

ENVIRONMENTAL SCAN 02 / RESEARCHER-FACING BRIEF

AI Can Act Across the Research Workflow. What Should Researchers Still Own?

Research agents can now search, analyze, run code, operate tools, and draft outputs across multiple steps. That makes one question more important: **what decisions are being delegated, and what responsibility remains yours?**

Main takeaway: Use research agents as assistants to scholarly judgment - not as substitutes for deciding what counts as adequate evidence, a sound method, or a defensible claim.

What changed

Recent systems show that AI can link together tasks that researchers once handled separately: literature search, hypothesis generation, experimentation, analysis, and writing. Some experimental systems can even interact with laboratory hardware.

But capability is not the same as dependable performance. Current research benchmarks still report major weaknesses in autonomous literature discovery and end-to-end scientific rediscovery.

A simple way to think about it

You frame the problem Question, stakes, assumptions, acceptable evidence	AI may act Search, select, analyze, code, operate tools, draft	You retain accountability Verify evidence, approve methods, interpret results, own the claim
--	--	--

Six checkpoints before you delegate research work to an AI agent

Research stage	What an agent may do	Human checkpoint
Discover / Search	Search, rank, follow references, rewrite queries	Choose the source environment; test coverage; inspect provenance and omissions.
Frame	Generate questions, hypotheses, variables, or plans	Own the problem definition, assumptions, stakes, and acceptable evidence.
Evaluate	Score relevance, credibility, or risk	Confirm authority, methods, limitations, and disciplinary fit.
Analyze	Run code, statistics, models, or instruments	Approve methods and protocols; verify data handling; inspect unexpected outputs.
Synthesize	Combine studies or results	Check whether the synthesis overstates certainty, misses counterevidence, or hides disagreement.
Write / Publish	Draft claims, figures, citations, or disclosure	Own accuracy, interpretation, scholarly voice, authorship responsibility, and disclosure.

Ask this at every stage: What decision is the agent making here? How is AI influencing that decision? What responsibility remains mine?

Five questions that make agentic research more trustworthy

1**What source environment is the agent using?**

Is it searching the open web, a bounded scholarly corpus, subscribed content you can access, or something unclear?

2**Can I reconstruct what happened?**

Keep enough traceability to identify sources, prompts/instructions, tools, code, parameters, and important intermediate outputs.

3**Is the evidence set adequate?**

A real citation is not the same thing as a complete or balanced evidence set. Check what another search strategy reveals.

4**Did I verify methods and unexpected results?**

When an agent runs code, models, statistics, or instruments, inspect the method, data handling, failure modes, and surprising outputs.

5**Can I defend the final claim myself?**

You remain responsible for accuracy, interpretation, citation traceability, scholarly voice, and appropriate disclosure.

What the evidence does - and does not - show

The evidence base is **Developing**. Peer-reviewed demonstrations show that linked research automation is technically feasible in bounded settings. At the same time, recent benchmarks show that complex literature discovery and end-to-end scientific rediscovery remain difficult.

What this scan supports

Agentic research environments are emerging and can take meaningful actions across the research workflow.

What this scan does not support

That autonomous research is broadly reliable, that one agent can replace disciplinary research judgment, or that humans can safely surrender accountability for methods and claims.

Source and access transparency

This environmental scan uses peer-reviewed research, open research benchmarks, official scholarly standards, and technology announcements. These sources are not weighted equally. Vendor announcements are treated as capability signals, while peer-reviewed studies and benchmarks carry the reliability assessment.

The scan is not restricted to open access, but automated scanning has a visibility advantage toward public and openly inspectable sources. It therefore does not claim that open-access and paywalled literatures reach the same conclusions.

Selected references

Lu, C., et al. (2026). Towards end-to-end automation of AI research. Nature, 651, 914-919. DOI: 10.1038/s41586-026-10265-5.

Ghareeb, A., et al. (2026). A multi-agent system for automating scientific discovery. Nature, 655, 497-505. DOI: 10.1038/s41586-026-10652-y.

Xiong, L., et al. (2026). AutoResearchBench: Benchmarking AI agents on complex scientific literature discovery. arXiv. DOI: 10.48550/arXiv.2604.25256.

Xu, W., et al. (2026). ResearchClawBench: A benchmark for end-to-end autonomous scientific research. arXiv. DOI: 10.48550/arXiv.2606.07591.

International Committee of Medical Journal Editors. (2026). Use of Artificial Intelligence in Publishing.

Anthropic. (2026, August 27). Previewing the Model Hardware Standard.

FOR RESEARCHERS: AI can help you move faster. Your job is to keep the scholarly decisions visible - what evidence counts, what methods are acceptable, what conclusions are justified, and what you are prepared to defend.