AI PROMPTS · PROMPT PACK

Research Summary and Literature Review Prompts

Read faster without reading worse — with verification built into every prompt, not bolted on at the end

FORMAT	LENGTH	VERSION
Prompt pack	12 pages	1.0

LAST REVIEWED
2026-08-01

Who this is for

You read research. Students, graduate students, faculty, and practitioners who want to stay evidence-based without pretending you have unlimited hours.

You've probably already noticed that AI tools will confidently tell you what a study found, and that some of those studies don't exist.

The one rule this whole pack is built on

> **NEVER LET A MODEL TELL YOU WHAT A STUDY FOUND. LET IT HELP YOU ORGANIZE WHAT YOU HAVE ALREADY READ.**

That sentence is the difference between a tool that makes you a faster researcher and a tool that ends your credibility.

Every prompt here is shaped by it. You'll notice most of them contain some version of *"work only from the text I gave you"* — that's not boilerplate, it's the load-bearing instruction.

What you'll walk away with

- **Eight prompts you can paste in with no editing** — they ask you for the source, then work only from it
- A verification protocol that takes about ninety seconds per citation
- Around twenty more prompts for synthesis, critique, and writing
- Three worked examples, including one where the AI got it wrong and how that was caught
- A pre-submission checklist

SCOPE

This is an educational resource about using a tool. It doesn't replace your own reading, your own judgment, or your institution's rules on research conduct, authorship, and AI use. Where this guide and your institution disagree, **your institution governs.** Check your programme's policy before you use any of this in submitted work.

Why this resource is more cautious than most

Large language models generate text that is *plausible*, and a citation is one of the most formulaic, predictable text patterns in existence. Author, year, title, journal, volume, pages, DOI. A model can produce a perfectly shaped citation for a paper that was never written, complete with real author names attached to a title they never published and a DOI that resolves to nothing — or worse, to something else entirely.

This is well documented, reproducible, and has already produced professional consequences for people who submitted work containing invented references.

It also isn't a bug that's about to be fixed. It's a consequence of how these systems work: they are extremely good at the *shape* of a citation and have no independent way to check whether a particular one exists.

Two things follow.

First, never ask a model to recall research. Not "what does the research say about X." Not "find me five studies on Y." Not "what did Smith 2019 find." Every one of those invites fabrication, and the output will look exactly as confident whether it's real or invented.

Second, you can safely ask it to work on text you supply. Paste an abstract, a paragraph, your own notes — and ask it to restructure, compare, question, or clarify *that text*. This is where the genuine time savings are, and the fabrication risk drops close to zero because there's nothing to invent.

Everything in this pack lives on the second side of that line.

What about tools that search and cite?

Some tools retrieve real documents and cite what they retrieved. These are meaningfully better and worth using.

They are not a reason to skip verification. A retrieval tool can still summarize a real paper incorrectly, attribute a finding to the wrong source among the ones it found, or surface a retracted or predatory-journal paper without flagging it. **Retrieval changes the failure mode from "invented" to "misattributed." Both end the same way if you cite it.**

Verify anything you will put your name to. It takes ninety seconds.

The ninety-second verification protocol

Do this for **every** citation before it enters anything you submit, publish, or teach — regardless of where it came from.

1. **Find it yourself.** Search the title in your library database, Google Scholar, or PubMed. Not the AI's link — your own search.
2. **Confirm the pieces match.** Authors, year, journal, title. A single wrong author or a year that's off by one is a signal the whole thing was assembled rather than retrieved.
3. **Open it.** Abstract at minimum. Full text if you're relying on a specific finding.
4. **Check it says what you think it says.** This is the step people skip. A real paper cited for a claim it doesn't make is a worse error than a fake one, because nobody catches it.
5. **Check it hasn't been retracted.** For anything load-bearing.

If a citation fails at any step, **do not repair it and keep going.** Delete the claim, or go find a real source that supports it. Patching a broken citation is how a fabrication ends up in print with a real-looking reference attached.

> THE TELL
>
> a DOI that doesn't resolve, or a paper that only appears in the AI's answer and nowhere else on the internet, is almost always invented. Real papers leave traces — citations, repository copies, author pages, preprint servers.

Sorting your sources

A quick tally before you start writing. If most of your citations sit in the bottom row, you don't have a literature review yet.

Status	What it means	Safe to cite?
Read	You have read it, at least the abstract and results	Yes
Retrieved	You have the actual paper open, not yet read properly	Not yet
Referenced	You've seen it cited by someone else	No — get the original
Suggested	An AI mentioned it	**No.** Verify first, then it becomes Retrieved

Part 1 · Paste these in as they are

No editing, no brackets. Each one asks you for the material, then works only from what you give it.

1. Is this worth reading properly?

PROMPT

```
I'm going to paste an abstract. Help me decide whether to read the full paper.

Work ONLY from the text I give you. Do not add anything you know about this
study, these authors, or this field. If something isn't in the abstract, say
"not stated" — don't fill the gap.

Tell me: the question it asks, what they actually did, what they found, and —
most importantly — what it does NOT establish.

Then tell me the single biggest reason this might not transfer to a different
population or setting.

Ask me to paste the abstract.
```

2. Explain this paper to me like I'm new to the method

PROMPT

```
I'll paste a section of a research paper that I'm struggling with.

Work only from the text I provide. Don't supplement it with outside knowledge,
and if the passage is unclear or assumes something it doesn't state, tell me
that rather than smoothing it over — the confusion may be the paper's, not mine.

Explain it in plain language, then tell me which technical terms I'd need to
understand properly to evaluate whether their approach was appropriate.

Ask me to paste the section.
```

3. What are the limitations they're not emphasising?

PROMPT

```
I'll paste a paper's methods and results.

Work only from what I give you. Based solely on that text, tell me what this
design cannot show — not flaws in their execution, but the boundaries of what
this kind of study can establish at all.

Cover: whether the design supports causal claims, who the sample does and
doesn't represent, what the outcome measure actually captures versus what it's
named, and anything measured but not reported.
```

```
If I ask you to speculate beyond the text, refuse and tell me what I'd need to
look up instead.

Ask me to paste the methods and results.
```

4. Organise my notes by what they disagree about

The most useful synthesis prompt in the pack.

PROMPT

```
I'll paste my own notes on several studies I have read.

Organise them by what they DISAGREE about, not by topic. For each disagreement,
tell me whether the difference is methodological, definitional, or a genuine
empirical conflict — and say which, explicitly.

Then tell me what question nobody in this set seems to have asked.

Work only from my notes. Do not add studies, findings, authors, or citations I
did not provide. If my notes are too thin to support a comparison, say so
instead of inventing the connective tissue.

Ask me to paste my notes.
```

5. Check my summary against the source

Use this before you cite anything.

PROMPT

```
I'm going to give you two things: a passage from a paper, and my own summary of
it.

Compare them. Tell me: anything in my summary that the passage doesn't support;
anywhere I've strengthened a hedge ("associated with" becoming "causes,"
"may" becoming "does"); anything important I left out; and any place I've
generalised beyond the population studied.

Quote the exact words from each where they diverge.

Don't rewrite my summary. I want to see the gaps myself.

Ask me for the passage first, then my summary.
```

6. Make this understandable to a non-research audience

PROMPT

```
I'll give you a research finding in my own words, along with what the study
actually did.

Rewrite it for a general audience at about a newspaper reading level. Keep every
qualification that matters. Where the original says "associated with," do not
```

```
drift into "causes." Where it studied one population, say so.

Then show me the version I'd be tempted to write instead — punchier and slightly
wrong — and point at exactly which word does the damage.

Ask me for the finding and the study details.
```

7. Stress-test my argument

PROMPT

```
I'll paste an argument I'm making in a paper or a talk.

Give me the three strongest objections a careful reader in my field would raise.
For each: is it fatal, serious, or answerable?

Then tell me what kind of evidence would close the biggest gap. Do NOT tell me
whether that evidence exists or name any studies — I'll go and look. If you name
a study I'll have to spend twenty minutes discovering it isn't real.

Ask me to paste the argument.
```

8. Audit my reference list

Run this before you submit anything.

PROMPT

```
I'll paste my reference list.

For each entry, tell me which parts you can and cannot assess from the text
alone — formatting consistency, internal contradictions (a journal that doesn't
match the field, a volume number inconsistent with the year, a DOI prefix that
doesn't match the publisher), and anything structurally odd.

Be explicit: you cannot confirm any of these exist. Flag the ones whose details
look internally inconsistent so I know where to check first.

Do not "correct" any citation. If something looks wrong, say so and let me
verify it against the actual source.

Ask me to paste the list.
```

That last instruction matters more than it looks. A model asked to fix a citation will confidently produce a *tidier* version of a reference that may still be fabricated — and now it looks more credible.

Part 2 · Prompts with more control

Reading and appraisal

PROMPT

```
Here is the methods section of a study on [TOPIC]: [PASTE].

Working only from this text, tell me: the design, the unit of analysis, what was
randomised or controlled if anything, the sample and its likely limits, and the
primary outcome measure.

Then tell me which of my questions this design can and cannot answer. My
question is: [YOUR QUESTION].
```

PROMPT

```
I'm comparing two papers. Paper A: [PASTE]. Paper B: [PASTE].

Working only from these texts, build me a comparison table: question, design,
sample, measures, main finding, and stated limitations.

Then tell me whether their findings actually conflict, or whether they're
answering different questions in a way that only looks like conflict.
```

Synthesis

PROMPT

```
Here are my notes on [NUMBER] studies about [TOPIC]: [PASTE].

Group them by the mechanism or explanation each one implies, not by outcome.
Some papers with similar findings may imply incompatible mechanisms — that's
what I want to see.

Work only from my notes. Flag anywhere my notes are too thin to place a study
confidently.
```

PROMPT

```
Here's my draft literature review section: [PASTE].

Tell me where I've written "the literature shows" or similar when I'm actually
describing one or two studies. Quote each instance.

Then tell me where I've presented a contested area as settled.
```

Writing

PROMPT

```
Here's a paragraph from my draft: [PASTE].

Check every claim against the citation attached to it. For each: does the cited
source, as I've described it, actually support the sentence? Flag any place
where the citation is doing more work than it can bear.

I'll verify the sources themselves — you're checking whether my sentence matches
my own description of them.
```

PROMPT

```
I need to describe this limitation honestly without undermining my whole paper:
[LIMITATION].

Give me three ways to state it — minimal, moderate, and fully transparent. For
each, tell me what a sceptical reviewer would say.

Don't help me hide it. Help me state it well.
```

Teaching with research

PROMPT

```
I want to teach [WHO] how to read a paper critically, using this one: [PASTE
ABSTRACT AND METHODS].

Give me five questions students should ask of this specific paper, ordered from
most to least obvious. For each, tell me what a good answer looks like and what
a superficial one looks like.

Work only from the text I gave you.
```

Three worked examples

> **EXAMPLE**
>
> *These examples are fictional, including the study details in example 3. They illustrate a process, not real findings.*

Example 1 — Triage that saved an afternoon

The situation. Amara, a second-year graduate student, has 40 papers from a database search and time to properly read maybe eight.

What she did. Pasted each abstract into prompt 1, one at a time. Roughly ninety seconds each.

What it gave her. The "what it does NOT establish" line did most of the work. Eleven papers turned out to be cross-sectional where she needed longitudinal evidence — visible from the abstract, but not something she'd have registered while skimming forty of them at 11pm.

What she still had to do. Read the eight. The AI didn't tell her what any paper found; it told her what each one was *shaped like*, which is exactly the right division of labour.

Example 2 — The synthesis she couldn't see

The situation. Ben has read fifteen papers for a review and can feel a pattern he can't articulate.

What he did. Pasted his own notes — his words, not abstracts — into prompt 4.

What it gave him. Organised by disagreement, the set split cleanly: studies using self-report measures found one thing, studies using behavioural measures found close to the opposite. He'd read every one of those papers and hadn't seen it, because he'd been grouping them by topic.

The "what has nobody asked" answer was less useful — fairly generic. He kept the first part.

Why it worked. The model had nothing to invent. Every study in the answer was one Ben had read and summarised himself. The only new thing was the arrangement, which is the one thing it's genuinely good at.

Example 3 — The one that was wrong

The situation. Priya asked a general AI assistant a question she now describes as "the obvious thing you're not supposed to do": *what does the research say about X?*

What came back. Six citations. Confident, well-formatted, plausible-looking. Real-sounding journals. DOIs.

What verification found. Following the ninety-second protocol:

- Two were entirely fabricated. The titles appeared nowhere except in that answer.
- One had real authors and a real journal — but the title belonged to a different paper by different people.
- One was real and correctly cited, but **said close to the opposite** of what the summary claimed.

- Two were real and correctly described.

The one that matters. Not the fabrications — those were caught in the first search, in about a minute. It's the fourth one. A real paper, real authors, real DOI, correctly formatted, cited for a finding it doesn't support. That one would have survived any check short of actually reading it.

Which is the whole argument for step 4 of the protocol. Confirming a paper exists is not the same as confirming it says what you were told.

Priya now uses prompt 1 with abstracts she has retrieved herself, and never asks what the research says.

Source tracking sheet

One row per source. Nothing gets cited until the last three columns are ticked.

#	Author / year	How I found it	Retrieved	Read	Says what I think
1			☐	☐	☐
2			☐	☐	☐
3			☐	☐	☐
4			☐	☐	☐
5			☐	☐	☐

"How I found it" matters. If the answer is "an AI suggested it," that source needs the full protocol before it goes anywhere near your draft.

SOURCES I CHECKED AND DROPPED — AND WHY

Keep this. It stops you re-finding the same dead end in three weeks, and it's evidence of a careful process if anyone ever asks.

Common mistakes

Mistake	What happens	Do instead
"What does the research say about X?"	Fabricated citations, delivered confidently	Retrieve papers yourself, then paste them in
"Find me five studies on Y"	Same, with better formatting	Search a real database
Asking it to summarise a paper by name	It summarises what a paper with that title *would probably* say	Paste the actual text
Trusting a retrieval tool without checking	Misattribution instead of invention — same outcome	Verify anyway. Ninety seconds.
Asking it to fix a citation	A tidier version of a possibly fake reference	Verify against the source, never patch
Checking existence but not content	A real paper cited for a claim it doesn't make	Step 4. Always step 4.

Pasting a whole paywalled PDF	A licence problem, and possibly an institutional one	Abstracts, your own notes, or check what your access terms allow
Letting it write the lit review	Prose that reads well and misrepresents the field	Let it organise *your* notes; you write the prose
Accepting "studies show" in its output	An unsourced claim in your voice	Ask which study, then verify — or cut it
Using it on work you'll submit without checking policy	An academic integrity problem	Read your programme's policy first

Before you submit — checklist

- [ ] Every citation was found by me, in a real database, independently of the AI's answer
- [ ] Authors, year, title, and journal all match the actual paper
- [ ] I have opened every source I cite — abstract minimum, full text where a specific finding matters
- [ ] Each source actually supports the specific sentence it's attached to
- [ ] Nothing load-bearing has been retracted
- [ ] No citation was "corrected" or reformatted by an AI
- [ ] Every hedge in the original survives in my version — no "associated with" became "causes"
- [ ] Where I've said "the literature shows," I have more than two studies
- [ ] Contested areas are described as contested
- [ ] The prose is mine
- [ ] I've read my institution's policy on AI use in submitted work, and this complies with it
- [ ] I could explain and defend every source in this list from memory

That last one is the honest test. If you couldn't discuss a source in a viva or a seminar, you don't know it well enough to cite it.

Your next step

Take a paper you've already read and think you understand. Write a two-sentence summary from memory. Then use **prompt 5** to compare your summary against the actual passage.

Most people find at least one place where they'd firmed up a hedge. That's the error this pack exists to prevent, and it's worth meeting it on a paper where the stakes are zero.

Sources and further reading

This pack asserts **no empirical claim and cites no study.** In a resource about citation integrity, that's the only defensible position — a fabricated or unverified reference here would discredit the entire thing.

Two factual statements are made, both widely documented and both stated without numbers:

- Language models can generate citations for papers that do not exist, including plausible authors, titles, journals, and DOIs.
- Retrieval-based tools reduce fabrication but introduce misattribution, so verification is still required.

For further reading, go to primary sources rather than any summary:

- Your own institution's or programme's current policy on AI use in research and submitted work. This is the one that actually governs you, and it's being revised almost everywhere.
- Your professional association's current published guidance. For psychology in the US, the American Psychological Association publishes on research and publication conduct.
- The Committee on Publication Ethics (COPE) publishes guidance on publication integrity for authors and editors.
- Retraction Watch maintains a public database of retracted papers.

> **RESEARCH VERIFICATION REQUIRED**
>
> The four resources above are named as starting points because they publish relevant material, **not** as citations — no claim here depends on them, and none was retrieved during drafting. Before publishing, confirm current URLs and add them with access dates.

Erica K. King · Psychology · AI · Build Something You Can Use

Research Summary and Literature Review Prompts · Version 1.0 · Created 2026-08-01 · Last reviewed 2026-08-01

Erica K. King

Psychology · AI · Build Something You Can Use

Research Summary and Literature Review Prompts · Version 1.0 · Created 2026-08-01 · Last reviewed 2026-08-01

This is a free educational resource. You're welcome to use it in your own work, teaching, and practice, and to share it with attribution. It isn't professional, clinical, legal, or financial advice.

> **SCOPE**
>
> This is an educational resource. It doesn't replace your own reading, your own judgment, or your institution's rules on research conduct, authorship, and the use of AI tools. Where this guide and your institution disagree, your institution governs.