

Agentic AI in academia: How to adopt for research, learning, and innovation



ELSEVIER

Advancing human progress together

Contents

Agentic AI in academia: How to adopt for research, learning, and innovation	4
What is agentic AI and how does it differ from generative and extractive AI?	5
How does agentic AI work?	7
Agentic AI in teaching and learning	9
Tips for choosing and using agentic AI	12
References	14

Agentic AI in academia

This high-level guide includes:

- A brief introduction to agentic AI
- A look at how it differs from extractive and generative AI
- Insights into how it can support core institutional and library goals
- Best practice tips for selecting and implementing the technologies in your institution

The rise of purpose-driven agents

It wasn't until late 2022 that ChatGPT burst onto the scene, bringing the exciting potential of generative AI (GenAI) to public awareness. Now, in the fast-paced world of technology, the focus is shifting to the next evolution of artificial intelligence – **agentic AI**.

While this new AI system draws on the large language models and systems that power GenAI, it is designed to move beyond content creation: An agentic AI solution can also set objectives, action them, iterate and interact with other tools and systems as it works to achieve the user's goal.

These capabilities have led some to predict the technology will radically change how we perform tasks and make decisions. And according to Yoav Shoham, professor emeritus at Stanford University: "The potential here is real. But we need to match the ambition with thoughtful design, clear definitions and realistic expectations. If we can do that, agents won't just be another passing trend; they could become the backbone of how we get things done in the digital world."¹


What is agentic AI and how does it differ from generative and extractive AI?

Agentic AI systems incorporate generative and extractive AI into their workflows, coordinating their use where best suited. While the table below highlights the differences between these concepts, they often work together and are not mutually exclusive.

Agentic AI	Generative AI	Extractive AI
Core use case: Solving complex, multifaceted problems.	Core use case: Creating new content.	Core use case: Retrieving and connecting information.
What is agentic AI? A sophisticated AI system that combines the data extraction and content creation abilities of GenAI with a reasoning engine, which simulates logical thinking and decision-making. This powerful blend enables agentic AI tools to break down complex requests (like booking a trip or writing a paper) into separate components and action each of them in turn, using the information surfaced to determine next steps. They can also connect to other systems (sometimes called “AI agents”) required to achieve the user’s goal.	What is generative AI? An AI system that generates fresh, original outputs in response to a user’s prompt. It does this by predicting what the most relevant response would be based on the knowledge and patterns it has ingested during its training. Large language models are an essential component of GenAI tools, enabling them to understand natural language requests and generate easy-to-understand responses.	What is extractive AI? In use for more than a decade, extractive AI technology acts like a helpful librarian, quickly scanning large volumes of data to identify the most relevant phrases, sentences or data points, which it then extracts and connects. It can also generate summaries using the extracted content. It is typically used to answer questions, highlight connections and provide concise summaries.
What kinds of queries is it designed for? While agentic AI tools can respond to any query, they really excel with those requiring action or multistep reasoning, such as open, complex or interdisciplinary queries. Academic example: “Show me emerging research trends in climate science and create a comprehensive literature review with methodology recommendations.”	What kinds of queries is it designed for? Generative AI is useful when the user wants new content created. Academic example: “Write a summary comparing the key findings from these 10 climate papers.”	What kinds of queries is it designed for? Extractive AI is great for “find me the answer to this question” queries. Academic example: “What are the main conclusions in this specific research paper?”
What form does the response take? This varies depending on the tool’s prompt engineering (guidelines) and the nature of the user’s query. It can range from real-world progress towards the user’s goal (e.g. booking of a flight) to an in-depth report.	What form does the response take? This depends on the tool’s prompt engineering (guidelines) and the user’s requirements. While the form of the output can vary – common options include text, images, video and audio – the content generated is always new and original.	What form does the response take? This varies depending on the design of the tool and the user’s query. Typically, users see a subset of the original input data pulled out, highlighted or reorganized: Extractive AI never generates new content.

How does agentic AI work?

Here's a high-level overview of how an agentic AI tool is designed to respond to a user query.



Increasingly however, agentic AI systems are becoming interactive, enabling the user to step in at key stages and tweak the strategy or direction. Many also make it possible to ask follow-up questions. The agentic AI system then determines whether it can respond to these questions from the sources already consulted, or whether a new search is required.



Agentic AI in academia

As we've seen, agentic AI is designed to process far more complex requests than other AI systems. And while human oversight remains essential from a quality, ethical and safety standpoint, in practice, agentic AI systems are designed to operate with minimal real-time supervision. For academic institutions, this offers fascinating opportunities. It enables librarians, academic leaders, researchers and students to move beyond just asking an AI system a basic query and receiving a straightforward response: It opens the door to solving core research problems and gaining deeper strategic insights. Agentic AI also has the potential to automate and accelerate more routine tasks, helping free up users to focus on work requiring critical thinking, complex problem-solving and emotional intelligence. And when it comes to teaching and learning, Kamalov, et al. (2025) see AI agents "as a promising new avenue for educational innovation."³

Crucially, agentic AI systems can allow for easier and more targeted human supervision of the processes that they support, offering users the possibility to supervise and intervene in the areas where they want to get involved. For example, in a literature review process, users could make changes to specific steps, such as excluding certain documents from consideration while including others, or modifying search parameters mid-process based on emerging patterns. This selective intervention means that researchers can maintain control over critical decision points while still benefiting from the AI's ability to handle the orchestration of complex, multistep tasks.

Although use of agentic AI is growing, many universities have yet to incorporate it into their digital strategies. But in a 2025 preprint, Gridach, et al. look at how the technology can aid standard academic tasks, such as hypothesis generation and literature reviews.⁴ They also point to its potential to "democratize access" to research tools. For example, the natural language queries and easy-to-digest responses enable users to explore highly academic topics, whatever their research capabilities or familiarity with the subject.

Agentic AI also offers improved support for specific language and learning needs. As scientist and venture capitalist Dr Gerald Chan notes: "Education is a perfect use case for agentic AI because educational resources around the world are both inadequate and unevenly distributed."⁵

But if institutions want to maximize the benefits of agentic AI, they are going to need to think carefully about how they use these systems. For example, it will be important to consider whether their query requires the computing power of agentic AI. And, if so, what kind of prompt(s) will elicit the information they are seeking?

As with all AI tools used in research and education, it will be important to *select systems thoughtfully and with intent*. This includes opting for tools that are transparent and grounded in peer-reviewed, scholarly content. And always checking to see how user data will be protected and used.

It will also be critical for users to ensure they remain the human-in-the-loop at every stage of the process, viewing the information generated as a stepping stone, rather than an end destination: As John Hopkins Sheridan Libraries advises its researchers: "Treat AI as a starting point, not a substitute for critical thinking and peer-reviewed sources."⁶

These are points we know librarians feel strongly about too. For example, when we asked information professionals what features they need in a generative AI tool, transparency, tool governance and human oversight topped the list.⁷

*"It may be helpful to consider the agent to be the virtual equivalent of a Ph.D.-holding research and teaching assistant working 24-7 on the project without break, every day of the year."*²

Ray Schroeder,

Senior fellow for UPCEA: the Association for Leaders in Online and Professional Education

Agentic AI and the library

Many believe that with their information expertise, technology fluency and teaching abilities, librarians are perfectly placed to guide the introduction and use of agentic AI on campus. But perhaps more importantly, emerging technologies like agentic AI offer librarians the opportunity to reinforce the value that they offer.

As Don Simmons, Assistant Professor at Simmons University's School of Library and Information Science commented in a recent *Library Connect* article: "Whether you perceive AI as a good or bad thing, it offers an opportunity to amplify our profession, so that we not only remain relevant, but become the source to go to when anyone wants to learn more about AI navigation and AI literacy."⁸

It's a view echoed by Marydee Ojala, an Editor-in-Chief for *Computers in Libraries* magazine: "Agentic AI is here to stay and it's worth the time and effort of librarians and information professionals to ascertain its relevance to the profession and to exploit the technology so that it serves the best interests of library clientele and librarians themselves."

Although she agrees that agentic AI is a "new competency" for the library community, Ojala doesn't view it as a new skill. "The empathy central to the work of human librarians is often missing from AI agents. Thus, the combination of emotion and technology fits perfectly into the role of AI librarian, whether it's called by that title or not..."⁹

Agentic AI offers librarians the opportunity to draw on their unique expertise and skillset to support areas such as:

- **AI literacy.** For example, they can educate users about the capabilities and limitations of agentic AI tools and address concerns about privacy, bias and misinformation. They can also champion the responsible and ethical use of these tools, including their broader societal impact.
- **AI consultancy.** This includes collaborating with other departments to develop institutional AI policies, and advising on academic integrity issues related to the use of AI in research and coursework. Members of the Association of Research Libraries (ARL) report that they are being approached by other departments to "be partners on campus and support institutions" with AI-related issues.¹⁰ In response, ARL has developed [seven principles](#) members can draw on when responding to these requests.

Importantly, **librarians can also take on the role of guide** – steering their library users toward specific agentic AI tools that meet their own and institutional standards. And, as a recent survey shows, that guidance will be valued: 58% of students told us they would be confident that a library-approved tool aligns to their institution's AI policy, making them more likely to use and adapt it.¹¹

Agentic AI can also play a role in enhancing library services, as Ojala explains: "... agentic AI could be instrumental in anticipating the research needs of library users, suggesting relevant resources and identifying any research gaps it perceives. For collection development, an AI agent could not only suggest materials to purchase but also initiate purchase orders independently. It might also take on weeding duties."⁹

Agentic AI in *teaching and learning*



As Kamalov, et al. note, education is an area in which the innovative potential of agentic AI can excel. It can help to make learning more personalized, efficient, engaging and inclusive for students. It can also equip them with the AI skills and literacy they need to succeed in their future careers.

Agentic AI can also support educators, aiding them with course content creation and delivery, as well as student engagement. For Chan, agentic AI only reinforces the value of the educator role. “Along with knowledge acquisition, the students need to make sense and be able to make use of what they are learning. This is where a great professor comes in.” But he believes that the rise of AI is going to require universities to “reimagine the classroom, reexamining each aspect of the educational process and redesigning the educational experience and outcome.”⁵

And that includes evaluation. Like the introduction of the calculator and the launch of Wikipedia, technologies like generative AI and agentic AI are leading educators to think carefully about how they should evaluate students and their work. With so much information so readily to hand, some universities are now shifting away from assessing the work produced and are focusing more on the process followed to complete it.¹³

“At Jefferson, we work hard to make sure our graduates are prepared for the future of work in an era of – I’ll call it – ubiquitous computation... We want them to be facile and comfortable working with new technological resources like AI; and teach them to balance technologies’ strengths and weaknesses by applying empathy, creativity, intuition and human caring.”¹²

Susan C. Aldridge,

PhD, President, Thomas Jefferson University,
United States of America

Moving toward an agentic AI future

With agentic AI evolving rapidly, no one can say with any certainty how it will change our lives. However, Adrian Raudaschl, a Principal Product Manager at Elsevier, does have some thoughts on what may happen in the short term.

“I can see a scenario where users, whether they are academic leaders, educators, researchers or librarians, will orchestrate a small team of these agents. The agents will work on four or five different requests in the background, bringing them the latest insights or preview considerations. The user can then take ownership of that information and use it to advance their goals; for example, identify the hypothesis they want to research next, craft a lesson plan or develop a strategy.”



“One early adopter told us that the report Deep Research generates could save their team 4-5 days of work.”

Adrian Raudaschl,
Principal Product Manager for Scopus AI, Elsevier

Adrian leads the Elsevier team behind a new agentic AI feature that has just been launched on Scopus AI – [Deep Research](#). **Just like Scopus AI, Deep Research draws exclusively on the peer-reviewed and curated content in Scopus** and it synthesizes its findings into a comprehensive report that Adrian and his team have intentionally designed to spark critical thinking.

And each step taken to research and prepare the report is shared on screen so users can follow the process in real time.

According to Adrian, Deep Research was developed in response to a user desire for more and deeper insights. He explains: “Since the launch of Scopus AI, one of the most popular elements of its response has been the automatically generated Go deeper questions. So, in April 2025, we released **Conversational Follow-Up**, enabling users to pose their own additional questions.

“But what we soon noticed is that up to half of the Conversational Follow-Up questions people asked weren’t related to their original query. Instead, they wanted to explore beyond the borders of that topic or reinterpret the information. And these were questions that Scopus AI couldn’t easily answer. With Deep Research that has changed. In fact, we’ve been asking some of our partners in life sciences to think of really hard questions we can test it with, and the results have been very promising.”

On a recent tour of universities in Asia, Adrian also got to see firsthand how Deep Research is helping students initiate deeper and more critical conversation with each other and their mentors.

Students especially valued the tool's ability to deliver comprehensive, multi-perspective analysis grounded in credible sources. Many highlighted how its transparent reasoning and easy access to full-text references empowered them to verify insights and explore topics more deeply. As one student put it, "It's better than what I would set up—that's for sure about research design." Another described the experience as feeling like a detective scene in a movie—where all the clues suddenly come together and the mystery is solved—thanks to the system's clarity, speed and visual presentation. Above all, they appreciated how Deep Research encouraged thoughtful engagement rather than overreliance on AI.

Elsevier's development teams are already busy looking at opportunities to further leverage agentic AI in Scopus AI and other Elsevier solutions. Adrian says: "We've got lots of ideas and are working to make them happen this year. Agentic AI offers so much potential – technology is now moving from a simple query/response format to providing capabilities for deeper intellectual work."

However, he and the team plan to introduce these new developments carefully, step by step.

"This is new technology, and it comes with risks. For us, ensuring that we introduce changes incrementally and take a test-and-learn approach is really important. It's not only the right thing to do, it also aligns with our [responsible approach to AI](#) at Elsevier."



[Learn more about Scopus AI](#)



Tips for choosing and using agentic AI

Four important questions to ask when selecting an agentic AI tool

Like every form of technology, agentic AI comes with risks attached, so it's helpful to bear these points in mind when assessing potential solutions.

1. Does it promote 'human-in-the-loop'?

There are understandable concerns that as the use of agentic AI increases, it could replace traditional academic roles or stifle creativity, originality and critical thinking. But if a tool is developed responsibly, that doesn't have to be the case – it can be structured intentionally to support, rather than replace, human critical thinking, acting as a catalyst for deeper reflection and reasoning. According to Hosseini and Seilani (2025), one of the ways that AI developers can ensure this happens is to design agentic AI systems that work with humans in collaborative partnership; for example, introduce “dynamic goal-sharing, negotiating in real time, shared decision-making, and adaptive task allocation.”¹⁴

2. Does it explain how it works?

With many AI systems, it's not always clear how they have reached their responses or the sources they have used. To combat this, Viswanathan (2025) points to the importance of making agentic AI transparent. “Systems must be designed with inherent explainability features that allow stakeholders to understand the reasoning behind autonomous decisions. This includes implementing mechanisms for tracking decision pathways and maintaining comprehensive audit trails of system actions.”¹⁵

3. What steps does it take to combat bias and hallucinations?

Users want AI to deliver responses that are accurate, bias-free, accountable and fair. But with non-academic grade tools, in which responses aren't grounded in verified scholarly content, the number of hallucinations (false or misleading AI outputs presented as fact) can be high.¹⁶ In an article in *Library Journal*, Nicole Hennig, eLearning Developer at the University of Arizona Libraries, said concerns over fabricated sources have led her institution to warn students to avoid non-academic grade AI tools when looking up articles.

She explains: “The articles sound very plausible because [the AI tool] knows who writes on certain topics, but it'll make up things most of the time because it doesn't have a way to look them up.”¹⁷

While eliminating bias and hallucinations in AI tools remains a challenge, there is much that can be done to minimize them. And here again, human input remains vital, according to Gridach, et al., including “robust oversight mechanisms, human-in-the-loop architectures, and frameworks to evaluate and mitigate these risks during training and deployment.”¹⁴

4. How does it handle and use your users' data?

Another major area of concern for users is how their data and the queries they enter will be stored and handled. Finding a reputable provider that uses secure, established systems and is transparent about its privacy policies can help to address these fears.

Using agentic AI to increase AI literacy at your institution

In his interview with [Library Connect](#), Former librarian Don Simmons, now Assistant Professor at Simmons University's School of Library and Information Science suggested five basic steps that libraries can take.⁸

1. Familiarize yourself with the technology.

2. Build your AI literacy skills.

3. Make the rules clear. This includes liaising with colleagues and the institution administration to ensure AI policies for users are up to date and clearly promoted. If none are currently in place, work together to establish them.

4. Don't reinvent the wheel. According to Simmons: “There are so many different examples of courses and trainings out there already.”

5. Launch your own AI trainings. “These can be as simple as one-shot workshops on how AI can help students build their resumes, or more complex programs. Don't worry if you are still fairly new to AI. In our profession, I firmly believe no one is truly an AI expert. We are all learning all the time” says Simmons.

Additional options include:

- **Developing exercises to nurture students' evaluation skills:** These can take the form of fact-checking challenges; for example, librarians can present students with fake news articles, biased summaries or plagiarized texts and guide them to critique and detect flaws using their AI literacy skills. These exercises can also be used to reinforce the importance of checking the credibility of sources when using non-academic grade tools. Other sessions can look at examples of ethics and bias in AI, with students encouraged to think about whether changes to data sources and system design could have helped to prevent these issues.
- **Holding training sessions for colleagues and faculty:** These can take the form of workshops, tutorials or online resources. And as Simmons says, they don't have to be elaborate. For example, Hennig has plans for an online "AI Tool Exploration Hour," during which colleagues will be able to spend time "individually or collectively playing with and exploring one or more [AI] tools," with breakout groups and in-person meetings optional.¹⁷
- **Using agentic AI to help develop lesson plans and information literacy materials:** Agentic AI can play a role in mapping out potential lesson plans, as well as contributing to the content of the lessons themselves. It can also help to craft AI FAQs and LibGuides.
- **Sharing learnings:** If there are other librarians, faculty or students actively using agentic AI, creating a space for them to share their experiences can help to guide responsible use of these tools on campus.



References

- 1 Shoham, Y. (July 2025). Don't Let Hype About AI Agents Get Ahead of Reality. MIT Technology Review.
<https://www.technologyreview.com/2025/07/03/1119545/dont-let-hype-about-ai-agents-get-ahead-of-reality/>
- 2 Schroeder, R. (February 2025). Setting a Context for Agentic AI in Higher Ed. Inside Higher Ed.
<https://www.insidehighered.com/opinion/columns/online-trending-now/2025/02/04/setting-context-agentic-ai-higher-ed>
- 3 Kamalov, F. et al. (April 2025). Evolution of AI in Education: Agentic Workflows. arxiv.
<https://doi.org/10.48550/arXiv.2504.20082>
- 4 Gridach, M. et al. (March 2025). Agentic AI for Scientific Discovery: A Survey of Progress, Challenges, and Future Directions. arxiv.
<https://doi.org/10.48550/arXiv.2503.08979>
- 5 Chan, G. (February 2025). Rethinking Clark Kerr: The Uses of the University in the Age of Generative AI. The Inaugural Dean's Distinguished Lecture, College of Computing, Data Science and Society, University of California, Berkeley, California.
<https://www.geraldchanspeeches.com/2025/02/10/1251/>
- 6 Using AI Tools for Research. (Accessed 01 August 2025). John Hopkins Sheridan Libraries. <https://guides.library.jhu.edu/c.php?g=1465762&p=10904518>
- 7 Insights: Librarian attitudes towards AI. (2024). Elsevier. <https://www.elsevier.com/insights/attitudes-toward-ai/librarians>
- 8 Willems, L. (July 2025). The role of AI in universities is growing – what does that mean for librarians? Library Connect.
<https://www.elsevier.com/connect/the-role-of-ai-in-universities-is-growing-what-does-that-mean-for-librarians>
- 9 Ojala, M. (January 2025). Agentic AI and AI Librarians: Future roles for the profession. Information Today Europe.
<https://www.infotoday.eu/Articles/Editorial/Featured-Articles/Agentic-AI-and-AI-Librarians-Future-roles-for-the-profession-167436.aspx>
- 10 Coffey, L. (May 2024). New AI Guidelines Aim to Help Research Libraries. Inside Higher Education.
<https://www.insidehighered.com/news/tech-innovation/libraries/2024/05/01/new-ai-guidelines-aim-help-research-librarians>
- 11 Librarian Futures Part IV. (2025). Technology from Sage.
<https://www.technologyfromsage.com/whitepapers/>
- 12 Thomas Jefferson University is Pioneering Professions-Focused Education for a Dynamic Century (accessed 01 August 2025). Inside Higher Ed – content sponsored and provided by Thomas Jefferson University.
https://narratives.insidehighered.com/thomas-jefferson-university-pioneering-professions-focused-education-dynamic-century/index.html?mc_cid=38004dd169&mc_eid=03820262d6
- 13 The Rise of Generative AI in Education. Embracing Digital Transformation. Podcast episode 158.
<https://embracingdigital.org/en/episodes/edt-158/index.html>
- 14 Hosseini, S. & Seilani, H. (2025). The Role of Agentic AI in Shaping a Smart Future: A systematic review. Array. Volume 26. <https://doi.org/10.1016/j.array.2025.100399>.
- 15 Viswanathan, P. S. (2025). Agentic AI: A Comprehensive Framework for Autonomous Decision-Making Systems in Artificial Intelligence. International Journal of Computer Engineering and Technology. Vol. 16 No. 01.
<https://ijcet.in/index.php/ijcet/article/view/223>
- 16 Dolan, E. (April 2024). PsyPost. ChatGPT Hallucinates Fake But Plausible Scientific Citations at a Staggering Rate, Study Finds
<https://www.psypost.org/chatgpt-hallucinates-fake-but-plausible-scientific-citations-at-a-staggering-rate-study-finds/>
- 17 Thornton, H. (April 2024). AI in Academia. Library Journal.
<https://www.libraryjournal.com/story/academiclibraries/ai-in-academia>