

AI for Scientific Research

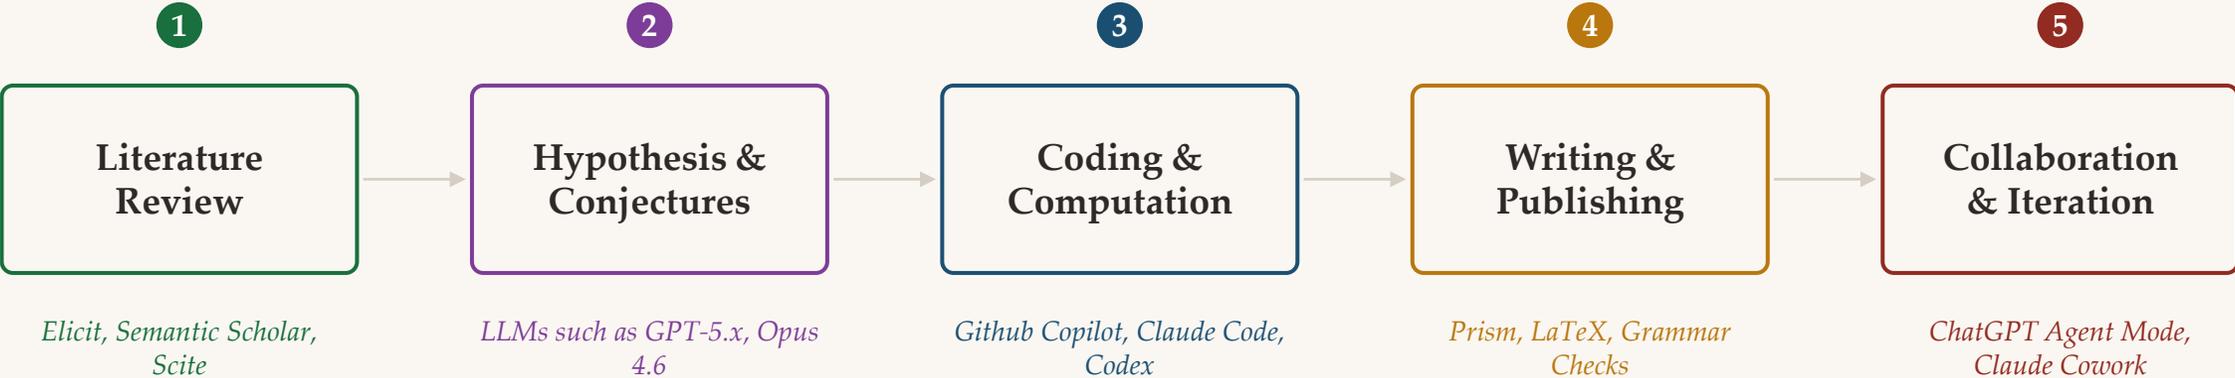
Examples and Limitations

Maria Crăciun & Rareș Răhăian
Tiberiu Popoviciu Institute of Numerical Analysis
Romanian Academy, Cluj-Napoca

11 March 2026



The Research Workflow & AI Tools



The AI Revolution in Science

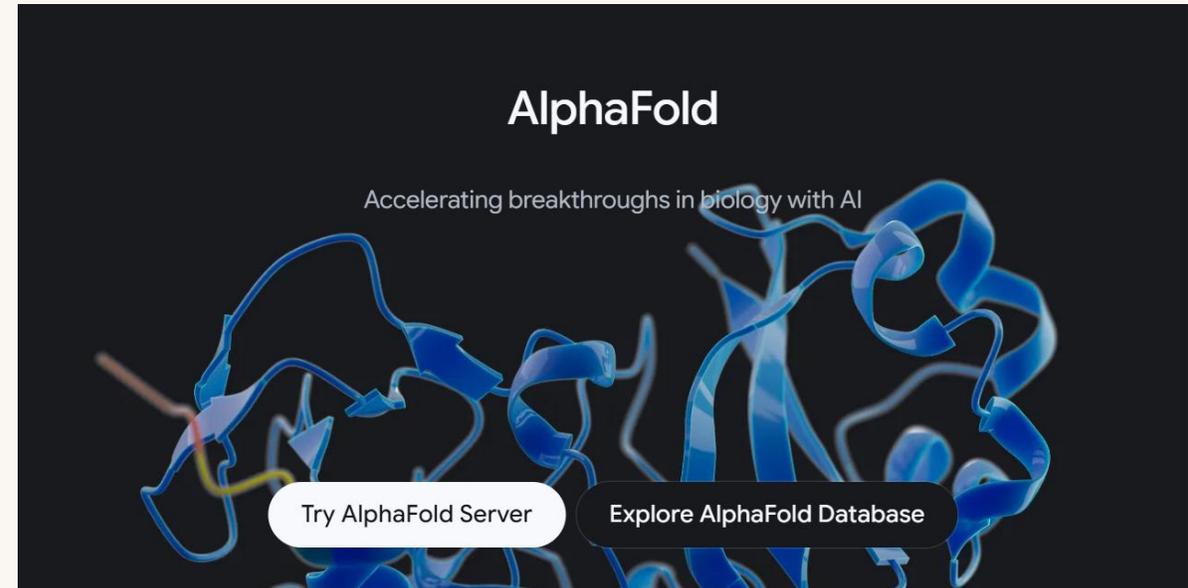
- AI is moving from a tool to a research partner

AlphaFold (2020) – the moment things changed

- Predicted 200M+ protein structures
- Solved a 50-year-old grand challenge in biology
- This was a purpose-built model for one specific problem

Now: general-purpose LLMs are entering the lab

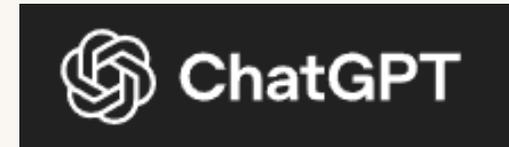
- Not built for one task – they can do many things at once
- Where does AI actually help? Where does it fail?



LLMs Enter the Lab

What are Large Language Models?

- Trained on huge amounts of text from the internet
- You talk to them in plain English — they respond in plain English
- The main ones: Claude (Anthropic), ChatGPT (OpenAI), Gemini (Google)
- *“We built Claude with Claude”* - Anthropic



What researchers use them for:

- Brainstorming ideas, debugging code, drafting papers
- Summarizing papers, explaining concepts, exploring data



Key limit: They remix existing knowledge — they do not discover new things

AI as a research assistant

Think of AI as a very fast, very well-read colleague

- It has read millions of papers but “understood” none of them
- Great at finding patterns, bad at knowing what they mean

Best used for:

- Speeding up tasks **you already know** how to do!
- Getting a first draft — **not a final answer**
- Exploring ideas before committing to a direction

“When solving a problem of interest, do not solve a more general problem as an intermediate step. Try to get the answer that you really need but not a more general one.” – Vladimir Vapnik

Not for: Replacing your judgment, verifying facts, or doing the thinking

AI accelerated science - examples

OpenAI — "Early experiments in accelerating scientific discovery"

- Scientists used GPT-5 for real research across many fields
- Good for exploring ideas, but not reliable enough to trust blindly

Donald Knuth and Claude Opus 4.6

- One of the greatest computer scientists, now testing LLMs
- Checking if AI can reason about math and algorithms

When top scientists take AI seriously, it is worth paying attention

I want to eventually tackle the full black hole problem with parameters M and J (the mass and spin of the black hole). But let's first set $M=0$ and tackle the flat-space case

+

🎤

↑

The Hallucination Problem

LLMs can fabricate references that look completely real

A real example:

- Prompt: "List 5 key papers on topological data analysis in neuroscience"
- AI returned 5 papers with plausible titles, valid journals, correct author formats
- But **3 of the 5 papers did not exist** — DOIs returned 404

How to protect yourself:

- Always verify citations via DOI, PubMed, or arXiv
- Prefer domain-specific tools (Semantic Scholar, Elicit, Consensus)
- Treat AI output as a first draft — always require human review

Takeaway: For literature search, use dedicated research tools — not general-purpose LLMs

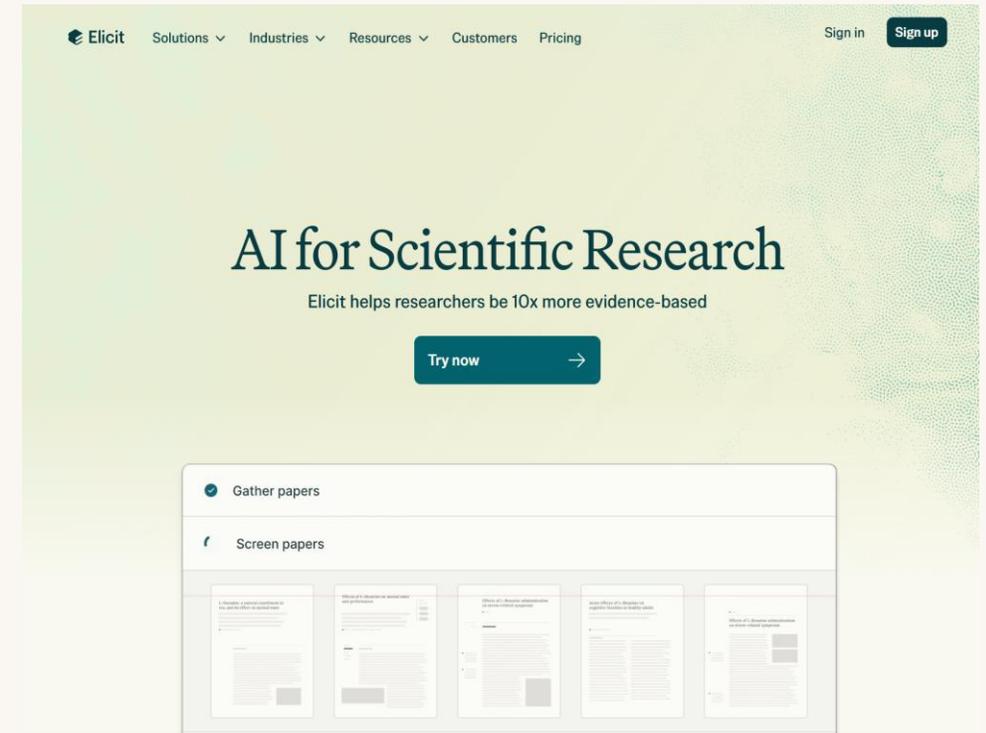
DOI NOT FOUND

10.1137/20M1378213

Literature Review & Discovery

Too many papers, too little time — AI can help you filter

- **Elicit** — ask a research question, get relevant papers, summaries
- **Semantic Scholar** — smart search, citation maps, summaries
- **Scite** — shows if a paper was supported or contradicted by others



Tip: Use these to discover papers, then read them yourself

Limit: They can miss very new or very niche papers

Writing Scientific Papers

Prism (OpenAI) — a LaTeX editor with AI built in

- Write, edit, and compile papers with AI help
- Good for getting past writer's block on boring sections

Any LLM can help with writing too

- Fix grammar, improve clarity, shorten long paragraphs
- Generate a rough abstract or related work section

The screenshot displays the Prism LaTeX editor interface. On the left, there is a sidebar with a 'New Project' button and a file list containing 'diagram.jpg' and 'main.tex'. The main area shows a LaTeX document with the following code:

```
1 \documentclass[11pt]{article}
2 \usepackage[margin=1in]{geometry}
3 \usepackage{amsmath}
4 \usepackage{graphicx}
5 \usepackage{tikz-cd}
6 \usepackage{multicol}
7
8 \setlength{\parindent}{0pt}
9 \setlength{\parskip}{1\baselineskip}
10
11 \begin{document}
12
13 \section*{What is Prism?}
14
15 \textbf{Prism} is an AI-powered \LaTeX{}
16 editor for writing scientific documents.
17 It supports real-time collaboration
18 with coauthors and includes OpenAI-powered
19 intelligence to help you draft and edit
20 text, reason through ideas, and handle
21 formatting.
22
23 \section*{Features}
24
25 \begin{multicols}{2}
26 Prism includes ChatGPT directly in the
27 editor and can access your project, so
28 you can ask it to do things like:
29
30 ``Add the equation for the Laplace
31 transform of  $\cos(at)$  to the
32 introduction.''
33
34 \[
35 \mathcal{L}\left\{t \cos(at)\right\}
36 = \frac{s^2 - a^2}{(s^2 + a^2)^2}
37 \]
```

On the right, a preview window shows the rendered LaTeX document. The title is 'What is Prism?'. The text describes Prism as an AI-powered LaTeX editor and lists features such as real-time collaboration and AI assistance. The preview includes a table with the following data:

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	16

The preview also shows a diagram of a vector space with points A, B, and C, and a table with the following data:

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	16

At the bottom of the preview window, there is a navigation bar with a 'Zoom to fit' button and a 'Tools' button.

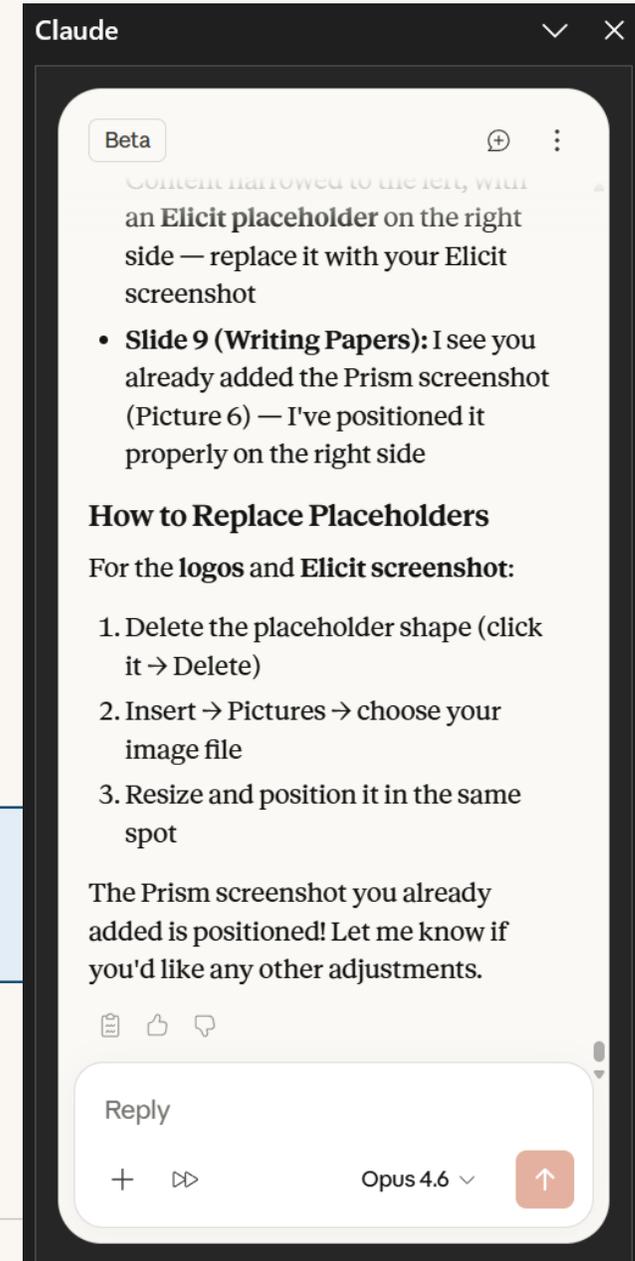
Limit: AI text often sounds generic; risk of all papers sounding the same

AI-enhanced Productivity Tools

AI is showing up inside the tools you already use

- **Claude in PowerPoint** — makes slides for you (like this one!)
- **Claude in Excel** — helps with formulas, data analysis, charts

The trend: AI is not a separate step — it is becoming part of the workflow.
This presentation was partly made with Claude in PowerPoint!



Coding & Computation

- **GitHub Copilot** — sits in your code editor and suggests code as you type
- **OpenAI Codex** — a coding agent: you describe a task, it writes and runs code
- **Claude Code** — works in the terminal, reads your whole project

Especially useful for:

- Writing boilerplate code, plotting graphs, data cleaning
- Explaining unfamiliar code, debugging error messages
- Everything can be used directly in your coding environment (e.g. VS Code)

Limit: AI code can look right but have subtle bugs — always test it

Agentic AI for Complex Tasks

What makes AI "agentic"?

- Instead of one question and one answer, it plans and takes steps
- It can browse the web, run code, search papers, write files

Tools you can try today:

- **ChatGPT Agent** — give it a task and it figures out the steps
- **Claude Cowork** — works on files and projects alongside you
- **Not yet very advanced!**

Limit: More autonomy means more chances to go wrong — check its work

Practical Matters: Cost & Access

Free versions exist — but they are limited

- ChatGPT Free, Claude Free, Gemini Free
- Good enough to explore, not enough for daily research use

Paid plans cost about \$20/month each

- ChatGPT Plus/Pro, Claude Pro, Google AI Pro
- Specialized tools (Elicit, Copilot) cost extra on top

 Free Try Claude \$0 Free for everyone Try Claude	 Pro For everyday productivity \$17 Per month with annual subscription discount (\$200 billed up front). \$20 if billed monthly. Try Claude	 Max Get the most out of Claude From \$100 Per person billed monthly Try Claude
---	---	---

Key Takeaways

AI is useful for research — but it is a tool, not an answer

Where it helps:

- Finding and summarizing literature
- Writing code, fixing bugs, making plots
- Drafting text, brainstorming, explaining things

Where it fails:

- Proving new things that nobody has proved before
- Being reliably correct — it makes confident mistakes

The most important skill: knowing when and how to use AI and when not to

Questions & Discussion

What has been your experience with AI in research?
