

# AI tools in primary care research

Workshop handout – key points and practical guidance

## Five principles for AI use

These principles should guide every decision you make about using AI in your work.

1. **You own the output.** You're accountable for accuracy and integrity. Would you put your name to this?
2. **Be transparent.** Disclose substantive AI use in research outputs.
3. **Protect data ruthlessly.** Match your tool to your data classification.
4. **Maintain academic integrity.** AI cannot be an author or make original contributions.
5. **Enable innovation responsibly.** Consider societal impact, equity, and sustainability.

## Critical restriction: personal data

**Currently, no AI processing of personal data is permitted with any commercial AI – including university-approved tools (ChatGPT Edu, Copilot, Gemini). This restriction applies until further guidance is issued.**

Personal data includes anything that could identify someone: names, emails, meeting attendees, student work, interview transcripts, or even descriptions like "the new research fellow in hypertension."

## Understanding what LLMs actually are

Traditional tools extend specific human capabilities in clear, bounded ways:

- A stethoscope extends your hearing
- A microscope extends your sight
- Statistical software extends calculation

**LLMs are something genuinely new.** They operate in the realm of language and meaning-making – territory that's always been uniquely human. They have **fluency without understanding** – they can discuss patient anxiety, draft care plans, and synthesise research, without ever experiencing worry, providing comfort, or being ill.

This makes them powerful *and* unreliable in ways no previous tool has been. Working with them requires a new mindset: not "use this tool" but "direct this conversation."

## Hallucinations: the confident fabricator

LLMs produce plausible, authoritative-sounding statements that are simply wrong. This is mathematically baked-in, not a bug that can be patched away. Newer "reasoning" models can actually hallucinate *more* because they talk themselves into longer stories.

### When hallucinations bite researchers:

- Niche or low-data topics (most of our work)
- Requests for exact facts, dates, or citations
- Long chats where early context drifts out of view

### Stay safe:

- Treat every "fact" as unverified until you check it
- Ask for brainstorms, structure, and rewrites – not data you can't verify
- Cross-check anything that matters

**The good news:** even with fabrications, LLMs are world-class at ideation, outlining, polishing prose, and teaching you the shape of a field. Verify the edges; exploit the speed.

## Prompting: simple and structured approaches

### For quick daily tasks: the three-part minimum

**[Role] → [Context] → [Task]**

**Good:** "Act as an experienced academic editor. I'm revising a paper about diabetes for publication in The BMJ. Please give feedback on and improve this paragraph's clarity and readability."

**Weak:** "Make this better."

### For reusable templates: the full framework

When you're creating prompts you'll use repeatedly – for plain language summaries, literature review assistance, or feedback on drafts – add more structure:

- **Role:** Who should respond (systematic reviewer, clinician, data analyst)
- **Context:** Your situation and background
- **Style:** Writing approach (academic, clinical, accessible)
- **Tone:** Attitude (professional, warm, technical)
- **Audience:** Who will read this (peers, patients, policymakers)
- **Response:** Specific format needed (table, report, bullet points, word limit)

## Why prompting works: attention and context

Inside LLMs are thousands of "attention heads" – think of them as mini-spotlights. Each one asks: "Given the word I'm looking at, which other words in the prompt are suddenly important?"

When you write "Act as a systematic reviewer..." the heads that learned statistics, PRISMA, or effect sizes light up. The ones trained on creative fiction stay dark. Better prompts mean brighter, better-chosen spotlights.

*It's like tuning a radio to the right frequency.*

### Context windows matter:

The context window is the LLM's working memory – how much of the conversation it can "see" at once.

- ChatGPT/Copilot: ~32,000–128,000 words
- Claude: ~150,000–200,000 words
- Gemini: up to 600,000–800,000 words

**Practical implication:** LLMs perform best in roughly the first 25–50% of their context limit. You can't dump every paper from a literature search into one chat. Break tasks into smaller chunks across multiple conversations, then combine the results.

## Key takeaways

**LLMs have fluency without understanding.** They're powerful collaborators in the space of ideas and language, but they lack the embodied experience that grounds human judgement.

**They're not databases.** Don't use them like Google. Use them for ideation, structure, drafting, and polishing – then verify anything factual yourself.

**Your prompts are the steering wheel.** How you cast your words decides how well the model serves your research. A simple Role → Context → Task structure beats 90% of vague requests.

**You remain accountable.** AI makes mistakes and can perpetuate biases. You're the expert, and you're responsible for everything you put your name to.

**Questions or feedback?** Contact the workshop team or visit the departmental AI guidance pages.