital technology that works with these representations in code. While their decisions had to account for the technical elements of digital technology, they were still making decisions about what a genre is, who it's for, when it's used, and what elements define it. They programmed by considering how to take their goals in defining a genre to generate new examples and make them work in the logic of digital technology. While the instructor does the actual coding, which is itself mediational, the focus is on how the students mediated human understanding and digital representation.

*Genre Generator Presentation (estimate: 15-20 minutes)*—

Show students the formatted text and code. Read through the name of the genre students chose, the names of the students who designed the generator, their description of the genre, and the key features they identified.

Relate a key feature to the code, which can generally be broken up for readability similar to the Vee implementation of the Love Letter Generator. For example, have all of the lists of terms in sequence, then the variables that choose from the lists, then the string template that makes use of the variables. If students identified that a salutation was important for an email to a professor, for example, then relate this key feature to the lists, variables, and place in the template that generate the salutation.

Show several sample outputs of the generator and ask students to identify which key features they can see in the outputs, and try to relate those outputs to what they can see in the code.

*Rhetorical Analysis of the Genre (estimate: 10-15 minutes)*—

Guide students through a rhetorical analysis of the genre as a mirror of the first session's discussion on love letters. You might ask about a number of specific features in the generator, but these general questions are useful starting points:

- What does this generator tell us about this genre?
- Why does this genre exist, and what does it do?
- How does this generator reflect these aspects of the genre?
- What aspects of the genre are important that you do not see included here? Why might they not be present in the generator?

The final question I like to ask to reiterate the sociotechnical nature of the genre generators is:

- Do these generators objectively identify what their respective genres *are*, or their key features?

### **Guiding Realizations**

1. Although the computer mechanically produces new instances of the genre based on pre-defined rules, there are still subjective decisions about what a genre is and how to represent it that are embodied in the code.
2. There is a distinction between *programming* as a conceptual process of shaping real-world things or ideas in terms that can be understood and processed by the computer, and *coding* as enacting those ideas in machine-readable language.
3. To define something like a genre in a way that can be interpreted by a computer does not mean that this definition is the only or most correct way to understand what that genre is. By virtue of making the genre fungible to the logic of the computer, certain understandings of the genre were prioritized and others, de-prioritized.



# Writing Against the Machine

## Debating with ChatGPT

Justin Lewis and Ted Wayland

Olympic College

*This assignment challenges students to use LLMs to map the counterarguments to their main claims through a process of counterclaiming. By writing prompts that position the AI as argumentative adversary, students can refine, extend and evolve their thesis over the course of the research, drafting and revision processes. This lesson is focused on 100- level composition courses but could be adopted to any undergraduate or graduate course focused on argument and claim making.*

---

*Learning Goals:*

- Discover new pathways and avenues for research
- Destabilize that existing research topics and reassess previous work by accounting for the intricate nature of an initial claim
- Introduce source credibility as an integral component of the research process.

*Materials Needed:* Large language model with linguistic analysis feature such as GPT-3.5 or GPT-4

*Original Assignment Context:* First-year writing course

*Timeframe:* ~2 course sessions

---

### Overview

AI chatbots that rely on LLMs, such as ChatGPT, provide a wealth of opportunities for instructors working in first year writing. When directed appropriately, a chatbot provides a collaborator for students in numerous invention activities (Anson 42). In our composition courses, we've used ChatGPT and other LLM wrappers to map the stakeholders or relevant audience(s) for a given argument, summarize key points from short course texts, provide alternative language to (re)shape thesis statements, and restructure drafts through the tool's imposition of outlines on draft-stage freewriting activities. In this short teaching reflection, we'll describe how we've used ChatGPT as a method of *counterclaiming*, or the process of mapping the complex counterarguments that work against the initial claim-making of students in first year composition. Inviting students to sketch counterclaims with ChatGPT not only provides them novel pathways to new research foci, it also invites them to resist overcommitting to a central claim too early, instead encouraging them to consider how their claims may evolve through a series of successive complications in the research process (Rosenwasser and Stephens 139; Nussbaum and Schraw 59).

### Context

Both authors teach first year writing courses at a mid-sized community college in the Pacific Northwest. Like most community college student populations, our demographics vary widely and include a diversity of ages, backgrounds, and

abilities. Through direct transfer agreements, many of our students move on to four-year universities after completing their Associates; however, students in our courses also pursue certificate and degree pathways in skilled trades or are completing pre-college requirements before applying to other schools. Because of the college's proximity to US military installations, we also provide educational opportunities to many active duty and recently discharged military personnel.

The counterclaiming activities that follow were developed and integrated into two first year writing courses. For Ted, the tool was integrated into his Winter 2023 section of ENGL101, a class dedicated to introducing students to academic literacy in reading and writing. For Justin, the tool was first utilized in Fall quarter 2022 but further refined in his Winter 2023 ENGL102, a course that emphasizes academic research and the development of source-supported academic arguments. While the broad course outcomes for ENGL101 and ENGL102 are similar, the major difference between the two courses is 102's emphasis on finding, evaluating, and integrating sources into student writing. Both courses focus on critical reading practices, rhetorical choices in academic writing, and the use of a recursive writing process.

Justin used the counterclaiming activities during the second week of Unit 1 (4 weeks total). The first project is an "Inquiry Memo" that invites students to submit a proposal for the research topic they aim to pursue over the course of the term. The main activities of the inquiry memo include defining research topics, defining researchable questions, and posing a main claim to guide research practices in Unit 2 (annotated bibliography – 2 weeks) before composing a researched argument in Unit 3 (4 weeks). The activities of this assignment occupied two class periods, or the third week of the first unit. After completing these activities, students will be introduced to questions of source credibility in the final week of the unit.

On the first day of his ENGL101 course, Ted showed students a [short video of Jacques Droz's 18<sup>th</sup> century writer automaton](#), sparking discussions on our collective definition of the act of writing and its relationship to cognition. With the memory of that uncanny, clockwork writer in mind, the course's final essay assignment introduced the use of ChatGPT as a digital interlocutor. The final assignment asked students to write a thesis-driven essay on a topic related to memory, synthesizing information from a film and essays relevant to the topic. Students were asked to use ChatGPT at two moments in the writing process for the assignment: when working on defining their topics, students used multi-stage queries with the bot to create more specific, narrow subtopics, and later, when working on crafting theses, students used the bots to generate counterclaims to test their arguments and refine their theses in response. Finally, students were asked to reflect on their experience using AI in the writing process and to revisit their ideas about the intersection of writing, technology, and thought.

### **Assignment Goals**

Over time, we've both attempted numerous invention activities to encourage students to map counterclaims to the initial positions they adopt early in the writing process. We advocate this kind of work because we want students to avoid overcommitting to a particular position until their research process is complete. Ideally, we want students to *evolve* their thesis over the course of their reading and writing activities. Unfortunately, as many students remind us, "We don't know what we don't know." While keyword association and topic relationship mapping tools provided some pathways for students to discover counterclaims, the authors found these technologies cumbersome and tied to library subscription services. Because ChatGPT's UI is intuitive and reflects a search-based mental model, we were hopeful

that developing counterclaim activities in this interface would encourage higher student participation due to usability and familiarity.

The goals for our use of ChatGPT to counterclaim include:

- Familiarizing students with the emerging role of AI in the writing process; and
- Improving the nuance and complexity of student claim-making and thesis evolution using the AI.

The specific learning outcomes of our use of ChatGPT to counterclaim include:

- Students will be able to define a subject position for ChatGPT;
- Students will be able to provide specific instructions to the AI to guide its counterclaiming;
- Students will be able to use ChatGPT to provide multiple counterclaims or alternative arguments to their initial argumentative thesis (claim); and
- Students will be able to evolve their initial argumentative thesis (main claim) in light of the counterclaims provided by the AI.

### **Required Materials**

At present, this lesson relies on the [free preview of ChatGPT](#); however, we expect access to this model will change and likely require a fee that could be paid by the student, instructor or the institution. Another prompt-based, generative AI interface would also work for this assignment. In addition to ChatGPT access, students were required to use the [prompts.chat Github repo](#) to shape language around defining the subject position of the AI. This lesson could be customized/extended by integrating additional subject roles for the AI using the prompts builder at [prompts.chat](#).

### **Assignment Results**

#### *Ted's ENGL101 Course*

ENGL101 students showed a great deal of enthusiasm and interest with using ChatGPT in the writing classroom. The majority of the students in Ted's two classes had not used an AI bot for writing before our activities in February 2023, and many had not even heard of ChatGPT. During the first activity, students were wowed by the capabilities of the AI, one student noting that it was "interesting how it imitates lifelike conversation." Their first activity centered on generating subtopics for an essay assignment on the theme of memory in a film they had selected, and the general feedback from students was that it was helpful to query the bot as a brainstorming tool, quickly generating ideas that suggested new possible topics for their essays; there was largely a consensus that the subtopics generated by ChatGPT were reflecting their own ideas back to them, albeit in a more focused manner. A student who fully embraced the tools said "it's hard to come up with a good topic, but with this, you put in your idea and it gets right to your point!" It was remarkable to them how well ChatGPT could generate subtopics with some degree of specificity on a particular film: it might mention, without a prompt, that *Rashomon's* director is Kurosawa, or seem to possess a knowledge of the disjointed narrative in *Memento*. On the other hand, students observed moments where the AI offered up fuzzy, incorrect, or incomplete information about films, especially less well known ones; it seems not to have "seen" Chris Marker's *La Jetée*, for example.

On this first day of working with ChatGPT, our discussions of the process had a consistent theme: amazement with what the AI could produce, mixed with some

uneasiness with the technology; one student explained that “it felt like cheating” to use AI this way, while another expressed a feeling that the AI “took the art out of writing.” In response, another student approvingly described the AI as offering up a skeleton that he would be able to flesh out with his own writing.

The second ENGL101 activity took students from topics to theses, and this time ChatGPT’s role was providing counterclaims to students’ first drafts of a thesis. Interestingly, on just the second day of using the bot, students were already past the novelty of AI writing and were more intent on stretching its capabilities and looking at it with a more critical eye. Students wrote a reflective note at the end of their assignments, and the general feedback was that using the AI was helpful and something they could imagine incorporating into their writing process—with, however, significant caveats. In class discussions and the students’ reflections, students returned again and again to the generic, boring, non-committal quality of the writing ChatGPT produced, and by the end of the activity, students were balancing their fascination with the possibilities of the technology with sharp observations on its limitations. Ultimately, this led us back to our class mascot, Droz’s writer automaton, and a recognition that thinking and writing remain, for now, in the sphere of the human mind.

### *Justin’s ENGL102 Course*

Overall, students enrolled in the research focused ENGL102 course responded positively to the use of ChatGPT in the invention process. Beyond students’ initial excitement over the “magic” of AI text generation, students noted that ChatGPT provided them with new pathways for research and discovery, leading them to question their initial reaction to a broad topic. For students with claims that were far too broad, ChatGPT tended to narrow through complication by identifying counterarguments that supported the counterposition to their own. For students whose initial claims were too narrow, the AI occasionally provided new vocabulary and new lines of research related to the counterposition. These new counterclaim subtopics tended toward less specificity, allowing students whose initial claim-making was too granular to zoom back out to identify a more appropriately scoped focus.

Yet, not all students reported that the counterclaiming activity was generative. For students with topics that were of the appropriate scope and specificity, the AI sometimes provided generic counterpositions that restated their own or that they had already identified in their own research and writing. For still others, the counterclaiming activity complexified their research questions and claim-making in negative ways by diverting the research process and creating too many opportunities for further exploration, especially in the context of a quarter-long research project.

Without equivocation, when challenged with evaluating the sourcing credibility of the AI, students found that ChatGPT fabricated evidence by inventing credible *sounding* sources. Interestingly, this didn’t invalidate the counterclaims themselves; rather, as one student put it, “The sources are bullshit, but the ideas seem legit.” In this way, we view ChatGPT – and other tools like it – in much the same way as Anson and Straume who see these AI as “writing assistants’ requiring some supervision of outputs” to ensure that the ideas and the sources aren’t both ‘bullshit’ (5).

### **Works Cited**

Anson, C.M. “AI-Based Text Generation and the Social Construction of Fraudulent Authorship: A Revisitation.” *Composition Studies*, vol. 50, no. 1, 2022, pp. 37-46.

Anson, C. M., and I. Straume. “Amazement and Trepidation: Implications of AI-Based Natural Language Production for the Teaching of Writing.” *Journal of Academic Writing*, vol. 12, no. 1, Dec. 2022, pp. 1-9, doi:10.18552/joaw.v12i1.820.

Nussbaum, E. Michael, and Gregory Schraw. “Promoting Argument-Counterargument Integration in Students’ Writing.” *The Journal of Experimental Education*, vol. 76, no. 1, 2007, pp. 59–92. *JSTOR*, <http://www.jstor.org/stable/20157471>. Accessed 13 Feb. 2023.

Rosenwasser David and Jill Stephen. *Writing Analytically*. 5th ed. Thomson Wadsworth 2009.

---

## The Assignment

Over the course of class this week we’ll be using artificial intelligence to assist us with discovering new pathways and avenues for research. During the first two weeks of the unit, you identified a main idea, or topic, that you’re invested in and that can sustain your research activities for the term. Further, you’ve also evaluated what’s *arguable* about that topic and developed an initial *arguable* main claim to guide your research. Now that you’ve found some stability around your topic, we’ll be conducting some activities to *destabilize* that topic. My hope is that in completing this process, you’ll reassess your existing previous work by accounting for the intricate nature of your initial claim. In addition to refining and evolving your claim, we’ll also use these activities to introduce source credibility as an integral component of the research process.

### Day 1: Writing Good Prompts for Chat GPT & Initial Queries

#### *Overview*

Today’s class will be dedicated to securing a ChatGPT account, crafting effective prompts for the AI and executing counterclaim queries.

#### *Learning Objectives*

After today’s class, students:

- will be able to define a subject position for ChatGPT;
- will be able to provide specific instructions to the AI to guide its counterclaiming;
- will be able to use ChatGPT to provide multiple counterclaims or alternative arguments to their initial argumentative thesis (claim);

#### *Activities*

1. Step One: Securing a ChatGPT Account.
  - a. First, visit [chat.openai.com](http://chat.openai.com) and click “Sign Up.” Follow the prompts to register your ChatGPT account.
2. Step Two: Writing the Script for Your Query
  - a. After you’ve secured your account, take a moment to play around a bit with the AI interface. Ask it some questions. Take note of the generative/

interesting responses you receive. Think expansively here and, more importantly, think playfully. [*Discuss interesting responses and brainstorm possible uses of the bot with students – 15 minutes*]

- b. Now that you're acquainted a bit with the tool, we're ready to begin our goal of generating counterclaims to your own research as a strategy for complicating and refining your research topic and thesis. To do so, you'll need to develop a prompt for your query. We'll be using language included on [prompts.chat](#).
  - i. *Activity 1: Finding your persona and prompt:* First, head over to [prompts.chat](#). Next, find the prompts that define personas that might work well for identifying counterclaims. I'm particularly fond of the persona that [acts as a debater](#). After you've selected your persona, look at the language included to generate a query with that persona. We'll use that language to generate your specific prompt. Copy+paste that language into a Word doc.
  - ii. *Activity 2: Customizing your prompt:* As I stated before, I like using the debater persona to generate counterclaims. Here's the prompt language for that persona:

I want you to act as a debater. I will provide you with some topics related to current events and your task is to research both sides of the debates, present valid arguments for each side, refute opposing points of view, and draw persuasive conclusions based on evidence. Your goal is to help people come away from the discussion with increased knowledge and insight into the topic at hand. My first request is "I want an opinion piece about Deno."

Your first writing assignment in this counterclaim activity is to rewrite this prompt so that it best fits your own research project and main claim/thesis. So, for example, if my topic is "Gifted programs for students in K-12" and my main claim was, "Gifted programs provide students with positive opportunities to excel beyond their peers," I would rewrite the prompt above this way:

I want you to act as a debater. I will provide you with a main topic and my claim about a topic. Your task is to research the claims that go against my topic and highlight the best arguments against my main claim. Your goal is to help me better understand alternative viewpoints on my topic and other claims that could be used to argue against my main claim. My main topic is: "Gifted programs for students in K-12." My main claim is, "Gifted programs for students in K12 provide students with positive opportunities to excel beyond their peers."

So, to begin, identify the persona you want to use to develop your counterclaims. Next, rewrite the prompt provided for that persona in a way that you believe will generate claims and/or arguments against your claim.

### 1. Step 3: Executing Your Query & Refining Your Prompt

- a. Now that you've written your prompt, it is time to see what ChatGPT can do! Execute your prompt using the text input at the bottom of the

ChatGPT homepage and sit back, relax, and see what counterclaims the AI creates in response to your query.

- b. *Optional Activity: Refining Your Prompt:* Some of you may have received a wealth of information related to your topic. If so, that's wonderful! However, if your prompt didn't initially generate a robust response, consider rephrasing or reworking your prompt so that it provides more detail and more information related to your main topic and main claim. Experiment and rerun your prompts multiple times to generate a diversity of counterclaims.

### *Homework*

Continue to refine your prompt and generate responses to your queries. For homework, refine your prompt at least three times, run the subsequent queries as new chats and copy+paste that information into a document to refer to in our next class meeting.

## **Day 2: Mapping Research Opportunities and Evaluating AI Credibility**

### *Overview*

Today's class will be dedicated to mapping the counterclaims provided by your ChatGPT queries to complicate your initial claim and identify new research opportunities. We'll also spend a bit of time considering the reliability and credibility of the AI responses.

### *Learning Objectives*

After today's class, students:

- will be able to evolve their initial argumentative thesis (main claim) in light of the counterclaims provided by the AI; and
- will be able to evaluate the accuracy and credibility of the results of their ChatGPT counterclaiming activities.

### *Activities*

1. Step One: Identifying Counterclaims and Research Clusters
  - a. At this point you've collected multiple ChatGPT outputs into a shared document. Now it's time to *distill* all those counterarguments into a usable collection of counterclaims you can use to complicate your initial main claim and reshape your research trajectory. Here's how:
    - i. First, look for patterns in the ChatGPT responses. The AI often frames responses in terms of "lenses" like "economic," "social," or "racial." First, identify if any "lenses" are present and group responses based on those lenses. If lenses aren't obvious, do your best to group like responses with like. After you've grouped the responses, name your groupings or response clusters.
    - ii. Next, reflect back on your initial topic and focus. Based on the ChatGPT clusters, how might you reshape your topic and your main claim based on the counterclaims provided by the AI? Rewrite your main claim in anticipation of addressing these counterclaims.
    - iii. Based on the research clusters identified in Step I, do you find any of these areas more compelling or interesting to address than

- others? Identify 1-2 clusters you're interested in investigating more in your research project. Rewrite your prompt from Day 1 with a focus on your new research cluster. Run the query again to identify and refine further the research questions and counterclaims identified in this cluster.
- iv. Based on the cluster grouping above, identify 4-5 counterarguments you believe you'll need to address in your own research paper. Google them. Wikipedia them. Do they seem credible? If so, how will you address them in your research? Why/why not?
  - v. In your working with ChatGPT, have you reconsidered the scope, focus or slant of your initial position? Why/why not?

*Homework*

Read [“Assessing Source Credibility for Crafting a Well-Formed Argument”](#) in *Writing Spaces Volume 3*.



# Using AI Text as Prompts for Critical Analysis

Juan Pablo Pardo-Guerra  
UC San Diego

*In this assignment, students are provided with an AI-generated text relevant to a course's topics and focus and then asked to comment, review, and expand on it using a feature such as "track changes." In engaging with the AI-generated text, students review their knowledge, offer critiques, modify theoretical and empirical claims, and provide concrete examples that illustrate or disprove the provided answer. Their critical and evaluative efforts for the course's topics are foregrounded, and additionally, they gain some AI literacy in evaluating the AI-generated text.*

---

*Learning Goals:*

- Contrast the AI-generated claims with those developed throughout the course
- Evaluate the quality of AI-generated texts in relation to formative and summative discussions, readings and exercises had during the course
- Modify the provided text in order to align its claims with those of the literature
- Give examples and counterexamples of claims made in the AI-generated text.

*Materials Needed:* Short texts (500-1000 words) generated by a LLM on the basis of prompts, Text-editing software which allows students to comment on the provided AI-generated texts

*Original Assignment Context:* upper division undergraduate class on economic sociology

*Timeframe:* ~1 course session

---

## **Introduction**

There is no getting around the fact that artificial intelligence has changed the landscape of higher education in critical ways. This is particularly clear in the case of Large Language Models that, trained on vast amounts of digital data, manage to produce intelligible texts from a user's prompt. Although Large Language Models do not "know anything at all" (Burrell 2023) and are merely "stochastic parrots" (Bender, et al.) echoing dead texts and archived digital interactions, their ability to create texts with a few keystrokes and a click is of concern.

A considerable part of the debate on the most recent generation of Large Language Models has centered on the detection of AI-generated text. This simply mirrors earlier discussions about plagiarism, where the emphasis was on locating instances of copying and paraphrasing existing texts without adequate attribution and with the clear intent of passing stolen materials as original contributions. Less time is spent understanding why students chose to use these shortcuts in the first place—placing focus not on surveillance and punishment but on the incentives that lead students to cheat. Students cheat more often than not because the assignments do not challenge their skills, lacking clear benefits and connections to learning outcomes against which they can measure their performance. Knowing how to summarize Georg Wilhelm Hegel, Karl Marx, Max Weber, and Michel Foucault is certainly a skill, but

one that has no clear value—specifically in how it is operationalized in practice—with the type of cognitive work that most students will have to engage in employment and civic life.

Given that Large Language Models are particularly effective at summarizing documents and producing credible accounts of frequently used historical texts (an LLM would likely generate a quite reasonable analysis of the Communist Manifesto), students may find that requesting them to do what, in their view, is the same kind of work as that performed by a machine is an invitation to disengagement. Students need to learn these skills, of course, but building assessment strategies around these as if they are true proxies of “critical thinking” is a mistake.

The solution to this puzzling situation is not avoiding LLMs but accepting them as extensions of our analytical and pedagogic toolkits. LLMs can become instruments around which we develop distinct critical and analytical competencies with our students. What matters here is crafting assignments that both develop interest *and* align directly with learning outcomes, providing students with a sense of development and empowerment rather than repetition and emulation. In what follows, I outline what one approach might look like.

### **Goals and Outcome**

I have been teaching an upper division undergraduate class on economic sociology since 2016, titled *Economy & Society*. The main purpose of the course is to impress upon students the argument that economic outcomes like wages, financial stability, and entrepreneurial success aren't simply products of individual efforts and merit but reflect the structural inequalities across race, gender, social class, and ability that shape most social situations.

In the past, I have assessed this class with short weekly write-ups on small empirical projects (for example, a discussion of the notorious case of Theranos, the biomedical company that was purely a sham, or a discussion of the consequences of redlining in San Diego using various maps and datasets of the region). These speak directly to one of the three learning outcomes for the course, namely, to provide students with the ability “to use sociological concepts and explanations to critically analyze how their own economic lives are shaped by broader social structures.” In addition to these small write-ups, students also submit a final, 1,500-word essay based on one of 5 possible prompts.

This last assessment is particularly exposed to potential cheating with LLMs. It would be possible to remove this element of assessment from course, but that would eliminate a final moment for integrative reflection about how social structures impinge on economic outcomes.

Participants of this exercise should be able to 1) contrast the AI-generated claims with those developed throughout the course; 2) evaluate the quality of AI-generated texts in relation to formative and summative discussions, readings and exercises had during the course; 3) modify the provided text in order to align its claims with those of the literature; 4) give examples and counterexamples of claims made in the AI-generated text.

### **Materials**

For this exercise or those that seek to emulate its design, instructors will require:

- Short texts (500-1000 words) generated by a LLM on the basis of prompts. Each text should correspond to one prompt. They should be selected by the instructor to be “reasonably competent” answers to the prompt while lacking the empirical and theoretical depth of a well-written response. Instructors may benefit from generating these texts through role-prompting, that is, by requesting the LLM to generate the texts “as if” written from a particular adversarial position (for example, a libertarian who does not believe in the welfare state, or a business person who is committed to individual merit).
  - Text-editing software which allows students to comment on the provided AI-generated texts. This software can be provided through a public service (like Google Docs) or a platform like Microsoft Office. It should provide commenting, strikethrough, editing, and markup features.
- 

## The Assignment

Given that the main objective of the course is to foster critical thinking among students in relation to real-world cases, the new assessment involves transforming the output of a LLM, ChatGPT, into the empirical object that students will have to engage with.

Because the course guides students to think about how expectations and assumptions about categories of class, gender, race, and ability play a role in shaping economic outcomes, asking them to identify and query these expectations and assumptions is a skill they are expected to have developed. Thus, rather than asking them to produce text, the new assignment asks them to criticize an AI-generated text, making it the kind of empirical material they analyze in their weekly write-ups.

In this assignment, students are provided with a 500-word computer-generated essay that responds to a prompt (for example, “Is economic inequality inevitable?”). They are then asked to use word processing software, such as Word or Google Docs, to make comments and edits on the AI-generated document. Specifically, students are asked to:

1. Identify and correct factual and theoretical mistakes in the provided text.
2. Add contextual information to the claims made in the text, referencing readings from the course.
3. Develop the themes of the essay through original cases and personal experience.
4. In one additional paragraph, describe how assumptions about class, gender, race and ability were propagated in the original text.
5. Students should use “track changes” and “comments” features to allow for the readers to see what was changed.

The result of this exercise is not an essay but corrections on an essay. This assessment exercises the critical abilities of students to evaluate and contextualize claims in relation to the themes of the course. This is harder to simulate with LLMs, given that it asks for a better understanding of contextual information (course readings and themes as well as personal experience) and novel empirical data (such as cases that illustrate the claims in the essay). Instead of reducing critical thinking to the production of (relatively predictable) texts, this exercise invites students to use their critical thinking skills to curate materials, connect topics, and propose changes

—a skill that better represents the type of cognitive tasks that they will confront in the future.

# Generate and Enact a Writing Style

## Examining Writing Style Through Generative AI

John J Silvestro  
Slippery Rock University

*This assignment asks students to develop a more critical and actionable understanding of the concept of writing style through the use of AI text generators. Writing style is a difficult concept to teach given the challenges of concretely presenting it to students. AI text generators offer tools for quickly generating multiple versions of sentences and paragraphs. Students can use multiple versions of a sentence to closely examine individual aspects of their own and others' writing styles. This assignment enables students to develop and define writing style more concretely.*

---

### *Learning Goals:*

- Develop and enact a critical definition of writing style
- Understand when to use AI text generators in the writing process
- Analyze others' writing, including that created by AI text generators, through definition of writing style

### *Materials Needed:*

- Students having reached the first draft phase of a writing project
- Essay "The Evolution of Imitation: Building Your Style" by Craig A Meyer (available free online in *Writing Spaces* textbook collection)
- Example paragraphs in a distinct writing style, such as Didion, Hemmingway, Morrison
- Access to an AI text generator, such as ChatGPT, Bing, or Sudowrite

*Original Assignment Context:* middle of academic research essay project in a first-year writing course

*Timeframe:* ~1 week

---

## Overview

Many writers have modified their writing approaches to accommodate technological systems like text correctors and text predictors. They continually check their grammar and spelling throughout the writing process because of the ways text correctors immediately scan and mark text and then efficiently provide lists of corrections. In turn, writers have modified their writing style to better accommodate the corrections and predictions of the aforementioned systems. The most advanced technological systems in this lineage, AI text generators like ChatGPT and Sudowrite, offer writers opportunities to expand their writing approaches, particularly their understanding of and applications of writing style.

Writing students in particular could benefit from some of the ways text generators operate to help them extend their approaches to writing style. Writing style can be

incredibly difficult to learn as it is a somewhat nebulous concept, unique in certain ways to each writer, and requires close considerations of multiple variants of similar sentences and/or paragraphs to be fully recognized. Text generators offer a tool that enables students to quickly generate multiple variations of the same sentence or paragraph. Students can then use readily-available variations to develop a sense of writing style and to experiment with their own writing style. Text generators thus offer an opportunity to help students better define and enact writing styles(s).

### **Focusing on Style**

The assignment outlined in this chapter proposes using the rhetorical concept of writing style as a framework for critically engaging and incorporating text generators into writing processes. Specifically, this assignment proposes that writing teachers work with their students to first develop in them a deeper sense of their own writing style and how their style shapes the meaning, value, and impact of their texts. From there, writing teachers can use text generators to help students expand their sense of writing style. Text generators can become a quick, readily available reference point for the structure and style of paragraphs that students can use to expand their own writing.

However, writing style is an under-used concept from a pedagogical perspective. Too often, writing style is approached without a critical framework. The teaching of writing style gets reduced to concerns for “flow” or “professional tone” (Aull and Lancaster 98). These simplified perspectives on writing style are particularly useless when applied to the paragraphs produced by text generator technologies. The paragraphs generated by most of these technologies have a coherent-yet-indistinct flow and abstractly professional tone. Writing teachers must work with students to develop in them both a sense of their own writing style and a more advanced concept of the rhetorical theory of writing style.

To enable students to develop a more advanced concept of writing style, this assignment draws from theories of rhetorical imitation and Laura I. Aull and Zak Lancaster's work on writing style as a critical stance. Imitation is a well-established way of introducing writing style: in most approaches to imitation students revise the structures of other writers' sentences to develop new sentences. Specifically, students keep the various parts of speech in the others' sentences, incorporating their own nouns, verbs, adjectives, and adverbs that work for their writing. For example, students replace the noun from a sentence with a noun relevant to their argument. As a writing activity, imitation offers a useful starting point for developing both an expanded understanding of writing style, as it offers a way to engage with the impact that writing style has on the meaning of writing, and it provides an obvious writing practice that students can apply to any text-generator-produced paragraphs.

Imitation is not enough, though. Students need to be able to do more than just potentially rewrite the nouns, verbs, adjectives, and adverbs from text generator-generated paragraphs. They need a framework that enables them to understand why those nouns, verbs, adjectives, and adverbs and the order those parts of speech are in matter. To help students develop a sense of writing style as a framework, I begin by introducing Laura L. Aull and Zak Lancaster's work on writing style as a stance. They define style as a stance that manifests through a writer's use of prepositions, adjectives, research, and/or their framing of readers. Put differently, writing style is a stance that writers gradually adopt through their revision process that shapes how they approach their argument, context, and audience. Writing stance offers dual uses for writing students. It offers a set of specific aspects of a text to analyze in their own

writing. It also offers a framework for analyzing the style of text-generator-generated text. It focuses students on specific elements of generated paragraphs to engage.

Thus, this assignment presents a week-long project that seeks to introduce students to a more complex critical framework for writing style, writing style as stance (Lancaster and Aull 100), then to use that framework to investigate their own writing style as well as the writing style of text generator technologies. A key part of this project is using text generator technologies to quickly generate versions of students' own paragraphs. Rapidly generated paragraphs offer students a clearer sense of other ways they could have worded and structured their sentences and paragraphs. Text generator technologies then offer a mirror of sorts that students can use to reflect their writing back at themselves, or at least alternative versions of their own writing. Once students gain a sense of their own writing style, they investigate the writing style of text-generator-produced writing. In examining text-generator-produced writing through a writing style lens, students can develop a more critical perspective on the writing generated by the technologies and gain a clearer sense of how they will likely want to rewrite and rework any text-generator-produced writing they decided to use.

If writing teachers provide students with critical frameworks for writing style, they can help students expand their writing processes so that they can productively incorporate text generator technologies. By deepening students' understanding of writing style, teachers can enable students to situate text generator technologies as resources, much like the spell-checker algorithm. If students can develop an expanded understanding of their own writing style and of writing style in general, they can more productively engage text generator technologies; they can maintain a sense of agency over their writing as they grapple with the streams of coherent, bland, semi-professional, culturally-specific paragraphs that these technologies generate.

---

## The Assignment

### Ideal Course

Any course engaging writing

### Requirements

- Access to writing software (Word, Google Docs, Pages, etc.) in and out of classroom
- Access to ChatGPT, or similar text-generator technology such as Lex
- Access to “The Evolution of Imitation” by Craig A. Myer and other writing samples that demonstrate a distinct writing style
  - Essay is open-source and readily available through the online writing textbook *Writing Spaces*
- Access to “Stance as Style: Toward a Framework for Analyzing Academic Language” by Laura Aull and Zak Lancaster.
  - Available in the *Style and the Future of Composition Studies* collection as well as online.

### Times Taught

Three as of publication

## General Outline

(Should be taught when students are revising a writing project)

1. After drafting a part of an essay or longer piece of writing, as an assignment, students read "[The Evolution of Imitation: Building Your Own Style](#)." The essay, by Craig A. Myer, outlines a broad definition of writing style and then presents a few imitation-centered writing activities that help students start to understand how writing style develops at the sentence level. Students then use the activities from the reading to rewrite a paragraph from a current draft into a different writing style based on an example.
2. In-class discussion of assignment work. Class articulates an initial understanding of the importance of style. The teacher introduces and outlines Aull and Lancaster's theory of writing style as stance, focusing on the role of prepositions, adjectives, research, and readers.
3. As an assignment, students summarize one of their own paragraphs and then put it into a text generator to have it create a version of the paragraph. Students write a comparison of the paragraphs.
4. In-class discussion about the paragraphs. Students discuss the differences between their take on a topic and the text generator's. They focus on the prepositions, research, and framing of readers.
5. As an assignment, students select a paragraph by an established author with a distinct writing style. Students then have a text generator create a version of the selected paragraph. Following this, students compare the two paragraphs, focusing on prepositions, adjectives, research, approach to readers, and other elements such as sentence length and word choice. rewrite the text-generated paragraph to be in their personal style. They focus on changing prepositions, adjectives, research, and framing of readers. They aim to rewrite every word of the text-generator's paragraph.
6. In-class discussion about writing style using the differences between established authors and the output of AI text generators as a focus. As a class, students write a class definition of writing style, articulating its rhetorical significance. The class ends with students using the definition of writing style to start revising their writing. an application of the class's definition of style to the writing produced by text generators.

### Out-of-Class Assignment #1

\*Should be assigned when students are deep into the draft of a longer piece of writing.

*Explanation for students:* You have reached a point in your draft where you should have a solid sense of what you are trying to do, what point(s) you are making, the argument you are articulating, and/or what you want to critically engage. This means that you have reached a point in your writing process where you should start to focus on *how* you are articulating those ideas. You should start to think about the words, phrases, and sentence structures you are using to manifest and shape your ideas. You should start thinking about your writing style.

With this in mind, you will read "The Evolution of Imitation: Building Your Own Style" by Craig A. Myer. This essay presents a series of tools that you can use – imitation tools – that can help you examine, rework, and/or broaden your writing style. The essay presents a few ways to alter your writing style.



After you finish the reading, you will enact one of the approaches from the essay on one of the paragraphs from your current draft. You will use the essay to revise your paragraph. You will rewrite your paragraph so that it is either a structural imitation or a contextual imitation. If you decide to do a structural imitation, you will need to find another paragraph, written by someone else, and then copy the structure of the paragraph. You are welcome to select any paragraph you want. The only requirement is that someone else authored the paragraph.

### **In-Class Lesson Plan #1**

\*Should follow Out-of-Class Assignment #1

#### *General Outline*

*Opening Mini-Lecture:* teachers discuss writing style and reading. They talk briefly about how writing style is something distinct and that plays an equally important role in the expression of the meaning and value of an idea, point, or argument.

*Discussion:* students are asked to review their revised paragraph and identify the most noteworthy or interesting change they made in the revised paragraph. Students then share their examples and discuss the changes they made. Discussion centers around drawing out ideas about how writing style is constructed – through prepositions, adjectives, use of research/information, and the framing of the reader.

*Mini-Lecture:* teacher introduces the concept of “writing style stance.” This concept forwarded by Laura A. Aull and Zak Lancaster presents writing style as a distinct stance that a writer takes toward their ideas as well as their readers. (Teachers can assign sections from Aull and Lancaster’s work if helpful. Otherwise, teachers will need to familiarize themselves with this approach and be prepared to introduce it to students). The teacher should focus on how stance is expressed, through prepositions, adjectives, use of research/information, and the framing of the reader. They should also focus on how writing style stance shapes the ideas, points, and/or arguments that a writer makes.

*End-of-class:* teacher should introduce Out-of-Class Assignment #2. They should quickly demonstrate how to use ChatGPT to students.

### **Out-of-Class Assignment #2**

\*Should be the work assigned after In-Class Lesson Plan #1

*Explanation for students:* Now that you have experience with writing style as an imitation practice as well as a stance that a writer adopts, we want to continue exploring your writing style. We want to start identifying what your writing style is. We also want to think about the stance you are taking.

To help you start examining your writing style, you will do a few things. First, you will write a brief one-to-two sentence summary of the paragraph that you revised for our previous class period. You want to have a clear understanding of what you are doing with the paragraph.

Second, you will go to ChatGPT (link). You will input the one-to-two sentence paragraph summary that you wrote. You will then ask ChatGPT to generate a paragraph on those points for you.

Third, you will copy-and-paste the paragraph that ChatGPT generates into a document. You want to save the paragraph that ChatGPT generates so that you can compare it to your own paragraph.

Fourth and finally, you will write a short comparison of your original, revised paragraph to the paragraph generated by ChatGPT. You want to focus on the aspects of writing-style-as-a-stance that we discussed during our previous class. You want to focus on how ChatGPT uses prepositions, adjectives and adverbs, how it uses information, and how it frames its readers (aka you). You want to think about what you are doing and how that differs from what ChatGPT does.

## **In-Class Lesson Plan #2**

\*Should follow Out-of-Class Assignment #2

### *General Outline*

*Opening Discussion:* students share their experiences using ChatGPT and the differences and similarities between their writing and the paragraph generated by the auto-writer. The teacher should keep a running list on the board or in a shared online writing document of all the points students make about the things that the auto-writer does.

*Activity:* students are then asked to define “writing style.” They individually write two-to-three sentence definitions of the rhetorical theory of writing style. The teacher then works with the students to develop a collective, negotiated definition of writing style.

*End-of-class:* teacher should introduce Out-of-Class Assignment #3.

## **Out-of-Class Assignment #3**

\*Should be the work assigned after In-Class Lesson Plan #2

*Explanation* *for*  
*students:* Hopefully, you now have an understanding of writing style and how it plays a co-facilitator role in writing. The writing style one uses influences their message, points, and/or arguments. The ultimate goal is for you to develop a sense of your own writing style. For you to be able to define what you want to do and why you want to do it.

However, before you can do that, you should learn to identify elements of a text’s writing style. You should learn to analyze and understand others’ writing styles. With that in mind, you will do two things for our next class period. First, you will select one of the paragraphs below. These paragraphs all provided different and distinct writing styles. You will select one paragraph that you feel is well written and that has elements that you want to copy. Following this, you will write a short summary of the paragraph, turning it into a prompt. You will enter that prompt into ChatGPT to

have it generate its own version of the paragraph. We are using ChatGPT to quickly get a quick sense of how else this paragraph could be written.

Finally, you will write a brief analysis of the writing style of each paragraph. For the comparison, you will explore the writing style of each individual paragraph. As part of this exploration, you will reference the other paragraph, using it as a point of comparison for the paragraph you are exploring. First, you will examine the writing style of the paragraph you selected. You will pay attention to how the paragraph uses prepositions, adjectives, research/information, and framing of the reader. As part of this, you will consider how the paragraph's writing style differs from the style ChatGPT generated for its paragraph. Following the first comparison, you will do the exact same analysis for the paragraph generated by ChatGPT. You want to analyze the style of the auto-generated paragraph and how it differs from the paragraph you selected.

We will then use these paragraphs to talk about how writing style is actualized. Following this, you will develop an outline of your own writing style which you will use to write your current essay.

*Example paragraphs:*

- “You think because he doesn’t love you that you are worthless. You think that because he doesn’t want you anymore that he is right — that his judgement and opinion of you are correct. If he throws you out, then you are garbage. You think he belongs to you because you want to belong to him. Don’t. It’s a bad word, ‘belong.’ Especially when you put it with somebody you love. Love shouldn’t be like that. Did you ever see the way the clouds love a mountain? They circle all around it; sometimes you can’t even see the mountain for the clouds. But you know what? You go up top and what do you see? His head. The clouds never cover the head. His head pokes through, because the clouds let him; they don’t wrap him up. They let him keep his head up high, free, with nothing to hide him or bind him. You can’t own a human being. You can’t lose what you don’t own. Suppose you did own him. Could you really love somebody who was absolutely nobody without you? You really want somebody like that? Somebody who falls apart when you walk out the door? You don’t, do you? And neither does he. You’re turning over your whole life to him. Your whole life, girl. And if it means so little to you that you can just give it away, hand it to him, then why should it mean any more to him? He can’t value you more than you value yourself.” —Toni Morrison, *Song of Solomon*
- “Hello babies. Welcome to Earth. It’s hot in the summer and cold in the winter. It’s round and wet and crowded. On the outside, babies, you’ve got a hundred years here. There’s only one rule that I know of, babies—“God damn it, you’ve got to be kind.” —Kurt Vonnegut, *God Bless You, Mr. Rosewater*
- “...I think we are well-advised to keep on nodding terms with the people we used to be, whether we find them attractive company or not. Otherwise, they turn up unannounced and surprise us, come hammering on the mind’s door at 4 a.m. of a bad night and demand to know who deserted them, who betrayed them, who is going to make amends. We forget all too soon the things we thought we could never forget. We forget the loves and the betrayals alike, forget what we whispered and what we screamed, forget who we were.” —Joan Didion, *Slouching Towards Bethlehem*
- “But I did not know how to make my apology. The words that had strung themselves so easily to make a blunder in the drawing room would not come now that I wished the blunder remedied. I stood there below her window, tongue-tied and ashamed. Suddenly I saw her turn and stretch behind her, and

then she leant forward once again and threw something at me from the window. It struck me on the cheek and fell to the ground. I stooped to pick it up. It was one of the flowers from her bowl, an autumn crocus.” —Daphne Du Maurier, *My Cousin Rachel*

### **In-Class Lesson Plan #3 Option A**

\*Should follow Out-of-Class Assignment #3

#### *General Outline*

##### *Opening*

##### *Activity:*

As a class, students vote on a two of the writers from (Out-of-class Assignment #3) that they found the most interesting. Next, students vote on two random things they want to have written about - could be the cancellation of class, reactions to a recent event, or thoughts on a particular individual. Following these votes, the teacher generates four different paragraphs. They generate two paragraphs for each “random thing” and each of the paragraphs is in the style of one of the two writers.

*Discussion:* students using the opening activity and their own out-of-class writing to share their ideas about writing style. They then explain the writing style that the paragraph generated. They also explain their understanding of writing style.

##### *Activity:*

students go through their own writing and start to generate an understanding of their own writing style. (Teacher circulates around the room to help students).

*Activity:* students start revising their current essay draft to be more in their “own” writing style

*End-of-class:* teacher sets up follow-up work which will focus on revising essays to be in the students’ own writing style

### **Works Cited**

Aull, Laura. L. Bandarage, Dineth, and Richardson Miller, Meredith. “Generality in Student and Expert Epistemic Stance: A Corpus Analysis of First-Year, Upper-Level and Published Academic Writing.” *Journal of English for Academic Purposes* 26 (2017): 29-41.

Aull, Laura L. and Lancaster, Zak. “Stance as Style: Toward a Framework for Analyzing Academic Language.” *Style: and the Future of Composition Studies*. Eds. Paul Butler, Brian Ray, and Star Medzerian Vanguri. U P of Colorado, 2020. 98-113.

Lancaster, Zak. “Expressing Stance in Undergraduate Writing: Discipline Specific and General Qualities.” *Journal of English for Academic Purposes* 23 (2016): 16-30.

Meyer, Craig. A. “The Evolution of Imitation: Building Your Own Style.” *Writing Spaces: Readings on Writing* vol 3. <https://writingspaces.org/past-volumes/the-evolution-of-imitation-building-your-style/>.