

Gate 4 Reference Card

Literature Review + Reference List

AI Writing Coach | Dr. Alianna J. Maren | Themesis, Inc. | themesis.com

What to Bring to Your Gate 4 Claude Session

- Your updated report document** — with all Gate 3 revisions incorporated
- Your current Literature Review draft** — at any stage — list of sources, outline, partial draft, or complete draft
- Your reference list** — however complete or incomplete
- This Reference Card** — open and in front of you

Why the Literature Review Matters

Before anyone invests in the hard work of reading your Data, Methods, and Results — they need to know: does this person belong in this conversation?

Your Literature Review is your credential flash. Think of a detective entering a room — the first thing they do is show their badge. Without it, they're a stranger claiming to belong. Your Literature Review tells the reader: I know this field, I know who did what, I know where the gaps are, and I know exactly where my work fits.

A sparse Literature Review, or one built on lightweight sources, signals the opposite — that you're a tourist in the field, not a resident.

The Three Red Flags of a Bad Literature Review

- **Too sparse** — not enough sources. You haven't done the work of finding what's already been said about your problem.
- **Too lightweight** — Medium articles, blogposts, Wikipedia. Fine for orientation, but not the load-bearing structures of your field. You need the canonical works.
- **Disconnected from the field's core** — missing the papers that every serious person in your community has read.

What a Literature Review Actually Is

Not a list. An argument. Its job is to build a case — in narrative form — that leads inevitably to your own contribution.

The architecture underneath every great Literature Review:

- **Problem** — stated broadly: here is the territory
- **Prior Work** — acknowledged generously: here are the people who mapped it
- **Limitation** — named precisely: here is where prior work falls short
- **Chain of Attempts** — narrative momentum: here is how the field tried to close the gap
- **Remaining Gap** — measured, not vague: here is exactly what is still missing
- **Your Solution** — earned, not just announced: and here is where my work lives

Three Key Factors in Reference Selection

- **Community Respect** — find the work with the most citations in your field. High citation count is the community's vote for importance.
- **Publication Source** — weakest to strongest: blogpost → institutional tech report → arXiv preprint → conference paper → journal article.
- **Recency** — this year = up to date; within two years = covering recent impacts; older work must be highly cited OR essential to your thinking. Always look for a more recent publication by the same author.

Read at minimum: Title & Authors → Abstract → Introduction → References → Conclusions.
Do not cite work you haven't read.

Claude as Literature Scout — and Why You Still Have to Read

Before Claude and other capable AI tools, researchers found literature three ways: keyword search into Google or Google Scholar; directed search starting from a few known landmark papers and following the authors' recent work forward; and lateral search through open-access repositories like arXiv for adjacent work. All three are still valid — and Claude can accelerate all three dramatically.

What changes:

Claude can surface the most-cited, most relevant works in a field you're just entering — if you give it the right dialogue. Don't just ask 'find me papers on X.' Tell Claude your argument. Tell it what point you need supported and why. Tell it what your audience cares about. That dialogue produces the right papers — often the best papers — faster than any keyword search.

What doesn't change:

You have to read them.

The Interviewer Scenario

Imagine sitting across a desk from an interviewer. Or having lunch with a prospective employer. They've read your report. They're asking you about the papers you cited — what Hofmann (1999) actually argued, where pLSI falls short, why Blei et al. chose the Dirichlet prior.

There is no quick AI check available. You either know it or you don't.

The student who used Claude to find great papers but never read them will not survive that conversation. The student who read every paper they cited will speak with the confidence of someone who actually belongs in the field. That confidence cannot be faked.

Read at minimum: Abstract → Introduction → References → Conclusions. Do not cite work you haven't read. Claude can find the sources. Only you can own them.

The Loop-Connection Question — Before You Close Gate 4

Before ending your Claude session, ask yourself:

Does my Literature Review lead inevitably to my problem statement?

If someone read only my Literature Review — would they arrive at exactly the gap my work fills?

- **If yes** — your argument is continuous. Gate 4 is complete.
- **If no** — either the problem statement needs adjustment, or the literature selection needs adjustment. Work with Claude to identify which.

Canvas Confirmation Statement

Gate 4 complete — Literature Review reviewed with AI Portfolio Coach. Source clusters identified. Gap articulated and positioned. Literature Review narrative argument drafted or revised. Conceptual map figure generated. Reference list checked for completeness and consistent formatting. Session logged in Appendix B.