AI Researchers
Embedding Responsible AI in Research

AI READINESS TOOLKIT

UC1-03 Use case 1: Literature review and synthesis

AI literature reviews: source enrichment and validation

Finding, validating and incorporating sources an author has not yet discovered

AT A GLANCE

A practical guide to finding, validating and adding sources that an author has not yet discovered. It shows how AI tools can match a draft paper against the published literature, suggest relevant work and help the author check that each new source is real, relevant and correctly cited. Use it when preparing or revising a literature review. Every source and interpretation is checked against the original.

TYPE
Case study

FORMAT
PDF

CONTRIBUTING PARTNER
UPF Barcelona School of Management

WHO IT IS FOR
Researchers at any career stage preparing or revising a literature review

FRAMEWORK BARRIERS
Trust in AI, Know-how, This task

HUMAN-ONLY BOUNDARY

Every source and interpretation is checked against the originals. The final analysis stays with the researcher.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author or authors only and do not necessarily reflect those of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the entity providing the grant can be held responsible for them.

Introduction: Using AI to Find What Your Draft Is Missing

Most literature reviews are built from the sources a researcher already knows: the papers read during the project, the authors cited in the field's key articles, and the results of a few database searches run months earlier. By the time a draft is written, the field may have moved on, and relevant work from neighbouring disciplines, other regions or other languages may never have appeared in those searches. Furthermore, peer reviewers may have opinions, requirements and requests of literature that should be included in our paper.

AI tools can help close that gap. They can read a draft, identify which claims need stronger support, turn those claims into searches, and surface candidate sources the author has not yet discovered. Used well, this makes a literature review more complete, more current and more balanced.

Used carelessly, it does the opposite. General AI assistants can produce references that look real but do not exist, attach real papers to claims they do not make, or find only the sources that confirm what the author already believes. A citation added without reading the source is a liability, not an improvement.

The guiding principle of this guide is:

> AI finds candidate sources → The researcher verifies each one against the original → The researcher decides whether and how it changes their argument. The Researcher always makes the decisions. The tool is there for assistance, for streamlining the word, and to provide efficiency, not to substitute human criteria.

New sources should be able to change the draft. If every source the AI finds confirms what the paper already says, the search has probably been too narrow. The aim is not to decorate existing claims with more citations, but to test those claims against a wider body of evidence.

This guide has two parts. Section 2 explains the main types of AI tools and what each brings to source discovery. Sections 3 to 7 set out a five-stage workflow (map, search, discover, validate, incorporate) with prompt templates for each stage. A companion template, Include, Adapt or Discard, provides a log for recording and justifying decisions about each source and is meant to be a guide for literature review.

The Tool Landscape: What Each Type of Tool Brings

No single tool does everything. Generalist assistants are good at reading and reasoning about your draft; specialised research tools are better at searching the scholarly record. The strongest workflow combines them. Features and plans change often: the descriptions below reflect the tools as of September 2026 and should be checked against each provider's current documentation and your institution's list of approved tools. This guide provides tools as examples and suggestions, but is not intend to endorse any.

GENERALIST AI ASSISTANTS (E.G., CLAUDE, CHATGPT, GEMINI, COPILOT)

This guide uses Claude (Anthropic) as its example, but the same techniques work with ChatGPT, Gemini, Microsoft Copilot or any comparable assistant your institution approves. What matters is not the brand, but which features are switched on.

> Without web search turned on, a generalist assistant works only from its training data and the material you provide. It is excellent at analysing your draft, but any references it produces from memory must be treated as unverified leads. Some will be accurate; some will be partly or entirely invented.
>
> With web search turned on, the assistant retrieves live pages and cites them, so you can click through to check each source. In Claude, this is enabled from the tools menu in the chat window, for example. Web searchers can

> be useful, but not all web searches will turn up sufficiently acceptable academic sources.
>
> With a deep research mode (called Research in Claude, deep research in ChatGPT and Gemini), the assistant runs a series of searches that build on each other and returns a longer report with citations. This is useful for mapping a topic, but it takes longer and is available mainly on paid plans.
>
> With projects or file uploads, you can keep your draft, notes and reference list in one place so the assistant can compare them across several conversations. Some scholarly tools (for example, scite) can also be connected directly to Claude or ChatGPT.

What generalists bring: close reading of your draft, identifying claims that need support, designing search strategies, comparing a claim with the text of a source you supply, and helping integrate sources into your argument.

Main risk: confident but unverified references, especially when search is off.

ANSWER ENGINES (E.G., PERPLEXITY)

A tool like Perplexity builds every answer from a live web search and shows numbered citations next to each statement. Its academic focus setting prioritises journals, preprint servers and scholarly databases; a deeper multi-step research mode handles more complex questions; and Spaces keep the context of a project across sessions.

What it brings: speed, recent material, and access to gray literature such as policy reports, working papers and institutional documents that academic databases may not index.

Main risk: sources of uneven quality appear side by side, and the summary may state a finding more strongly than the source does. Always open the cited source.

SCHOLARLY AI SEARCH ENGINES (E.G., CONSENSUS, ELICIT)

These tools search only academic literature, drawing mainly on large open databases such as Semantic Scholar and OpenAlex (Consensus reports more than 200 million papers). Every answer is tied to specific papers.

> Consensus is designed for research questions. For yes/no questions (“Does X increase Y?”), its Consensus Meter shows how the most relevant studies divide between yes, possibly and no. Filters allow you to narrow results by study design, journal quality, field and year, and a deep search mode produces a longer report across more papers.
>
> Elicit finds papers through semantic search and presents them in a table, with columns you define (for example, sample, method, main finding). This makes it easier to compare studies and supports systematic-review-style screening and data extraction.

What they bring: a quick view of the empirical evidence on a question, including studies that disagree.

Main risks: they work best for empirical, quantitative questions and less well for law, humanities or theory; summaries are often based on abstracts; and they do not cover subscription databases such as Scopus or Web of Science.

CITATION INTELLIGENCE (E.G., SCITE)

scite shows how later papers have cited a given paper, sorting each citation as supporting, contrasting or mentioning, and showing the sentence where the citation appears. It also flags retractions and editorial notices and can check a whole reference list.

What it brings: validation. It helps answer whether a source still stands, and whether later work has challenged it.

Main risk: the classification is automated; a high number of "supporting" citations is not the same as strong evidence.

CITATION MAPPING (E.G., RESEARCHRABBIT, LITMAPS, CONNECTED PAPERS)

These tools start from one or more seed papers you already trust and show the papers around them: the works they cite (backward), the works that cite them (forward) and papers with similar content. Results appear as networks or timelines.

What they bring: discovery of the "neighbourhood" of your key sources, including recent papers that build on them.

Main risk: they reproduce existing citation clusters, and coverage of books, non-English publications and regional journals is uneven. They show connections, not quality.

HOW THE TOOLS FIT TOGETHER

TASK	BEST SUITED	WHY
Analyse the draft and identify claims needing support	Generalist assistant	Reads and reasons about long texts you provide
Turn claims into search queries	Generalist assistant	Generates synonyms, Boolean strings and counter-questions
Get a quick view of the evidence on an empirical question	Consensus, Elicit	Search only scholarly papers; show agreement and disagreement
Find recent work and gray literature	Perplexity; generalist with web search	Live web search with citations
Expand from papers you trust	ResearchRabbit, Litmaps, Connected Papers	Backward, forward and similarity links
Check whether a source still stands	scite; publisher and retraction databases	Supporting/contrasting citations and retraction notices
Confirm that a source exists and details are correct	Library databases, Google Scholar, DOI lookup, reference manager	Authoritative bibliographic records
Integrate verified sources into the text	Generalist assistant	Revises structure and synthesis within your constraints

> Before uploading an unpublished draft to any tool, check that the tool is approved by your institution for confidential or unpublished work, and review its data and training settings.

Stage 1 — Map Your Draft

Start with your own text, not with a search. Before looking for new sources, identify exactly which claims need support and where the current reference list is thin. A generalist assistant is well suited to this stage because it works only with the material you provide.

Prompt 1 — Extract the claims that need sources Purpose: Build an inventory of claims and their current support.

Prompt template: Act as a critical academic reviewer. Below is a section of my draft paper on [RESEARCH TOPIC]: [INSERT DRAFT SECTION WITH IN-TEXT CITATIONS] List every substantive claim in a table with these columns: claim (quoted verbatim); type of claim (empirical, theoretical, definitional, contextual); source currently cited, if any; and the kind of evidence that would be needed to support it adequately. Flag claims that are uncited, supported by only one source, or stated more strongly than the cited source is likely to allow. Do not suggest specific references. Describe the type of evidence needed, not particular authors or titles.

Researcher check: Do the flagged claims match your own sense of where the argument is weakest? Add any claims the AI has missed.

Prompt 2 — Identify blind spots in the reference list Purpose: Detect narrowness in time, region, method or perspective.

Prompt template: Below are my draft section and its reference list: [INSERT DRAFT SECTION] [INSERT REFERENCE LIST] Analyse the reference list and identify possible blind spots, such as: heavy reliance on a few authors, journals or research groups; sources that may be outdated for a fast-moving topic; a single method, discipline, region or language dominating the evidence; no sources that challenge or qualify the main argument. For each blind spot, describe what type of literature might be missing. Do not name specific studies.

Researcher check: Blind spots are hypotheses to test through searching, not conclusions.

Stage 2 — Build Search Strategies

A good search strategy separates what you are looking for from where you look. The same claim should be searched in different ways for different tools: keywords and Boolean strings for databases, natural-language questions for scholarly AI search engines, and seed papers for citation mapping. Every strategy should include a search for evidence that could contradict the claim.

Prompt 3 — Translate claims into searches Purpose: Produce ready-to-use queries for several tools.

Prompt template: For each of the following claims from my paper, prepare a search strategy: [INSERT 3–5 CLAIMS FROM PROMPT 1] For each claim, provide: key concepts, with synonyms and related terms used in [DISCIPLINE(S)]; a Boolean search string suitable for Scopus or Web of Science; a natural-language research question suitable for Consensus or Elicit; one question designed to find evidence that contradicts or limits the claim; terms in [OTHER LANGUAGE, e.g., Spanish], if relevant literature may exist in that language. Do not answer the questions or suggest sources.

Researcher check: Are the terms the ones used in your field? Adjust jargon before running the searches.

Stage 3 — Discover Candidate Sources

Run the searches in more than one type of tool. Treat everything found at this stage as a candidate: a lead that must still be verified. Record where each candidate came from, as you will need this for the log in the companion template.

Prompt 4 — Search with a generalist assistant Purpose: Find candidates with live search and traceable evidence.
Best with: Claude or ChatGPT with web search or deep research on

Prompt template: Using web search, find peer-reviewed studies, published since [YEAR], that address the following question: [INSERT RESEARCH QUESTION FROM PROMPT 3] For each source, provide: full reference; DOI or stable link; one sentence on what the study actually found; and the phrase from the abstract or text that supports your summary. Only list sources you have retrieved in this session. Do not list sources from memory. If you cannot find a DOI or link, say so. Include studies whose findings disagree with or qualify the question's premise, and label them.

Researcher check: Open every link. A citation that looks right in the chat can point to a different paper or to nothing at all.

Prompt 5 — Search for recent and gray literature Purpose: Capture new work and policy material outside journals.
Best with: Perplexity (academic focus) or a generalist with web search

Prompt template: Using academic sources where possible, what has been published on [TOPIC] since [YEAR]? Separate the results into: (1) peer-reviewed articles; (2) preprints and working papers; (3) reports by public bodies, international organizations or think tanks. For each item, give the publisher or journal, the date and a direct link. State clearly when an item has not been peer reviewed.

Researcher check: Is each gray-literature source authoritative enough for the claim you want to use it for?

USING SCHOLARLY SEARCH AND CITATION MAPPING TOOLS

Scholarly AI search engines work best with short, precise questions rather than long prompts. Examples:

Consensus: "Does AI literacy training increase researchers' use of generative AI?" Then filter by study design (e.g., randomized trials, meta-analyses) and year, and open the papers classed as "no" and "possibly" as well as "yes."

Elicit: "What factors predict academic staff adoption of generative AI?" Then add columns such as population, country, method and main finding to compare studies side by side.

ResearchRabbit, Litmaps or Connected Papers: add three to five papers you have read and trust as seeds. Review the later works that cite them (to find recent developments) and the similar works (to find parallel literatures).

Stage 4 — Validate Every Source

Validation is the stage that cannot be delegated. The AI can help organise checks, but the researcher must confirm each one against the original record and, for any source that carries a claim, read the relevant part of the text. Besides, if the machine makes decisions for the author, at what point do you cease to be the author of your own paper? Could this be understood as academic dishonesty? It is the sole job of the researcher to assess and validate every source.

Apply five checks to every candidate:

CHECK	QUESTION	HOW TO VERIFY
1. Exists	Is this a real publication?	Resolve the DOI; find the record in a library database or Google Scholar; import it into your reference manager.
2. Accurate	Are authors, year, title, journal and pages correct?	Compare every field with the publisher's record. AI tools often mix details from different papers.
3. Relevant	Does it say what you need it to say?	Read at least the abstract and the relevant section; read the full text for any source that supports a key claim.
4. Credible	Is it good enough for this claim?	Check venue, peer review, design, sample and limitations. Watch for predatory journals.
5. Standing	Has it been retracted, corrected or challenged?	Check publisher notices, the Retraction Watch database and scite's supporting and contrasting citations.

PROMPT 6 — CHECK CLAIM–SOURCE FIT PURPOSE

Test whether a source supports the exact claim you are making.

Prompt template: Here is a claim from my paper: [INSERT CLAIM] Here is the relevant passage from the source I intend to cite, copied from the original: [INSERT PASSAGE, WITH PAGE NUMBER] Assess whether the passage supports the claim fully, partially, qualifies it, contradicts it, or is not relevant. Consider differences in population, context, method, time period and strength of language (for example, association versus causation). If the fit is partial, propose a rewording of my claim that matches what the source actually shows. Base your assessment only on the passage provided.

Researcher check: The AI sees only the passage you paste. Make sure it is representative of the source's overall findings.

PROMPT 7 — VERIFY A REFERENCE LIST PURPOSE

Detect invented or inaccurate references before submission.
Best with: a generalist with web search

Prompt template: Using web search, check each reference in the list below: [INSERT REFERENCE LIST] For each one, report: whether you found the publication; a link to the publisher's or database record; and any mismatch in authors, year, title, journal, volume, pages or DOI. Do not correct a reference by guessing. If you cannot find a record, mark the reference as "unverified."

Researcher check: Treat the AI's check as a first pass. Confirm every "unverified" or mismatched item yourself.

Prompt 8 — Assess the strength of a study Purpose: Judge how much weight a source can carry.

Prompt template: Below are the abstract and methods section of a study I plan to cite: [INSERT TEXT] Summarise the research design, sample, setting and main limitations. Then explain what kind of claim this study can support (for example, descriptive, correlational, causal) and what claims it cannot support. Do not assess the quality of the journal or the authors' reputation.

Stage 5 — Incorporate Sources into the Argument

A verified source is not yet part of the argument. Decide what role it plays: does it support a claim, extend it to a new context, qualify it with conditions, or contradict it? Sources that qualify or contradict a claim are often the most valuable additions, because they show that the review is balanced and they sharpen the contribution. Arguments

should be made by researchers, not by machines.

> Prompt 9 — Position a new source in the argument Purpose: Decide where and how a verified source fits.
>
> Prompt template: Below is a paragraph from my literature review and my notes on a new source I have read and verified: [INSERT PARAGRAPH] [INSERT YOUR NOTES ON THE NEW SOURCE: FINDINGS, CONTEXT, METHOD] Explain whether the new source supports, extends, qualifies or contradicts the paragraph's argument. Then propose two ways of integrating it into the paragraph. Preserve my existing claims and citations unless the new source requires a claim to be qualified. If it does, say so explicitly and explain why. Use only the information in my notes. Do not add details about the source.
>
> Researcher check: If the new source weakens a claim, have you changed the claim, or just added the citation?

> Prompt 10 — Synthesize instead of stacking citations Purpose: Turn a list of sources into an argument.
>
> Prompt template: The paragraph below now cites several additional sources: [INSERT PARAGRAPH WITH CITATIONS] Identify where sources are simply listed one after another. Propose a revised structure that groups them by agreement, disagreement and the conditions under which findings differ. Keep every citation attached to the claim it currently supports. Do not introduce new claims or references.

> Prompt 11 — Final citation consistency check Purpose: Catch orphan and mismatched citations.
>
> Prompt template: Compare the in-text citations in the manuscript below with its reference list. [INSERT MANUSCRIPT] [INSERT REFERENCE LIST] List: citations missing from the reference list; references never cited in the text; and inconsistencies in author names or years between the two. Do not change the text.

Record, Disclose and Avoid Common Pitfalls

Keep a record of each source found with AI assistance: which tool found it, how it was verified, and what you decided. The companion template, Include, Adapt or Discard, provides a log for this purpose. Check whether your target journal, institution or funder requires you to disclose the use of AI in the literature search and follow its wording.

PITFALL	WHAT IT LOOKS LIKE	HOW TO PREVENT IT
Invented references	A plausible title, real-sounding authors and a DOI that leads nowhere or to a different paper.	Never use references generated without search. Resolve every DOI.
Citation drift	A real paper cited for a claim it does not make.	Read the relevant passage; use Prompt 6.
Evidence inflation	A correlational study cited as proof of causation, or a small sample cited as a general finding.	Match the claim's strength to the design; use Prompt 8.
Confirmation bias	Every new source supports the draft.	Always run a counter-question (Prompt 3).
Abstract-only reading	The source is cited on the basis of an AI summary or abstract.	Read the full text for any source carrying a key claim.
Coverage gaps	Books, non-English work and regional journals are missing.	Combine AI tools with library databases and searches in other languages.
Secondary citation	Citing an original study found only through another paper's description of it.	Locate and read the original, or cite it as a secondary source.

QUICK CHECKLIST

☐ I mapped my own claims before searching.

☐ I searched in more than one type of tool, including for counter-evidence.

☐ Every new source has been located in an authoritative record.

☐ I have read the part of each source that supports my claim.

☐ Each claim is no stronger than its sources allow.

☐ I checked for retractions and contrasting evidence.

☐ I have recorded how each source was found and verified, and disclosed AI use where required.

> **AI CAN WIDEN THE SEARCH**
> Only the researcher can decide what the evidence shows. Only the researcher is the author of the paper.

AI-assisted source discovery and validation log

> **AI MAY SUGGEST SOURCES**
> The researcher remains responsible for verifying, reading and correctly using every source cited.

This reusable template helps researchers evaluate and document sources found with the help of AI tools when enriching a literature review. It accompanies the guide AI Lit Reviews: Source Enrichment and Validation. It is designed for finding and checking sources, not for delegating the reading of the literature or the judgement about what it shows.

RESEARCHER	CLICK OR TYPE HERE
Project / manuscript	Click or type here
Section reviewed	Click or type here
AI tools used	Click or type here
Date / version	Click or type here

HOW TO USE THIS TEMPLATE

> Define which claims you want to support, test or update, and complete the checks in Section 1.
>
> Search with AI tools and record every candidate source you consider using.
>
> Verify each candidate against the original record and read the relevant part of the text.
>
> Decide whether to INCLUDE, ADAPT or DISCARD each source, and record why.
>
> Complete the final integrity check before submitting or sharing the text.

This log does not need to capture every search result. Prioritise sources you intend to cite, and sources that could change a claim, such as studies that contradict or qualify your argument.

Before you begin

Use this page to set up a safe and transparent search before asking an AI tool to find sources.

FIELD	DETAIL	
Check	Yes / No / N.A.	Notes or action required
Authorised environment The tools are approved by your institution for the material you will share.	☐ Yes ☐ No ☐ N.A.	
Confidentiality No unpublished, confidential or personal data will be uploaded to an unapproved tool.	☐ Yes ☐ No ☐ N.A.	
Publisher and institutional rules Relevant journal, funder, university and project policies on AI use have been checked.	☐ Yes ☐ No ☐ N.A.	
Claims mapped first The claims to support, test or update have been identified before searching.	☐ Yes ☐ No ☐ N.A.	
Counter-evidence included The search will look for studies that contradict or qualify the argument, not only for support.	☐ Yes ☐ No ☐ N.A.	
Reading commitment Every source carrying a claim will be read, at least in the relevant section, before it is cited.	☐ Yes ☐ No ☐ N.A.	

DEFINE THE PURPOSE OF THE SEARCH

☐ Support uncited or weakly supported claims

☐ Test claims against contradicting evidence

☐ Update the review with recent work

☐ Broaden coverage (other disciplines, regions or languages)

☐ Find methodological or theoretical references

☐ Other: _____

Risk map for AI-found sources

Use the risk level to decide how much scrutiny a source requires. Risk depends on how the source was found and checked, and on the claim it is asked to carry, not on how relevant the AI's summary makes it sound.

GREEN: Usually low risk after review

Located in a library database or on the publisher's site, with all details matching

Full text read; the cited passage supports the claim as written

Peer-reviewed, with a design suited to the claim

No retraction, correction or major contrasting evidence found

AMBER: Substantive review required

Known only through an AI summary or the abstract

Preprint, working paper or gray literature

> Different population, context or period from the claim
>
> Found through another paper's description (secondary citation)
>
> Contrasting citations exist and have not yet been assessed

> RED: Do not cite until resolved

> Cannot be located, or the DOI leads to a different paper
>
> Authors, title, year or journal do not match the record
>
> Retracted, or subject to an expression of concern
>
> Cited for a claim the source does not make
>
> Added to support a claim without being read

WATCH FOR THESE SOURCE PROBLEMS

PROBLEM	TYPICAL SIGN
Invented reference	Plausible title and authors; no record exists
Hybrid reference	Real authors attached to a wrong title, year or journal
Citation drift	Real paper; the claim it is cited for is not in it
Evidence inflation	correlational study → cited as causal evidence
Scope inflation	one country or sample → cited as a general finding
Summary distortion	AI summary stronger or broader than the paper's own conclusion
Outdated standing	Later work has retracted, corrected or overturned the finding

Suggested instruction to the AI

Adapt this instruction to the tool and the task. It is intentionally restrictive.

> Search for sources relevant to the claims below using live search only. For each source, give the full reference, a DOI or stable link, and the passage that supports your summary. Do not list sources from memory. If you cannot verify a detail, say so. Include studies that contradict or qualify the claims, and label them.

FIVE CHECKS FOR EVERY SOURCE

CHECK	QUESTION	EVIDENCE
1. Exists	Is it a real publication?	DOI resolves; record in a database or on the publisher's site
2. Accurate	Are all bibliographic details correct?	Match against the publisher's record
3. Relevant	Does it say what I need it to say?	Relevant passage read; page number noted
4. Credible	Is it strong enough for this claim?	Venue, design, sample and limitations checked
5. Standing	Has it been retracted or challenged?	Retraction notices and contrasting citations checked

DECISION DEFINITIONS

INCLUDE: The source passes all five checks and supports the claim as written.

ADAPT: The source is valid, but the claim or the source's role must change (for example, the claim is narrowed, or the source is cited as qualifying or contradicting evidence).

DISCARD: The source fails a check, cannot be verified, or does not support any claim in the paper.

Source log

Record each source you consider citing. Add rows as needed. Use short, evidence-based reasons rather than general statements such as "relevant."

FIELD	DETAIL							
No.	Claim / section in draft	Candidate source (full reference)	Found via	Verified	Role	Risk	Decision	Reason / evidence checked	Final use in text
1		☐ Exists ☐ Details ☐ Read ☐ Standing	Supports / Extends / Qualifies / Contradicts		Include / Adapt / Discard				
2		☐ Exists ☐ Details ☐ Read ☐ Standing	Supports / Extends / Qualifies / Contradicts		Include / Adapt / Discard				
3		☐ Exists ☐ Details ☐ Read ☐ Standing	Supports / Extends / Qualifies / Contradicts		Include / Adapt / Discard				
4		☐ Exists ☐ Details ☐ Read ☐ Standing	Supports / Extends / Qualifies / Contradicts		Include / Adapt / Discard				
5		☐ Exists ☐ Details ☐ Read ☐ Standing	Supports / Extends / Qualifies / Contradicts		Include / Adapt / Discard				

Source log (continued)

FIELD	DETAIL								
No.	Claim / section in draft	Candidate source (full reference)	Found via	Verified	Role	Risk	Decision	Reason / evidence checked	Final use in text
6									
7									
8									

FIELD	DETAIL
9	
10	
11	
12	

Worked example

Illustrative example. The sources below are fictional and are used only to show how the log works.

Claim in the draft: "AI literacy training increases researchers' adoption of generative AI."

CANDIDATE A — SUGGESTED BY A GENERALIST ASSISTANT WITHOUT WEB SEARCH

FIELD	DETAIL
AI suggestion	[Author], [Year]. "AI literacy and faculty adoption of generative AI: a longitudinal study." [Journal], with a DOI.
Checks	DOI leads to an unrelated article. No record of the title in Scopus, Google Scholar or the journal's archive.
Risk level	RED
Decision	DISCARD
Why	Invented reference. The assistant produced it from memory, not from a search.

CANDIDATE B — FOUND WITH A SCHOLARLY AI SEARCH ENGINE

FIELD	DETAIL
AI summary	"AI literacy training increases adoption of generative AI."
Checks	Exists and details match. Full text read: cross-sectional survey of doctoral students in one country; finds that self-reported AI literacy is associated with intention to use generative AI. No training intervention was tested.
Role	Qualifies the claim: supports an association, not a causal effect of training, and in a narrower population.
Risk level	AMBER
Decision	ADAPT
Final approved version	"Among doctoral students, higher self-reported AI literacy has been associated with greater intention to use generative AI [Source B], although evidence on the effect of training interventions remains limited."

> A source that exists is not necessarily a source that supports your claim.

FINAL INTEGRITY CHECK

EXISTENCE AND ACCURACY

☐ Every cited source has been located in an authoritative record.

☐ All bibliographic details match the publisher's record.

☐ No reference generated from an AI tool's memory remains unverified.

FIT AND STRENGTH

☐ I have read the passage that supports each claim.

☐ No claim is stronger, broader or more causal than its sources allow.

☐ Sources that qualify or contradict the argument are acknowledged.

STANDING

☐ Retractions, corrections and expressions of concern have been checked.

☐ Contrasting evidence for key sources has been considered.

TRANSPARENCY AND COMPLIANCE

☐ The tools used were appropriate for the material shared.

☐ Institutional, publisher and funder rules were followed.

☐ This log has been retained where appropriate.

☐ AI use has been disclosed or acknowledged where required.

RESEARCHER APPROVAL

FINAL DECISION	☐ APPROVED ☐ FURTHER VERIFICATION REQUIRED
Researcher name	
Signature / confirmation	
Date and version	

> **FINAL RESPONSIBILITY CANNOT BE DELEGATED**
> the named author remains accountable for every source cited.

OPTIONAL AI-USE RECORD AND DISCLOSURE

Use this section only where relevant. Adapt the wording to the journal, institution, funder or disciplinary guidance that applies.

FIELD	DETAIL
Tools and versions	
Date(s) used	
Purpose of use	Claim mapping / search strategy / source discovery / verification / integration / other

FIELD	DETAIL
Material provided to the tools	
Human verification performed	
Relevant policy checked	
Location of retained log	

EXAMPLE DISCLOSURE STATEMENT

> AI tools were used to help identify candidate sources for the literature review. The authors verified each source against the original publication, read the relevant sections, and remain fully responsible for the selection, interpretation and citation of all sources.

Do not use this example automatically. Disclosure requirements differ across publishers, institutions and disciplines. Always follow the rules that apply to the specific output.

About this resource

TITLE	AI literature reviews: source enrichment and validation
TOOLKIT REFERENCE	UC1-03, use case 1: Literature review and synthesis
DEVELOPED BY	UPF Barcelona School of Management, Rodrigo Cetina
SUGGESTED CITATION	AI Researchers (2026). AI literature reviews: source enrichment and validation. AI Readiness Toolkit, UC1-03. UPF Barcelona School of Management, Rodrigo Cetina. airesearchers.eu
QUALITY ASSURANCE	Reviewed by the AI Researchers consortium for quality, categorisation and alignment with the AI Engagement Framework. External peer review by senior academics is under way, and the item may be updated as a result.
VERSION	1.0, November 2026
PROJECT	AI Researchers: Embedding Responsible AI in Research. Erasmus+ Cooperation Partnership in Higher Education, KA220-HED-66835D2B. October 2025 to March 2028.
CONSORTIUM	University of Lodz (coordinator), FEP University of Porto, UPF Barcelona School of Management, ACEEU, NIPI (OPI PIB), European E-learning Institute and Momentum Marketing Services

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author or authors only and do not necessarily reflect those of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the entity providing the grant can be held responsible for them.

AI Researchers
Embedding Responsible AI in Research