Generative AI in Higher Education

The Product Landscape

Claire Baytas
Dylan Ruediger

March 7, 2024

ITHAKA S+R

Ithaka S+R provides research and strategic guidance to help the academic and cultural communities serve the public good and navigate economic, demographic, and technological change.

Ithaka S+R is part of ITHAKA, a not-for-profit with a mission to improve access to knowledge and education for people around the world. We believe education is key to the wellbeing of individuals and society, and we work to make it more effective and affordable.

Introduction

Understanding the value of individual products in a now-crowded marketplace is a major challenge for end users and for university CIOs, IT departments, and others involved in decision-making.

Generative AI (GAI) has quickly gained a significant foothold in academia, and is now used widely for teaching, learning, and research purposes.[1] While national trends in student and faculty adoption are unclear, surveys conducted by individual institutions have found that approximately 50 percent to 65 percent of both students and faculty have used ChatGPT or one of its commercial competitors.[2] If current trends continue, in the near future GAI use will be ubiquitous, fully integrated into the core mission of colleges and universities.

While the full effects of this transformation are, at best, clear only in outline, 2023 saw both well-established vendors and start-ups racing to bring GAI applications optimized for use in educational contexts to market. As we write this brief, new products are appearing so rapidly that just keeping up with them is difficult, and understanding the value of individual products in a now-crowded marketplace is a major challenge for end users and for university CIOs, IT departments, and others involved in decision making about which products will be supported and/or licensed for campus users.

---

[1] For overviews, see Yiheng Liu et al, "Summary of ChatGPT-Related Research and Perspectives Towards the Future of Large Language Models," *Meta-Radiology* 1 (2023), https://doi.org/10.1016/j.metrad.2023.100017; Michelle Kassorla, "Teaching with GAI in Mind," *EDUCAUSE Review*, 14 December 2023, https://er.educause.edu/articles/2023/12/teaching-with-gai-in-mind; Sinan Onal and Derya Kulavuz-Onal, "A Cross-Disciplinary Examination of the Instructional Uses of ChatGPT in Higher Education," *Journal of Educational Technology Systems*, September 2023, https://doi.org/10.1177/00472395231196532; Rishab Jain and Aditya Jain, "Generative AI in Writing Research Papers: A New Type of Algorithmic Bias and Uncertainty in Scholarly Work" *arXiv preprint: 2312.10057*, December 2023, https://doi.org/10.48550/arXiv.2312.10057.

[2] An institution-wide survey from the University of Baltimore found that 67 percent of faculty and staff and 54 percent of students had used ChatGPT; see https://drive.google.com/file/d/1ufdagea0Xm8TpiKsyvbr1Kp-kpez3z6Z/view. An institution-wide survey from the University of Michigan found that 56 percent of undergraduates, 58 percent of faculty, and 66 percent of graduate students had used a GAI tool; see https://genai.umich.edu/committee-report. Michigan also reported that students had greater familiarity with the tools than faculty. Baltimore, on the other hand, found that faculty had greater familiarity with ChatGPT and Bard/Gemini, while students had greater familiarity with Grammarly.

Since last fall, Ithaka S+R has been partnering with 19 colleges and universities from the US and Canada to assess GAI's impact on higher education and make evidence-based, proactive decisions about how to manage the far-ranging effects of GAI.[3] As part of this project, Ithaka S+R has been cataloging GAI applications geared towards teaching, learning, and research in the higher education context. Today, we are excited to make our Product Tracking tool (https://sr.ithaka.org/our-work/generative-ai-product-tracker/) publicly available.

The Product Tracker includes a basic description of each tool, as well as information about the pricing model, key features, and other relevant details such as the large language model (LLM) or datasets behind the tool or background on the vendor. As it would be impossible to track every GAI application that might conceivably be used in higher education contexts, we have limited the Tracker to a) products marketed specifically towards faculty or student use, and b) products that appear to be in active use in teaching, learning, or research activities. At present, the Product Tracker contains data on over 100 GAI tools and applications. While we make no claims to including every relevant or potentially relevant product, the Tracker includes data on the most visible individual products on the market and is comprehensive enough to provide a landscape perspective on the market itself. The Tracker is also a living document which we update regularly to include new products or add new data about existing products. We will continue to do so for at least the remainder of 2024 if the tool retains value to the community.

This issue brief is designed to enrich the descriptive data captured in the Product Tracker. In the brief's first section, we provide a typology of existing products and value propositions. In the second, we offer observations about what the product landscape suggests about the future of teaching, learning, and research practices, and speculations on the near-term future of the academic GAI market.

[3] Danielle Miriam Cooper and Dylan Ruediger, "Making AI Generative for Higher Education: Announcing the Partners for a Multi-Year Research Project," *Ithaka S+R*, 24 May 2023, https://sr.ithaka.org/blog/making-ai-generative-for-higher-education-2/.

Use Cases and Value Proposition

At the most fundamental level, most products included in our tracker are marketed as ways of speeding up and/or enriching three key phases of the research lifecycle: discovery, understanding, and creation. Some products are designed to perform hyper-specific tasks such as providing analysis of a single PDF; others are designed to support students and researchers from the conception of a project to its completion.

In nearly every case, it is important to recognize that individual products are essentially apps running on the "operating system" of a commercial LLM, and specifically on OpenAI's GPT models. With a few notable exceptions, primarily coming out of biomedical fields, ground-up, custom-built LLMs are not part of the current product landscape.

Discovery

GAI powered discovery tools are at this point perhaps the most mature application of the technology to higher education contexts, driven in part by the integration of GAI into existing search and discovery tools used by major scholarly publishers and aggregators. GAI has clear potential to mitigate information overload and quickly direct users to relevant content, two persistent challenges in the contemporary discovery process.[4] GAI's ability to engage users in dialogue and its use of natural language processing over keywords and metadata intend to simplify and improve the specificity of the discovery process. Users benefit from these tools' capacity to gain a deep semantic understanding of both natural language queries and the content of academic papers in databases to have an increasingly conversational search experience. For instance, Consensus (https://consensus.app/), a startup with an AI-powered search engine,

[4] Mark Glickman and Yi Zhang, "AI and Generative AI for Research Discovery and Summarization," *arXiv preprint:2401.06795* (2024), https://doi.org/10.48550/arXiv.2401.06795; Richard Van Noorden, "ChatGPT-like AIs are Coming to Major Science Search Engines," *Nature*, 2 August 2023, https://www.nature.com/articles/d41586-023-02470-3.

now has a Copilot feature in its main product as well as a Consensus GPT in OpenAI's store, both of which bring a chat dimension to discovery.[5]

Moreover, the discovery products aimed at higher ed users address a clear and significant problem with using consumer LLMs such as ChatGPT for research and educational purposes, namely the poor quality and incorrect data included in their training data and their limited ability to cite specific sources. The potential of GAI powered tools to simultaneously address the limitations of both existing academic discovery tools and of general purpose LLMs provide a compelling value proposition that improves an identifiable use case applicable to undergraduate student assignments and specialist research questions.

Major scholarly publishers and content aggregators have the advantage of being gatekeepers to large corpora of such material and have spent the past year racing to integrate GAI into their discovery platforms.

High quality discovery depends on high-quality content, and one of the core ways vendors are differentiating their offerings is based on the value of the datasets that their tools provide access to. Major scholarly publishers and content aggregators have the advantage of being gatekeepers to large corpora of such material and have spent the past year racing to integrate GAI into their discovery platforms. The growth of open access publishing has created opportunities for start-ups to enter this discovery space as well. Consensus, for example, pulls its source material from the over 200 million scientific papers in the open Semantic Scholar database (https://www.semanticscholar.org/), with particular strengths in queries related to "medical research and physics to social sciences and economics."

[5] Consensus introduced its Copilot feature and commented on the popularity of the "chat experience" in search and discovery on its blog in February 2024. See: Eric Olson, "Introducing: the Consensus Copilot," *Consensus*, 7 February 2024, https://consensus.app/home/blog/introducing-the-consensus-co-pilot. Perplexity, another popular startup, also has a Copilot feature for its search engine that asks follow-up questions to specify user queries, as recently discussed in *The New York Times*. See: Kevin Roose, "Can This A.I.-Powered Search Engine Replace Google? It Has for Me," *The New York Times*, 1 February 2024, https://www.nytimes.com/2024/02/01/technology/perplexity-search-ai-google.html.

Understanding

With partial exceptions such as ResearchRabbit (https://www.researchrabbit.ai/), "the Spotify for papers," and Keenious (https://keenious.com/), few vendors are investing in products designed primarily to perform traditional search better. GAI's ability to help users understand relevant material through summarization and synthesis is the core value proposition of many products on the market.

This is invariably the case with the tools released by large publishers and aggregators. For example, ITHAKA's JSTOR GAI assistant (https://www.jstor.org/generative-ai-faq), currently in beta release, allows users to ask questions about, generate summaries of, and find other content similar to the content they are viewing.[6] Elsevier's Scopus AI (https://www.elsevier.com/products/scopus/scopus-ai) search can provide summaries from abstracts of documents relevant to the user's search within the Scopus database, suggest "go deeper" questions, and identify other top researchers in the same field. Digital Science's Dimensions AI Assistant (https://www.dimensions.ai/products/all-products/dimensions-ai-assistant/), in beta, generates summaries from the Dimensions dataset based on user's queries. Commercial start-ups like Consensus, and not-for-profit vendors like Ought's Elicit (https://elicit.com/welcome) or the Allen Institute for AI's Semantic Scholar (https://www.semanticscholar.org/), are building products with similar functionality that use millions of open access research publications. Scite (https://scite.ai/), a startup purchased by Research Solutions in November 2023, draws from open repositories such as PubMed and Unpaywall, but is also signing indexing agreements with commercial publishers.

These tools point towards a future in which the distinction between the initial act of identifying and accessing relevant sources and the subsequent work of reading and digesting those sources is irretrievably blurred if not rendered irrelevant. For organizations providing access to paywalled content, it seems likely that many of these new tools will soon become baseline features of their user interface and presage an era where that content is less "discovered" than queried and in which secondary sources are consumed largely through tertiary summaries.

[6] Like JSTOR, Ithaka S+R is a service of ITHAKA (https://www.ithaka.org/).

Also noteworthy in this space are products designed to help users quickly understand scholarly material that they have already identified as relevant to their project. Scholarcy (https://www.scholarcy.com/), ChatPDF (https://www.chatpdf.com/), Adobe's AI Assistant (https://www.adobe.com/acrobat/generative-ai-pdf.html), and the aptly named TLDR This (https://www.tldrthis.com/) and Explainpaper (https://www.explainpaper.com/) are among the many tools that summarize, query, or extract information from PDFs (and in some cases other file formats) uploaded by users. The fundamental value proposition for many offerings in this space is their promise of allowing quicker and better comprehension of lengthy and complex material.

The most comprehensive products in this category combine various different features into a workflow platform. Such products present themselves as one-stop shops for the different steps of the research process. SciSpace (https://typeset.io/), for instance, combines the capability to query uploaded documents with a search engine that can also help identify related papers. It also contributes to the writing process by paraphrasing the user's text in different styles, and checks texts for AI presence. Products such as genei (https://www.genei.io/), Notion (https://www.notion.so/product), and Iris.ai (https://iris.ai/) present themselves as workspaces that help users organize their research data in addition to having GAI features. Genei, for example, offers a search engine and the ability to extract keywords or summaries from search results or user uploaded documents that can be stored with other files on the platform. It additionally includes a citation generator and a tool that rephrases or expands users' writing. As we will discuss further below, these multi-function platforms may have an edge in the market because of their ability to combine the capabilities of multiple tools being marketed separately into one workspace.

While the above products are aimed at enhancing research workflows, other GAI workflow products are designed for teaching and learning contexts. An example is Kortext Premium (https://www.kortext.com/premium-live/), an enhanced version of the Kortext study platform. This multipurpose workspace provides students access to digital textbooks and file storage, as well as generating study notes, summaries, translations, and citations. Another student learning-oriented product is Clarivate's Alethea (https://clarivate.com/products/books/alethea/), a reading assistant that asks students questions and creates tasks for them to facilitate

engagement with their reading. Alethea thus advertises itself as using GAI to make students think more critically, rather than producing work for them. In essence, products geared towards students aim to facilitate students' processes of gathering important information from course materials, as well as verifying and reinforcing comprehension.

Creation

Content generation is central to the very definition of GAI, and while content generation is a part of nearly all GAI products, we include in this category products whose primary function for teachers, researchers, or students is generating text, images, or code.

Text-generating tools market themselves as simplifying the process of moving from unformed ideas to polished text. Some products in this space use GAI to improve functions that have been part of word processing software for decades. For example, the popular platform Grammarly's (https://www.grammarly.com/) GAI features utilize user inputs to generate lists of ideas during brainstorming, as well as produce written drafts. Products such as Jenni (https://jenni.ai/) and Quillbot (https://quillbot.com/) auto-complete users' sentences. Citation generators and plagiarism detectors are other examples of GAI products serving well-established higher ed use cases.

Perhaps the most exciting developments in this space are several products designed specifically to facilitate scientific communication. Prominent products in this space include Springer-Nature's Curie (https://www.aje.com/curie/), a Microsoft Word Plugin that provides suggestions about how to improve the flow and structure of academic writing, Digital Science's Writefull (https://www.writefull.com/), and Trinka (https://www.trinka.ai/). All three products are designed to better understand the terminology and structure of academic writing than general purpose writing tools, thanks to their specialized training on scientific publications. Several tools in this category include features designed to assess what Trinka calls the "publication readiness" of manuscripts. Trinka includes tools to cross-check a manuscript against the scope of a particular journal, generate abstracts and keywords, and verify compliance with individual journals' technical and ethical standards. HeyScience's Intelligent Review tool (https://reviewer.heyscience.ai/) generates feedback on manuscripts tailored to the journal that a

researcher has identified as a potential venue for publication. Such tools have the potential to level the playing field for academics publishing in a non-native language, as well as to simplify the process of making a manuscript more likely to be accepted for publication.

Beyond text generation, GAI is known for its capacity to generate images, popularized by tools such as OpenAI's DALL-E (https://openai.com/dall-e-3), Midjourney (https://www.midjourney.com/home), Stability AI's Stable Diffusion (https://stability.ai/stable-diffusion), and Adobe Firefly (https://www.adobe.com/sensei/generative-ai/firefly.html). The baseline function of this family of tools is to allow users to create and manipulate images based on text inputs. Certain products have additional offerings: Stability AI, for example, has also released GAI-powered tools that create music and sound effects and video. While these tools that generate images and other media are not marketed exclusively for teachers, researchers, or students, they are often included on lists of products that may find use in educational contexts.

Researchers and students are also finding products that generate code useful. Meta's Code Llama (https://ai.meta.com/blog/code-llama-large-language-model-coding/), for example, produces code and natural language about code from both code and natural language inputs. Github Copilot (https://github.com/features/copilot), along similar lines, also allows users to chat about their code base and receive code suggestions as outputs. It also includes a "code completion" feature. Startup coding products with similar capabilities include Source AI (https://sourceai.dev/), Replit AI (https://replit.com/ai), and Tabnine (https://www.tabnine.com/). Often, such products can help debug code too.

Tools aimed specifically at teachers and students are more common in the creation space than in the discovery or understanding space. For example, Anthology's AI Design Assistant for Blackboard Learn Ultra (https://www.anthology.com/ai-design-assistant) and the startups Teachermatic (https://teachermatic.com/) and Curipod (https://curipod.com/) generate materials teachers can use in the classroom, such as lesson plans, worksheets, tests, or rubrics. These teacher-oriented content generators market themselves as saving teachers valuable time by facilitating the process of creating course materials and sometimes offer options to tailor the content to meet student needs. Other content generators are geared towards learning

reinforcement for students. Wolfram's Problem Generator (https://www.wolframalpha.com/problem-generator/), for instance, creates practice math problems and worksheets. The startup Wisdolia (https://www.wisdolia.com/) creates study flashcards based on content uploaded by users. Such products offer students customizable ways to practice what they have been learning in the classroom on their own time.

What's Next for Higher Education?

Products that cannot differentiate themselves or their value propositions are likely fated to fold, be acquired by bigger players, or limp along at the margins of the market.

As the sheer size of our Product Tracker indicates, the marketplace for GAI products useful to students and faculty is growing more crowded daily. The maturation of the GPT Store as a venue for reaching customers and the ease of building custom versions of ChatGPT indicate that we will continue to see new entries coming to market for some time. At the same time, many of the products in our Tracker are essentially interchangeable in terms of their core features and functionality. Products that cannot differentiate themselves or their value propositions are likely fated to fold, be acquired by bigger players, or limp along at the margins of the market. Because of the low barrier to entry involved in customizing an instance of ChatGPT, it is reasonable to expect products and players to continue to proliferate and the market to eventually consolidate around a small number of companies.

In academic spaces, the control that large commercial publishers have over large and unique portions of the most important form of scholarly content (particularly in research contexts) is a significant competitive advantage. This content adds considerable value to the discovery tools they are developing and can serve as a foundation for unified platforms with functionality across the discovery, understanding, and writing processes. What is less clear to us is how the significant investments these publishers have made in open publishing models may compete with the increased value of closed collections over the longer term. Already, the sizable corpus of OA publications has given not-for-profit and commercial competitors access to the content necessary to compete in this space: GAI may dramatically increase the commercial value of being a gatekeeper to otherwise inaccessible scholarly content.

Tools embedded within larger platforms are likely to gain advantages reaching higher education users in contexts beyond research and scholarly publishing. Learning management systems could function similarly in teaching and learning contexts. One example is the aforementioned AI course design assistant for Blackboard; another is D2L's beta features that use GAI to create quizzes and practice questions (https://www.d2l.com/newsroom/d2l-rolls-out-generative-ai-program-practice-quiz-questions/). Given the central role that learning management systems play in college courses, these organizations seem well positioned to build out suites of GAI applications aimed at course instructors and their students. The large office suites offered by Microsoft and Google are other potential vessels for the platformization of GAI in higher education.

Lurking behind both the diversity of the current marketplace and possible future of a smaller pool of major platforms is OpenAI, whose dominance of GAI in higher education—and other major economic sectors—is difficult to overstate. The vast majority of the products mentioned in this issue brief and included on our Product Tracker run on OpenAI's GPT models. Other significant LLMs such as Meta's Llama or Google's Gemini are bit players and, if anything, may be losing ground to OpenAI.[7] The slow but steady stream of universities entering into licensing agreements with OpenAI to provide GAI services to campus communities are another key measure of OpenAI's rapidly deepening entanglement with the IT infrastructure of higher education. If OpenAI and Microsoft (thanks to its close relationship to OpenAI) continue to consolidate their current advantage, they will be the true "winners" in the market.

OpenAI's dominance raises important questions for the future of GAI in higher education, particularly as universities inch towards fully incorporating the technology into their teaching, learning, and research missions. One significant unknown is what the market price for OpenAI's services will be once OpenAI shifts from pursuing market share to profit maximization. Universities are likely also unsure how many tokens might be required to fully integrate GAI across campus users and use cases,

[7] For example, a recent campus survey at the University of Baltimore showed a significantly higher percentage of students, faculty, and staff were using ChatGPT than Google Bard. See: Jessica A. Stansbury, Sarah Lausch, Nima Zahadat, David Kelly, "White Paper: AI Perceptions at the University of Baltimore," *University of Baltimore Center for Excellence in Learning, Teaching and Technology*, https://drive.google.com/file/d/1ufdagea0Xm8TpiKsyvbr1Kp-kpez3z6Z/view.

making it doubly difficult to understand the financial aspects involved in going all in on GAI.

A second important question relates to the opportunities and costs of normalizing GAI use in pedagogical and research contexts. Technology profoundly shapes the cognitive aspects of learning and how institutions articulate and measure it, and OpenAI has begun taking steps to plan an active role in this process. Its recent licensing agreement with Arizona State University included provisions that go beyond provision of services to create a "partnership" in which OpenAI will be involved in "designing, supporting and ensuring the effective use of its tools at ASU."[8] Universities, of course, have significant interest in ensuring effective use of GAI, and it is easy to see how collaboration with OpenAI could benefit the university. Even so, universities and other stakeholders in higher education will need to consider how to make sure that they retain some measure of influence over how, when, and why to incorporate a technology with so much potential to disrupt institutional practices and norms, not to mention the meaning of learning and knowledge itself.

> What we can see in the product landscape are glimpses of the individualized activities intrinsic to research and learning that the market is identifying as ripe for transformation.

Our Product Tracker has little to contribute to questions of this magnitude. However, what we can see in the product landscape are glimpses of the individualized activities intrinsic to research and learning that the market is identifying as ripe for transformation. Ethical debates in higher education about GAI usage are likewise often focused on assessing the pros and cons of very specific actions—whether using GAI to, say edit a paragraph for clarity is more acceptable than using it to generate a paragraph for subsequent editing. We will continue to benefit from engagement with action-level conversations about GAI usage to develop academic and research integrity standards and to pursue an

[8] Olivia Sanchez, "A New Partnership Paves the Way for Greater Use of AI in Higher Ed," *The Hechinger Report*, 26 January 2024, https://hechingerreport.org/a-new-partnership-paves-the-way-for-greater-use-of-ai-in-higher-ed/. See also John Warner, "ChatGPT Can't Teach Writing," *Inside Higher Ed*, 22 January 2024, https://www.insidehighered.com/opinion/blogs/just-visiting/2024/01/22/arizona-state-announces-plan-give-education.

understanding of which activities have unique value to learning and knowledge creation and which are “busy work.”

At the same time, focusing too much attention on individual steps risks losing sight of learning and insight as holistic processes. We may be very comfortable with the widespread use of GAI to help us identify information, summarize it, or bring important passages to our attention, to generate comprehensive literature reviews capable of giving us the lay of the land in a fraction of the time it would take to map it ourselves, or to improve the internal organization and clarity of a research paper or assignment. But it is also necessary to think about the cumulative impact of mediating each of these activities together to assess how best to make GAI support teaching, learning, and research fit for the purpose of producing the insights required to address complex social and scientific challenges and foster critical and engaged students, citizens, and communities.