

This is an expanded version of an article published in *Psychology Today*

Citation: Martin, A.J. (May, 2023). Educational psychology meets generative AI: A look at the impact of ChatGPT, Bard, and others on education and beyond. *Psychology Today*, <https://www.psychologytoday.com/us/blog/psyched/202305/educational-psychology-meets-generative-ai>

Educational Psychology Meets Generative AI

A Look at the Impact of ChatGPT, Bard, and Others on Education and Beyond

Andrew J. Martin

School of Education, University of New South Wales, Australia

Generative artificial intelligence (AI) is kicking-up a storm in universities, colleges, and K-12 schools. Chatbot applications such as OpenAI's "ChatGPT", "GPT-4" and "DALL-E", as well as Google's "Bard" and "Socratic" are among those receiving the most interest in the education space at present. But there will be more as the AI field advances at breakneck speed. AI applications use "machine learning" techniques to generate human-like text, including academic text, as well as visual images and audio. Whereas traditional AI was designed to recognize patterns in large data sets and make predictions to inform decisions, generative AI uses massive amounts of data to create new, original, content. Educational institutions the world over are wrestling with the pedagogical, ethical, personnel, and logistical implications of generative AI. Currently, the technology harbors worrisome problems, like inventing fake sources and sometimes drawing questionable conclusions. However, there is little doubt that the technology will continue to develop and "learn", and that it will be hugely disruptive and transformative in education.

In this article I share early thoughts on the potential role of generative AI in the educational psychology research landscape. I label my ideas "early" because in light of the rapid evolution of this technology, I anticipate that some of these ideas might be soon outdated, and new ideas would need to be developed.

I approach this in two complementary ways. The first examines how educational psychology research can help us better understand the ways that generative AI impacts learning and learners. The second examines how generative AI is used to advance research methods in educational psychology. In this discussion, I use GPT-4 as the illustrative generative AI, in light of the attention it has been receiving; but, the ideas hold for other existing AI and new generative applications that will be created.

How Generative AI Impacts Learning and Learners

In broad terms, educational psychology focuses on the factors and processes involved in learning and development. Therefore, educational psychology research has much to contribute in how generative AI such as GPT-4 impacts learners. Here I will point to three possible contributions.

1. *How generative AI affects learning*: In what ways is learning impacted by the use of generative AI? This research might involve, for example, an investigation into how using

GPT-4 to develop a first draft of a history essay impacts students' persistence, recall of key facts, and understanding of the arguments in that essay. Alongside its potential impacts on people's engagement, learning, and achievement, researchers might also investigate generative AI as a resource for "interpersonal" interactions and support. For example, given that instrumental help is one pillar of a quality teacher-student relationship, what is the potential role of GPT-4 to support teachers' efforts in this domain? Taking this a step further, a more confronting reality will be the emotional relationships that people develop with AI products designed precisely for this purpose. Researchers will need to examine the implications of this for young people's social-emotional development.

2. *How generative AI progresses key ideas about learning:* In what ways does generative AI transform our understanding of how and what people learn? For example, educational psychology researchers specializing in cognitive science point to the development of students' long-term memory as a primary purpose of education. What are the implications for educational psychology and educators if knowledge that would previously have been encoded in long-term memory is increasingly offloaded to the likes of GPT-4? Or, if high-quality use of generative AI relies on students asking it "good" questions, what do we want students to learn on a particular topic so they can generate "better" questions?
3. *How generative AI can be optimized to enhance learning:* To what extent can educational psychology research inform the development of generative AI so it can best support learners? As noted above, generative AI raises challenging questions for educators about what is important for students to learn. Researchers may provide guidance for AI developers to adapt the technology to ensure that students harness opportunities to learn important content that is vital for their optimal functioning through school—and beyond.

How Generative AI Impacts Research and Researchers

We can also consider how generative AI impacts the ways that educational psychology research is conducted. Following are two potential contributions.

4. *How generative AI is used to improve research methods:* To what extent can generative AI be used to progress research methods in educational psychology? This endeavor might entail using GPT-4 to assist in developing research instruments, such as adaptive text in a motivation survey. Here, researchers could use the AI to generate unique item wording based on the way respondents answer questions in the survey. Or, researchers may use GPT-4 to generate new material for students to respond to as they participate in an educational intervention. Indeed, generative AI may also be used to produce data for educational psychology research. If students are using GPT-4 in their everyday learning contexts, there could be opportunities to capture in-the-moment academic engagement and learning data by this technology. It is also possible that AI research and development may produce data capture methodologies that educational psychologists had not previously considered and that can be integrated into educational psychology research to garner novel insights into student learning.
5. *How generative AI is used to perform scholarly tasks:* What role does generative AI play in performing key aspects of the research and scholarship process? Reviews of literature are perhaps the lowest hanging fruit, but the technology can also be used to conduct many other tasks such as managing data and generating syntax for statistical analysis. Attention is also turning to the use of generative AI to address labor shortages in academia—such as using GPT-4 as an additional reviewer of a manuscript or research grant application. That said,

generative AI is prone to error and human oversight will be vital for some time to come—suggesting a supportive rather than fully performative role for generative AI at present.

What is the “Right” Balance?

Realistically, the debate is not about “whether” students and researchers should use generative AI, since it is clear that they will. Rather, the debate needs to be about “how much”, “in what ways”, “when”, and “for what purposes” they should use it. This is where the contested and vexed territory lies and researchers, administrators, educators, and the public would benefit from clear, timely, and ongoing guidelines from policymakers, institutions, and respective administrators (e.g., for researchers, the journals in which they publish; for educators, the lesson plans they generate). It is also where unpacking and considering these matters in terms of points #1 to #5 above may provide some concrete structure for ways forward. I will take #5 (using generative AI to perform scholarly tasks) as a case in point: “How much”, “in what ways”, “when”, and “for what purposes” should a researcher use GPT-4 to develop a review of literature for a research paper? What “rules” need to be developed to guide us here?

For example, would it be helpful to apply the classic 80-20 “rule” as appropriate to the researcher’s status as a novice (e.g., research student) or an expert (e.g., a mid- to late-career scholar). Here, a novice researcher would dedicate substantial time to searching for and reading the scope of literature, summarizing it, and then crafting it into a narrative as appropriate to a research question. Then GPT-4 would be used to identify further literature and ideas not captured by the novice’s extensive efforts. Here, the balance of the work is 80% novice and 20% AI. Through this balance, the novice has the opportunity to systematically build their own body of long-term knowledge and understanding using GPT-4 to augment this as appropriate.

However, for the expert who already knows much of the field and has likely drafted similar ideas in many previous papers, it may be appropriate for much of the initial information to be generated by GPT-4. Then, the expert corrects the errors that are known to arise in generative AI text, includes additional information and research that may have fallen outside AI’s algorithmic parameters, and finesses the argument to the specific research question posed. Here, the balance of the work is 80% AI and 20% expert.

Of course, this is not prescriptive or definitive. In fact, it quite possibly is naïve. Ultimately, it may not be possible or desirable to fully moderate generative AI, compelling us to land on something like an 80 (AI) - 20 (researcher) mix across the board, irrespective of a researcher’s status as novice or expert. If this is the case, the task for universities and research supervisors may be to consider how research activities need to change to incentivize the desirable use of the AI by research students—much like, for example, how educational psychology researchers will need to guide educational policy-makers and practitioners on how to shape learning and learning activities to incentivize students to use generative AI in desirable ways.

But even in an 80 (AI) - 20 (researcher) split, the use of generative AI can be highly scaffolded for the novice. Here, for example, a supervisor would encourage the novice to use GPT-4 to generate,

review, edit, and master small chunks of subject matter in an iterative way. Approaching a larger task through smaller chunks may make it easier for the novice to encode what is generated by GPT-4 in long-term memory. The supervisor would also strongly encourage verification through additional reading of published sources. Here, for example, the novice would chase up and consult references and ideas identified by GPT-4 so that as they use generative AI, their own knowledge and skills are building in long-term memory (and they catch errors produced by the AI). As long-term memory builds, the novice learns to use GPT-4 in a more expert way through better questions, search terms, prompts, etc.

Taken together, I believe that researchers who learn to effectively harness generative AI, whilst continuing to build and maintain their own body of knowledge and skill, will flourish in the world of scholarship ahead. As tempting as it may be for researchers (given the performance pressures they are under) to lean heavily or exclusively on generative AI to do their thinking and work for them, this will ultimately leave them falling short as scholars.

The Future of Scholarship: “New Tech” Meets “Old School”

There is no question that generative AI is rapidly becoming exceedingly accurate, fast, productive, efficient, and comprehensive, and that it is helping research and researchers to be the same. In this regard, this technology will eventually expedite and improve scholarship, even if it will also involve less desirable uses of AI. At the same time, we must not neglect the crucial yields of developing extensive domain-specific knowledge in long-term memory (i.e., expertise) through years of engagement in research, review, and authorship using traditional means and technologies. We should not underestimate the deep and authentic learning that comes from reading the rich and complex ideas laid out in a seminal or pioneering journal article, or an exquisitely crafted chapter or monograph—not just a synthesis of it generated by AI. So, in addition to the extractive power of generative AI, “old school” scholarship will also continue to generate yields as researchers themselves continue to develop novel ideas and answers—and, in turn, provide generative AI with new highly informed prompts and “better” questions—all with a view to advance understanding of human experience, learning, and development.

The ideas presented here are drawn from an online presentation by the author at the UNSW Educational Psychology Research Group Roundtable on March 27, 2023. Thanks are extended to Keiko Bostwick, Rebecca Collie, Colin Evers, M’lane Field, Helena Granziera, Kieran Martin, Rory Martin, and Avi Kaplan for comments on an earlier draft. No part of the text was developed using generative AI.

Andrew Martin is Scientia Professor of Educational Psychology in the School of Education, UNSW, Australia (<https://www.researchgate.net/profile/Andrew-Martin-22>)