

AI tools cheat sheet

How to use AI in your research · COACT 2026 · Dr Luke Philbrooks · l.a.philbrooks@qmul.ac.uk

1 OF 2 · TOOLS

01 Chatbots

Claude

My main one: writing, analysis, long documents. Projects keep files and instructions together.

ChatGPT

All-rounder with voice, images and a large app and connector directory.

Gemini

Search grounding and Google Workspace. Gems are its saved instructions.

Google AI Studio

Very large documents and new models. Free tier trains on your data.

Microsoft Copilot

QMUL's official tool. Use it for anything sensitive.

02 Finding and reading papers

Elicit

Search and pull methods and sample sizes into a table. Check a few rows.

Consensus

"What does the evidence say about X?" Imports your Zotero library.

Perplexity

Quick searches and deep research. Open the sources: they don't always say what it claims.

Research Rabbit

Add one paper, get a map of related work.

Scite

Shows whether later papers support or contradict a paper.

03 Knowledge and notes

Gemini Notebook

WAS NOTEBOOKLM

Questions about papers and manuals you upload; points to where it found each answer. Only knows what you give it.

Zotero

Reference manager. [Zotero MCP](#) lets an AI search your library. It can write too: back up first.

Obsidian

Markdown notes. [mcp-obsidian](#) connects them to an AI.

04 Agents and coding

Claude Code · Codex · Gemini CLI

Agents that work inside a project folder: read files, run code, edit. Permissions are yours to set.

MCP **CONNECTORS**

Open standard for plugging AI into other software. Official ones first; check who made the rest.

Find skills: [skills.sh](#) · [anthropics/skills](#) · [K-Dense scientific skills](#) | **Awesome lists:** [MCP servers](#) · [main list](#)

Antigravity · Cursor · VS Code

Editors with agents built in. More setup, easier to miss edits.

Agent Skills **REUSABLE KNOW-HOW**

Corrected it twice? Make it a skill. Read any skill's code before installing.

05 Local models

Ollama

Command-line, one-step hosting of open models.

LM Studio

Point-and-click app for finding and running them.

Hugging Face

Where open models and datasets live.

Nothing leaves your machine, but a laptop only runs small, weaker models.

06 Landmark papers

Word2vec: embeddings

InstructGPT: training from human feedback

Attention Is All You Need

Direct Preference Optimisation

Chain-of-thought prompting

Found in the Middle: long context

Habits that pay off

CRCF: Context, Request, Constraints, Format

Interview me: ask it for its questions before it starts

Show, don't tell: two or three examples beat ten rules

Say what not to do: list the exclusions

Fresh chats: long chats degrade, so restart and hand over a note

Links, then open them: a real paper can still not say that

Check memory: read what it has saved about you and delete what's stale

Markdown beats PDF: convert before you feed it in

One task per chat: break big jobs into steps

Plan, check, do: make agents show a plan before they act

Least access: read-only and one folder, with backups

You're the boss: it's the intern, and you sign the claims

AI in your research: cheat sheet

COACT 2026 · Dr Luke Philbrooks · l.a.philbrooks@qmul.ac.uk

2 OF 2 · METHOD

Three questions, every time

1. **Could I explain and defend this without the AI?**
2. **Is this something my PhD is meant to teach me?** Then use AI to learn it while you do it.
3. **Would I notice if it were wrong?**

Offload the mechanics. Never the judgement.

Where each task sits

- **Automate:** formatting, references, boilerplate code, file conversion
- **Assist:** debugging, editing your draft, summarising papers you've read
- **Augment:** critiquing your design, tutoring in a new field
- **Own:** your question, design, interpretation and every claim you sign

A good prompt (CRCF)

- **Context:** who you are, what it needs to know
- **Request:** exactly what you want
- **Constraints:** sources, dates, length, what to avoid
- **Format:** table, headings, bullets, code only

Add: "Before you start, ask me any questions you need answered."

By research task

Task	Where AI helps	The trap	Tools
Reading	Finding and triaging papers; mapping a field; questioning papers you've read	Summaries replace reading; invented or mis-attributed citations	Elicit, Consensus, Perplexity, Research Rabbit
Crossing fields	A patient tutor; loading manuals; translating jargon	Fluent answers in a field you can't yet judge. Check with a person in it	Learning modes, Gemini Notebook
Designing	Critic: confounders, missing controls, pre-mortem	Letting it design; confidently wrong maths; not understanding your stats	Claude, ChatGPT, Gemini (reasoning mode)
Analysing	Writing and running code; agents with checkpoints	Silent data changes; debugging loops; results you can't explain	Claude Code, Codex, Antigravity
Writing	Drafts from well-structured notes; structure; "reviewer 2" feedback	Losing your voice; AI-written discussion; sending unreviewed output	Claude, skills, style guides

Prompt: teach me, don't tell me

I'm new to [topic]. Teach me [concept]. Don't give me the full answer. Ask me one question at a time, check my answer, and explain only when I'm stuck. Finish by asking me to explain it back.

Prompt: critic for your protocol

Here is my protocol. Don't rewrite it. List the three most serious threats to validity, most serious first. For each, say what a sceptical reviewer would ask. Be specific; skip generic advice.

Prompt: your discussion

My results: [...]. My interpretation: [...]. What are the strongest objections? Which alternative explanations have I not addressed? Don't rewrite my text.

Give your projects a memory

- **AGENTS.md / CLAUDE.md:** what it is, where the data is, conventions, what not to touch. Keep it short and current
- **Projects** (Claude, ChatGPT), **Gems** (Gemini)
- **Skills:** anything you've corrected twice
- **Connectors (MCP):** only what you need

Prompt: hand-off note

Write a hand-off note for a fresh chat: project goal in two sentences; what we did and which files changed; decisions and why; open problems; next steps. Under 300 words.

Data and privacy

- **Public papers, own drafts:** most tools, with care
- **Unpublished results, IP:** QMUL Copilot, or a paid tool with training switched off
- **Participant data:** approved systems only; check your ethics approval
- **Local models** keep data on your machine

Integrity

- Disclose AI use as your journal, funder and thesis regulations require
- You're responsible for all of it (UKRI puts this on the researcher)
- Keep a short log of how you used AI on each piece of work
- Never send anyone unreviewed AI output
- Agents: least access, backups and Git, approve deletions yourself

Read more

- Bastani et al. (2025) *PNAS*. [doi:10.1073/pnas.2422633122](https://doi.org/10.1073/pnas.2422633122)
- METR (2025). [arXiv:2507.09089](https://arxiv.org/abs/2507.09089)
- [QMUL Library guide](#) · [JRMO](#)