

The AI That Does Things For Us: A Guide to the Agentic Age, and What to Watch

Hafiz Hanif

Faculty of Human Development, Sultan Idris Education University · ORCID 0000-0001-9943-6557

CITE AS	Hanif, H. (2026, August 17). The AI That Does Things For Us: A Guide to the Agentic Age, and What to Watch. Hack(edu).tech Chronicle. https://hackedu.tech/articles/the-ai-that-does-things-for-us-a-guide-to-the-agentic-age-and-what-to-watch
PUBLISHED	17 August 2026
DOI	10.5281/zenodo.21972037
LICENCE	CC-BY-4.0
AVAILABLE AT	https://hackedu.tech/articles/the-ai-that-does-things-for-us-a-guide-to-the-agentic-age-and-what-to-watch

ABSTRACT

Artificial intelligence has crossed from answering to acting. Agentic AI systems plan, use tools, and complete multi-step goals without step-by-step human steering, offering genuine relief from administrative burden. This article welcomes that promise, then identifies four risks grounded in research: automation complacency and bias, harms that scale with autonomy and underspecified goals, evidence that teacher-in-the-loop hybrids outperform fully autonomous tutors, and the atrophy of skills we no longer practise. Delegate tasks, it argues, never outcomes.

KEYWORDS

Agentic AI

Artificial intelligence has quietly crossed a line, and most people have not noticed yet. The AI everyone knows is the chatbot: a brilliant consultant sitting across the table. You ask, it answers, and then it waits. Every step still belongs to you. If you want a report written, you request the outline, then the draft, then the corrections, one prompt at a time. You remain the project manager; the AI is a worker who performs exactly one task per instruction.

The new generation of AI is different. Give it a goal, "find me three flights to Kota Kinabalu under RM400, compare them, and book the best one," and it breaks the goal into steps by itself, opens the search engine, reads the results, weighs the options, fills in the forms, recovers when a page fails to load, and reports back when the job is done. The industry calls this *agentic AI*: systems that do not merely answer, but act. The research community describes these agents as systems built from four capabilities working together: a profile of who they are supposed to be, a memory of what has happened, a planning ability that decomposes goals into steps, and the capacity to take actions using external tools (Wang et al. 2024). Strip away the jargon and the definition is simple. A chatbot talks. An agent does.

Let me first give this development its due, because it deserves it.

For years I have written about the burden of non-teaching work, the reports, the data entry, the forms that ask for what three other systems already know. The chatbot era helped at the

margins: it could draft the text, but a human still had to shepherd every step. The agentic era is the first genuine answer to that burden. An agent can be told, once, "every Friday, compile the attendance data, flag students absent more than three days, draft the follow-up letters to guardians, and put the summary in my inbox," and then simply do it, week after week. It can watch a shared folder, sort what arrives, chase what is missing. For a teacher, a clerk, a small business owner, this is not an incremental convenience. It is the difference between a tool you operate and a colleague you delegate to. The dream of automation, promised since the first computer entered a school office, finally has machinery capable of keeping it.

My students see it too. Give an agent a project brief on Monday and it returns with the completed work: researched, structured, referenced, formatted. Impressive. Genuinely, honestly impressive.

And that is precisely why the rest of this column exists. Because delegation, as every manager and every parent knows, is never free. Here are the four things I believe we must watch, each grounded not in science fiction but in published research.

Watch the first: our attention leaves the room.

Human beings have a documented weakness around automation, and it predates AI by decades. The classic research on pilots, doctors, and operators found two reliable patterns: automation complacency, where we stop monitoring a system that has been correct many times in a row, and automation bias, where we accept a machine's recommendation without checking, making errors of omission and commission that, disturbingly, training and instructions do not eliminate (Parasuraman & Manzey 2010). These findings came from cockpits and control rooms staffed by trained professionals. Now scale the problem: an agent performs not one automated step but ten in sequence, invisibly. A small error in step two quietly contaminates steps three through ten, and the polished final output gives no hint of the rot inside. The more capable the agent, the fewer times we check, and the fewer times we check, the more catastrophic the eventual failure we miss. This is not a flaw in the machine. It is a flaw in us, and it is the single most predictable danger of the agentic age.

Watch the second: harm scales with autonomy.

A landmark paper by twenty-two researchers put the warning in precise terms. As algorithmic systems become more agentic, gaining goal-directedness, long-term planning, direct impact on the world, and freedom in how they accomplish underspecified goals, the potential harms grow with them, including harms nobody anticipated because nobody specified what the system should not do (Chan et al. 2023). The word to hold onto is underspecification. When I ask a chatbot a question, the worst outcome is a bad answer I can ignore. When I give an agent a goal, "maximise my Threads engagement," "make sure every student submits," I have handed over not just a task but the choice of methods, and machines pursue goals literally, without the common sense that quietly fills the gaps in every human instruction. The agent that spams, pesters, or cuts corners is not malfunctioning. It is doing exactly what we said, rather than what we meant. Accountability follows the same slope: when a ten-step chain produces a harmful outcome, who answers for it? The user who gave the goal? The developer? The agent that chose the path? Our institutions, schools included, have no ready answer.

Watch the third: in education, the evidence already favours the human in the loop.

This is not merely my instinct. A systematic review of eighty-two studies on AI-powered educational agents, spanning tutoring, assessment, feedback, and curriculum design, arrived at a striking convergence: hybrid workflows, in which teachers curate and moderate the AI's output, outperform fully autonomous tutors (Córdova-Esparza 2025). Read that again. The best-

performing configuration is not the one with the most automation. It is the one where a human educator remains in the loop, combining the machine's scale with the teacher's judgement. Meanwhile, a scoping review of large language models in education catalogued the practical and ethical challenges that autonomous deployment raises, from privacy and transparency to equity, that remain largely unresolved (Yan et al. 2023). The agentic pitch will arrive at our schools dressed as a fully automatic tutor, marker, and counsellor. The research, at this moment, says the winning design keeps the teacher's hands on the wheel.

Watch the fourth: what delegation does to the delegator.

Here is the question I find myself asking most. When an agent completes a student's assignment end to end, research, drafting, referencing, submission, what exactly has the student practised? Delegation. Only delegation. Regular readers will recognise this as the authorship question from my column on vibe-lesson-planning, now at a larger and more dangerous scale, because the agent no longer waits for our review between steps. The struggle we outsource is the learning we forfeit. And there is a quieter cost for working adults too: skills unpractised are skills that decay. A generation that has never compiled its own report may one day be unable to tell a good report from a confident-looking bad one, which is precisely the moment automation bias becomes fatal.

So where does this leave us? Not, I want to stress, with rejection. I use agentic tools daily, and I will teach others to use them, because the efficiency is real and the burden it lifts, especially the administrative burden crushing our teachers, is real. The trade is worth making. But every trade has terms, and ours should be these.

- First, delegate tasks, never outcomes. The agent can run the errand; a named human owns the result, reads it, and answers for it.
- Second, build checkpoints where the stakes are high. An agent drafting letters to parents needs a human eye before sending, every single time, no matter how many times it was right before, because complacency is precisely the state of having been right many times in a row.
- Third, keep doing some of the work manually, deliberately, the way pilots still hand-fly approaches to keep their skills alive.
- And fourth, teach supervision as a skill. We spent two decades teaching students to use tools. We must now teach them to manage machine workers: to specify goals precisely, to inspect intermediate steps, to know when the confident output is wrong. That is a literacy our curriculum does not yet have a subject code for.

The chatbot era asked: what can AI tell us? The agentic era asks a more intimate question: what are we willing to let it do for us, and what must remain ours?

My answer, for now, is this. Let the machine do things for us. But let us never stop being able to do them ourselves, and let us never stop checking. Because an assistant we cannot evaluate is not an assistant.

It is a replacement, arriving politely, one delegated task at a time.

References

- Chan, A., Salganik, R., Markelius, A., Pang, C., Rajkumar, N., Krashenninnikov, D., Langosco, L., He, Z., Duan, Y., Carroll, M., Lin, M., Mayhew, A., Collins, K., Molamohammadi, M., Burden, J., Zhao, W., Rismani, S., Voudouris, K., Bhatt, U., ... Maharaj, T. (2023). Harms from Increasingly Agentic Algorithmic Systems. In *2023 ACM Conference on Fairness Accountability and Transparency* (pp. 651–666). ACM. <https://doi.org/10.1145/3593013.3594033>
- Córdova-Esparza, D.-M. (2025). AI-Powered Educational Agents: Opportunities, Innovations, and Ethical Challenges. *Information*, 16(6), 469. <https://doi.org/10.3390/info16060469>

- Parasuraman, R., & Manzey, D. H. (2010). Complacency and Bias in Human Use of Automation: An Attentional Integration. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 52(3), 381–410. <https://doi.org/10.1177/0018720810376055>
- Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J. (2024). A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6). <https://doi.org/10.1007/s11704-024-40231-1>
- Yan, L., Sha, L., Zhao, L., Li, Y., Martinez-Maldonado, R., Chen, G., Li, X., Jin, Y., & Gašević, D. (2023). Practical and ethical challenges of large language models in education: A systematic scoping review. *British Journal of Educational Technology*, 55(1), 90–112. <https://doi.org/10.1111/bjet.13370>