Thinking Smarter, not Harder? Google NotebookLM's Misalignment Problem in Education

Christa Albrecht-Crane
Department of English
Utah Valley University
Orem, UT, USA
christaa@uvu.edu

Abstract

This paper examines Google's NotebookLM as a case study of how consumer-facing generative AI technologies misalign with educational values and user needs. While marketed as an "AI-powered research assistant," NotebookLM exemplifies the disconnect between AI industry promises and actual capabilities. Through technical analysis of large language model mechanisms, this paper reveals how NotebookLM's statistical compression methods fundamentally differ from human cognitive processes of reading and analysis. The paper argues that despite claims of source-grounding, NotebookLM produces outputs that compress rather than comprehend texts, often missing crucial arguments and generating confabulated content. Drawing on examinations of "smart" technology rhetoric and extreme usability design, the analysis demonstrates how the tool's frictionless interface obscures computational limitations while potentially undermining cognitive development. The paper concludes by advocating for critical AI literacy in writing studies and technical communication, proposing pedagogical approaches that demystify AI hype and preserve the essential friction necessary for meaningful learning.

CCS Concepts

• **Human-centered computing** → Human computer interaction (HCI); HCI design and evaluation methods; Usability testing; Human computer interaction (HCI); HCI design and evaluation methods; Walkthrough evaluations; Interaction design; Interaction design process and methods; User interface design; Interaction design; Interaction design process and methods; User centered design; Human computer interaction (HCI); Interaction paradigms; Natural language interfaces; • **Computing methodologies** → Artificial intelligence; Natural language processing; Information extraction; Artificial intelligence; Natural language processing; Natural language generation.

Keywords

extreme usability, usability, NotebookLM, AI chatbot interface, cognitive offloading, critical AI literacy, AI summarization, large-language models, misalignment, education, technical communication

SIGDOC '25, Lubbock, TX, USA

ACM ISBN 979-8-4007-1444-3/2025/10
https://doi.org/10.1145/3711670.3764628

ACM Reference Format:
Christa Albrecht-Crane. 2025. Thinking Smarter, not Harder? Google NotebookLM's Misalignment Problem in Education. In *The 43rd ACM International Conference on Design of Communication (SIGDOC '25), October 24, 25, 2025, Lubbock, TX, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3711670.3764628

1 Introduction

The current artificial intelligence (AI) boom presents magnificent economic and productivity promises while implicating our students and educators in complicated and concerning ways. The tech industry's claims about AI tools do not match what these tools actually do or how they affect education. One the one hand, for instance, a recent article in The Journal of Marketing recommends that companies and marketers expand the adoption of AI-based products by targeting naïve users who do not know much about the technology and feel awe about its purported capabilities, therefore being highly receptive to using it [1]. The authors note that "businesses may benefit from targeting those with lower AI literacy, designing products to meet the needs of this target segment, and tailoring their marketing messages to highlight the perceived magicalness of AI technology (p. 2) [1]. And just in time for final exams in May 2025, OpenAI and Anthropic simultaneously launched efforts to capture such naïve users, namely college students, offering free access to their premium AI services [2]. On the other hand, it's becoming clear that overreliance on generative AI tools risks delegating users' thinking tasks to generative AI tools. For instance, a study from Microsoft's research division calls attention to the perceptions of cognitive skill loss among professional knowledge workers who use AI tools in their workflows [3]. Similarly, in educational settings, research suggests a negative relationship between the frequency of students' use of AI tools and their critical thinking abilities, a relationship "mediated by increased cognitive offloading," particularly pronounced in younger, less-educated students (17-25 years), who reported substantial dependence on AI tools [4].

These introductory observations illustrate the mismatch between the objectives of consumer-facing generative AI technologies and their mismatch with the goals of education. I will examine a particular LLM-powered application, Google's popular NotebookLM, to argue that this program is symptomatic of a broad push of AI products representing an ideological value system which compromises educational values, including in our work as writing scholars and teachers. NotebookLM stands as a particularly instructive example of the dominant "intelligent interactive machine" [5] exploiting the intelligence angle that has fascinated computer science and the popular imagination for a long time. But as I argue in this article, it turns out that what these programs actually deliver does not match the alluring claims; NotebookLM's underlying computational mechanism remains misaligned with the tasks it wants users to perform (especially when considered in educational settings). My argument contributes to the important project in technical communication pedagogy to analyze the integrated human skills, competencies, and habits of mind needed as AI technologies proliferate [6].

2 What Is NotebookLM?

Marketed as an "AI-powered research assistant" [7], NotebookLM's popularity centers on its alleged ability to summarize and engage with user-uploaded source texts, in a process called "source-grounding" [8]. Users can upload PDFs, Google Docs, websites, YouTube videos, or Google slideshows. The three-panel interface layout reflects conventional research software design, emulating a three-step process of gathering materials, analyzing them, and presenting the findings; thus, on the left of the program's interface, the Sources panel serves as a document repository; the central Chat panel provides a conversational interface for "analysis," and the third panel on the right, aptly called Studio, generates derivative content that is entirely pre-determined, including study guides, briefing docs, timelines, FAQs, and most mesmerizingly, podcast-style audio conversations of uploaded documents between two synthetic hosts. This interface design, as Tiffany DeRewal notes, "encourages users to imagine that the tool is only working with selected sources, and that it progressively achieves an increasing depth of engagement with those sources" [9]. Initially launched in 2023 under the name "Project Tailwind" [10], as of July 2025, a basic version of the program is available at no cost for Google account holders, and an upgraded version is sold as part of a Google One AI Premium subscription [11]. The audio overview feature is also available as a stand-alone experimental app called Google Illuminate [12].

The promotional language on NotebookLM's landing page lauds the program's benefits as a research assistant, where the chatbot presumably assumes human-like cognitive qualities and responsibilities. Until May 2025, the main tag line on the landing page read, "Think Smarter, Not Harder" (with the subtitle "The ultimate tool for understanding the information that matters most to you, built with Gemini 1.5") [13] and has since been replaced with the phrase, "Understand Anything" (with the subtitle "Your research and thinking partner, grounded in information you trust, built with the latest Gemini models") [14]. NotebookLM's target audience appears to be college students, which is dramatized by the example animations on the landing page, showing course numbers and titles of a literature course (featuring one of the most notoriously "difficult" literary texts in the English language, James Joyce's *Ulysses* [15]). This clever marketing strategy hits on pressure points in the public imagination, enticing potential student users to see a computer program as a research assistant that can, finally, save them loads of time, provide reprieve from drudge homework, and make accessible impenetrable texts, thus supplying what the teacher wants–i.e., the text's meaning and outputs toward a research paper with the required analysis and documentation.

Many enthusiastic users, such as the two writers of the Substack "AI Supremacy," celebrate NotebookLM as "a game changer for education" [16]. The popular science writer Steven Johnson, who collaborated with Google on designing the application, touts the program's promise as a Socratic, interactive learning tool [17]. He also lauds NotebookLM's source-grounding principle, positing that it's "capable of answering questions and tracking down references based entirely on the source materials you've uploaded" [17]. In his advice blog for marketers, Tom Lawrie proclaims, "You can ask questions, to summarize and pick out key points across source material. It's a mindblowing way to learn faster, organize your knowledge or use it to generate new ideas. The really cool shit is what you can do once you've established this knowledge base" [13]. Clearly, these perspectives position NotebookLM not just as useful, but as transformative through their strategic claims that the program excels across multiple dimensions.

Two specific elements are remarkable in Google's marketing and design choices: first, the rhetoric of "smart" technology, and second, the extreme usability principle that's used across digital consumer products of low-friction interfaces. First, the landing page's tag line, "Think Smarter, Not Harder," connects to a broader pattern in tech marketing that deliberately cultivates a mystification about the inner workings of digital products. As David Berry explains, the "smartness" lies not in making users actually smarter, but in making them not need to understand how smart devices work. This creates a paradox: the less you know about how the technology operates, the "smarter" it appears to be. For Berry, "ignorance of computational processes is, under this epistemology, celebrated as a means to an end of smartness" (p. 219) [18]. That is to say, simply using NotebookLM is a smart move in itself. Using this "smart" program means we're being smart.

The strategy of smart technology connects to the second element, which is the design philosophy of usability and user experience (UX). Traditionally, UX designers have advocated accessibility, consistency, clarity, user control, and error prevention, privileging functionality without frustration or interruption. The foundational UX principle is, "everything starts with the end user" [19]. However, a decisive shift has occurred since WWII in the deployment of consumer products, bringing user-centered concerns into tension with profit goals. As Bradley Dilger explains, rather than foregrounding the user's experience, designers and product developers have increasingly exploited UX principles to prioritize profits through transactional uses of digital applications [20]. The design techniques for this goal center on "extreme usability, meaning that products need to exhibit almost no friction, intuitive interfaces, convenience, and speed [20]. Nearly frictionless online interactions have become the norm in UX design, often directly aiming for the psychology of addiction [21]. Connecting to the rhetoric about smart technology, extreme usability design creates an online milieu in which users consume AI-powered products as commodities with minimal friction and easy, enjoyable interfaces [22]. And above all, users are not offered any insight into how these systems actually operate and what their abilities and limitations are.

3 LLM Mechanism and Appropriate Use cases

From its inception in the 1940s, computer science has had a fascination (if not obsession) for computational methods to emulate and outperform the human brain. As Thomas Haigh puts it, "AI was born in hype" in 1955 (p. 35) [23]. The biological metaphor of "neurons" became the founding lens of the field through which to articulate computational logic [24]. This alluring space of brains and intelligence dominates the computer science and pop culture imaginary to this day, with terms like "memory," "deep learning," or "reasoning" occupying the public imaginary.

Until neural networks took off in the 2010s, computer engineers used a "symbolic" approach to language processing in which they manually programmed a system with grammatical rules and syntactic relationships to enable the software to parse English sentences [25]. For example, when I type, "The small child walks into the room," the software breaks down this sentence using formal grammar rules into components like [action: walks] [subject: the child] [location: the room]. The computer can create outputs that strictly mirror the encoded rules. Use cases were and still are basic question/answer systems, word search capability, machine transcription and language translation software. Symbolic systems are accurate and dependable within their hard-coded rule architecture, and they are transparent, as their code is explicitly written and verifiable [25]. But they require lots of manual coding, are unable to deal with real-world linguistic material that was not programmed into their software, and they struggle with language ambiguity [25]. For instance, a hard-coded program has trouble deciding how to accurately transcribe the spoken words "red" or "read" in the following two sentences: "The girl picked up the **red** ball," and "The girl **read** a book." Because of these drawbacks, symbolic language processing stagnated until the 2010s, when several developments occurred simultaneously: a dramatic increase in computational power through fast, parallel chips; availability of massive human data sets from the Internet; the creation of large investment capital by a few large tech companies; and the development of a novel neural network architecture called the transformer [26].

Developed in 2017 by a team of eight Google machine translation researchers [27, 28], the generative pre-trained transformer (GPT) can compute statistical probability scores during training to provide a good guess about which word is used in the two different linguistic contexts [29]. Transformers don't need explicitly programmed rules but can observe from their training data that, in the example used above, "picked up" is statistically more likely to be followed by the word "red" than the past tense verb "read." Interconnected processing layers contain computational units ("neurons") that observe linguistic patterns from large sets of examples by adjusting the strength of connections between them, allowing the network to develop internal mathematical representations (i.e., statistical regularities) of a language [30]. What's important to capture here is that these machine learning efforts were narrow in their immediate utility, favoring probability scoring as the decisive metric of linguistic success [25]. But large tech companies took this computational lens and applied it wholesale to many other language and communication acts through their commercial products.

Since November 2022, when OpenAI released its first generative AI chatbot, these statistical systems have been presented as intelligent replicas in communication scenarios. These systems "reduce intelligence to that which their machines can do and then claim their machines are intelligent" [31]. Statistical GPT systems can process huge amounts of language data but are fundamentally not built for precision; they sometimes miss the "correct" next word in what is called a "hallucination," as if the error is an aberrance in an overall effort to be correct (p. 217) [32]. Hallucinations remain endemic to LLMs' statistical modeling of language. OpenAI researchers express exactly this predicament in their 2020 paper when they explain that "the language modeling objective used for many large LMs–predicting the next token on a webpage from the internet–is different from the objective 'follow the user's instructions helpfully and safely.' Thus we say that the language modeling objective is misaligned" [33]. A research team at Anthropic also notes that LLMs are not intrinsically aligned with users' needs for "helpful, honest, and harmless" results (p. 3) [34]. To help mitigate this consequential misalignment, tech companies employ fine-tuning methods and reinforcement learning through human feedback (RLHF) techniques, which are time-consuming and highly exploitative of human labor [32, 35]. In the end, these techniques cannot alter the underlying statistical mechanism that powers GPT systems.

4 LLMs and Summarization Tasks

Strictly speaking, a large language model like Gemini, which powers NotebookLM, does not summarize a text. Instead, LLMs use statistical calculations to compress texts [36]. Source texts are essentially shortened to their most frequently used words, using extractive techniques in which the model employs a statistical analysis to calculate the most important sentences in a given text and pulls them, often verbatim, rearranging them into a condensed, coherent-sounding output [37]. Extractive methods "are not capable of identifying the underlying contextual meaning or the additional information that is often implicit in the text, thus providing summaries that are not very detailed" [37]. To complement this rather mechanical process, abstractive methods have been added. They use the transformer technology to pull from patterns a language model has learned during training, computing longer-range dependencies across a text. [37]. The problem with the abstractive method is that it has to use a language model's entire system, and as a result, the output can include too much content pulled from the training data rather than the source text, resulting in "hallucinations," i.e., content that does not exist in the source text [37]. Any new output will be the result of an unpredictable and hard-to-control combination of the LLM pulling from the immediate context and the larger parameters of the training data. Summarization is never the result of an LLM seeing information solely from a source text.

To improve the accuracy of summarization tasks, a technique called Retrieval Augmented Generation (RAG) was developed by Meta researchers in 2021 [38]. To make sure that its next-word predictions draw in some way from specific texts, the RAG technique adds access to external content through a dynamic retrieval mechanism. Based on the words of the prompt, the system searches and retrieves what it considers relevant information from source texts, and only those segments will be provided to the LLM to help generate its output [38]. That is to say, the LLM only has access to the segments that are returned based on word relevancy. All in all, RAG does provide an improvement of the data sets an LLM sees and processes into a summary, but the basic limitations of transformer-based text generation remain. As Miriam Reynoldson puts it in a comment on her blog about RAG, "once we understand there's no magic in the base model, they just tell us there's magic powder in the wrapper" [39]. RAG cannot eliminate LLMs' confabulation problem.

NotebookLM uses RAG to augment the Gemini language model [40]. But to say that it tracks "down references based entirely on the source materials you've uploaded," as Steven Johnson claimed on his blog, is inaccurate. NotebookLM's context window is 1 million tokens per query and 2 million tokens across 50 sources [41], with 100 tokens being roughly 75 words; this translates to about 1,500 pages per query. In theory, the program is capable of summarizing extremely long texts; however, as we just saw, the RAG technique allows only chunks of text from uploaded sources that match the user's prompt to be used. The inline citations NotebookLM provides in the Chat panel "are mere pretenses," as they are guesses about what words or sentences might be relevant from the chunks of text the program can see [9]. In this way, then, LLMs and RAG systems do not "read" text in some sort of holistic way as they don't even have access to the text themselves. A language model's summary is pattern-driven, compression-focused, and programmed for statistical optimization. It is not reasoning, or thinking, or being intelligent [42, 43].

5 NotebookLM Summarization analysis

Despite enthusiastic endorsements and hyperbolic claims that NotebookLM can serve as a research partner, it is essential to clarify this program's affordances and limitations as a summarization tool. Here are a few significant aspects:

- NotebookLM provides statistical word compression of source texts that sometimes misses a text's actual arguments [36]. One very unsettling example of such a "misreading" occurred in my own experience with NotebookLM, when it misconstrued Dilger's argument about extreme usability, confidently claiming that the author argued for "the growing importance of usability and suggests that extreme usability is crucial for navigating the complexities of modern technology and ensuring effective communication." Even my follow-up prompt to include the critical elements of Dilger's argument did not dissuade the model from its confident declarations about Dilger's embracing the concept of extreme usability. Other researchers have reported similar problems [9].
- Another layer in NotebookLM is the extensive use of invisible system prompts, which developers include to steer the model towards providing output that better matches what users are told to expect [9]. Thus, all the options in the Chat window, such as automatic summaries and all the pre-designed artifacts in the Studio window, are powered by extensive instructions that run in the background. Users cannot see the elaborate instructions, nor can they turn them off. These unseen system prompts also make outputs more susceptible to confabulation because they add substantial additional parameters to prompts. Users report that audio overviews can include made-up content [44] and show gender stereotyping [45]. Reddit threads [46] and NotebookLM's Discord page [47] address a substantial number of user-reported issues.
- Summaries remain limited in their capacity to synthesize source content and provide real analysis. NotebookLM generates flat and generic output [48]. For example, it mostly regurgitates the exact language of source documents as I observed with my research writing class. The program failed to provide any sort of interpretive or analytical depth, which the students were able to generate in a lively and engaging class discussion about the same text.
- NotebookLM's summaries certainly look and sound like legitimate summaries. Gemini's attention mechanism during training picks up on the form of a conventional summary and can convincingly reproduce that form. If nothing else, LLMs are genre machines, and NotebookLM can generate in a formulaic manner the summary genre that is highly represented in Internet content [49]. But form is not the same as content, a distinction that matters.
- Google claims not to use personal data to train the model [14], (which is why I felt it was okay to upload Dilger's essay in the first place). When I repeated the experiment and asked NotebookLM to summarize the chapter's argument again, it did much better. I do not know whether the language model was updated between the two interactions, or whether Google does perhaps train their models on user data and input.
- NotebookLM does not save a user's chat history between sessions or screen refreshes [50]. Users are not alerted to this limitation on the interface (ironically this coming from a company that saves all other user data). In addition, the "Save to note" button in the Chat window is unreliable, which I discovered when I wanted to review the saved note about the program's first summary of Dilger's chapter, only to find that it had been deleted.

6 Learning and Friction

Statistical methods of text extrusion are misaligned with users' needs for precise, consistent results. They sloppily fulfill a transactional purpose. In this paradigm, writing becomes a commodity (p. 510) [51]; "writing is reduced to nothing more than documentation," as John Gallagher puts it [52]. But even in this transactional task, statistical text summary lets users down. But because of the industry's refusal to face this misalignment, students (and users in general) are asked to carry the burden of spotting and correcting confabulations and weaknesses [53].

It's worth contrasting machinic summary with how a human approaches summarization:

> "Human conceptual systems, in contrast, while appearing 'suboptimal' by these statistical measures, are likely shaped by a broader array of functional imperatives. These include the demands of adaptive generalization, rich causal and functional inference,

> the constraints of neural embodiment, and the requirements of nuanced communication–pressures that favor representations less statistically 'tidy,' but ultimately more flexible and powerful for navigating a complex world" (p. 8) [42].

Summarization is a decidedly rhetorical and embodied activity with a different objective. It views texts as embedded in specific contexts with particular audiences, and emphasizes relationships, implications, and consequences that may not be statistically prominent but are contextually crucial for our understanding of the world. Gallagher stresses, "written summaries aren't meant to be valuable themselves. Rather, summarizing is meant to indicate having thought through some ideas, putting them through your mind's meat grinder. It's an exercise designed to make you think [54]. Programs like NotebookLM remove enough friction from their easy, consumer-driven interfaces to thwart cognitive effort [20, 21].

In educational literature, friction is discussed as "productive failure," favoring problem-solving pedagogies that start with students grappling with a problem, before they learn to cope with it [55]. Inquiry-based writing pedagogy also focuses on "growth or belonging or productive uncertainty or the pleasure of wrestling with difficult ideas" (p. 178) [56]. Writing teachers know that reading and writing are hard and painful, acknowledging the "embodiment and the materiality of the affective wound" (p. 42) [57]. Good education supports the vulnerable learner in courses that ought to focus on uncertainly and friction as building blocks for learning. And so it goes for reading and summarization. As Stone, Goodlad, and Sammons explain, large language models do not "frame linguistic exchange through any of the many theorizations of human communication as an interactive, dialogic, and potentially emancipatory practice" [58]. These intersubjective affordances take time and effort. NotebookLM can be said to fail its target audience in three ways: (1) we need more time, not a computer program that removes this necessary component, (2) the program cannot even do a good job working with sources, and (3) its frictionless interface obscures the computational mechanism that is misaligned with the user task.

One can go as far as to say that to be human is to encounter and cope with productive friction; Goodlad and Stone posit that the ideology of "frictionless knowing" in generative AI products obscures that "it is only through active engagement that human beings acquire the learning that helps them to enrich a plural world of other situated people, places, objects, and recorded ideas" [25]. They refer to "human poiesis" as the process of transformative creation that includes artistic engagement and social interactions [25]. While large language models are trained precisely on such human textual corpora, their predictive, algorithmic calculations stand to replace the core human faculties of reflection, thinking, and invention. Shah and Bender urge us to remember that digital spaces should support user-centered, learner-focused environments that do not replace cognitive work [59], but include friction as a valuable and necessary component of learning [9].

7 Conclusion

Commercial generative AI systems take advantage of students (and faculty) who know little about machine learning. Demystifying AI hype from conceptual and methodological confusion that bombards us from industry, educational, and entertainment sources becomes part of the project of critical AI literacy with socio-political goals [60]. While NotebookLM promises to occupy the role of "research assistant" with expert knowledge of sources, it risks depriving users of agency, of control over their cognitive labor, and of knowledge about how exactly the program is misaligned with user goals. My advocacy here is not outright AI refusal, but a refusal of the dominant narratives about these technologies. Reframing the AI hype narrative can take many forms in our classrooms. I suggest some preliminary ideas:

1. Exposing students to accessible and critical materials about the inner workings of LLMs and their fundamental difference to human cognitive processes. Engaging examples might include Casey Fiedler's approachable short videos demystifying generative AI on social media [61], multi-modal essays [62, 63], or Bender and Hanna's "ridicule as praxis" podcast [64].
2. Inviting students to test programs like NotebookLM for their limits, in a sort of jail-breaking experimentation. I cannot think of more revealing example than a Reddit user's clever idea to upload a PDF that repeated just two words, "poop" and "fart" one thousand times; the resulting AI podcast delivers an earnest, utterly meaningless nearly 10-minute synthetic conversation about the two-word subject, complete with the requisite genre components of turn-taking and verbal banter [65]. Is there a better example of utter AI slop?
3. Offering students exposure to academic and professional organizations and networks which frame AI technologies through ethical and responsible perspectives. Examples include the grass-roots organization Civics of Technology, DAIR (the Distributed AI Research Institute), the Algorithmic Justice League, or ACM's FAcct conference (Fairness, Accountability, and Transparency).
4. Exploring justice-oriented practices for language model applications. Machine learning is not new. Its branding as "artificial intelligence" is. It is possible to design LLM-based applications that align their capabilities with specific user tasks, transparency, and social responsibility. In technical communication, see, for example [66, 67].
5. Reflecting with students on broader industry trends they will face as future professionals. While AI proficiency is touted as a necessary professional skill, we can refuse this simplistic framing in favor of robust conversations about AI design, implementation, and governance [25]. Jason Tham is clear about the challenges: "This requires a deep understanding of the capabilities and limitations of AI and the ability to critically evaluate its impact on user experience" [68]. Lucy Suchman's analysis of changes in the industry is remarkably prescient in this context [69].

I opened this paper by referencing a pro-business article, in which the authors suggest "that until capability considerations outweigh AI receptivity fueled by perceptions of AI as magical, there may be unintended consequences of policymakers' efforts to educate the public about AI" (p. 18) [1]. That is to say, companies benefit when consumers don't fully understand how AI works because mystery breeds enthusiasm. It's clear that to naively accept AI hype remains an exercise in normalization and acquiescence. As writing teachers, we must continue to destabilize the dominant AI rhetoric because at the current late capitalist moment, big tech AI hype is infiltrating the public arena and institutions (including education) to restructure resources, information, and human agency [70]. Let's think harder, with our students, to retain cognitive labor as categorically human and inalienable.

References

[1] Stephanie Tully, Chiara Longoni, and Gil Appel. 2023. Lower artificial intelligence literacy predicts greater AI receptivity. Journal of Marketing 89, 5 (January 2025), 1-20. https://journals.sagepub.com/doi/10.1177/00222429251314491

[2] Hayden Field. 2025. OpenAI, Anthropic target college students with latest education AI announcements. Retrieved July 26, 2025 from https://www.cnbc.com/2025/04/03/openai-anthropic-target-college-students-with-education-ai-services.html

[3] Hao-Ping (Hank) Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The impact of Generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In CHI Conference on Human Factors in Computing Systems (CHI '25), April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 23 pages. https://doi.org/101145/3706598.3713778

[4] Michael Gerlich. 2025. AI tools in society: Impacts on cognitive offloading and the future of critical thinking. Societies 15, 6 (January 3, 2025), 1-28. https://doi.org/10.3390/soc15010006

[5] Lucy Suchman. 2021. Talk with machines, redux. Interface Critique 3, 69-80. DOI: 10.11588/ic.2021.3.81328

[6] Stambler, Danielle Mollie, Nupoor Ranade, Daniel L. Hocutt, Stephen Fonash, Jessica Lynn Campbell, Ann Hill Duin, Isabel Pedersen, Jason Tham, Saveena (Chakrika) Veeramoothoo, and Gustav Verhulsdonck. 2024. Beyond digital literacy: Investigating threshold concepts to foster engagement with digital life in technical communication pedagogy." Technical Communication Quarterly 34, 2, 218–38. doi:10.1080/10572252.2024.2388038

[7] NotebookLM Help. 2025. Google.com. Retrieved April 17, 2025 from https://support.google.com/notebooklm/?hl$=$en#topic$=$14287611TUG 2017. Institutional members of the LaTeX Users Group. Retrieved May 27, 2017 from http://wwtug.org/instmem.html

[8] Training Data: EP17. Google NotebookLM's Raiza Martin and Jason Spielman on Creating Delightful AI Podcast Hosts and the Potential for Source-Grounded AI. Sequoia.com. Retrieved July 26, 2025 from https://www.sequoiacap.com/podcast/training-data-google-notebooklm/

[9] Tiffany DeRewal. Ask an expert: Evaluating LLM "research assistants" and their risks for novice researchers. Critical AI Blog. Retrieved July 26, 2025 from https://criticalai.org/2025/03/26/ask-an-expert-evaluating-llm-research-assistants-and-their-risks-for-novice-researchers/

[10] James Vincent. 2023. Google teases Project Tailwind–a prototype AI notebook that learns from your documents. The Verge. Retrieved July 26, 2025 from https://www.theverge.com/2023/5/10/23718707/google-io-2023-project-tailwind-personal-ai-notebook

[11] Vikas Kansai. 2025. GoogleLM Plus is now available in the Google One AI Premium subscription. Google Blog. Retrieved April 17, 2025 from https://blog.google/feed/notebooklm-google-one/

[12] Illuminate. Google. Retrieved July 26, 2025 from https://illuminate.google.com/explore

[13] Tom Lawrie. Google Gemini's NotebookLM: what marketers need to know. Threerooms Blog. Retrieved July 26, 2025 from https://www.threerooms.com/blog/google-gemini-notebooklm-for-marketers~

[14] NotebookLM. Google. Retrieved July 26, 2025 from https://notebooklm.google/

[15] Sian Cain. 2016. Is James Joyce's Ulysses the hardest novel to finish? The Guadian. January 13, 2016. Retrieved July 26, 2025 from https://www.theguardian.com/books/booksblog/2016/jan/13/is-james-joyce-ulysses-the-hardest-novel-to-finish~

[16] Michael Spencer and Nick Potlakistsky. 2024. Google NotebookLM: A game-changer for education and beyond. Substack. September 25, 2024. Retrieved July26, 2025 from https://www.ai-supremacy.com/p/googles-notebooklm-a-game-changer~

[17] Steven Johnson. 2024. How to use NotebookLM as a research tool. Substack. February 19, 2024. Retrieved July 26, 2025 from https://stevenberlinjohnson.com/how-to-use-notebooklm-as-a-research-tool-6ad5c3a227cc~

[18] David Berry. 2021. Explanatory publics: Explainability and democratic thought. In Fabricating Publics: The Dissemination of Culture in the Post-Truth Era. Eds. Bill Balaskas and Carolina Rito. DATA Browser 07, Open Humanities Press, 211-232. https://openhumanitiespress.org/books/download/Balaskas-Rito_2021_Fabricating-Publics.pdf

[19] Ginny Redish and Carol Barnum. 2011. Overlap, influence, intertwining: the interplay of UX and Technical Communication. Journal of User Experience 6, 3 (May 2011). https://uxpajournal.org/overlap-influence-intertwining-the-interplay-of-ux-and-technical-communication/~

[20] Bradley Dilger. 2006. Extreme usability and technical communication. In Critical Power Tools: Technical Communication and Cultural Studies. Eds. J. Blake Scott, Bernadette Longo, and Katherine V. Wills. State University of New York Press, New York, 47-69.

[21] Marcus Tomalin. 2023. Rethinking online friction in the information society. Journal of Information Technology 38, 1 (2023). DOI: 10.1177/02683962211067812

[22] Safiya Umoja Noble and Sarah T. Roberts. 2016. Through Google-colored glass(es): Design, emotion, class, and wearables as commodity and control. In Emotions and Technology: Communication of Feelings for, with, and through Digital Media. Eds. Sharon Y. Tettegah and Safiya Umoja Noble. Elsevier, London, 187-212.

[23] Thomas Haigh. 2023. There was no 'first AI winter.' Commun. ACM 66, 12 (Novermber 2023), 35-39. https://doi.org/10.1145/3625833

[24] Thomas Haigh. 2023. Conjoined twins: artificial intelligence and the invention of computer science. Commun. ACM 66, 6 (May 2023), 33-37. https://doi.org/10.1145/3593007

[25] Matthew Stone, Lauren M. E. Goodlad, and Mark Sammons. 2024. The origins of generative AI in transcription and machine translation, and why that matters. Critical AI 2, 1. https://doi.org/10.1215/2834703X-11256853

[26] Tomas Haigh. 2025. Historical reflections: Artificial Intelligence then and now. Communications of the ACM 68, 2 (February 2025), 24-29. DOI: 10.1145/3708554

[27] Steven Levy. 2024. Eight Google employees invented modern AI: here's the inside story. Wired, March 20, 2024. Retrieved July 26, 2025 from https://www.wired.com/story/eight-google-employees-invented-modern-ai-transformers-paper/

[28] Stephen Marche. 2024. Was linguistic AI created by accident? The New Yorker, August 23, 2024. Retrieved July 26, 2025 from https://www.newyorker.com/science/annals-of-artificial-intelligence/was-linguistic-ai-created-by-accident

[29] Ashish Vaswami, *et al.* 2017. Attention is all you need. 31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf

[30] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In Conference on Fairness, Accountability, and Transparency (FAccT '21), March 3–10, 2021, Virtual Event, Canada. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3442188.3445922

[31] Nicholas Carr. 2025. Against compresison. New Cartographies, July 11, 2015. Retrieved July 26, 2025 from https://www.newcartographies.com/p/against-compression?utm_source$=$post-email-title&publication_id$=$2992012&post_id$=$161162756&utm_campaign$=$email-post-title&isFreemail$=$true&r$=$22pnoc&triedRedirect$=$true&utm_medium$=$email~

[32] Karen Hao. 2025. Empire of AI: Dreams and nightmares in Sam Altman's OpenAI. Penguin Press, New York, USA.

[33] Long Ouyang *et al.* 2022. Training language models to follow instructions with human feedback. OpenAI, March 2022. https://arxiv.org/pdf/2203.02155~

[34] Amanda Askell *et al.* 2021. A general language assistant as a laboratory for alignment. Anthropic, December 9, 2021. https://arxiv.org/pdf/2112.00861~

[35] Emily M. Bender and Alex Hanna. 2025. The AI con: How to fight big tech's hype and create the future we want. Harper Collins, New York, USA.

[36] René Maatman, *et al.* 2024. When ChatGPT summarizes, it actually does nothing of the kind. R&A IT Strategy & Architecture, May 27, 2024. https://ea.rna.nl/2024/05/27/when-chatgpt-summarises-it-actually-does-nothing-of-the-kind/

[37] Aniket K. Shahade *et al.* 2024. Text summarization: Exploring classical, machine learning, and deep learning methods. Nanotechnology Perceptions 20, S9, 360-376. https://nano-ntp.com/index.php/nano/article/view/1573/1286

[38] Patrick Lewis, *et al.* 2021. Retrieval-augmented generation for knowledge-intensive NLP tasks. April 12, 2021. https://arxiv.org/pdf/2005.11401~

[39] Miriam Reynoldson. 2025. Why RAG is not the answer. The Mind File, May 22, 2025. Retrieved July 26, 2025 from https://themindfile.substack.com/p/rag-is-not-the-answer.

[40] Amanda Iglesias Moreno. 2025. Step-by-step guide to building a RAG system with NotebookLM. Data Science Collective. Medium, March 21, 2025. https://medium.com/data-science-collective/step-by-step-guide-to-building-a-rag-system-with-notebooklm-81688b9f516f~

[41] Yu Chen Kuang. 2025. Google gace its AI a bigger brain: a deep dive into NotebookLM's superpowered context window. Yu Chen Kuang, July 9, 2025. https://www.yuchenkuang.com/ai-tools/google-notebooklm/04-google-gave-its-ai-a-bigger-brain-a-deep-dive-into-notebooklms-superpowered-context-window~

[42] Chen Shani, *et al.* 2025. From tokens to thoughts: How LLMs and humans trade compression for meaning. Preprint, under review, June 30, 2025. https://arxiv.org/pdf/2505.17117~

[43] Ivy B. Grey. 2025. Why you're thinking about "reasoning" all wrong. Wordrake, May 2025. Retrieved July 26, 2025 from https://www.wordrake.com/blog/youre-thinking-about-reasoning-wrong#_ftn2

[44] Yann Gourvennec. 2024. NotebookLM by Google: Artificial voices, real concerns. Visionary Marketing (October 2024). Retrieved April 17, 2025. https://visionarymarketing.com/en/2024/10/21/notebooklm-by-google-artificial-voices-real-concerns/~

[45] Craig Van Slyke. 2024. Is NotebookLM Biased? Substack, November 22, 2024. Retrieved July 26, 2025 from https://aigoestocollege.substack.com/p/is-notebook-lm-biased

[46] Jess_askin. 2025. I now understand NotebookLM's limitations–and you should to. Reddit, June 2025. Retrieved July 26, 2025 from https://www.reddit.com/r/notebooklm/comments/1l2aosy/i_now_understand_notebook_llms_limitations_and/

[47] Join the NotebookLM community. Google NotebookLM Help. Retrieved July 26, 2025 from https://support.google.com/notebooklm/answer/16268096?hl$=$en&co$=$GENIE.Platform%3DDesktop~

[48] Philippa Hardman. 2024. Are ChatGPT, Claude & NotebookLM *really* disrupting education? Substack, October 17, 2024. Retrieved July 26, 2025 from https://drphilippahardman.substack.com/p/are-chatgpt-claude-and-notebooklm

[49] Sarah Brommer, *et al.* ChatGPT and its genre competence: an exploratory study. Weizenbaum Journal of the Digital Society 4, 4 (2024), 1-33. DOI: https://doi.org/10.34669/wi.wjds/4.4.6

[50] Frequently asked questions. Google NotebookLM. Retrieved July 26, 2025 from https://support.google.com/notebooklm/answer/16269187?hl$=$en#zippy$=$%2Chow-to-report-offensive-responses-generated-in-notebooklm%2Climits-for-notebooklm%2Cfile-size-limit-for-sources-in-notebooklm%2Csave-chat-history-between-sessions~

[51] Matthew Kirschenbaum and Rita Raley. 2024. AI and the university as a service. PMLA 139, 3, 504 - 515.

[52] John R, Gallagher. 2025. AI writing is not a calculator. So what is a good analogy? Substack, April 25, 2025. Retrieved July 26, 2025 from https://meresophistry.substack.com/p/ai-writing-is-not-a-calculator-so

[53] Daniel Cryer. 2024. The use of not to use AI? A student's burden. Inside Higher Ed, December 9, 2024. https://www.insidehighered.com/opinion/views/2024/12/09/ai-shifts-responsibility-academic-integrity-opinion#

[54] John R. Gallagher, 2025. A future high school teacher asked me about AI. Here's what I said. Substack, May 17, 2025. Retrieved July 26, 2025 from https://meresophistry.substack.com/p/a-future-high-school-teacher-asked

[55] Tanmay Sinka and Manu Kapur. 2021. When problem solving followed by instruction works: evidence for productive failure. Review of Educational Research 91, 5 (October 2021), 761-798. https://doi.org/10.3102/00346543211019105

[56] Annette Vee. 2023. Large language models write answers. Composition Studies 51,1 (2023), p. 176-181. https://compstudiesjournal.com/wp-content/uploads/2023/06/vee.pdf

[57] Edward Comstock. 2023. Toward a new neurobiology of writing: plasticity and the feeling of failure. CCC 74, 4 (June 2023), 695 - 730. DOI:10.58680/ccc202332523

[58] Lauren M. E. Goodlad and Matthew Stone. 2024. Beyond chatbot-k: On large language models, "generative AI," and the rise of chatbots–an introduction. Critical AI 2, 1. https://doi.org/10.1215/2834703X-11205147

[59] Chirag Shah and Emily M. Bender. 2022. Situating search. CHIIR '22, March 14–18, 2022, Regensburg, Germany. https://doi.org/10.1145/3498366.3505816

[60] Meredith Whittaker. 2021. The steep cost of capture. Interactions Magazine (November 2021). Interactions.ACM.org, 51-55. https://doi.org/10.1145/3488666

[61] Casey Fiedler, Information Science Professor, https://www.instagram.com/professorcasey/

[62] Angie Wang. 2023. Is my child a stochastic parrot? The New Yorker, November 15, 2023. https://www.newyorker.com/humor/sketchbook/is-my-toddler-a-stochastic-parrot

[63] Seán Clarke, Dan Milmo, and Garry Blight. 2023. How AI chatbots like ChatGPT or Bard work–visual explainer. The Guardian, November 1, 2023. https://www.theguardian.com/technology/ng-interactive/2023/nov/01/how-ai-chatbots-like-chatgpt-or-bard-work-visual-explainer

[64] Emily M. Bender and Alex Hanna. Mystery AI Hype Theater 3000. https://www.dair-institute.org/maiht3k/

[65] Sorryaboutyourcats. 2024. Analyzing poop & fart written 1,000 times–creating meaning from the meaningless. Reddit. https://www.reddit.com/r/notebooklm/comments/1fs9wlf/analyzing_poop_fart_written_1000_times_creating/

[66] S. Scott Graham and Hannah R. Hopkins. 2022. AI for social justice: new methodological horizons in technical communication. Technical Communication Quarterly 31, 1 (2022), p. 89-102. https://doi.org/10.1080/10572252.2021.1955151

[67] Scott R. Gallagher. 2025. Presentation at Computer and Writing, Athens, Georgia, USA.

[68] Jason Tham. 2024. Teaching UX: amid the hype of generative AI. Journal of User Experience 20, 1 (November 2024), 1-18. https://uxpajournal.org/wp-content/uploads/sites/7/2024/11/JUX_Tham_Nov-2024.pdf

[69] Lucy Suchman. 2002. Located accountabilities in technology production. Scandinavian Journal of Information Systems 14, 2 (2002), 90-104. https://uxpajournal.org/wp-content/uploads/sites/7/2024/11/JUX_Tham_Nov-2024.pdf

[70] Matteo Pasquinelli. 2019. How a machine learns and fails–a grammar of error for artificial intelligence. Spheres: Journal for Digital Cultures 5 (2019), 1-17.