



FUTUREBANKER

YOUR TOMORROW STARTS HERE

AI RESEARCH TOOLKIT

A Practical Guide to Gemini, ChatGPT, Claude & Perplexity for Research

COMPREHENSIVE WORKFLOW & REFERENCE GUIDE

YOUR TOMORROW STARTS HERE

Version 4.0 • August 2026

CONTENTS

This guide brings together four of the leading AI research tools — Claude, ChatGPT, Google Gemini with Gemini Notebook, and Perplexity — into a single, practical reference for research work. The method is taught once, tool-agnostically, in Part Two; Parts Three to Six show exactly how to run it in each tool; Part Seven covers what the workflow alone does not — data protection and classification, a pre-use sign-off checkpoint, ethics clearance, bias, reproducibility, and disclosure.

PART ONE EXECUTIVE DECISION MATRIX

- Which tool and mode for which job, at a glance

PART TWO THE UNIVERSAL AI RESEARCH METHODOLOGY

- Define the question, build a plan, search & gather
- Evaluate & verify sources (SIFT), resolve contradictions
- Synthesise, build a research brief, audit the final answer
- Prompt library & master workflow checklist

PART THREE CLAUDE IN PRACTICE

- Modes, Chat, Projects, Cowork, Design, Code
- Worked example: a regulatory research task

PART FOUR CHATGPT IN PRACTICE

- Search, Deep Research, uploaded documents
- Worked example: a competitor landscape

PART FIVE GEMINI & GEMINI NOTEBOOK IN PRACTICE

- Gemini Notebook: the 2026 rebrand of NotebookLM
- Deep Research, Studio tools, worked example

PART SIX PERPLEXITY IN PRACTICE

- Source-first search, Pro Search, Deep Research, Projects
- Source hierarchy, tool-comparison workflow, worked example

PART SEVEN RESPONSIBLE & ETHICAL USE IN RESEARCH

- POPIA data classification and pre-use sign-off checklist
- Ethics clearance, bias & African relevance, reproducibility
- AI-use disclosure template, escalation contacts

PART ONE — EXECUTIVE DECISION MATRIX

Choose the right tool for the job in seconds — then turn to the relevant part for the step-by-step.

At a Glance: Which Tool for Which Job

Tool	Core strength	Best for
Claude Chat — Web Search	Fast, conversational first-pass research	Quick lookups and single-topic web search.
Claude Chat — Extended Thinking	Deep reasoning without new information	Maths, logic, code debugging, argument analysis.
Claude Chat — Research	Multi-source, documented investigation	Complex questions needing a cited report.
Claude Projects	Persistent research workspace	Reusing the same standards, tone, and knowledge across engagements.
Claude Cowork	Agentic, multi-file workspace	Synthesising several source files into a finished Word/Excel/PPT deliverable.
Claude Design	Prompt-to-canvas visual output	Turning finished research into client-ready decks and one-pagers.
Claude Code	Terminal/IDE technical research	Data-heavy or repeatable, script-based analysis.
ChatGPT Search	Web search with citations	Current facts, recent developments, source discovery.
ChatGPT Deep Research	Multi-step, documented research	Complex questions needing structured, cited synthesis.
ChatGPT (files)	Document-grounded analysis	Researching your own reports, PDFs, spreadsheets, or notes.
Gemini	Web-connected discovery	Finding initial sources, trends, hypotheses, draft outlines.
Gemini Deep Research	Multi-step, documented research	Complex questions needing structured, cited synthesis.
Gemini Notebook	Source-grounded synthesis	Reading dense PDFs, cross-analysing documents, audio/video overviews.
Perplexity Search	Fast, source-cited default answer engine	Quick, current, citation-backed lookups.
Perplexity Pro Search	Deeper multi-source research pass	More complex queries needing broader sourcing or model/code options.

Tool	Core strength	Best for
Perplexity Deep Research	Multi-step, documented research agent	Complex questions needing a long, structured, cited report.
Perplexity Projects	Persistent, shareable research workspace	Reusing context and files across an ongoing research project or team.

Research Mode Comparison

Claude modes

Mode	Best for	Typical output
Web Search	Straightforward, current facts — one or two searches	A direct answer with source links
Extended Thinking	Complex reasoning that doesn't need the web	A more thorough, multi-angle answer
Research	Complex questions needing five or more searches	A structured, cited research report
Uploaded files / Projects	Researching your own reports, PDFs, or standing knowledge	Document-grounded analysis

ChatGPT modes

Mode	Best for	Typical output
Normal chat	Explaining concepts, brainstorming, outlining a topic	Brief explanation or working notes
Search	Current facts, recent developments, source discovery	Answer with web citations/sources
Deep Research	Complex questions requiring multiple sources and synthesis	Structured, documented research report
Uploaded files	Researching your own reports, PDFs, spreadsheets, or notes	Document-grounded analysis

Gemini modes

Mode	Best for	Typical output
Gemini (chat)	Explaining concepts, brainstorming, quick current facts	Brief explanation, with web citations

Mode	Best for	Typical output
Gemini Deep Research	Complex questions requiring multiple sources and synthesis	Structured, documented research report
Gemini Notebook	Researching your own reports, PDFs, notes, or links	Grounded, citation-backed analysis

Perplexity modes

Mode	Best for	Typical output
Search	Straightforward, current facts answerable in one search	A cited, synthesised answer
Pro Search	More complex queries needing broader sourcing, or model/code options	A deeper, multi-source cited answer
Deep Research	Complex questions needing a long, multi-step investigation	A structured, cited research report
Projects	Reusing the same research context, files, or team across a project	A persistent, shared research workspace

Gemini vs. Gemini Notebook

Feature / Function	Gemini App	Gemini Notebook
Primary Source Data	Entire Public Web + Real-Time Google Search	User-Uploaded Files (PDFs, Docs, Web)
Hallucination Risk	Moderate — always requires independent verification	Lower when answers are constrained to the uploaded sources — still requires verification
Citations	Web search link cards	Inline exact passage highlights & quotes
Context Window	Large context (Gemini Advanced)	Up to 50 sources, 25 million words total
Data Privacy	Subject to workspace privacy controls — verify current terms	Provider states personal data isn't used to train global models — verify current terms and your POPIA basis before uploading confidential data

HOW TO USE THIS MATRIX

Pick your tool and mode above, then jump to its part: Claude (Part Three), ChatGPT (Part Four), Gemini & Gemini Notebook (Part Five), or Perplexity (Part Six). Every part assumes you already know the shared research method in Part Two — it covers only what's specific to that tool.

Note: Feature names, plan availability, and context-window limits change frequently across all four tools — Google renamed NotebookLM to Gemini Notebook in July 2026, and Perplexity has since renamed Spaces to Projects. Confirm current specifications against each provider's official documentation before relying on them in a client-facing session.

PART TWO — THE UNIVERSAL AI RESEARCH METHODOLOGY

One disciplined process, usable in Claude, ChatGPT, Gemini, or Perplexity — taught once here, then applied tool by tool in Parts Three to Six.

Good AI-assisted research does not depend on which tool you open. It depends on running the same disciplined sequence every time: define the question, gather evidence, verify it, resolve disagreement, synthesise with an honest confidence level, structure the deliverable, and audit before you rely on it. This part teaches that sequence once. Parts Three to Six show exactly how to execute each step inside Claude, ChatGPT, Gemini / Gemini Notebook, and Perplexity — without re-explaining the method itself.

1. Define the Research Question

Weak research begins with a vague prompt. Strong research starts by specifying the topic, purpose, audience, timeframe, geography, evidence standard and desired output.

Use the 6-part research question formula:

TOPIC What?	PURPOSE Why?	WHO? For whom?
TIME When?	SCOPE Where / how much?	OUTPUT What format?

Why this is stronger.

- It sets a clear subject and geography.
- It defines a time window.
- It tells the AI which source types of matter.
- It requests separation of evidence from interpretation.
- It specifies the final deliverable.

EXAMPLE PROMPT

“Act as a research analyst. Help me investigate the adoption of generative AI in South African banking. Focus on 2024–2026, identify major adoption drivers, risks, regulatory considerations, and practical use cases. Prioritise official regulators, banks, academic research, and reputable industry reports. Separate facts from interpretation. End with 5 research gaps and a source list.”

THIS PROMPT IS TOOL-AGNOSTIC

Every prompt in this part works unchanged in Claude, ChatGPT, Gemini, or Perplexity — paste it in as written. Parts Three to Six show the platform-specific mechanics (which button, which mode, which upload panel) for running it.

2. Build a Research Plan

Before asking for a long report, break the problem into sub-questions. This reduces the risk of overlooking an important dimension of the topic.

EXAMPLE PROMPT

“Turn my research question into a research plan. Create 6–8 sub-questions. For each, state: (1) what evidence is needed, (2) ideal source types, (3) keywords, (4) bias or limitation, and (5) how the finding will contribute to the conclusion.”

Level	Question
Research question	What is happening and why?
Sub-question 1	What does the evidence show?
Sub-question 2	Who is affected?
Sub-question 3	What are the risks or constraints?
Sub-question 4	What do authoritative sources say?
Sub-question 5	Where do sources disagree?
Sub-question 6	What remains unknown?

TIP

If the plan is too broad, ask the tool to reduce the scope. If it is too narrow, ask it to add missing dimensions such as geography, stakeholder groups, counterarguments, costs, or implementation barriers.

3. Search & Gather Evidence

Whichever tool you are in, prioritise the same tier of source — and tell the tool explicitly to do the same.

Tier	Description
1 — Primary	Legislation, official statistics, original studies, regulator publications

Tier	Description
2 — Strong secondary	Systematic reviews, reputable research institutions, established industry reports
3 — Useful secondary	Quality journalism, professional analysis
4 — Discovery only	Blogs, forums, unsourced summaries, social posts

EXAMPLE PROMPT

“Search for recent evidence on [TOPIC]. Prioritise primary sources, government/regulator publications, peer-reviewed research, and official organisational reports. For each important claim, show the source and explain exactly what it supports. Flag sources older than [X] years unless they are foundational.”

4. Evaluate and Verify Sources

Never treat a citation as proof by itself. The source must support the claim.

A SIFT-style verification sequence.

Step	Action
STOP	Pause before accepting a surprising or consequential claim.
INVESTIGATE	Who published it? What is the organisation's authority?
FIND BETTER COVERAGE	Look for an original study, official dataset, regulator, or primary document.
TRACE	Ask the tool to quote the exact passage the claim comes from and check the wording matches.

EXAMPLE PROMPT

“Audit the following claims. For each, rate the evidence as strong / moderate / weak / unsupported. Identify the exact source that supports it, explain any mismatch between the claim and the source, and tell me what additional evidence would be needed to establish it confidently.”

5. Compare Evidence and Resolve Contradictions

Good research does not hide disagreement. Ask the tool to map competing claims and explain whether the disagreement comes from different definitions, populations, dates, methods, samples, or assumptions.

Contradiction map

Claim A	Claim B	Why they differ?	What would resolve it?
Source says X	Source says Y	Different samples / dates / definitions	Primary dataset or newer study

Claim A	Claim B	Why they differ?	What would resolve it?
Study reports association	Commentator says causation	Evidence is being interpreted too strongly	Causal design / stronger evidence

EXAMPLE PROMPT

“Compare these sources on [QUESTION]. Create a claim-by-claim evidence matrix. For each disagreement, determine whether it is caused by different dates, definitions, samples, methods, source quality, or interpretation. Do not force consensus. State which conclusion is best supported and why.”

6. Synthesise Findings

Synthesis is more than summarising each source. Ask the tool to combine evidence across sources and explain the pattern that emerges.

A strong synthesis has four layers:

WHAT? What does the evidence show?	SO WHAT? Why does it matter?	HOW SURE? How strong is the evidence?	NOW WHAT? What should happen next?
--	--	---	--

Confidence language

- High confidence: multiple strong, independent sources converge.
- Moderate confidence: evidence is consistent but has limitations.
- Low confidence: evidence is sparse, indirect, or conflicting.
- Unknown: insufficient evidence to make a responsible conclusion.

EXAMPLE PROMPT

"Synthesize the evidence into 5–7 key findings. For each finding provide: finding → evidence → interpretation → confidence → implication. Include counter-evidence where it exists. Do not introduce claims that cannot be traced to the evidence set."

7. Build a Research Brief

Once the evidence has been verified, ask the tool to produce the final deliverable. Give it a structure so the output is useful to your audience.

#	Section	Content
1	Executive summary	5–8 bullets covering the answer and implications.
2	Research question	The exact question, scope, and definitions.
3	Method	Sources searched, date range, inclusion/exclusion criteria.
4	Findings	Evidence-backed findings with citations.
5	Contradictions	Where evidence disagrees and why.
6	Implications	What the findings mean for the audience.
7	Limitations	What cannot be concluded.
8	Research gaps	Questions still requiring investigation.
9	References	Source list with links/citations.

EXAMPLE PROMPT

“Write a professional research brief for [AUDIENCE] based only on the verified evidence. Use the structure above. Put citations immediately after claims they support. Clearly label inference, uncertainty, and recommendations. End with a concise list of the most important sources.”

8. Audit the Final Answer

Before using a research output, run a final audit. This is one of the most important steps because polished prose can conceal weak evidence.

Audit question	Pass?
Does every major factual claim have evidence?	☐
Does the source support the wording used?	☐
Are important sources primary or authoritative where possible?	☐
Are dates current enough for the research question?	☐
Are contradictory findings represented fairly?	☐
Are facts separated from interpretation and recommendations?	☐
Are limitations and evidence gaps stated?	☐
Have high-stakes claims been independently verified?	☐

EXAMPLE PROMPT

“Perform a final research audit. Identify any factual statement that is weakly supported, uncited, outdated, overstated, or based on an indirect source. Return a table with claim, source, support level, problem, correction required.”

9. Prompt Library

Copy, adapt and reuse these prompts — everyone works unchanged in Claude, ChatGPT, or Gemini.

Topic exploration

EXAMPLE PROMPT

"Give me a research landscape for [TOPIC]. Identify the major themes, debates, key terms, leading organisations/researchers, and unanswered questions. Do not draft the final report yet."

Literature / source map

EXAMPLE PROMPT

"Map the sources on [TOPIC]. Group them by theme, methodology, findings, and limitations. Identify where they agree and disagree."

Source discovery

EXAMPLE PROMPT

"Find high-quality sources for [QUESTION]. Prioritise primary sources, official statistics, regulators, peer-reviewed research, and original datasets. Explain why each source is relevant."

Source audit

EXAMPLE PROMPT

"Audit these references for existence, relevance, date, authority and whether they support the associated claim. Flag anything that needs replacement."

Counterargument

EXAMPLE PROMPT

"Challenge my current conclusion: [CONCLUSION]. Find credible evidence that supports, weakens, or contradicts it. Separate strong counter-evidence from speculation."

Evidence matrix

EXAMPLE PROMPT

"Create an evidence matrix with columns: claim, source, evidence, method, date, limitations, confidence and implication."

Research gap

EXAMPLE PROMPT

"Based on the evidence, identify the most important unanswered questions. Rank them by importance, feasibility, and potential impact."

Executive summary

EXAMPLE PROMPT

"Turn the verified findings into a one-page executive summary for [AUDIENCE]. Keep claims evidence-backed and include the most important implications."

10. Master Research Workflow Checklist

Use this page as a repeatable process for every future research project, in any tool.

- Define the research question and the decision or purpose it serves.
- State the audience and purpose up front (e.g. "for a facilitator briefing" vs "for a client-facing deck").
- Set the date range, geography, and scope.
- Break the question into 6–8 sub-questions before requesting a long report.
- Specify the format you want back — Word document, comparison table, bullet summary, slide outline.
- Choose the right mode or surface for the job (see Part One's Executive Decision Matrix) — or a combination.
- Prioritise authoritative and primary sources.
- Ask for citations and inspect the underlying sources — never treat a citation as proof by itself.
- Compare multiple sources rather than relying on one answer.
- Look specifically for contradictions and counter-evidence.
- Separate facts, interpretations, assumptions, and recommendations.
- Record limitations and evidence gaps.
- Combine evidence across sources and state a clear confidence level for each finding.
- Structure the final deliverable for its audience.
- Audit important claims before publication or decision-making.
- Keep a source list / research log for reproducibility.
- Review and verify before use — treat all AI output as a strong first draft for your professional review.
- Reuse what works — save prompts, Projects or Notebooks that produce a good result, so future research starts from a consistent baseline.

GOLDEN RULE

Move from question → evidence → verification → synthesis → decision. The quality of AI-assisted research depends less on one perfect prompt and more on running this disciplined, iterative process — in every tool.

PART THREE — CLAUDE IN PRACTICE

How to run the Part Two methodology inside Claude Chat, Projects, Cowork, Design and Code.

This part covers what is specific to Claude: its working surfaces, its three Chat modes, and the mechanics of each. For the research method itself — defining a question, verifying sources, resolving contradictions, synthesising, building a brief, auditing — see Part Two, which applies here without any changes.

The working surfaces at a glance

- Claude Chat — the conversation window at claude.ai, with three research modes built in: Web Search, Extended Thinking, and Research. Best for quick lookups through to a full multi-source investigation of a single topic.
- Claude Projects — a persistent workspace inside Chat that holds custom instructions and a standing knowledge base of your own documents, so every chat in the Project starts from the same context.
- Claude Cowork — an agentic workspace (desktop, web or mobile) that reads and writes real files. Best for multi-step research that pulls together several sources into a finished document, spreadsheet, or deck.
- Claude Design — a prompt-to-canvas product for turning findings into polished visuals: decks, one-pagers, prototypes.
- Claude Code — a terminal/IDE tool built for technical work: data pipelines, scripted analysis, repository-based research.

Which surface should you use?

If your research need is...	Use...
A quick, current fact or first-pass exploration of a topic	Claude Chat — Web Search
Careful reasoning on a problem that doesn't need new information	Claude Chat — Extended Thinking
A full multi-source investigation with a cited report	Claude Chat — Research
The same research structure, standards or tone reused across many engagements	Claude Projects
Synthesising several source files into one finished document	Claude Cowork
Turning finished findings into a client-ready visual	Claude Design
Data-heavy or repeatable, script-based analysis	Claude Code

1. Choose the Right Mode Within Claude

Claude Chat is not a single mode — Web Search, Extended Thinking and Research are three distinct tools you can call on depending on what the question needs. Anthropic's own guidance is using Web Search for quick, current

facts; use Extended Thinking for complex reasoning that doesn't need the internet; and use Research when a question needs five or more searches synthesised into a documented report.

Mode	Best for	Typical output
Web Search	Straightforward, current facts answerable in one or two searches — news, figures, quick lookups	A direct answer with source links
Extended Thinking	Complex reasoning that doesn't need the web — maths, code debugging, argument analysis	A more thorough, multi-angle answer
Research	Complex questions needing five or more searches and 1–3 minutes of synthesis — competitor analysis, updating a document with current data	A structured, cited research report
Research + Extended Thinking	Large projects needing both strategic planning and broad evidence gathering	A comprehensive, documented report

NOTE

Feature names and plan availability (Web Search, Extended Thinking, Research, Projects) change as Anthropic updates Claude. Confirm current names and access on your account before training clients on this material.

2. Claude Chat as a Research Tool

Step-by-step: running a research query in Chat

Step	What to do
1	Open claude.ai and start a new chat.
2	State your research question clearly, including the audience and purpose (e.g. "Summarise recent changes to FAIS Board Notice requirements for a facilitator briefing").
3	If the topic is time-sensitive, ask Claude to search the web, or enable Web Search if you do not see it search automatically.
4	For a broad or multi-source question, ask Claude whether Research would suit it better — this runs a longer, multi-source investigation and returns a documented report with citations.
5	Review the answer and, importantly, check any cited sources before relying on the content for client-facing material.
6	Ask targeted follow-up questions to narrow, deepen, or reframe the research for your specific workshop or curriculum need.

EXAMPLE PROMPT

"Search for the latest guidance on FAIS Board Notice 194 of 2017 and summarise the key compliance obligations for financial services facilitators, with sources."

COMPLIANCE NOTE

Claude Chat should verify current facts, regulations, and figures with the search tool rather than relying purely on its training data. Always verify regulatory citations (FAIS Board Notices, NQF levels, SAQA standards) against the official source before using them in client materials.

Working with your own documents in Chat

Upload PDFs, notes or spreadsheets directly into the chat and ask Claude to treat them as a defined evidence set, rather than pasting long text into the conversation — this is the fastest way to apply Part Two's methodology to material you already have.

Claude Projects: a persistent research workspace

A Project keeps custom instructions and a standing knowledge base — your brand guidelines, standards documents, or a library of prior briefings — available to every chat inside it. Set it up once, and every future research task in that Project starts from the same baseline.

Step	What to do
1	Create a new Project and give it a clear name (e.g. "NQF & SAQA Regulatory Research").
2	Add custom instructions describing the audience, tone, and output format you want by default.
3	Upload standing reference material to the Project's knowledge — standards documents, previous briefings, style guides.
4	Start each new research task as a chat inside the Project, rather than a standalone chat, so it inherits that context automatically.
5	Update the Project's knowledge as standards or standing material change, so future chats stay current.

3. Claude Cowork as a Research Tool

Cowork is Claude's agentic workspace for multi-step, file-based research. It reads a folder of source material — PDFs, notes, transcripts, spreadsheets — and produces finished deliverables: a formatted Word briefing, an Excel comparison table, or a synthesised report, without you copying and pasting content between windows.

Step-by-step: a Cowork research workflow

Step	What to do
1	Create a dedicated folder for the research task (e.g. "NQF-Level6-Curriculum-Research").

Step	What to do
2	Add your source files to the folder: existing curriculum documents, SAQA unit standards, prior workshop material, client briefs, articles.
3	Open Cowork and grant it access to that folder.
4	Describe the outcome you want in plain language — be specific about format, audience, length, and tone.
5	Let Claude work through the task. For longer research tasks, you can step away — Cowork can continue running and you can check back later, including from your phone.
6	Review the draft output in the folder, and give Claude follow-up instructions to refine tone, structure, or depth.

A Cowork document-analysis workflow

CONNECT Grant access to a source folder	EXTRACT & COMPARE Cross-reference every file	DRAFT Produce the finished deliverable	REFINE Review and give follow-up instructions
---	--	--	---

Where Cowork is strongest

- Synthesising multiple regulatory or standards documents into a single briefing.
- Turning raw workshop notes or transcripts into a polished curriculum draft.
- Building comparison tables (e.g. current vs updated NQF unit standards) directly into a working Excel file.
- Running the same research structure repeatedly via a saved Project, so formatting and depth stay consistent across engagements.

EXAMPLE PROMPT

“Point Cowork at the folder [FOLDER NAME]. Produce a [LENGTH]-page Word briefing for [AUDIENCE] comparing [CURRENT MATERIAL] against [NEW/UPDATED MATERIAL]. Flag any gaps or contradictions, cite the specific source document for each claim, and end with an action list.”

DATA CARE NOTE

Cowork can browse the web and connect to tools such as email or cloud storage when authorised. Treat any client-identifying or confidential material with the same care you would in any other software — check what the task needs access to before granting it.

4. Claude Design as a Research Presentation Tool

Claude Design is not a research-gathering tool in itself — it is where research findings become client-ready visuals. Use it once your research is settled (from Chat, Cowork, or Part Two's brief structure), and you need to present it as a deck, one-pager, or infographic in your organisation's brand identity.

Step-by-step: turning research into a visual

Step	What to do
1	Finish your research and synthesis first — Design works best from a clear brief, not raw source material.
2	Open Claude Design and describe the deliverable: audience, key messages, structure, and any brand guidelines (colours, fonts, logo) to apply.
3	Paste in your synthesised findings, key statistics, or talking points as the content for the visual.
4	Review the generated layout and iterate — ask for changes to structure, hierarchy, or specific slides/sections.
5	Export the finished visual for use in a workshop, proposal, or social content.

Giving Design a strong brief

Brief element	What to specify
Audience	Who is reading or viewing this, and in what setting.
Key message	The one finding or recommendation the visual must land.
Structure	Sections or slides, and the order they should appear in.
Brand assets	Colours, fonts, logo, and any layout conventions to follow.
Evidence	The specific statistics, quotes, or findings to include — sourced from your finished research.

NOTE

Confirm current plan availability and feature names for Claude Design before relying on this section for a client-facing session, as Anthropic's products evolve quickly.

5. Claude Code as a Research Tool

Claude Code is built for technical work — but it is a genuinely useful research tool whenever the material is data-heavy, lives in a repository, or needs scripted, repeatable analysis rather than one-off synthesis.

Step-by-step: research with Claude Code

Step	What to do
1	Create or open a project folder for the research task.

Step	What to do
2	Open Claude Code in the terminal, desktop app, or your IDE, pointed at that folder.
3	Optionally run an initialisation step so Claude reads the folder contents and creates a working notes file it can reuse across the session.
4	Describe the analysis you need — for example, processing a batch of exported survey responses, cross-referencing multiple data files, or scripting a repeatable check against a dataset.
5	Review the code and output together; ask Claude to explain its method so you can verify the logic before trusting the result.
6	Save the working notes file so the same analysis can be re-run or extended for a future engagement.

A Claude Code research workflow

INGEST Load the raw data or documents	SCRIPT Write the repeatable analysis	VERIFY Check the method and output together	DOCUMENT Save notes so it can be re-run
---	--	---	---

Where Code fits

- Processing exported assessment or survey data from a workshop into summary statistics.
- Building a repeatable script to check curriculum documents against a structured standards list.
- Any research task that needs to be re-run on updated data without redoing the analysis from scratch.

6. Applying the Part Two Methodology in Claude

Every step in Part Two — verifying sources, mapping contradictions, synthesising with confidence language, building a brief, auditing — runs the same way in Claude: ask for it directly, in plain language, in whichever surface you're using. Claude does not need different phrasing than Part Two's prompts; the only thing that changes are where you're working.

EXAMPLE PROMPT

"Run the SIFT verification sequence and the 8-point audit checklist from Part Two on the answer above. Return the results as two tables."

7. Worked Example: A Regulatory Research Task

This example shows several surfaces working together on a single research task — checking whether training material remains aligned with current FAIS Board Notice guidance.

Surface	Action
Chat — Web Search	Ask Claude to search for any recent regulatory communication on the relevant FAIS Board Notices and summarise what has changed.
Chat — Research	For a broader scan across several regulators or standards bodies at once, run a Research task instead of a single search.
Cowork	Point Cowork at a folder containing the current workshop materials and the regulatory summary and ask for a gap-analysis briefing document.
Design	Once the briefing is reviewed and approved, turn the key findings into a one-page visual summary for internal sign-off or client circulation.
Code	If the review involves checking many documents at scale (e.g. an entire course library), script a repeatable check that flags documents referencing outdated clauses.

COMPLIANCE NOTE

Every regulatory citation, client name, and figure produced through this workflow should be verified against the primary source before it is used in a compliance-facing deliverable. Treat AI-drafted output as a strong first draft for your review — not a final answer.

8. Best Practices for Research Prompts in Claude

- State the audience and purpose up front (e.g. "for a facilitator briefing" vs "for a client-facing deck").
- Specify the format you want back — Word document, comparison table, bullet summary, slide outline.
- Ask for sources whenever the research involves current facts, regulations, or statistics.
- Break large research tasks into a folder of source material rather than pasting long text into chat.
- Review and verify before use — treat all output as a draft for your professional review, in line with your organisation's standards.
- Reuse what works — save prompts or Project instructions that produce a good result, so future research tasks start from a consistent baseline.

Official resources

- Anthropic Help Center — Web Search, Extended Thinking & Research: support.claude.com/en/articles/11095361
- Claude Docs: docs.claude.com

KEEPING THIS GUIDE CURRENT

This part is a working reference. As Claude's interfaces evolve, revisit the official Anthropic documentation (docs.claude.com and support.claude.com) to confirm menu names and available features before training clients on this material.

PART FOUR — CHATGPT IN PRACTICE

How to run the Part Two methodology inside ChatGPT — Search, Deep Research, and file uploads.

Updated for the current ChatGPT research workflow — August 2026.

LEARNING OUTCOME

This part covers what is specific to ChatGPT: choosing between Normal chat, Search, Deep Research and uploaded files, and the mechanics of each. For the research method itself — defining a question, verifying sources, resolving contradictions, synthesising, building a brief, auditing — see Part Two, which applies here without any changes.

1. Choose the Right Research Mode

A useful rule is Search for quick/current answers; use Deep Research for complex, multi-step investigations. OpenAI's current guidance distinguishes ChatGPT Search from Deep Research: Search retrieves timely information and links, while Deep Research plans, researches and synthesises a documented report from selected sources.

Decision diagram

- Is the answer likely to change? → Yes → Search
- Does it require many sources / several steps? → Yes → Deep Research
- Is the evidence in your files? → Yes → Upload / analyse files.

IMPORTANT

ChatGPT is a research assistant, not the final authority. For important academic, legal, financial, medical, or business claims, verify the underlying source before relying on the answer. Use ChatGPT to accelerate research — not to replace source evaluation; a fluent answer can still contain an error, an outdated claim, or a weak source.

2. Search the Web

ChatGPT Search can retrieve current information from the web and provide links/citations to sources. In the current ChatGPT workflow, Search can be selected from the tool's/source controls, and ChatGPT may also search automatically when a question benefits from current web information.

Step-by-step

- 1. Open ChatGPT.
- 2. Enter a focused question rather than a generic topic.
- 3. Use Search when current or source-specific information is required.
- 4. Inspect the citations/sources returned.

- 5. Open important sources and read the underlying evidence.
- 6. Ask follow-up questions to compare, challenge, or narrow the evidence.

EXAMPLE PROMPT

“Search the web for recent evidence on [TOPIC]. Prioritise primary sources, government/regulator publications, peer-reviewed research, and official organisational reports. For each important claim, show the source and explain exactly what it supports. Flag sources older than [X] years unless they are foundational.”

3. Use Deep Research for Complex Questions

Deep Research is designed for complex online tasks. OpenAI describes a workflow in which you define the outcome, select permitted sources, review a proposed research plan, monitor progress, and receive a structured report with citations or source links.

Recommended workflow

1. DEFINE State the decision or report you need	2. CONSTRAIN Specify sources, geography, dates	3. REVIEW PLAN Check sub-questions before running	4. MONITOR Interrupt or refine if scope drifts
---	--	---	--

Deep Research quality controls

- Specify which websites or source categories should be prioritised.
- Tell it what to exclude — for example, promotional pages or sources outside the date range.
- Ask for contradictions and evidence gaps, not only supporting evidence.
- Review the plan before the research proceeds.

EXAMPLE PROMPT

“Conduct deep research on [TOPIC]. The goal is to help me decide [DECISION]. Research [DATE RANGE] and [GEOGRAPHY]. Prioritise [SOURCE TYPES / WEBSITES]. Compare competing viewpoints, identify evidence quality, quantify findings where possible, and clearly distinguish established facts from expert interpretation and uncertainty. Include a concise executive summary, methodology, findings, contradictions, limitations, research gaps, and references.”

4. Work with Uploaded Documents

Research is often based on your own evidence: policy documents, annual reports, academic papers, survey data, meeting notes, or spreadsheets. Upload the relevant files and ask ChatGPT to analyse them as a defined evidence set.

Document analysis workflow.

UPLOAD Add source files	EXTRACT Find relevant evidence	COMPARE Identify patterns & conflicts	SYNTHESISE Create findings
-----------------------------------	--	---	--------------------------------------

EXAMPLE PROMPT

"Analyse the uploaded documents as my primary evidence set. First create a source inventory with title, date, organisation/author, and purpose. Then identify the passages or data relevant to [QUESTION]. Compare the documents, flag contradictions, and distinguish what is explicitly stated from what is inferred. Do not introduce outside facts unless I explicitly ask for external research."

For spreadsheets or datasets, ask for:

- Variable definitions and missing values.
- Descriptive statistics and notable patterns.
- Outliers and possible data-quality problems.
- Alternative explanations.
- Charts/tables that directly answer the research question.

5. Applying the Part Two Methodology in ChatGPT

Every step in Part Two — verifying sources, mapping contradictions, synthesising with confidence language, building a brief, auditing — runs the same way in ChatGPT: ask for it directly, in plain language, in Search, Deep Research, or a file-analysis chat. The prompts in Part Two need no ChatGPT-specific rewording.

EXAMPLE PROMPT

"Run the SIFT verification sequence and the 8-point audit checklist from Part Two on the answer above. Return the results as two tables."

6. Worked Example: A Competitor Landscape

This example shows Search, Deep Research and file uploads working together on a single task — building a competitor landscape for a new AI-training offering before pricing and positioning it.

Mode	Action
Search	Run quick lookups on each named competitor — recent launches, pricing pages, public course catalogues — to build an initial list.
Deep Research	Once the list is confirmed, run a Deep Research task: "Research the AI-skills-training market for financial services in South Africa, 2024–2026, covering [COMPETITOR LIST]. Compare positioning, pricing model, and NQF/SETA accreditation status."

Mode	Action
Uploaded files	Upload your own past proposals and pricing sheets and ask ChatGPT to compare your positioning against the Deep Research findings and flag any gap.
Audit (Part Two)	Run the Part Two audit checklist on the final comparison before it goes into a client-facing deck.

COMPLIANCE NOTE

Competitor claims, pricing, and accreditation status change quickly. Verify every figure in a competitor landscape against the competitor's own current public material before using it in a proposal or client-facing deck.

7. Official Resources

- OpenAI Help Center — ChatGPT Search: help.openai.com/en/articles/9237897-chatgpt-search.
- OpenAI Help Center — Deep Research: help.openai.com/en/articles/10500283-deep-research.
- OpenAI Academy — Research with ChatGPT: openai.com/academy/search-and-deep-research

KEEPING THIS GUIDE CURRENT

This part is a working reference. As ChatGPT's interfaces evolve, revisit OpenAI's official documentation to confirm menu names and available features before training clients on this material.

PART FIVE — GEMINI & GEMINI NOTEBOOK IN PRACTICE

How to run the Part Two methodology inside Gemini and Gemini Notebook.

NAME CHANGE: NOTEBOOKLM IS NOW GEMINI NOTEBOOK

In July 2026, Google renamed NotebookLM to Gemini Notebook as part of closer integration with the wider Gemini ecosystem. The core product, your existing notebooks, and the notebook.google.com address remains unchanged — only the name and logo changed. The rebrand also introduced code execution as a new capability, and the Studio panel now additionally offers Flashcards, Infographics, and Slide Decks. This guide uses "Gemini Notebook" throughout; older material, search results and menus may still say "NotebookLM" for some time.

This part covers what is specific to Gemini and Gemini Notebook: the two-engine architecture, Deep Research, source grounding, and the Studio tools. For the research method itself — defining a question, verifying sources, resolving contradictions, synthesising, building a brief, auditing — see Part Two, which applies here without any changes.

Gemini vs. Gemini Notebook — dual-engine architecture

Gemini (Web-Connected LLM)	Gemini Notebook (Grounded AI)
Role: Discovery, Search & Brainstorming Strengths: Live web search, broad topic exploration, code execution, multi-modal input. Best for: Finding initial papers, trends, generating hypotheses, draft outlines.	Role: Synthesis, Analysis & Verification Strengths: Strongly grounded in uploaded sources, minimal web bias, strict inline citations — still requires verification for high-stakes claims. Best for: Reading dense PDFs, cross-analysing documents, building audio/video overviews.

Decision diagram

- Is the answer likely to change, or does it need current web information? → Yes → Gemini
- Does it require many sources and several steps of independent research? → Yes → Gemini Deep Research
- Is the evidence already in your own files? → Yes → Gemini Notebook
- Do you need an audio/video explainer, mind map, or written report generated from your sources? → Yes → Gemini Notebook's Studio tools

(See Part One's Executive Decision Matrix for the full mode-by-mode comparison table.)

1. Literature Discovery Using Gemini

Start in Gemini to explore the breadth of your research domain. Use structured prompts to gather specific, high-value source documents.

RECOMMENDED PROMPT FOR INITIAL MAPPING

“Act as a senior academic researcher in [Field/Topic]. Identify the top 5 seminal papers and top 3 recent review articles (2023–present) regarding [Specific Query]. Provide key findings, methodologies used, and direct URLs or DOI references where available.”

2. Using Gemini Deep Research for Complex Topics

Deep Research is Gemini's agent for complex, multi-step investigations: it plans a research approach, searches and reads a large number of sources, and returns a structured, cited report — well suited to broad landscape questions before you narrow into Gemini Notebook.

1. DEFINE State the decision or report you need	2. REVIEW PLAN Check the proposed research plan	3. MONITOR Let it research, or refine the plan	4. REPORT Receive a structured, cited report
---	---	--	--

Deep Research quality controls

- Specify which websites or source categories should be prioritised.
- Tell it what to exclude — for example, promotional pages or sources outside the date range.
- Ask for contradictions and evidence gaps, not only supporting evidence.
- Review the proposed plan before the research runs and refine it if it is too broad or too narrow.

EXAMPLE PROMPT

“Conduct deep research on [TOPIC]. The goal is to help me decide [DECISION]. Research [DATE RANGE] and [GEOGRAPHY]. Prioritise [SOURCE TYPES / WEBSITES]. Compare competing viewpoints, quantify findings where possible, and clearly distinguish established facts from expert interpretation and uncertainty. Include an executive summary, methodology, findings, contradictions, limitations, research gaps, and references.”

3. Curating and Uploading Sources

Collect the relevant assets discovered through Gemini. Gemini Notebook supports a variety of file formats and inputs:

- PDF Documents: Academic papers, whitepapers, financial reports (up to 500,000 words per source [verify against current documentation]).
 - Google Docs & Slides: Existing meeting notes, research drafts, and presentation decks.
 - Web URLs: Articles, documentation, and web sources.
 - Text Files & Copied Text: Transcripts, interview notes, or quick excerpts.
1. Navigate to notebook.google.com and click New Notebook.
 2. Name your notebook after your project (e.g., ‘Quantum Computing Hardware Comparative Analysis’).

- 3. Drag and drop your curated sources into the source box.
- 4. Wait for indexing to complete. Gemini Notebook will build an automatic summary for each source.

KEY FEATURE: SOURCE GROUNDING

Unlike standard chatbots, Gemini Notebook is designed to answer only using the information present in your uploaded sources, and most claims carry a clickable inline numeral that highlights the source passage. Treat this as strong grounding, not a guarantee — spot-check citations and use this as your fastest way to run Part Two's SIFT "trace" step.

4. Advanced Gemini Notebook Techniques

Audio & Video Overviews

Gemini Notebook can transform your uploaded research papers into a dynamic two-person AI podcast discussion (Audio Overview) or a narrated visual walkthrough (Video Overview) — ideal for quickly absorbing complex papers or opening a workshop with a summary of the source material.

- Go to the Studio panel on the right.
- Click Generate under 'Audio Overview' or 'Video Overview.'
- Use the Customize option to instruct hosts to focus on specific aspects (e.g., 'Focus specifically on methodology and limitations').

Mind Maps

Gemini Notebook can generate a visual mind map of the concepts and relationships across your source set. Use it early, before drafting, to check you have not missed a branch of the topic — and again at the end, to sanity-check the structure of your final report.

Generating structured assets

Asset Type	Primary Use Case	How to Trigger
Study Guide	Extracts short-answer questions, key terminology, and essay prompts.	Select 'Study Guide' in the Studio panel.
Briefing Doc	Creates an executive summary categorised by key themes.	Select 'Briefing Doc' in the Studio panel.
FAQ	Anticipates common questions and provides grounded answers.	Select 'FAQ' format from the template list.
Timeline	Extracts chronological events or experimental milestones.	Select 'Timeline' for historical/process papers.
Report	Produces a longer, structured written report drawing on the full source set — the fastest route to Part Two's research-brief structure.	Select 'Reports' in the Studio panel.

NOTE

Studio tools, code execution, and export formats are being added quickly following the Gemini Notebook rebrand. Confirm current feature names and plan availability against your own screen before training clients on this material.

5. Applying the Part Two Methodology in Gemini Notebook

Every step in Part Two — verifying sources, mapping contradictions, synthesising with confidence language, building a brief, auditing — runs the same way here: ask for it directly, and let Gemini Notebook's grounding do the heavy lifting on citation-tracing. The prompts in Part Two need no rewording for Gemini or Gemini Notebook.

EXAMPLE PROMPT

“Run the SIFT verification sequence and the 8-point audit checklist from Part Two on this notebook's synthesis above, quoting the exact source passage for every claim. Return the results as two tables.”

6. Worked Example: Academic Literature Synthesis

This example shows Gemini and Gemini Notebook working together to turn a batch of research papers into a workshop-ready summary.

Stage	Action
Gemini	Use the Recommended Prompt for Initial Mapping (Section 1) to identify the 5–8 most relevant papers on the topic.
Gemini Notebook	Create a new notebook, upload the selected PDFs, and wait for indexing and the automatic per-source summary.
Studio — Mind Map	Generate a Mind Map first; to confirm the source set covers every branch of the topic before drafting.
Studio — Audio/Video Overview	Generate an Audio or Video Overview to brief a co-facilitator or client quickly, ahead of the written brief.
Part Two, Section 7	Ask Gemini Notebook to produce the final research brief, in the 9-part structure from Part Two, grounded only in the uploaded papers.

VERIFICATION CHECKLIST FOR FINAL REPORTS

1. Check inline citations in Gemini Notebook to verify the original source context.
2. Use Gemini's Double-Check Response button to verify any web facts.
3. Export core findings into your document editor and finalise your report structure.

7. Official Resources

- Gemini Apps Help — Deep Research: support.google.com/gemini/answer/15719111
- Gemini Notebook Help: support.google.com/gemininotebook

KEEPING THIS GUIDE CURRENT

This part is a working reference, and Google renamed this product mid-2026 — expect more change. Revisit the official Gemini and Gemini Notebook help centres to confirm menu names and available features before training clients on this material.

GOLDEN RULE

Move from question → evidence → verification → synthesis → decision. The quality of AI-assisted research depends less on one perfect prompt and more on running this disciplined, iterative process — in every tool.

PART SIX — PERPLEXITY IN PRACTICE

How to run the Part Two methodology inside Perplexity's search, Pro Search, Deep Research and Projects.

This part covers what is specific to Perplexity: its source-first, citation-beside-claim style, Deep Research, Projects, and how to use it well for source-backed comparisons. For the research method itself — defining a question, verifying sources, resolving contradictions, synthesising, building a brief, auditing — see Part Two, which applies here without any changes.

The working surfaces at a glance

- Perplexity Search — the default answer engine. A query returns a synthesised answer with inline citations to the sources used. Best for quick, current, well-sourced lookups.
- Perplexity Pro Search — an advanced research mode that goes beyond a single search pass, drawing on more sources and offering model selection and code-interpretation for more complex queries.
- Perplexity Deep Research (Advanced Deep Research) — a multi-step research agent that searches, cross-references, and synthesises across many sources to produce a longer, structured, cited report.
- Perplexity Projects (formerly Spaces) — a persistent, shareable workspace that keeps a project's search conversations, files, custom instructions, and context together, similar in purpose to Claude Projects.

Which surface should you use?

If your research need is...	Use...
A quick, current, citation-backed fact	Perplexity Search
A more complex query needing broader sourcing, model choice, or code interpretation	Perplexity Pro Search
A full multi-step investigation with a long, structured, cited report	Perplexity Deep Research
The same research context, files or team reused across an ongoing project	Perplexity Projects

1. Perplexity's Source-First Approach

Perplexity places a citation beside each factual claim rather than at the end of a long answer, which makes it easy to check individual statements as you read. Treat this as a strong starting point, not a finished verification: Perplexity can search and synthesise results, but it should not be treated as the original source. Open the cited pages before relying on any consequential claim.

VERIFY THE CITATIONS

A citation beside a sentence confirms that Perplexity found a source saying something similar — it does not confirm the claim is accurate, current, or stated the way the summary implies. Open every source behind a claim that affects money, safety, legal compliance, security, or a strategic decision before you rely on it.

2. A Source Hierarchy for Perplexity Research

Part Two's 4-tier source hierarchy applies here without change. Perplexity's own mix of results also makes two additional categories worth calling out explicitly — vendor marketing and forum/social content — since Perplexity readily surfaces both alongside stronger sources.

Tier	What it covers
Primary sources	Legislation, regulatory guidance, research papers, datasets, company filings, and official documentation.
Independent secondary sources	Reputable journalism, systematic reviews, and established industry analysis.
Expert commentary	Useful for interpretation, but not definitive evidence on its own.
Vendor marketing	Useful for product claims — but these still need independent verification before use.
Forums & social media	Useful for discovering problems or edge cases, not for establishing general facts.

SOURCE-TIERING PROMPT

“For every major claim, identify whether its evidence is primary, independent secondary, expert commentary, vendor-produced, or anecdotal. Flag conclusions supported only by weak evidence.”

3. Applying the Part Two Methodology in Perplexity

Every step in Part Two — planning, source-backed search, evidence extraction, contradiction-checking, synthesis, brief-building, auditing — runs the same way in Perplexity: ask for it directly and explicitly tell Perplexity to search the current web when the question touches products, prices, laws, statistics, leadership, APIs, or recent developments.

RESEARCH PLAN PROMPT

“Create a research plan for this question: [research question]. Identify: the key sub questions; the evidence needed for each one; the best source types; important assumptions; and criteria for evaluating the findings. Do not draw conclusions yet.”

SOURCE-BACKED RESEARCH PROMPT

“Research the following question using current web sources: [question]. Requirements: prioritise primary and official sources; cross-check important claims; cite each factual claim beside the relevant sentence; separate verified facts, expert opinions, and estimates; note disagreements or missing evidence; do not infer conclusions that the sources do not support.”

EVIDENCE-EXTRACTION PROMPT

“Create an evidence table with research sub question, finding, supporting source, publication date, source type, evidence strength, limitations, and confidence level. Do not combine multiple unsupported claims into one row.”

4. Comparing AI Tools and Products in Perplexity

Perplexity's citation-beside-claim style makes it particularly well suited to structured, evidence-checked comparisons of tools or products for an AI toolkit — provided the evaluation criteria are set before the comparison runs, and product claims are verified against each provider's own current documentation.

TOOL-COMPARISON PROMPT

“Compare [tool A], [tool B], and [tool C] using primary function; research and web-search capability; citation quality; context limits; API and integration options; data retention and privacy; supported models; pricing; strengths; limitations; and best use case. Verify current details from official documentation. State the evaluation criteria before recommending anything.”

A practical research toolkit stack

Stage	Includes
Discovery	Web search, news databases, and academic search.
Collection	Reference managers, bookmarks, document stores, and APIs.
Analysis	An AI assistant, spreadsheets, notebooks, or qualitative-analysis tools.
Verification	Original documents, source comparisons, and fact-checking.
Production	Reports, briefs, presentations, or knowledge-base entries.

5. Interrogating and Verifying Findings

Do not accept the first synthesis. Ask Perplexity to challenge its own findings, then manually verify anything that affects money, safety, legal compliance, security, or a strategic decision.

SKEPTICAL-REVIEWER PROMPT

“Critique the findings as a sceptical reviewer. Identify: unsupported assumptions; missing stakeholders or evidence; selection bias; outdated information; conflicts of interest; correlation presented as causation; plausible alternative explanations; and claims that require manual verification.”

OPPOSING-CASE PROMPT

“Construct the strongest evidence-based case against the current conclusion. Explain which additional evidence would resolve the disagreement.”

CITATION-AUDIT PROMPT

“Audit the following draft claim by claim. For each factual claim, label it: supported, partially supported, unsupported, contradicted, or requires current verification. Explain the label and identify the best available source. Do not silently rewrite unsupported claims.”

Manual verification steps for high-impact claims

- Open the cited source.
- Confirm that the source says what the summary claims.
- Check the publication and update dates.
- Examine the methodology, sample size, and jurisdiction.
- Look for newer or contradictory evidence.
- Record the exact page, section, or table that supports the claim.

6. Worked Example: Evaluating AI Agents for Client Onboarding

This example shows the full plan → search → extract → compare → challenge → verify → synthesise workflow running in Perplexity, for a specific and decision-oriented question: how small accounting firms can use AI agents to automate client onboarding.

Stage	Action
Define & plan	State the objective precisely: tools available in 2026, implementation costs, security risks, and three real examples — not just "research AI agents." Ask Perplexity for a research plan and remove irrelevant branches.
Source-backed search	Run the source-backed research prompt, requiring current web sources, prioritised primary/official sources, and citations beside each claim.
Extract evidence	Build the evidence table: sub question, finding, source, date, source type, evidence strength, limitations, confidence.
Compare	Where the toolkit is also comparing specific vendors' agent products, run the tool-comparison prompt with pricing, security, and integration as explicit criteria.
Challenge	Run the sceptical-reviewer prompt, then the opposing-case prompt, to surface unsupported assumptions and unresolved disagreement.
Verify	Manually open the sources behind the cost, security, and compliance claims — the highest-impact claims in this question — before relying on them.
Synthesise	Produce the final research brief for the intended audience, with confidence levels assigned to each major finding.

7. Perplexity-Specific Prompt Library

Copy, adapt, and reuse these prompts — save them in your toolkit alongside the ones from Part Two.

FINAL RESEARCH BRIEF PROMPT

“Write a research brief for [audience] based only on the verified evidence. Structure: executive summary; research question and scope; method; key findings; comparison or analysis; risks and limitations; recommendations; open questions. Distinguish facts from interpretations. Place citations immediately after the claims they support. Assign confidence levels to major conclusions.”

MASTER PROMPT — COMPLETE RESEARCH WORKFLOW

“Act as a rigorous research assistant supporting an AI toolkit. Research objective: [insert objective]. Audience and decision: [insert who will use the research and what decision it supports]. Scope: [insert geography, industry, products, and time period]. Research process: break the objective into answerable sub questions; search for current evidence; prioritise official and primary sources; cross-check consequential claims using independent sources; build an evidence table; distinguish facts, estimates, opinions, and inferences; identify conflicting evidence and information gaps; challenge the preliminary conclusions; assign confidence levels; produce a concise final report. Output requirements: cite claims next to the supporting evidence; include publication dates and source types; explain material limitations; do not invent missing facts or citations; flag information requiring manual verification; finish with actionable recommendations tied directly to the evidence.”

THE CORE WORKFLOW

Plan → search → extract → compare → challenge → verify → synthesise. The most important safeguard is keeping Perplexity's synthesis separate from the underlying evidence — read the source, not just the summary of it.

8. Official Resources

- Perplexity Help Center: perplexity.ai/help-center/en
- What is Pro Search?: perplexity.ai/help-center/en/articles/10352903-what-is-pro-search
- What's New in Advanced Deep Research: perplexity.ai/help-center/en/articles/13600190-what-s-new-in-advanced-deep-research
- What are Projects (formerly Spaces)? perplexity.ai/help-center/en/articles/10352961-what-are-spaces?

KEEPING THIS GUIDE CURRENT

This part is a working reference. Perplexity has already renamed Spaces to Projects and refreshed Deep Research once — expect more change. Revisit the official Perplexity Help Center to confirm current feature names and plan availability before training clients on this material.

PART SEVEN — RESPONSIBLE & ETHICAL USE IN RESEARCH

What this toolkit does not cover on its own — data protection, sign-off, ethics clearance, bias, reproducibility, and disclosure — and how to close those gaps.

This part follows the order a researcher actually needs it in first work out whether material may be entered into an AI tool at all, then pause at a sign-off checkpoint before starting, then run the research responsibly, and finally disclose the AI assistance transparently. POPIA requires lawful, purpose-limited, safeguarded processing of personal information; responsible AI practice also requires human oversight and accountability throughout.

1. This Guide Is a Workflow, Not an Ethics Clearance

Parts Two to Six teach a disciplined research method: define the question, verify sources, resolve contradictions, synthesise with an honest confidence level, and audit before you rely on the output. That method is compatible with good research practice — but following it is not, by itself, proof of ethics approval, POPIA compliance, acceptable authorship practice, or institutional academic-integrity compliance. Those sit alongside this toolkit, not inside it. The human researcher retains responsibility throughout.

CONFIRM SEPARATELY

Before relying on AI-assisted research for a thesis, dissertation, journal submission, client deliverable, or funded project, confirm current requirements with the relevant authority — your institution's ethics committee, Information Officer, supervisor, client, or any publisher/funder AI-use policy. This part is orientation, not a substitute for that sign-off.

2. POPIA and Data Protection

Before using any AI system, classify every document, dataset, excerpt, prompt, or attachment you intend to enter. Do not upload personal, special personal, children's, confidential, restricted, or contractually protected information into a public or unapproved AI tool.

Where personal information is involved, researchers must first confirm the lawful basis for processing, the defined research purpose, appropriate security safeguards, institutional approval, and any applicable ethics or consent requirements. De-identification can reduce risk but does not necessarily mean that information falls outside POPIA, particularly where re-identification remains possible.

When in doubt, pause and consult the supervisor, principal investigator, Research Ethics Committee, Information Officer, data steward, Research Office, or authorised institutional support function.

POPIA data-classification table

Classification	Typical material	AI-tool rule	Conditions before use
Green: Public / non-confidential	Published articles, legislation, public reports, public web content.	May use in approved or standard AI tools.	Check original sources; observe copyright, licensing, and citation requirements.
Amber: Internal / non-personal	Unpublished drafts, de-identified aggregate data, internal strategy, or policy material.	Use only with approval.	Check institutional policy, contracts, platform terms, access controls, and data-retention settings. Remove unnecessary confidential information.
Red: Personal information	Names, contact details, student/staff records, identifiable transcripts, email correspondence.	Do not upload to public or unapproved tools.	Use only in an approved environment after lawful basis, safeguards, permissions, and institutional requirements have been confirmed.
Red: Special personal information	Health, race or ethnicity, religion, politics, trade-union membership, biometric, sex-life, or criminal-information data.	Do not upload unless explicitly authorised.	Require approved secure infrastructure, documented lawful basis, access controls, appropriate ethics clearance, and institutional guidance.
Red: Children's information	Learner records, data from minors, school surveys, images, recordings, identifiable classroom data.	Do not upload unless expressly authorised.	Confirm authority, consent/permission requirements, data minimisation, secure approved systems, and specialist institutional guidance.
Red: Restricted / contractually protected data	Bank data, partner datasets, NDA material, proprietary code, funder-restricted material.	Do not upload without written approval.	Check the applicable contract, data-sharing agreement, confidentiality terms, intellectual-property requirements, and platform approval.

POPIA's lawful-processing conditions include accountability, processing limitation, purpose specification, and security safeguards. Rule of thumb: de-identification reduces risk but does not automatically remove POPIA obligations, particularly if re-identification remains possible in your context.

3. Pre-Use Sign-Off Checklist

Complete this checklist before starting an AI-assisted research task. It is mandatory where Amber or Red material is involved (see Section 2) and strongly recommended for all substantive AI-supported research. If any item cannot be confirmed, pause and seek guidance from the supervisor, principal investigator, Research Ethics Committee, Information Officer, data steward, or authorised institutional support function before proceeding.

Checkpoint	Researcher confirmation / record
1. Research purpose	I can state the research task, output, audience, and reason for AI assistance. Task: _____
2. Permitted use	I have checked university, faculty, supervisor, funder, publisher, partner, and assessment rules. Policy checked: Yes / No
3. Data classification	I have classified all intended inputs using the Green / Amber / Red table above. Classification: _____
4. Personal-data screen	No personal, special personal, children's, confidential, restricted, or contractual data will be entered without documented approval. Cleared / approval attached: _____
5. Ethics status	I have confirmed whether ethics approval, an amendment, consent, or data-owner permission is needed. Ethics approval no.: _____
6. Approved platform	The selected tool is institutionally approved, or permission to use it has been documented. Tool/platform: _____
7. Privacy and security	I have checked current provider settings, retention/training terms, access controls, and cross-border processing implications. Reviewed on: _____
8. Data minimisation	I will use only the minimum necessary data and remove identifiers and unnecessary confidential material. Confirmed: Yes / No
9. Research integrity	I will independently verify claims, citations, quotations, calculations, and interpretations against original evidence. Confirmed: Yes / No
10. Bias and context	I will assess whether local, South African, African, indigenous, and relevant language perspectives are represented. Confirmed: Yes / No
11. Disclosure	I have identified whether and where material AI assistance must be declared. Location: _____

Checkpoint	Researcher confirmation / record
12. Audit trail	I will retain prompts, tool/model details, dates, outputs relied on, researcher edits, and verification notes. Storage location: _____
13. Accountability	I accept responsibility for the legality, ethics, accuracy, and final use of the research output. Researcher name/signature: _____
14. Escalation	If a risk or uncertainty arises, I will stop and seek guidance. Contact consulted: _____

4. Human-Subjects & Ethics-Committee Research

The workflow in Parts Two to Six is built for literature and document research. It does not, on its own, cover research involving people — interviews, surveys, health records, student records, or work with vulnerable groups.

- AI use does not replace ethics-committee clearance.
- AI use does not replace informed consent processes.
- AI use does not replace a data-management plan.
- AI use does not replace permission for secondary use of existing data or records.

IF PEOPLE ARE INVOLVED

If your research involves human participants, their identifiable data, or secondary use of records about them, obtain ethics clearance and follow your institution's research-ethics policy before using AI on any part of that data or those materials — regardless of how the AI tool's own privacy terms read.

5. Bias, Representation & South African/African Relevance

The source hierarchy in Part Two protects against weak sources, but it does not by itself correct for bias embedded in the AI models themselves, or for uneven representation of African scholarship, local languages, and indigenous knowledge in what these tools return. UNESCO's guidance on AI ethics specifically highlights inclusion, cultural diversity, and the risk of AI widening existing inequalities — worth keeping in view for any South African research context.

A contextual bias check

- Whose evidence is missing? Ask explicitly for South African or African primary sources, not only international ones.
- Are findings imported from a Global North context assumed to transfer directly to South Africa? Ask the tool — and yourself — whether they do.

- Is local-language or indigenous-knowledge material represented, or is it systematically absent from what the tool returns?

EXAMPLE PROMPT

“Review the sources used so far on this topic. Identify where South African or African primary scholarship, data, or perspectives are missing. Flag any conclusion that assumes findings from another context transfer directly to South Africa and explain why it may or may not.”

6. Reproducibility: Keeping an Audit Trail

AI outputs can change between sessions, model versions, and providers. For any research that will be examined, published, or relied on by someone else, keep enough of a trail that a supervisor, examiner, ethics committee, or peer reviewer can see how the AI shaped the work.

Record	Why it matters
Tool, model/version, and date used	Answers can differ meaningfully across model versions.
Prompts used, verbatim	Lets someone else see exactly what was asked.
Search terms and retrieved sources	Supports independent verification of what was found.
Verification notes (Part Two, Sections 4 & 8)	Shows what was checked, and how.
Final human-edited version	Distinguishes the AI draft from your own judgement and wording.

Claude Projects, a saved ChatGPT conversation, or a Gemini Notebook all keep a version of this trail automatically — but only for as long as you keep the workspace. For anything going into a formal submission, export or copy the trail into your own records.

7. AI-Use Disclosure

Disclose material AI use honestly and proportionately in the location required by the university, funder, conference, assessment brief, or journal. Disclosure does not replace ethical approval, POPIA compliance, source verification, or human responsibility. AI tools must not be listed as authors because they cannot take responsibility for scholarly work.

DECLARATION OF GENERATIVE-AI USE

Generative AI was used as a research-assistance tool in this project.

Tool(s) used: *[Tool name, provider, model/version where available].*

Dates of use: *[Date range].*

Purpose and scope: *The tool was used for [for example: brainstorming search terms; preparing an outline; language editing; summarising researcher-supplied, approved sources; organising themes; identifying counterarguments; or drafting code].*

Data handling: *No identifiable, special personal, confidential, restricted, or unapproved participant data were entered into the tool. / The approved institutional platform and data-governance controls used were [details].*

Human oversight and verification: *The researcher reviewed and revised all AI-assisted material. All factual claims, references, quotations, calculations, and interpretations were independently checked against original sources.*

Accountability: *The researcher accepts full responsibility for the accuracy, integrity, ethical compliance, and final content of this work.*

Authorship: *No AI tool has been listed as an author.*

Research record: *Prompts, outputs relied upon, source-verification notes, and material revisions have been retained in accordance with applicable institutional or project record-retention requirements.*

MINIMAL-USE WORDING

For minimal use — such as grammar correction only — the "Purpose and scope" line can instead state: "The tool was used only for language and readability suggestions; it was not used to generate substantive claims, analyse research data, or produce citations."

8. Escalation Contacts & Institutional Support

Escalation contacts vary by institution — confirm and record the current names or roles for your context before you need them:

- **[Supervisor / Principal Investigator]** — day-to-day guidance on research design and AI use within the project.
- **[Research Ethics Committee]** — approval and guidance for human-participant research, sensitive data, or special personal information.
- **[Information Officer]** — POPIA compliance, lawful basis, data classification, and cross-border transfer questions.
- **[Research Office / Postgraduate Office]** — institutional AI-use policy, approved-platform lists, and funder or publisher requirements.

ESCALATION RULE

Stop and seek guidance if the task includes Red-classified material, human-participant data, uncertainty about consent or ethical approval, third-party confidential information, or a platform that is not institutionally approved. This reflects POPIA's safeguard obligations and UNESCO's expectation of human oversight, accountability, and appropriate risk controls.

THE ACCOUNTABILITY RULE

AI may assist your research. It may not be the unverified origin of evidence, the undisclosed writer of substantive scholarly claims, or a destination for identifiable or confidential data without approval. You — not the tool — remain responsible for the trustworthiness of your research.

A NOTE ON HOW THIS GUIDE WAS PRODUCED

Transparency on the AI-assisted process behind this document.

Consistent with the disclosure standard set out in Part Seven, this guide was produced through a multi-stage, human-supervised AI-assisted process, summarised here for transparency:

Stage	What happened
1. Tool-specific drafting	The workflow content specific to each platform (Claude, ChatGPT, Gemini, and Perplexity) was originally generated by that respective AI tool.
2. Consolidation & formatting	All tool-generated material was provided to Claude, which consolidated, restructured, and formatted it into this document.
3. Moderation	The consolidated draft was reviewed in Gemini Notebook (formerly NotebookLM) for moderation, and its suggested revisions were incorporated.
4. Final triage	The revised document was submitted to Perplexity for a final triage review, and its resulting comments were incorporated into this final version.

HUMAN OVERSIGHT

The following human researchers directed, reviewed, and approved each stage of this process and remain responsible for the accuracy, integrity, and final content of this guide, consistent with the accountability standard set out in Part Seven.

Dr Susheila Moodley (PhD)

Academic Director | FutureBanker

Professor of Practice | College of Business & Economics, University of Johannesburg

Jacqui Shawe

EDTP Practitioner

