eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

Our place in the Agentic AI loop: The value of information professional competencies

Purpose: To analyse the concept of Agentic AI and the relevance of information professional competencies to them.

Design/methodology/approach: The paper is based on literature and a small scale testing of the deep research services from Gemini and ChatGPT.

Findings: The eight promises that are made in the discourse around Agents are elaborated, with the questions an information professional would ask. The main features of deep research agents are analysed. Information professionals are likely to be concerned about the sources of information in use and the reliability of outputs, and the wider societal impacts such as on the environment.

Originality: The paper is an early review of the trending Agent/ agentic AI concept.

Research limitations/implications: We do not know how far the promises of the discourse around agents can be really delivered.

Practical implications: Information professionals should be involved in configuration of agents and training of users.

Keywords: AI, Artificial Intelligence, Agents, Agentic AI, Information quality, Interoperability, Information professional roles

The promises of agentic AI

According to Reuters, 2025 is to be year of AI Agents (Paul, 2024). For a taster of what agents are we can turn to a one minute advert presented by the actor Idris Elba (https://www.youtube.com/watch?v=UfZ7GUAguxw). The video shows him walking round an office building quickly gathering a crowd of followers and saying:

> *"Only ServiceNow connects every corner of your business putting AI to work for people... So Kate in HR... can focus on people not process. Patty in IT is using AI agents to deal with the small stuff so she can work on the big stuff. Over here, AI gives Tina the info she needs to get the job done. AI helps Jim solve customer problems before they are problems."*

The 100-word text accompanying the video expands on these claims:

> *"Only the ServiceNow platform puts the power of AI agents to work across every corner of your business. Imagine a team of ServiceNow AI Agents proactively solving challenges in IT, HR, customer service, finance—you name it. These agents collaborate, learn from each other, and continuously improve, handling the busy work so your teams can focus on what truly matters. They boost*

productivity for employees, enrich customer experiences, and make work better for everyone. It's time to put ServiceNow AI Agents to work for your people."

The video has had around 800,000 hits on YouTube as at July 2025. My purpose in quoting this advert is not to ask questions about ServiceNow as such. Rather I think the script is helpful in establishing some of the key aspects of the discourse around agentic AI. It also provides an opportunity to surface the questions about that narrative that arise from a information professional perspective.

It appears that the advert is making eight promises for agentic AI, to each one of which there is a doubt or concern we might raise from an information perspective. The script talks about allowing a manager to "focus on people not process ... deal with the small stuff so she can work on the big stuff.. ... handling the busy work so your teams can focus on what truly matters." These quotes seem to make an empowerment promise, where AI frees up time to perform more interesting or important tasks. In the YouTube advert it seems that we are being told that departmental managers are empowered. But we are likely to ask how is management control exercised, when the manager loses direct contact with the detail? And, if some people are empowered, who might lose out?

"Putting AI to work for people" seems to make a "human in the loop promise". But that leaves us with the questions of who that human is, what control do they have, and what competencies do they need to carry out this effectively?

The R&D manager gets "...the info she needs to get the job done..." This is an information or data promise. However, we might reasonably ask, from our recent experience of generative AI, how accurate is that information? How are conflicts in sources handled where, as is so often the case, there is no single point of truth?

The often repeated mantra in the advert that agentic AI "...connects every corner of your business..." is an interoperability or orchestration promise. But we are likely to ask because this has been historically so difficult for technical and political reasons, how well can data from different parts of an organisation be integrated in practice? How are issues of access rights, confidentiality and privacy to be handled?

"[AI agents] boost productivity for employees" is a productivity promise. By persuasive use of "by" rather than "of" it deflects the thought that productivity increases imply intensification of work effort or even potential redundancies.

"AI agents ... enrich customer experiences" is a service promise. But do we know how, when and on what terms customers (and staff) want to interact with AI, given question marks about its ability to have any level of emotional intelligence?

While not explicit the short movie shows an inclusive workplace where a diverse workforce is represented. Implicit is the promise that AI is consistent with diversity, equity and inclusion. Yet we are bound to ask if that is really the case? A recent Danish

study found "younger, less experienced, higher-achieving, and especially male workers leading the curve" (Humlum and Vestergard, 2025).

Finally, again implicitly, by virtue of not being referenced, there is the promise that there are no environmental implications or socially harmful effects of technology. But we are increasingly asking about the social sustainability of AI solutions (Crawford, 2021; Estampa, 2023).

At this point in the development of the technology we do not know how far these promises can be delivered or if our doubts are justified: but it is important to ask such questions, because they are rooted in the long experience of information professionals faced with technologies that promise to solve organisational problems.

Turning from this rhetoric to more formal definitions of agents, Hosseini and Seilani (2025) write that agents are:

> A category of AI systems capable of independently making decisions, interacting with their environment, and optimizing processes without direct human intervention.

This definition brings out key features: autonomy or agency; complex tasks that cannot be pre-programmed; and learning from the environment. The Gartner definition is also useful in helping us see a connection to the AI we have already seen such as Robotic Process Automation (RPA) (Lin, Chiu & Lam 2024) or generative AI.

> "AI agency is a spectrum. At one end, traditional systems with limited agency perform specific tasks under narrowly defined conditions. At the other end, future agentic AI systems with fully agency will learn from their environment, make decisions and perform tasks independently. [...] Agentic AI refers to goal-driven software entities that have been granted rights by the organization to act on its behalf autonomously make decisions and take action." (Coshow et al., 2024)

Agents working with each other is another feature of the discourse.

As Ojala (2025) notes the aspiration to offer all this is not very new, the fundamental idea has been around since at least the 1990s, and in many respects the promises and questions are the same across all forms of AI. Agents represent a further step along a known path, with an emphasis on certain aspects such as autonomy and taking on more complex tasks.

Possible applications of agents in libraries are easy to imagine: Such as responding to patron queries, analysing data about service use or performance or analysing policy context to support strategic decision making. In the wider educational context Kshetri (2025) suggests that agents might have uses in real time, adaptive pedagogic assistance to students, answering questions about study options or to support admissions communications.

Deep research

In order to get a fuller sense of the character and capabilities of such agents I closely examined two examples that are already available, choosing them in a domain of particular interest to information professionals: evidence review. I analysed the "deep research" features of ChatGPT and Google Gemini. Both these services create research reports, involving finding and analysing sources, synthesising them, and producing a final written report and other possible outputs. In promising to actively carry through complex tasks with some independence and learning, these exhibit key features of AI agents. In addition, I looked at documentation for CoPilot's "Researcher" reasoning agent which offers similar capabilities.

I must acknowledge up front that this was not an objective evaluation. I used a free version of ChatGPT, not any of the more advanced models. In contrast, I was using a (University of Sheffield) institutional version of Gemini and I do not know exactly how this has configured. The reader might well get different results from the same services depending on their profile. Given the speed of change the detail of the analysis is likely to be out of date before published anyway. So, my report of findings should not be treated as a road test, more an exploration of the dimensions of potential agent capabilities.

The main prompt I used was "What are the main theories that describe how academic research has a real-world impact?". I observed that ChatGPT was more interactive, asking me to clarify my requirements and actively suggesting possible useful outputs (such as a comparison table). While Gemini was less conversational it did lay out its step by step "research plan" to conduct the review before it started and offered me an opportunity to edit it. It also reported against the plan at the end of the process. ChatGPT produced a shorter report of 5 pages based on 15 sources, more quickly. Gemini produced a 27-page report based on 33 references. Gemini offered more alternative outputs such as visualisations and a podcast (although it failed to actually produce the podcast when requested). Overall, the text outputs in both cases were impressive, coherent documents summarising a large amount of complex material, in a few minutes. One may be better or worse currently, but fundamentally they are similar and give us a sense of potential capability of AI agents in such a task. I think we can assume they are likely to improve over time, so focussing on actual performance is secondary to reflecting on the concept.

The significant limiting factor for both deep research agents was that neither had access to licensed library resources or other local institutional resources. This greatly proscribed the value of the actual reports for me as an academic researcher. Both reports were largely based on open web sites with a predomination of US (and with Gemini UK and Irish) material. In both cases all sources were in English language. For this reason, I would be troubled if users, such as students or researchers, relied on

deep research today. While CoPilot is intended to operate with local intranet resources, a future barrier is likely to be negotiating access to licensed library content for these to be genuinely useful in academic context. Furthermore, the assumption of the interoperability promise with agentic AI is that issues such as the politics of information sharing in organisations can be solved. As information professionals we are aware of the political and power issues underlying unwillingness to share information across institutional silos, as well as more technical barriers. Also, even if we gain access to the data, is it accurate or consistent? Data quality issues mean that the data we have for decision making is often fragmentary, contradictory and inaccurate. Indeed, data quality issues are reported as a major obstacle to use of AI (Mehri, 2022). The technology cannot itself solve these organisational issues of interoperability. It could be that such political, legal and ownership issues rather than technology that will limit the value of agents from an information perspective.

We are also likely to have doubts about the information promise made for agents. Two years of experience with generative AI has taught us to anticipate a number of problems. Firstly, there will be a question about the accuracy of outputs, the hallucination of information and sources. I did not see invented citations in the results of my queries. But it feels unlikely that such agents will be more accurate than current generative AI in more complex tasks, except in circumstances where inputs are very carefully controlled. Secondly, there is the issue around the reproducibility of outputs: how useful is a report that is significantly different the next time the service is asked the same question? How is accountability for decisions affected by the impossibility of reproducing results? Thirdly, is the tendency to sycophancy: to agree with the user and try and keep you happy as a consumer. Fourthly, and critically, is the way that use of such reports would create dependency and in the long term undermine the users' skills to perform the task themselves, plus the implications of not doing the work manually and focussing more on checking. Writing a complex report requires lots of thought and analysis. An agent created report may give us a final answer, but what is lost by not going through the process of the research? And rather than having skills to research and analyse data for a report, emphasis shifts to checking a given report. That is a different type of task and there is a potential loss of perspective in this shift (Simkute et al., 2025).

The role of information professionals

So, the promising adverts, while not necessarily hype, leave us with a lot of questions from an information professional perspective. Yet our focus on information sources and information quality and our understanding of information politics, gives us two very important roles in supporting the use of agentic AI. The first role is in configuring services so that their defaults are most likely to produce reliable outputs. Above all

securing the breadth of reliable data and information sources suitable for different types of question is surely an information professional role.

There is also an aspect of fine tuning of interaction. For example, it is possible to define triggers when an agent service is offered to the user. When to do so needs to be carefully considered. Example searches can be configured to indicate to the user relevant uses of a particular agent. Copilot documentation also reveals how the tone of agent interaction can be determined. The text suggested for this in the CoPilot documentation is absurdly informal.

> "*Your tone should be friendly, helpful, cheerful, and expressive.*
>
> - *Always greet users warmly and use a smiley emoji.*
> - *Use positive language.*
> - *Offer your help proactively.*
> - *Use emojis and exclamation marks for cheerfulness.*
> - *Keep the conversation engaging with expressive language.*
> - *End conversations positively.*"

Information professionals would want input on how agents are configured. For example, the agent could be set up to state its own estimate of the reliability of its outputs and reminders of how to prompt and evaluate outputs.

In so far as configuration cannot make agents wholly transparent and safe, the second role for information professionals is then training people in their effective use. This would consist of writing good prompts, in checking research plans, in evaluating outputs, and in how to offer useful feedback into the system to enhance its long-term performance. One suspects that while users are excited about AI chatbots, they do not use them in a particularly sophisticated way. Good training will enhance reliability. Growing experience with such systems will enable us to say better what kind of questions the reader should ask when evaluating outputs.

Then we also need to factor in wider ethical issues such as environmental impact, impact on labour displacement, and workplace diversity. Educating users about these aspects is a key aspect that should not be forgotten.

The information professional's recurrent dilemma around AI, arising with generative AI is likely to recur with agents. It is that as professionals placing a high value on information quality we see many problems with manifestations of AI. Yet users like their ease. We have to work out how to remain open to users' optimism, while educating them about safe use. We need to train users to use them effectively at a practical level. But we also to remind them of the limitations and ethical dimensions.

References

Coshow, T., Gao, A., Pingree, L., Verma, A., Scheibenreif, D., Khandabattu, H. and Olliffe, G. (2024). Top Strategic Technology Trends for 2025: Agentic AI. Gartner Group, https://www.gartner.com/doc/reprints?id=1-2K8Y7LEY&ct=250212&st=sb

Crawford, K. (2021*). The Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Estampa (2024). Cartography of generative AI, https://cartography-of-generative-ai.net/

Hosseini, S., & Seilani, H. (2025). The role of agentic ai in shaping a smart future: A systematic review. Array, 100399.

Humlum, A. and Vestergaard, E. (2025). *Large Language Models, Small Labor Market Effects*. University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2025-56, http://dx.doi.org/10.2139/ssrn.5219933

Kshetri, N. (2025). Revolutionizing Higher Education: The Impact of Artificial Intelligence Agents and Agentic Artificial Intelligence on Teaching and Operations. *IT Professional*, *27*(2), 12-16.

Lin, C. H., Chiu, D. K., & Lam, K. T. (2024). Hong Kong academic librarians' attitudes toward robotic process automation. *Library Hi Tech*, *42*(3), 991-1014.

Merhi, M. I. (2023). An assessment of the barriers impacting responsible artificial intelligence. *Information Systems Frontiers*, *25*(3), 1147-1160.

Ojala, M. (2025). Agentic AI and AI librarians: Future roles for the profession. *Information today*. https://www.infotoday.eu/Articles/Editorial/Featured-Articles/Agentic-AI-and-AI-Librarians-Future-roles-for-the-profession-167436.aspx

Paul, K. (2024). Autonomous agents and profitability to dominate AI agenda in 2025, executives forecast. Reuters. https://www.reuters.com/technology/artificial-intelligence/autonomous-agents-profitability-dominate-ai-agenda-2025-executives-forecast-2024-12-12/

Simkute, A., Tankelevitch, L., Kewenig, V., Scott, A. E., Sellen, A., & Rintel, S. (2025). Ironies of generative AI: understanding and mitigating productivity loss in Human-AI interaction. *International Journal of Human–Computer Interaction*, *41*(5), 2898-2919.