

Chapter 1. The Current State of Artificial Intelligence in Writing

Research Ensemble of Humans, Robots, Texts, and Computers

As scholars and researchers, we increasingly find ourselves in collaborative spaces where humans and artificial intelligence (AI) tools work side by side. This section is a reflection on our experience of working within such an ensemble, composed of human researchers, undergraduate interns, and AI agents like Otter.ai, all of which contributed to a more streamlined, productive, and less tedious research process. Our purpose is to demystify the alien-seeming partnership with our virtual tools and normalize the technology by emphasizing its mundane use rather than highlighting its more exotic, emergent qualities. By leveraging the capabilities of AI to handle repetitive tasks, we were able to focus our attention on analysis, writing, and organizing the core elements of our work. An earlier version of this section is published as an experience report paper published by the ACM SIGDOC and presented at the 2024 conference (Salvo & Sherrill, 2024).

We assembled a collaborative research team comprising a senior researcher, a junior faculty member, and six undergraduates. Alongside our usual technology tools—computers, software, smartphones, tablets, internet access, email, student backchannels, scheduling apps, and file sharing—we used AI to generate transcripts. Each team member, including our AI collaborators, contributed uniquely. In particular, Otter.ai, an AI-powered transcription tool, played a key role by providing quick, rough transcripts of interviews. These transcripts weren't publishable but offered a reliable foundation for further analysis. By handling transcription, Otter.ai saved us considerable time, allowing more focus on interpretation. It belongs to a more traditional class of AI tools, distinct from emergent generative AI—a distinction discussed in the next section.

Our interviewees were compelling figures in their own right. But were they team members? Chosen for their experience with both traditional and generative AI, they had no stake in this text but generously offered time, energy, and access to otherwise unavailable insights. Their interviews affirmed and extended observations made by Michael and John, offering vital triangulation and strengthening our findings. While many points aligned with expectations, each conversation—including a final exchange between John and Michael—offered surprising, insightful turns. Perhaps the most fitting way to describe the interviewees is by their proximity to the center of concentric circles: human and technological, primary and secondary researchers, informants and collaborators. These distinctions highlight varied participant roles, yet all contributed meaningfully to the collective.

The novice researchers on our team took these initial AI-generated transcripts and, using Robert S. Weiss's *Learning from Strangers* (1995) as a guide, edited them into coherent, readable prose. The students learned to transform the raw transcripts into formal writing that stayed true to the spirit of the spoken word while making them readable as text, a task that required both human intuition and attention to rhetorical nuance. In short, while Otter.ai gave us a head start, it was the human researchers who refined and enriched the material, ensuring accuracy, engagement, and readability (or so we hope).

What we found in this process is that AI need not displace humans but rather enhanced our ability to focus on more creative and intellectually demanding work. In this regard, we are aligning ourselves with the concept of *Human-Centered Artificial Intelligence* (HCAI) as outlined by Ben Shneiderman (2022) as well as echoed here both in Bridget and Kate's interviews. AI tools can give researchers what Shneiderman describes as "superpowers"—amplifying our abilities to complete tasks faster and with less effort. The tedious work of transcription, for example, was completed much more efficiently by our AI partner, leaving our human team with the more satisfying task of crafting the research narrative.

This is not to say that the role of AI in research is without challenges or concerns. A prevailing worry in many academic and professional circles is that AI will replace human jobs, taking over the very work that defines our expertise. We, however, experienced the opposite. AI was a collaborative partner, not a replacement for human labor. The technology relieved us of some of the more monotonous aspects of the work—like transcription—without taking away from the core research tasks that required human insight. Instead, we were able to discuss word choice and replacement, argue over what the word meant on the recording, and try to decipher competing threads of language when participants spoke over each other or referred to shared knowledge not explicitly stated. Each is an example of content beyond the ability of AI to ponder, as pondering is quite beyond the robots' capacity. It's important to acknowledge that while AI tools might multiply our productivity, they do so under the guidance and supervision of human researchers. We, humans, remain the driving force behind the research agenda, making critical decisions, interpreting data, and generating the scholarly content.

This collaboration between humans and machines led us to reflect on the concept of technological determinism—the idea that technology drives progress, independent of human will or agency. In our work, we consciously avoided falling into this trap. While AI undoubtedly sped up certain parts of our research process, it did not drive the project. Instead, it supported us. The research ensemble—a collective of human and non-human agents—functioned because we designed it to, always with a human-centered focus. Without our direction and critical oversight, the AI tools would have produced transcripts, but they would not have resulted in readable prose.

One of the standout moments in our collaboration came from Teah, an undergraduate member of our team. Teah quickly emerged as a leader, using

AI-enhanced tools to support the research in innovative ways. Specifically, she worked with image-generation technologies to develop visual content that represented video interview data in still images. She produced the screenshot images that appear in Kate's transcript. The AI generated numerous draft images, each offering different visual possibilities, but it was Teah's human expertise that selected and refined the images for inclusion in the final manuscript. The AI, in this case, did not replace the labor of visual design but accelerated the process of generating alternatives, allowing Teah to concentrate on the more complex work of refining and editing the images to meet the project's needs.

Teah's screenshots remain an important touchstone when John and Michael discuss the role of AI on the research team. Kate had a wonderfully structured presentation prepared and walked us through important steps in her and her teams' use of automated tools in three different environments (see Kate's interview for more). As video, the presentation and steps she took helped us understand her teams' working processes. The challenge was to translate the expert presentation, following John and Michael's conversational cues, and determining what a different audience—the readers of this text, *you*—would need visually to comprehend the details of Kate's presentation. As a team, we watched and rewatched the relevant sections of the interview, determining key elements and advising Teah on details she might want to represent for clarity with team members commenting and asking for refinement and clarity. The AI tool enabled rapid prototyping and Teah sketched a number of possible redesigns, sharing each draft not as contenders for inclusion but as visual strategies to represent fluid video in static text. We think we succeeded in capturing key elements of Kate's expert presentation, but ultimately it is the reader's role to determine whether the final product conveys the concepts Kate, Michael, and John seek to share.

This experience reveals how AI can support human creativity, rather than undermine it. AI tools are not autonomous agents with independent desires or capabilities. Here, they are extensions of our human research teams, assisting us with tasks that would otherwise slow us down and focus attention elsewhere than on the research. The human-AI partnership we established within this research ensemble allowed us to direct our energy toward more meaningful and challenging aspects of the project. This is, in part, why our undergraduate team was able to remain productive and engaged throughout the process.

Our collaborative team—the mixture of interview subjects, senior and junior faculty, undergraduate researchers, and AI agents—was successful because of shared goals and collaborative, inclusive design. The AI agents handled the initial legwork, such as transcriptions, but their role was always to support the human researchers. As we worked, we were reminded of a critical distinction: AI tools, no matter how advanced, are limited by the instructions and intentions of their human creators, dependent on human direction to derive outcomes that hold significance for the humans directing action. The AI agents in our project were effective because we directed them and infused them

with our intellectual goals. In other words, the AI did not dictate the direction of the research; it facilitated it.

There is, however, a broader lesson to be drawn from this experience. The integration of AI into research ensembles has the potential to reshape how we work, but only if we maintain a human-centered approach. AI cannot and should not replace the uniquely human elements of research: interpretation, creativity, and critical thinking. Instead, AI offers us an opportunity to offload routine tasks, allowing us to focus on higher-order intellectual work. So long as humans retain control over the research process, AI can serve as a powerful partner.

The research ensemble we created represents one future of scholarly collaboration. The partnership between humans and AI allowed us to create a workflow that was not only efficient and effective but also deeply satisfying. By leveraging AI tools for routine tasks, we were able to concentrate on more engaging, intellectually stimulating work: should we conclude with pizza metaphors, and will the audience understand or care about John's struggles with reporting trail-blocking barriers, and do these metaphors illuminate complicated relationships between humans and our devices? Pondering these questions are worthy of our attention while formatting and transcribing audio to text shrinks back as a distraction rather than a burden. This experience highlights the potential of AI to enhance work. As we look to the future, we believe that AI will continue to play an increasingly important role in research, but its success will depend on maintaining a balance where human insight and direction remain at the center of the process.

Distinguishing Generative and Traditional Artificial Intelligence

Artificial intelligence (AI) has evolved dramatically over the past few decades, encompassing a variety of systems and applications. In a pair of 2018 *MIT Technology Review* articles, Karen Hao published two flowcharts to help individuals quickly identify when systems are using AI, and when AI is using machine learning (Hao, 2018a, 2018b). We similarly draw two distinctions here: A key differentiation within the realm of AI is between *generative AI*, which is designed to create new content, and *traditional AI*, which is primarily focused on analyzing and making predictions based on existing data. While both types of AI have transformative potential, they operate in fundamentally different ways, with distinct goals and methodologies. Understanding the differences between these two approaches is crucial for appreciating the full spectrum of AI's capabilities and its implications for various fields, including technology, art, business, and science.

Traditional AI, often referred to as *discriminative AI*, is grounded in the task of recognizing patterns and making decisions based on those patterns. This type of AI excels at classification, prediction, and optimization tasks, all of which are contingent upon large datasets that have been pre-labeled for specific purposes. Examples include facial recognition systems, credit fraud detection algorithms,

and recommendation engines used by platforms like Netflix and Amazon. These systems are designed to process input data, compare it against pre-existing patterns, and output a predefined result based on that analysis.

At its core, traditional AI is built around supervised learning methods, where a model is trained on labeled datasets. The AI learns to map specific inputs to their corresponding outputs, refining its understanding as it is exposed to more data. For example, a traditional AI model trained to detect spam emails will rely on a dataset of emails labeled as either “spam” or “not spam.” As the AI analyzes the features of these labeled emails (such as certain keywords, patterns, or sender behaviors), it learns to identify future emails that fit the characteristics of spam. Its purpose is not to generate new emails but rather to classify existing data according to learned rules.

This type of AI is task-specific and excels at automating processes that require repetitive decision-making. A well-known example is AI-powered diagnostic systems in healthcare, which can assist doctors in identifying diseases by analyzing medical scans. These systems have proven to be highly efficient at processing vast amounts of data and making accurate predictions based on known patterns. However, they are inherently limited in that they do not create new knowledge or content—they simply act on what they have learned from existing data. This is also referred to as machine learning.

In contrast, generative AI represents a more dynamic and creative form of artificial intelligence. Rather than simply identifying or classifying patterns, generative AI is designed to produce entirely new content. This new content can take many forms, including text, images, music, or even videos, and is generated based on the underlying structures and patterns the AI has learned from training data. Notable examples of generative AI include GPT-x (a language model capable of generating human-like text), DALL-E (an image generator that creates visuals from text prompts), and StyleGAN (a model that generates realistic human faces and other images).

The foundation of generative AI lies in techniques such as unsupervised learning or self-supervised learning, where the AI does not rely on explicit labels or predefined outputs. Instead, the model learns to recognize patterns and relationships within the data on its own. A common method used in generative AI is the Generative Adversarial Network (GAN), which consists of two neural networks—a generator and a discriminator—working in tandem. The generator creates new data, while the discriminator evaluates the data’s authenticity compared to real data. Over time, the generator improves its ability to produce increasingly realistic outputs, whether it is generating artwork, audio files, or synthetic datasets.

A significant aspect of generative AI is its ability to simulate creativity. For instance, when a generative language model like GPT-3 is trained on large corpora of text, it does not merely replicate what it has seen. Instead, it generates new text based on the patterns of language, grammar, and context it has learned.

This ability to create content that resembles human creativity represents a shift from traditional AI's analytical role. Generative AI is not limited to responding to input in predictable ways; it can generate novel outputs that may not have been explicitly programmed into it.

Fundamental Differences Between Generative AI and Traditional AI

- **Objective:** One fundamental difference between these two types of AI lies in their goals. Traditional AI focuses on analyzing existing data and providing accurate classifications or predictions. It is task-specific, aiming to improve efficiency in tasks like sorting, diagnosing, or recommending. Generative AI, on the other hand, is focused on creating new content, often attempting to mimic human-like creativity. Its objective is to innovate and generate outputs that are novel and previously unseen.
- **Methodology:** Traditional AI typically employs supervised learning, requiring large, labeled datasets to function meaningfully. Generative AI often utilizes unsupervised or self-supervised learning, which does not depend on labeled data. Instead, generative models are designed to uncover hidden patterns within the data, allowing them to generate new and original content. Techniques like GANs and transformer-based models (such as OpenAI) enable these systems to push the boundaries of creativity.
- **Output:** The output of traditional AI is often deterministic and predictable. Given a certain input, traditional AI models will provide a predefined outcome, such as identifying whether an email is spam or not, or determining if a medical scan shows signs of a disease. Generative AI, by contrast, produces novel outputs—text, images, or music that did not previously exist. This creative output is not a replication of the data it has been trained on but a synthesis of patterns that allow it to generate something new, at least probabilistically.
- **Use Cases:** Traditional AI is used in scenarios where repetitive, accurate decision-making and classification are critical. This includes applications like fraud detection, recommendation systems, or autonomous driving. Generative AI, however, is used in contexts where new content creation is the goal. It is employed in fields such as text and image generation, where AI creates articles, stories, or marketing copy, and in visual and aural applications, where it generates images or sound. Additionally, generative AI is being used in synthetic data generation, which helps in creating artificial datasets for training other AI models without exposing real, sensitive data. John and Michael have published on this question in usability and whether data generated through user participation can be automated. While our answer, in short, is no, the expense of usability testing will make synthetic feedback increasingly attractive, if ultimately rendering usability useless in an entropic feedback loop. See the Conclusion for more on information entropy.

Working with Both Traditional and Generative AI

The emergence of generative AI marks a milestone in the development of artificial intelligence. By enabling machines to generate novel content, we enter an era where machines can simulate creativity—a domain traditionally reserved for humans. The potential applications range from automating creative processes to generating synthetic data for training other AI systems. The industry claims this process preserves privacy because no discernable link can be made between the data the AI used to produce its personas and real people.

While AI appears to preserve privacy, the rationale sidesteps the issues—ethical and otherwise—regarding whether this data is authentic or reliable. Further, generative capabilities raise a host of ethical concerns and challenges. The ability to generate convincing synthetic content—such as deepfake videos or AI-written articles—raises questions about intellectual property, misinformation, and the ethical deployment of these technologies. It becomes essential to ensure that generative AI is used responsibly, balancing innovation with safeguards against misuse.

The primary distinction between generative AI and traditional AI lies in their objectives and methods. Traditional AI excels at analyzing and classifying existing data, while generative AI pushes the boundaries of creativity by generating new content. Both forms of AI are transformative in their own ways, but generative AI introduces novel possibilities by enabling machines to mimic human creativity, potentially revolutionizing industries from the arts to technology. As generative AI continues to evolve, it will be essential to explore its applications thoughtfully and ensure ethical guidelines are in place to manage its far-reaching impacts.

Distinguishing between different kinds of artificial agents is important: while new generative AI applications have captured the attention of the technology press, investors, and the general public, traditional forms of AI have been in unremarkable and mundane use for quite some time. The implication of creativity in novel expressions of language and image both make creatives nervous and promise access. In this text, John and Michael have been careful to point to promises of the new technologies, and we agree that the emergent agents open a new realm of creativity not available previously, and agree with Bridget that AI deserves the coverage it receives as a “big deal.” What we find ourselves concerned about, though, differs from many critical angles. We are less worried about power consumption: all electronic technologies require electricity to function and numerous investigative journalists are revealing some truly horrifying abuses, and we leave it to them and cheer them on. Another thread of criticism has been intellectual property, and we rely on Peter Schoppert’s important work cataloguing the ISBNs included in various language models. Interestingly, Michael was actually disappointed that the ISBNs of neither of his books were included, but the third wave included the edited collection. And with Routledge’s announcement that they have a collaboration with Microsoft, the monograph has been assimilated into the Borg (Dutton, 2024; Schoppert, 2023a, 2023b).

Latour used the novelist Richard Powers to explain *scientifiction*: not space opera or alien horror but science fact stretched just slightly into the future—plausible and based in emergent laboratory findings, stretched to imagine what *might* or *could* happen. Powers has written about genomic intervention, articulating a new Frankenstein’s Monster and imagined how trees communicate over eons. That last book, *Overstory*, won the Pulitzer Prize in 2019. And he has written about artificial intelligence and saving the planet from environmental devastation in 2024 by “colonizing” oceans. Powers famously revealed his use of speech-to-text technology only after winning the National Book Award in 2006 for *EchoMaker*, a book simultaneously about a fractured human mind and the equally fragmented environment. In 2007, National Book Award in hand, he described how he was no longer able to sit at a desk and type for hours every day. He was using technology—speech to text—to write. He didn’t want to be known as the dictating novelist, so did not reveal his technology usage profile until after winning the prestigious prize. He published a piece in the *New York Times Review of Books* in 2007 describing his process (Powers, 2007).

We have no delusions about winning prizes for our work, but did want to describe our technology practice, particularly how we are using AI, both traditional and generative, in the interests of transparency. Michael’s practice has changed through the process of writing this text, starting tentatively with Otter.ai for transcription. Besides creating the transliterated transcripts described elsewhere in this book, Otter.ai also provides summaries of the scripts. These summaries supported writing the timestamped overviews of the interviews, again reducing tedious tasks by creating structure and then we rewrote and revised the summary text. We started using the transcriptions to navigate the interviews, remembering a keyword or pithy phrase and then not being able to locate it in a timely manner, so the AI-generated summary became a tool for our own use. When we considered the reader and the accessibility of the transcript texts, we decided to leave them as part of the printed artifact. As reader, you will ultimately decide whether this is a helpful gloss accompanying the transcripts and, as an open-source text, having two copies, side by side, one open to the transcript and the other open to the summary may provide a further use case for scholarly open-source electronic texts. The most “generative” of AI use was in scraping text from existing slide decks and pasting them into AI prompts, and the chatbot compiled messy notes into a zero draft by putting related content together and stripping any formatting from the text. Editing and revision was all done by its human authors. This process is consistent with the descriptions of technology support and elimination of repetitive tasks in writing, while acknowledging that every word was not selected and typed by the credited authors. Like Powers before us, we confess to using technology.

Similarly, Michael and John used AI to help clarify long, complex sentences and organize confusing sections of text. While algorithmic interventions in the book were modest, it does signal a shift in the way we think about text and

the fear of the empty page, which is where most of Michael's writing frustrations arise.

What am I trying to say here? It is a question I ask myself during revision, and the literal prompt I used when interacting with the generative AI chatbot, asking the AI to locate subjects, verbs, and nouns, independent and dependent clauses, and to untangle subordinated clauses and sentences, to highlight fragments.

With something already on the page, the task of revision somehow seems less daunting. For John, generative AI is helpful in stitching chunks of writing together into the confines of linear text (writing for the web vs. writing for print). Though the generative output is often corny or trite, it can help reveal implicit connections that require signposting. Is it important whether the author is revising bad prose written by a human or that provided by an algorithm? As Michael asked his class, *how much* AI provided content is an ethical amount? Twenty-five percent? And what does it mean to have 25 percent or 50 percent of the content provided algorithmically? And what of the meme we've all seen online asserting: "why would I bother reading something no human could be bothered to write?" We wonder: Are these questions that arise because generative AI assistance is new, and will these questions fall away as we generate a new sense of normalcy with our AI tools? Right now, because of hype and marketing, the most mundane of assistive applications is labelled AI. If there is a backlash, will the marketing change without altering the assistive component? All these questions remain asked but unanswered.

Meanwhile, John has discussed how he has used AI to generate routine, repetitious, and tedious communications with students and to address institutional demands. He has described the value in using the tireless robotic eyes to edit dozens of fraught end-of-semester emails with students without concern that a tired human would miss an unedited reference to someone else's name or a simple mistake that would allow a student to make an appeal based on a claim of an erroneous deadline or mathematical error. At a time of semester when all participants are tired, stressed, and error-prone, generative AI can help facilitate clearer, more empathetic communication. Similarly, John has used generative AI to quell the tedium of writing point-based rubrics—institutionally required not only for every assignment, but even for daily class activities, which in other institutional contexts would commonly be graded as complete/incomplete. Manually establishing clear performance criteria for each category of a rubric, John still does the rhetorically and pedagogically nuanced work of aligning grading criteria used to evaluate assignment outcomes with the respective assignment goals and means. John manually drafts a baseline description of excellent student work. But ChatGPT can quickly adjust the wording of descriptions within each category to fit performance levels of "Excellent," "Competent," or "Needs Improvement," requiring John to make just

minor revisions. For example, compare the following lines from a longer rubric section evaluating the organization of proposals:

Competent: Proposal is generally organized for readers and follows a logical progression of ideas, but may have some issues with navigation or skimmability.

Needs Improvement: Proposal is poorly organized and impedes skimming or easy navigation.

Such interventions do not replace expertise in education or writing pedagogy, but maximize the AI's labor-saving potential, allowing John to concentrate on the writing of new parts of this manuscript, mentally and emotionally complex aspects of teaching, and other related and unrelated research.

We imagine scenarios. We report on existing collaborations. And we gather available means of persuasion, some of which we might have overlooked without effective, if synthetic, partnerships. The integration of both traditional and generative AI into academic and professional workflows demonstrates the potential for effectiveness and innovation. Traditional AI, with its ability to automate routine, rule-based tasks, complements the creative capabilities of generative AI, which produces new content based on patterns and data. Together, they form a hybrid model that supports educators, researchers, and professionals by streamlining administrative duties and enhancing productivity. Traditional AI tools are invaluable in checking emails, proofreading documents, and providing drafts of timely feedback, as John experienced when addressing student communications. The benefit of such tools lies in their ability to reduce errors, save time, and enable more focused intellectual engagement, as well as to create reflective moments when humans can rest weary brains and let artificial agents grind through routinized tasks.

Meanwhile, generative AI offers rough new material, ideal in early ideation in research to the initial drafting of content, fostering creativity that might otherwise be hindered by more time-consuming tasks. The thoughtful use of both systems enhances human potential rather than replacing it, allowing educators to focus on high-level tasks such as pedagogy and research development: infrastructural decision-making. As we consider these tools' emerging roles, it's crucial to evaluate their effectiveness in balancing repetitive tasks with the pursuit of innovative, thought-provoking work. This synergy between traditional and generative AI serves as a promising model for how technology can assist rather than detract from academic and professional growth.

Responsible AI Use in Writing Courses

In the discussion of student experiences with generative AI, classroom insights offer a valuable perspective on students' expectations and usage of AI tools. We provide some of our experiential insights here. Our survey findings and

observations align broadly with classroom observations shared by many in the field: while most students are familiar with the concept of generative AI, fewer have engaged with it. The policies governing AI use in academic settings remain inconsistent, ranging from outright bans to vague warnings, with Penn State's extensive site the clear exception, for which we are grateful in part to Stuart Selber's efforts to push the institution to articulate its reliance on and interwoven relationship with development of AI (The Pennsylvania State University, n.d.). Students tend to utilize AI primarily for editing and smoothing text, but seldom for substantive revisions or text improvements. Interestingly, both students and faculty express shared concerns about plagiarism and the fear that reliance on AI might reduce the depth of student learning. Recent technical communication research suggests that students use LLMs in ways similar to how Kate describes her use of Writer in Chapter 4, as a helpful brainstorming and editing tool. Aligning with our own observations and studies, Gustav Verhulsdonck and Jialei Jiang reported that students used AI-based Figma plugins to brainstorm ideas, rapidly prototype, and review design checklists in a user experience (UX) design program (2025); Aditya Johri et al., described similar uses of AI tools in the workplace to generate ideas and refine output (2025); and such experiences resemble Teah's use of AI tools to efficiently prototype images for this book.

Furthermore, the ethical considerations around AI use are being actively discussed. Students are cognizant of the implications and express a collective anxiety about the possibility of academic dishonesty. Many fear that some students exploit AI to complete assignments, undermining the integrity of those who labor without such assistance. This frustration was well articulated by students who did not want to feel like the "rube" working diligently while others leveraged AI to "automagically" produce well-formatted papers with minimal effort.

Bits and Bots: Students Using AI

The 49-page report titled "Bits and Bots: A Guide to Ethical Artificial Intelligence, by & for Students" (Blunt, et al., 2023) was the product of an advanced professional writing course Michael recently taught. In this class, students engaged with large language model (LLM) AI tools to critically analyze and reflect on the ethical implications of AI-generated text. The report is set to be archived in Purdue University's digital library, and we look forward to it becoming publicly accessible. Once available, it will offer insights into students' collaborative work with AI, allowing readers to trace its impact (Blunt et al., 2023).⁵

Jeffrey was the first student to push the boundaries of AI use within our classroom, prompting ChatGPT to write in the first person—a result that surprised

5. The Bits & Bots report is available upon request: <https://asco3d.lib.purdue.edu/repositories/2/accessions/2773>. There is a large "Request" button in the upper-right hand corner. The Purdue Libraries trace downloads.

both him and his classmates. Prior to working on the report, Michael asked students to produce three to five pages of writing within two weeks, with one catch: they had to “collaborate” with AI to varying degrees—ranging from 0 percent to 100 percent. Jeffrey, who was assigned a 100 percent AI collaboration, initially expressed concern, wondering if the assignment was a trap. However, after reassurances, he embraced the challenge and interacted with multiple AI platforms including ChatGPT, Jenni.ai, and writer.com, exhausting their free trials over a weekend. His most notable achievement came when he prompted ChatGPT to rewrite his draft in the first person, an outcome that no one in the class had anticipated. His classmates estimated that about 50 percent of his text had been AI-generated, though the true figure was 100 percent. Zero percent was the baseline of complete authorial control and revision—bespoke text, and what authors created prior to the availability of generative AI tools. The class discussed what it meant to use 25 percent as a measure of AI text production, defining it as providing minimal structural help but adding, changing or drafting a quarter to a third of the text. The class defined the 50 percent threshold as the human author providing some minimal parameters and making changes to sentence structure and order, and leaving some of the AI-provided sentences largely intact surrounded by human-authored text. At 75 percent, students defined the text as largely generatively algorithmically structured with human authorial interventions where the provided text was unclear, vague, awkward, or just wrong (as described in the description of an AI “hallucination” below). Jeffrey was working at 100 percent AI generated, which he and others defined as having no human intervention besides providing chatbot prompts, and this usually resulted in the least readable and least convincing texts. However, Jeffrey’s persistent re-submission of outputs as inputs, moving from platform to platform, and the addition of asking for the output in first person resulted in an awkward but universally mislabeled “human” production. This mislabeling indicates there are ways AI produced texts evade detection. As the technology improves, predictions are that errors and erroneous information—hallucinations—will likely decrease, which has already been borne out in the results of ever-increasingly sophisticated output from freely available AI chatbots.

Another remarkable incident occurred when Olivia discovered an AI hallucination while working with ChatGPT. The AI erroneously claimed that a medieval French philosopher named Franciscus Niger had been the first to document the myth that the Moon is made of cheese. Despite Olivia’s diligent research, she was unable to find any credible reference to such a figure. Franciscus Niger, it turned out, was a fictional character generated by the AI.

Olivia, an assiduous and ethically driven student, invested significant time into probing ChatGPT for inconsistencies, and she ultimately uncovered this significant error. Although she initially expressed reluctance to use AI for academic work, she embraced the opportunity to expose its limitations. Her experience mirrors broader concerns expressed by students like those in John’s forthcoming

study—students who are apprehensive about generative AI’s potential to facilitate academic dishonesty while simultaneously undermining their own hard work.

Olivia’s encounter with the Franciscus Niger hallucination sparked a deeper inquiry into AI reliability, prompting students to question where else the AI might have gone wrong. Olivia’s unique perspective contributed significantly to the class discussions, and her proactive approach led her to contact the Purdue archives about preserving the class report. On the last day of the semester, just as the class was winding down and reflecting on its experiences with generative AI, an email from the archivist confirmed the report would be archived—a testament to Olivia’s initiative, as well as the interest in recognizing 2023 as the year of generative AI’s emergence as an important disruptive technology, worthy of capturing students’ concerns and attitudes towards the new *technê*.

The report is available through Purdue’s digital archives, allowing both on-campus and global access. The document will be persistently findable, ensuring its contributions to discussions on ethical AI use remain accessible to a wide audience. As we reflect on this age of (mis)information, we must ask ourselves: does something truly exist if it cannot be found?

History Redux: AI & Rhetoric Intertwined

Rhetoric, artificial intelligence, and automation have been deeply intertwined from the late 20th century into the 21st. The roots of rhetoric—an ancient *technê* for communication, argumentation, and persuasion—are closely linked to mechanistic views of language, a connection explored by Walter Ong in his studies of Ramus (1958, 2004) and the shift from medieval to early modern rhetoric. More recently, Lynette Hunter’s “Rhetoric and Artificial Intelligence” (1991) traced the rhetorical grounding of AI, positioning it as a technological resource integral to persuasion. Hunter’s historical analysis demonstrates rhetoric’s long-standing relationship with AI, independent of the distinctions between functional rhetoric and its romanticized forms.

Technology, especially digital technology, has profoundly shaped rhetoric and literacy. Since Ong’s *Orality and Literacy* (1982), scholars have examined media’s implications for argumentation and persuasion. Hunter’s work provides a historical foundation for understanding rhetoric and AI at the dawn of the networked digital age. She explores how rhetoric transformed during the 16th and 17th centuries, when logic was separated from rhetoric, leading to an emphasis on pure reasoning (1991). This shift pushed rhetoric toward mere ornamentation, yet it remained essential for persuasion. Ong’s studies of Ramus remain instructive here.

Hunter argues that AI’s development, grounded in formal logic and heuristic procedures, mirrors Aristotle’s distinction between logic and dialectic. While AI has significantly advanced rhetorical strategies, it has largely overlooked “stance”—the dynamic relationship between rhetor, audience, and text (1991).

Modern science, with its focus on rationality, often frames rhetoric as unnecessary, aligning with Aristotle's "demonstrative" argument, which dismisses persuasion in favor of self-evident reasoning. Empirical bias isolates rhetoric from its social and contextual foundations.

AI's focus on problem-solving and knowledge representation within isolated systems limits its ability to engage with context—central to rhetoric. By prioritizing exact representation over persuasion, AI denies the necessity of rhetoric in navigating uncertainty. Instead, Hunter suggests that AI should acknowledge its limitations and contribute to evaluating the multiple realities shaped by modern technology. Ultimately, her work underscores the importance of integrating rhetorical awareness into AI, allowing for a more nuanced understanding of its role in shaping discourse. Recognizing rhetoric's complexity may better equip AI to address ethical and communicative challenges, or will continue to define its limitations.

As early as 1991, scholars recognized context as key to AI's development. That demand for contextual understanding remains, yet generative AI in 2025 performs remarkably well despite limited context. The intersection of Platonic and Aristotelian rhetoric with AI expectations reveals rhetoric as *technê*—a structured, almost algorithmic set of principles designed to persuade. AI's reliance on probabilistic word prediction exposes rhetoric's foundational mechanisms: increasing the likelihood of convincing an audience of credibility, authority, and emotional connection, even when the "author" lacks corporeal existence. One can almost hear Roland Barthes' (1967) sardonic laugh in response to his famous essay *The Death of the Author*.

Attila Hallsby's "A Copious Void: Rhetoric as Artificial Intelligence 1.0" (2024) moves beyond the "stochastic parrots" critique (Bender et al., 2021) to examine the conceptual void between human authorship and AI-generated text. Hallsby argues that rhetoric and AI share core concerns: managing information overload, addressing social inequalities, and mitigating biases. Despite their historical separation, both disciplines navigate similar challenges.

The shared terminology of "stochastic" and "artifice" further highlights their connection. In ancient rhetoric, "stochastic" referred to probabilities and everyday occurrences, while in AI, it describes the structured randomness of algorithms. Similarly, "artifice" encapsulates both skilled invention and deception, reflecting AI's dual role as a tool for creativity and misinformation. This overlap underscores the ways AI, like rhetoric, relies on probabilistic construction and strategic persuasion.

Hallsby introduces the concept of "zero-agency" to describe AI's oscillation between presence and absence in rhetorical action. Such action parallels how rhetorical agency emerges not solely from an individual but from the interplay between text and context. AI can empower or disempower, manipulating data and representation. It extracts information from marginalized communities without consent while simultaneously excluding their vernaculars from AI

models. Agency, in this sense, becomes a function of access and control.

The article further argues that rhetorical tropes function similarly to algorithms: both establish rule-like structures that generate meaning. *Copia*, the rhetorical trope emphasizing abundance, reflects AI's vast data processing capabilities—both illuminating and overwhelming. Hallsby's analysis positions rhetoric as an underlying framework within AI, a conceptual bridge for understanding its potential and limitations. The "copious void" metaphor captures AI's dual nature: an expansive source of possibility and a space of overwhelming complexity.

Hallsby's work benefits from 43 years of technological advancements since Hunter's analysis, particularly the widespread adoption of generative AI chatbots. However, to fully explore agency, we turn to Carolyn R. Miller's "What Can Automation Tell Us About Agency?" (2007), which examines rhetorical agency in the context of automation. Positioned between Hunter's and Hallsby's work, Miller's research offers crucial insights into how automation reshapes rhetorical action.

Miller explores the tension between human rhetorical agency and automated assessment systems in education. While such systems are praised for efficiency and consistency, they struggle to capture the nuances of human communication, including creativity and emotional resonance. She argues that rhetorical agency is not an inherent trait of the speaker or writer but emerges dynamically through interaction with an audience. Agency, she suggests, functions as the "kinetic energy of performance," reliant on mutual attribution between rhetor and audience.

Surveying instructors of writing and public speaking, Miller found widespread skepticism toward automated assessment's ability to evaluate communication effectively. Instructors emphasized the importance of a live audience, particularly in public speaking, where engagement is central to rhetorical agency. She contrasts writing and speaking along performance and interaction dimensions: speaking is inherently interactive, requiring real-time feedback, while writing is temporally dislocated.

These insights raise critical questions about whether agency can be attributed to machines. Automated systems challenge our traditional understanding of rhetorical agency, pushing us to reconsider how meaning and persuasion emerge. Miller ultimately concludes that agency, though constructed, is essential to meaningful work. It is an attribution granted by one agent to another, carrying moral and pedagogical significance. Recognizing co-construction, educators must remain committed to fostering rhetorical agency, even (or *especially*) in an age of automation.

Miller's exploration aligns with the broader concerns of this book: the ideological construction of agency, its evolving meaning, and its role in defining work. Autonomy remains a crucial element of meaningful labor—determining what work is valued, how it is accomplished, and what tools are deemed necessary. As AI reconfigures work, these questions become increasingly urgent.

Alain de Botton offers a striking historical contrast between coercion in early industry and contemporary forms of motivation. In the past, productivity was

enforced through brute force. By the early 21st century, however, many jobs relied on satisfaction rather than obedience:

In the earliest days of industry, it had been an easy enough matter to motivate a workforce, requiring only a single and basic tool: the whip. Workers could be struck hard and with impunity to encourage them to quarry stones or pull on their oars with greater enthusiasm. But the rules had had to be revised with the development of jobs—by the early twenty-first century comprising the dominant sector of the market—that could be successfully performed only if their protagonists were to a significant degree satisfied rather than resentfully obedient. (de Botton, 2009, p. 32)

Creativity, initiative, and ambition are integral to meaningful work. The historical evolution of agency illustrates that while agency is socially constructed, it is no less real. Understanding its changing definition over time is essential, particularly as AI challenges long-held assumptions about autonomy and decision-making.

As workplaces become increasingly automated—not just in physical labor but also in mental and organizational processes—many who once believed themselves impervious to automation now face uncertainty. These concerns are not unfounded; AI-driven mass layoffs are already occurring. However, the best workplaces will integrate AI thoughtfully, just as digital technology in the 1990s and 2000s reshaped productivity and engagement. AI will redefine meaningful work, just as past technological revolutions did.

Hunter, Hallsby, and Miller document rhetoric's deep engagement with automation and AI, revealing its long-standing role in shaping technological discourse. Their work underscores the necessity of continuing to advocate for meaningful work in an AI-driven world. Feenberg's *right of refusal*—the rejection of certain technologies—remains a powerful stance. Yet outright refusal is an abdication of engagement. To ignore AI is to forfeit before the game begins. The challenge ahead is not just resisting automation but shaping it to preserve and enhance rhetorical agency—however constructed, partial, incomplete, and alienated our labors are in the latest round of late capitalism's post-industrial shenanigans.