

Educational Research and AI-Generated Writing: Confronting the Coming Tsunami

Tamara Tate¹, Shayan Doroudi¹, Daniel Ritchie¹, Ying Xu², Mark Warschauer¹

¹ University of California, Irvine

² University of Michigan

Author Note

Tamara Tate <https://orcid.org/0000-0002-1753-8435>

Shayan Doroudi

Daniel Ritchie

Ying Xu

Mark Warschauer <https://orcid.org/0000-0002-6817-4416>

We have no known conflict of interest to disclose.

Correspondence concerning this article should be addressed to Tamara Tate, University of California, Irvine, 3200 Education, Irvine, CA 92697-5500, (949) 431-6540, tatet@uci.edu

**Educational Research and AI-Generated Writing:
Confronting the Coming Tsunami**

Abstract

The public release and surprising capacity of ChatGPT has brought AI-enabled text generation into the forefront for educators and academics. ChatGPT and similar text generation tools raise numerous questions for educational practitioners, policymakers, and researchers. We begin by first describing what large language models are and how they function, and then situate them in the history of technology's complex interrelationship with literacy, cognition, and education. Finally, we discuss implications for the field.

On November 30, 2022, the technology company OpenAI released for public use the latest version of its AI-based large language model for text generation, ChatGPT (Schulman et al., 2022). Within a week, ChatGPT had taken the world by storm, with articles in the *The New York Times* (Metz, 2022), *Washington Post*, *The Atlantic* (Bogust, 2022; Herman, 2022), *Nature* (Stokel-Walker, 2022), *Wired* (Katwala, 2022), and elsewhere celebrating its remarkable capacities, decrying its biases, and pondering its impact on the future of education, especially writing instruction. All told, more than one million people accessed this web-based text interface in the first week (Altman, 2022), marking one of the fastest diffusions of new technology in history.

ChatGPT raises numerous questions for educational practitioners, policymakers, and researchers, each of which we will discuss. We begin by first describing what large language models are and how they function, and then situate them in the history of technology's complex interrelationship with literacy, cognition, and education.

Large Language Models¹

ChatGPT and other large language models are a type of artificial intelligence that are trained to generate text similar to human-generated text. They are called "large" because they are typically trained on a corpus of text data ranging from millions to billions of words; ChatGPT in particular is based on a model that was trained on several hundred billion words from a combination of curated corpora consisting of a vast crawl of websites on the Internet, books, and English-language Wikipedia (Brown et al., 2020). This allows the model to learn the nuances of natural language and generate text that is both coherent and natural-sounding.

¹ This section is based in part on the text output of ChatGPT.

The specific way that large language models work can vary, but they typically use a technique called deep learning, which involves training a network of artificial neurons on a large dataset. The network is then able to generate text by predicting the next word in a sequence based on the words that came before it. This allows the model to produce text that flows naturally and is similar to human-generated text. However, it is important to note that the way such a model learns language is vastly different from how humans learn language; simply the amount of text it is trained on is more than a human would read in a lifetime². Similar uses of AI allow automated generation in response to prompts of music and art.

Large language models have become increasingly popular in recent years and have been used for a variety of applications, including language translation, text summarization, and content generation. ChatGPT, produced by a company called OpenAI, has caught so much attention not because it is brand new—it builds on an earlier OpenAI product, GPT-3—but rather because it was the first released for free to the broad public requiring no coding or specialized skills to use; its user interface is styled as “chat” and allows for users to ask questions in natural language. As of this writing, anyone can simply log on to its public website, type in a prompt, and wait very briefly for the results. If the user asks a follow up question or comment, ChatGPT takes the earlier dialogue into account in its response.

The range of tasks that ChatGPT and other large language models can carry out is indeed remarkable and includes answering complex factual questions; writing an essay, story, play, or poem on almost any subject described³; writing a critique of that same text from the point of

² The amount of text GPT-3 has access to is around 100 times more than reading a 90,000-word book every day for 80 years.

³ For example, here is a poem ChatGPT wrote based on the abstract of this paper:

Ah, sweet muse, inspire thy servant's pen
To craft a verse in Shakespeare's den
Of ChatGPT, a tool most grand
Whose release doth shake the literate land

view of a teacher, professor or literary scholar; further revising the work based on that critique; writing computer code based on natural language descriptions; summarizing or expanding on texts; and changing texts to reflect a different perspective or author (e.g., rewrite the Star-Spangled Banner as Bob Dylan might). ChatGPT produces remarkably fluent text in many structured language genres and in various programming languages and will continue to learn as it receives feedback from users. In spite of its startling capacities, its serious flaws are also evident. In answering questions, it may give completely wrong answers, failing to note the use of negation in the text corpora it draws on (Katwala, 2022), and in many other cases it simply makes up what it does not know. It can display horrific racial or gender biases, again by distilling biases from the corpora it draws on, and its capacities change from day to day as the for-profit company that owns it, OpenAI, tweaks the model. And although it is free to access as of this writing, no one knows when and how it might be taken behind a paywall. However, one thing is certain: no matter what happens with ChatGPT, it will not be the only or last large language model we hear about, as other major technology companies will certainly release their own models for a wide variety of purposes and tasks, while OpenAI is well on its way to releasing its next iteration, GPT-4, which is said to be 500 times as large and powerful as current models (Altman, 2021, 2022)

Technology and Literacy

How then should we think about the pedagogical use of large language models? It is helpful to consider this in the context of the historical relationship between technology, literacy,

With power to generate text with ease
Leaving educators, policymakers to seize
Upon its history and complex interplay
With literacy, cognition, and education's sway
And ponder its implications, fair and true
For the field, ah what shall we do?

cognition, and education. Philosophers and scholars have been pondering this relationship for millenia, starting with Plato's consternations about writing in Phaedrus: "[Men] will cease to exercise memory because they rely on that which is written, calling things to remembrance no longer from within themselves, but by means of external marks" (Plato, n.d.). The invention of the printing press completely transformed literacy, though this process took several hundred years (Eisenstein, 1980). Over this period, writing became associated with original creation rather than the copying of manuscripts by hand; reading became a silent act by an individual rather than public oration; and scholarship grew to include expertise in a broad range of content areas rather than in-depth mastery of a small number of religious texts. These vast changes were facilitated by small technological reforms along the way, such as the now seemingly obvious step of putting spaces between words, which greatly aided the process of silent reading (Saenger, 1997).

The invention and diffusion of digital media brought about what Harnad (2003) called a fourth revolution in the means of production of knowledge, together with earlier revolutions of language, writing, and print. The impact of this fourth revolution is still unfolding, but it has already had a transformative impact on how people read, write, and communicate. Educators at first resisted these changes (Cuban, 1986), under the belief that digital tools, such as spell checkers and even word processors represented a form of cheating, and that students should thus focus exclusively on writing with pencil or pen and paper. Over time, minds have gradually changed with digital literacies recognized as their own unique and invaluable skills, both overlapping and distinct from print literacies (Warschauer, 1998; Lankshear & Knobel, 2008). More recently, standardized testing firms have come on board, both by allowing test-taking on digital devices and by trying to measure the distinct characteristics of digital literacy. Even the

Nation's Report Card, the National Assessment for Educational Progress, tested writing digitally (including allowing automated spell checking and grammar checking) beginning with the 2011 assessment because the Governing Board found digital writing skills important (Driscoll et al., 2010).

Debates about the use of digital technologies in the classroom often fall between two extremes, from the radical techno-optimists such as Nicholas Negroponte (Negroponte, 2006) who assume that simply passing out computers will solve all educational problems to the pessimists like Neil Postman (Postman, 2003) who warn that misinformed uses of technology will rob us of essential human value in education. We instead agree with Kranzberg (1986) who argued that "technology is neither good nor bad; nor is it neutral" (p. 545). Every technology poses both affordances and challenges, and also comes with its own biases.

Our research on classroom uses of technology, and in particular artificial intelligence, bears this out. The last author of this paper conducted early classroom research on what he coined "automated writing evaluation" (AWE; Warschauer & Ware, 2006), broadening out from the narrower term in use at the time of automated essay scoring. His classroom research suggested that AWE software was both fallible and useful (Grimes & Warschauer, 2010). Though its scoring and guidance were often faulty, and it was far from the silver bullet of improving writing that its cheerleaders claimed, it could still be a helpful tool for improving writing instruction by engaging and motivating students and assisting teachers with managing large classes, thus providing them more time for individual feedback. Its value could be realized in understanding both its strengths and limitations, and deploying it in a way that maximized the former and minimized the latter. This included educating teachers and learners on how AWE software functions, and using it in a critical, reflective way that also incorporated more social

forms of writing and assessment. Similar results have been found from our research on other forms of AI for language development, including visual-syntactic text formatting (Tate et al., 2019) and conversational agents (Xu et al., 2021, 2022). We thus take our cues on how to approach large language models from this body of research, seeking out, as Grimes and Warschauer (2010) expressed, “utility in a fallible tool.”

Implications

Instructors

Though the long-term impact of automated content generation on writing instruction is difficult to foresee, we are not persuaded by calls to ban it completely from student use due to its capabilities, faults, and biases. Our reasoning is based on two beliefs. First, its weaknesses are matched by powerful affordances, some of which favor the less privileged. These include non-native speakers of English around the world, who are forced to use the language in international academic and professional communication, as well as those with language or learning disabilities who struggle to write well—both groups that could benefit substantially from AI-enabled text generation. Secondly, the proverbial cat is already out of the bag, and if we do not teach people in marginalized communities to use these tools well, it will once again be the more tech-savvy elite who disproportionately benefit from them.

What then would a pedagogy that includes AI-based text generation look like? We argue it should include four elements. First, students need to understand large language models and other AI communication tools’ functions, strengths, weaknesses, and biases. What was the body of material on which the AI was trained? Who and what were excluded? What assumptions and biases might be implicit in the tool? Students can generate their own content and then work together to reflect on it. Second, students need to be able to find and access a range of AI tools

appropriate for specific communication tasks, such as writing papers or emails, creating slide presentations, or gathering background information. A number of such tools, beyond ChatGPT, exist, including Grammarly, a writing and grammar-checking tool that uses AI to check texts for grammar, spelling, punctuation, and other writing-related issues, and offers suggestions for how to fix errors or improve the writing; and Elicit, a tool built using GPT-3 and other language models, which searches hundreds of millions of research papers to answer questions. Third, students need to develop expert prompting strategies and search optimization to generate the most helpful content. In that sense, large language models are like search engines: garbage in, garbage out. The right prompts are required to hone in on what is needed. AI-generated text can be wrong and reference nonexistent studies; students need to understand that and include a corroboration step in their writing process if they are using AI-generated text. Finally, and most importantly, students need to learn how to incorporate AI-generated texts in their own writing ethically and effectively, noting and citing their use of AI in the authoring process. The standards for how to do so have not yet emerged, but will certainly do so over time, and we as educational researchers will have an important say in what they are. For now, OpenAI's policy (OpenAI, n.d.) that all written content created in part with text from ChatGPT should be "clearly disclosed in a way that no reader could possibly miss, and that a typical reader would find sufficiently easy to understand," is a good starting point. OpenAI has attempted to create a watermark or digital fingerprint within its AI-generated content (Wiggers, 2022), but that is challenging to do in a way that cannot be defeated through workarounds. Other third parties (e.g., <https://writer.com/ai-content-detector/>) have created tools to detect AI-written content; these could be used by teachers to detect plagiarism or used by students to modify their text to avoid being caught.

However, these tools do not bypass the importance of teaching students how to ethically and responsibly use and cite AI in their writing.

Of course any use of large language models in writing instruction should be introduced in a balanced and age-appropriate manner. Just like a young child should first learn arithmetic before later learning how to use a graphing calculator in high school math courses—and then doing assignments with and without the calculator—many believe that children should learn to write without these AI tools at a young age, before being introduced to them in secondary school or college. At-home writing assignments where use of these tools may be either allowed or difficult to control can be accompanied by in-class writing without them. Teachers will have to balance the teaching of effective writing with AI and writing without AI to ensure that students build up the necessary “muscle tone” to write and do not move too quickly to AI-generated text. Writing is hard; thinking is hard. But there is a time and a place for practicing hard things in order to become more proficient at them.

While much of the discussion around the use of large language models in education focuses on writing-based courses and assessments, we note that such tools can also have powerful affordances for other subjects as well. For example, students can ask such tools for explanations of complex ideas or step-by-step reasoning as they solve mathematics problems. Given ChatGPT’s tendency to make mistakes when reporting factual information or solving problems (tasks which it was never directly trained on), students should be cautious to carefully read over and possibly challenge its responses, rather than take them to be true. But the fallibility of such tools offers another affordance: large language models can serve as teachable agents. Taylor (1980) described three promising uses of the computer in education: tutor, tool, and tutee. While the uses of the computer as tool and tutor (e.g., intelligent tutoring systems) have

proliferated over the decades, its use as a tutee has remained very limited. As Taylor described how to use the computer as tutee, “the student or teacher doing the tutoring must learn to program, to talk to the computer in a language it understands” (p. 245). Until recently, this has remained limited mainly to students learning programming languages (like Logo or Scratch). However, perhaps for the first time, it is reasonable “to talk to the computer in a language it understands” without programming the computer; students can now learn by teaching the computer. If the computer answers a question incorrectly the student must identify either a mistake in their own thinking or an error that the tutee made. Although several teachable agents have been created before (Blair et al., 2007; Matsuda et al., 2013; Michie et al., 1989), these were developed to work for specific educational tasks. AI models like ChatGPT can potentially serve as general-purpose teachable agents with dialogue that feels like a natural tutor-tutee relationship, though research is needed to see how to effectively repurpose such models to best serve as tutees.

All of this also means that teacher pre-service education needs to provide the necessary pedagogical information and practical skills to incorporate AI in their lessons. Not only English teachers, but also history teachers and others will need to be aware of the challenges and opportunities for their students so that teachers can successfully navigate the (inevitable) use of AI in their classrooms. Teachers already in the classroom will need similar opportunities to learn and practice. Although traditionally teacher preparation has siloed technology aspects of education into a single course, we recommend incorporating this tool more organically into the various relevant courses.

Policymakers

Researchers, educators, and technology experts need to meet to discuss and clearly articulate the affordances and challenges of AI-generation tools like ChatGPT for education. Resources are needed to educate teachers at all levels about AI in education; curriculum should be developed to create good consumers of the technology. Educators should work with researchers and technology experts to address the four pedagogical points above and come up with best practices for effective and ethical use of AI-generation tools. The ensuing curriculum and best practices should be widely disseminated, to educators, students, and parents, harnessing the reach of nonprofit organizations already active in the digital-education field, such as Digital Promise, the International Society for Technology in Education, and Common Sense Media. Similarly, the professional organizations are quickly working to provide professional guidance for instructors and departments figuring out how to create both pedagogy and use policies for courses.

Large-scale assessments also need to consider whether and how to respond to the new AI-text generation tools. While spell check is routinely included as a tool available in the NAEP writing test and annual state English language arts assessments such as the California CAASPP, does text generation have any role in assessment? Should students be assessed on their ability to prompt, revise, or analyze AI-generated text? We believe that at the moment it is too early to do so, since students will have had limited if any experience on such tools and teachers have likely not begun to include them in the routine curriculum. However, we can imagine a time in the not-so-distant future when such tools should be included so that students can be assessed in a real-world, authentic writing context. More immediately, what does this tool do for the assessments in many AP writing courses, which encourage a fairly formulaic approach to essays? If ChatGPT

can produce an essay that could score a 4 or a 5 on an AP prompt in seconds, what is the real value of that type of writing and why are we giving students college credit for it?

Finally, the existence of powerful text-generation tools will exacerbate so-called “digital divides” and bring about new ones. At the time of this writing, ChatGPT is free—at least in many countries. In some countries, though, it is not even accessible. And where accessible it requires a mobile phone number to register a login. As large language models are possibly put behind paywalls, differential access will become even more important, let alone differences in knowledge and skill in using these tools, such as how to write the best prompts. It also can be expected that those who may benefit from large language models the most, including non-native speakers of national languages, people with learning disabilities, and simply those with weak literacy skills, will be most subject to accusations or punishment for seeking to use these tools. This exacerbates an already present dilemma, in which English learners are instructed to speak and write like a native speaker, but then get noticed most and thus stigmatized most when they master tools that allow them to do just that.

Researchers

On an almost meta level, we need to understand how AI-generated text changes the nature of academic writing and research itself. This summer, researchers queried OpenAI’s GPT-3 and caused it to draft an academic paper about itself which was submitted for publication (Thunstrom, 2022; preprint, GPT et al., 2022). They found that the AI generated clear and concise descriptions of its capabilities and features, but warned that any such writing must be closely monitored by researchers in order to mitigate any potential negative consequences. Publications, research organizations, and affinity groups such as APA and AERA need to gather experts and create a common agreement on how AI-generated text can ethically be used in

research, including how to make use of AI in research transparency and how to cite its use and contribution.

On an applied level, there needs to be more research on curriculum and effective teaching with and about large language models. The public release of ChatGPT has thrust onto the public's attention the challenges and affordances that AI presents for teaching students how to write and communicate. The urgent concerns and attention this issue is receiving call for rapid exploratory research on the topic. We recommend focusing on higher education and high school uses, such as studying the best ways of incorporating ChatGPT and other AI tools into writing and communication instruction or using it as a tool that can teach or be taught complex STEM topics. Educators throughout the US and the world are hungry for guidance on how to adapt their teaching to these new tools.

There needs to be more research on what learning is like with large language models and the impact of using AI tools. We need to observe what happens when students utilize various prompting and editing strategies, whether they revise more when using AI tools, which students use the tool more and what such use predicts for their future writing skills or understanding of content areas. Does using these tools on an ongoing basis make people better writers when the tools are removed (similar to visual-syntactic text formatting, see., e.g. Tate et al., 2019) or worse writers?

Finally, there needs to be more conceptual scholarship about the nature of cognition, writing, and tools. Thinking of Gregory Bateson's questions: Bateson (1973, p. 318) asks us to consider a blind man with a stick. "Where does the blind man's self begin? At the tip of the stick? At the handle of the stick? Or at some point halfway up the stick?" At what point does it become more important for us to know what a person can do *with* tools rather than without

them? From a sociocultural perspective, and Activity Theory in particular (Engeström, 1987), social interaction plays a role in creating an environment to learn language, learn about language, and learn “through” language (Vygotsky & Cole, 1978)—how does AI impact this environment? If human learning and development are bound up in activity, that is, purposeful action mediated by various tools (Vygotsky & Cole, 1978; Wertsch, 1979), what will this new tool mean? Differences in modality matter and the process of writing is shaped in part by the available tools (Bazerman et al., 2018; Graham 2018; Wertsch 1991). How will AI-enabled tools change the process, and product, of writing?

This is a pivotal moment in literacy. We must make the most of it.

References

Altman, S. (2022, Dec. 4). Twitter.

<https://twitter.com/sama/status/1599668808285028353?s=20&t=j5ymflUeTpeQuJKIWAKaQ>

Bazerman, C., Applebee, A.N., Berninger, V.W., Brandt, D., Graham, S., Jeffery, J.V., Kei, P., Matsuda, S., Rowe, D.W. Schleppegrell, M., & Wilcox, K. C. (2018). *The Lifespan Development of Writing*. National Council of Teachers of English.

<https://wac.colostate.edu/books/ncte/lifespan-writing/>

Birhane, A., Kalluri, P., Card, D., Agnew, W., Dotan, R., & Bao, M. (2022, June). The values encoded in machine learning research. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 173-184).

Blair, K., Schwartz, D. L., Biswas, G., & Leelawong, K. (2007). Pedagogical agents for learning by teaching: Teachable agents. *Educational Technology*, 47(1), 56-61.

<https://www.jstor.org/stable/44429380>

Bogost, I. (2022, Dec. 7). ChatGPT is dumber than you think: Treat it like a toy, not a tool. *The Atlantic*. <https://www.theatlantic.com/technology/archive/2022/12/chatgpt-openai-artificial-intelligence-writing-ethics/672386/>

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, Shyam, P., Sastry, G., Askell, A., Agarwal, A., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., ... Amodei, D. (2020). Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)* (pp. 1877-1901). <https://proceedings.neurips.cc/paper/2020>

Cuban, L. (1986). *Teachers and machines: The classroom use of technology since 1920*.

Teachers College Press.

Driscoll, D. P., Alukonis, D. J., Fabrizio, L. M., Flores, A., Friedman, A. J., Gordon, D. W., Hicks, D. R., King, K., Kozbial-Hess, K., Kranendonk, H., Miles, T., Paine, S.L., Perdue, S., Pimentel, S., Popham, W.J., Porter, A.C., Smith, W.T., Taymans, M.F., ... Easton, J.Q. (2010). Writing Framework for the 2011 National Assessment of Educational Progress.

National Assessment Governing Board.

<https://www.nagb.gov/content/dam/nagb/en/documents/publications/frameworks/writing/2011-writing-framework.pdf>

Eisenstein, E. L. (1980). *The printing press as an agent of change*. Cambridge University Press.

<https://doi.org/10.1017/CBO9781107049963>

Gpt, G.P.T., Thunström, A. O., & Steingrímsson, S. (2022). Can GPT-3 write an academic paper on itself, with minimal human input? <https://hal.archives-ouvertes.fr/hal-03701250/>

Graham, S. (2018). A revised writer (s)-within-community model of writing. *Educational Psychologist*, 53(4), 258-279.

Grimes, D., & Warschauer, M. (2010). Utility in a fallible tool: A multi-site case study of automated writing evaluation. *The Journal of Technology, Learning and Assessment*, 8(6).

Harnad, S. (2003). Back to the oral tradition through skywriting at the speed of thought.

Interdisciplines.

Herman, D. (2022, Dec. 9). The end of high-school English. *The Atlantic*.

<https://www.theatlantic.com/technology/archive/2022/12/openai-chatgpt-writing-high-school-english-essay/672412/>

Katwala, A. (2022, Dec. 9). ChatGPT's fluent BS is compelling because everything is fluent BS.

Wired. <https://www.wired.com/story/chatgpt-fluent-bs/>

Kranzberg, M. (1986). Technology and History:" Kranzberg's Laws". *Technology and culture*, 27(3), 544-560. <https://doi.org/10.2307/3105385>

Lankshear, C., & Knobel, M. (2008). *Digital Literacies: Concepts, Policies, and Practices*. Peter Lang.

Matsuda, N., Yarzebinski, E., Keiser, V., Raizada, R., Stylianides, G. J., & Koedinger, K. R. (2013). Studying the effect of a competitive game show in a learning by teaching environment. *International Journal of Artificial Intelligence in Education*, 23(1), 1-21. <https://doi.org/10.1007/s40593-013-0009-1>

Metz, C. (2022, Dec.11). The new chatbots could change the world. Can you trust them? *The New York Times*. <https://www.nytimes.com/2022/12/10/technology/ai-chat-bot-chatgpt.html>

Michie, D., Paterson, A., & Michie, J. H. (1989). Learning by teaching. In Jaakkola, H., & Linnainmaa, S. (Eds.), *Scandinavian Conference on Artificial Intelligence 89: Proceedings of the SCAI'89* (pp. 307–331). IOS Press.

Negroponte, N. (2006). *One laptop per child*. [Video]. TED Conferences. https://www.ted.com/talks/nicholas_negroponte_one_laptop_per_child?language=en

OpenAI (n.d.). Sharing & publication policy. <https://openai.com/api/policies/sharing-publication/>

Plato, P. (1952). *Phaedrus* (Vol. 275, p. 276A). Bobbs-Merrill.

Postman, N. (2003). *Technopoly*. Bollati Boringhieri, Torino.

Saenger, P. (1997). *Space between words: The origins of silent reading*. Stanford Univ. Press.

- Schulman, J., Zoph, B., Kim, C., Hilton, J., Menick, J., Weng, J., Uribe, J. F. C., Fedus, L., Metz, L., Pokorny, M., Lopes, R. G., Zhao, S., Vijayvergiya, A., Sigler, E., Perelman, A., Voss, C., Heaton, M., Parish, J., Cummings, R. N., ... Ryder, N. (2022, Nov. 30). *ChatGPT: Optimizing language models for dialogue*. OpenAI.com. <https://openai.com/blog/chatgpt/>
- Stokel-Walker, C. (2022, Dec. 9). AI Bot ChatGPT writes smart essays--should professors worry? *Nature*. <https://www.nature.com/articles/d41586-022-04397-7> doi: <https://doi.org/10.1038/d41586-022-04397-7>
- Taylor, R. (1980). The computer in the school: Tutor, tool, tutee.
- Tate, T. P., Collins, P., Xu, Y., Yau, J. C., Krishnan, J., Prado, Y., Farkas, G., & Warschauer, M. (2019). Visual-syntactic text format: Improving adolescent literacy. *Scientific Studies of Reading, 23*(4), 287-304.
- Thunstrom, A. O. (2022, June 30). We asked GPT-3 to write an academic paper about itself--then we tried to get it published. *Scientific American*.
<https://www.scientificamerican.com/article/we-asked-gpt-3-to-write-an-academic-paper-about-itself-mdash-then-we-tried-to-get-it-published/>
- Vygotsky, L. S., & Cole, M. (1978). *Mind in society: Development of higher psychological processes*. Harvard University Press.
- Warschauer, M. (1998). *Electronic literacies: Language, culture, and power in online education*. Routledge.
- Warschauer, M., & Ware, P. (2006). Automated writing evaluation: Defining the classroom research agenda. *Language Teaching Research, 10*(2), 157-180.
- Wertsch, J. V. (1979). From social interaction to higher psychological processes. A clarification and application of Vygotsky's theory. *Human Development, 22*(1), 1-22.

Wiggers, K. (2022, December 10). OpenAI's attempts to watermark AI text hit limits.

TechCrunch. <https://techcrunch.com/2022/12/10/openais-attempts-to-watermark-ai-text-hit-limits>

Xu, Y., Aubele, J., Vigil, V., Bustamante, A. S., Kim, Y. S., & Warschauer, M. (2022). Dialogue with a conversational agent promotes children's story comprehension via enhancing engagement. *Child Development, 93*(2), e149-e167. <https://doi.org/10.1111/cdev.13708>

Xu, Y., Vigil, V., Bustamante, A. S., & Warschauer, M. (2022). Contingent interaction with a television character promotes children's science learning and engagement. *Journal of Applied Developmental Psychology, 81*, 101439. <https://doi.org/10.1016/j.appdev.2022.101439>