

DR

Kristyn

Sommer

● DEVELOPMENTAL SCIENTIST

FULL GUIDE EDITION · VERSION 1.0 · 6 AUGUST
2026

Understanding & using **AI** as a scientist.

A practical, plain-language guide to using LLMs and Generative AI for academic researchers. Using the skills you already have to help you work and teach with more freedom and less ambiguity.

About this guide.

If you feel like the conversation about AI is already happening at a level too technical for you, I understand. Why? Because a year ago, that was me too.

I walked back onto a university campus after two maternity leaves feeling completely lost. Even Copilot felt beyond my reach. I told myself I wasn't smart enough, wasn't technical enough, didn't have the hours to learn what sounded like a new field and a new language all at once.

Now, I probably don't need to tell **you** this because you have been living it alongside me, but twice in the last five years, our jobs as academics have been remade.

First in 2020, when the pandemic forced every academic to rebuild teaching and research from a kitchen table. Then in late 2022, when ChatGPT landed and a year later it felt like our students had quietly stolen the answer sheet to every essay and exam we'd spent a decade perfecting.

We are not 'behind' because we are not capable. We are 'behind' because the rules of our jobs keep changing and no one has taken the time to catch us up.

And that is why this guide exists. I caught myself up in a flurry of necessity and the kind of deep intense focus that only an autistic special interest can derive. So, with that said, this guide assumes some things about you as the reader.

It assumes you've been asked to redesign your teaching, research, supervision, and grants more often in the last five years than the previous twenty. It assumes you are tired and the prospect of doing it all again, for a tool you can't quite see the value of, feels punitive. It assumes the fear of getting it wrong is louder than the curiosity to try.

It also assumes you got into research not to become a puppet master of computers, but for the thinking and questioning. **The part only you can do.**

This guide will not give you all the answers, but it will translate the jargon, walk you through what these tools actually are and aren't, and offer small adjustments you can test in your own work.



 DR KRISTYN SOMMER · 2025

AI USE

Written with Claude's assistance: drafting, flagging inconsistencies, and generating the disclosure templates in Chapter 7. Every output reviewed, edited, and approved by the author. My voice, my standards, my responsibility. Full disclosure in Chapter 7.

© & DISCLAIMER

© 2026 Dr Kristyn Sommer. All rights reserved. Provided for educational and professional development purposes; reflects the author's personal practice as at August 2026. Not institutional policy or legal advice. Always check your institution's current AI policy before use.

What's in here.

FOUNDATIONS

Chapters 1–4

1	Why this guide	5
2	AI, untangled	11
3	Supervision, not automation	17
4	Your data, your rules	25

TRUST & ACCOUNTABILITY

Chapters 5–7

5	When it goes wrong	33
6	Checking the work	41
7	Your name on it	48

GET SET UP

Chapter 8

8	Setting up Claude	62
---	-------------------------	----

RESEARCH & INTEGRITY

Chapters 9–13

9	Privacy and consent	79
10	Hidden assumptions	85
11	Open science	90
12	Writing the paper	95
13	Copyright and risk	101

What's in here.

PEER SERVICE

Chapters 14–16

14	Reviewing papers	111
15	Editorial work	117
16	Grant review	123

TEACHING

Chapters 17–21

17	The teaching load	128
18	Better teaching	133
19	AI-safe assessment	137
20	AI in the curriculum	142
21	Judging students	147

SUPERVISION & DUTY

Chapters 22–24

22	Supervising PhDs	155
23	Ethics approval	161
24	The energy cost	167

REFERENCE

Chapter 25 & back matter

25	Prompting	180
·	Glossary	194
·	Afterword	196

The four-check **loop**.

Run every output through these four checks before you accept it. If you only do these, you are already ahead of most AI use in the wild. It is a loop because you repeat it for every new output, not once, every time. **Any "no" means stop.** Revise, re-prompt, or write it yourself, then re-enter at *Source*.

FIGURE 6.1

Source · Scope · Voice · Ownership

The loop, in order

- 1 SOURCE**
Can you verify every claim against a real source?
- 2 SCOPE**
Is it claiming more than the sources actually support?
- 3 VOICE**
Does this sound like something you would write?
- 4 OWNERSHIP**
Would you put your name on this exact text?

ANY "NO" → STOP Revise, re-prompt, or rewrite it yourself, then re-enter at *Source*. It is a loop: run all four on every output, every time.

The loop in **action**.

The same Source check, run on two literature-review outputs. Same prompt, same tool. The only difference is whether the loop was actually run.

✓ SOURCE CHECK, RUN

- Every citation searched by title and first author in Scholar
- Two DOIs failed to resolve, so both claims were pulled before the draft moved on
- One real paper said the opposite of the summary, so the sentence was rewritten
- Result: a synthesis you can defend line by line

✗ SOURCE CHECK, SKIPPED

- Citations read as plausible, so they were accepted on sight
- Two of nine references did not exist. The DOIs were invented
- A real paper was cited for a finding it never reported
- Result: a reviewer found it first, and the whole section lost trust

That's five pages of a hundred and ninety-six.

What you just read is the chapter that does the most work: the four-check loop, and the same literature review run with the check and without it.

The other twenty-four chapters follow the same shape. Every one opens with the question you actually arrived with. Every worked success is paired with a worked failure, because the failure is the same tool on the same task with the supervision slipped, and you need both to see the difference.

The full guide comes as five files

The guide

196 pages, 25 chapters. Foundations, trust and accountability, a complete setup walkthrough, research and integrity, peer service, teaching, supervision, and a prompting reference.

Start here

What to open first, and a one-hour route in if you don't have a week.

Working tools

The things meant to live next to your keyboard, printable on their own. The four-check loop, the data classification tree, the 11-point de-identification checklist, a blank session log, and the disclosure templates.

Starter files

The context files Chapter 8 teaches you to build, ready to fill in rather than retype, plus the same four already filled in so you can see what finished looks like.

A changelog

The guide keeps changing, and updates are free once you've bought it.

Get the guide

drkristynsommer.com.au/ai-for-academics

Instant download. Not useful? Email me within 5 days for a refund. PhD and grad students, email me and ask, and I'll send you a code for 80% off. No proof needed.

Dr Kristyn Sommer · developmental scientist, autism advocate, science communicator. Written with Claude's assistance and fully supervised, reviewed, and owned by the author. The guide practices what it teaches, full disclosure inside.