



Founders AI

founders-ai.dev

PRACTICAL AI FIELD GUIDE

AI Research System

How to Find Better Answers Without Trusting AI Blindly

A practical research guide for founders, business owners, creators, consultants, and professionals who want AI to find, verify, and synthesize better answers before they act.

- Define the decision
- Build source ladders
- Verify claims
- Track evidence
- Write uncertainty
- Turn research into action

Most people do not need AI to make them more confident. They need AI to help them ask better questions, find stronger sources, compare evidence, and slow down before they act on a bad answer.

This guide gives you a practical research system you can use with ChatGPT, Claude, Perplexity, Gemini, Google, YouTube, Reddit, official documentation, industry reports, customer interviews, and your own notes.

The rule is simple:

AI can help you research. AI is not automatically the source of truth.

That distinction matters because AI can sound confident when it is wrong. Deep-research tools can produce useful, cited reports, but OpenAI's own help documentation frames deep research as a tool for multi-step online research that should be reviewed through its citations and source links ([OpenAI Help Center](#)). Researchers and product builders still warn that these systems can produce inaccurate claims, weak citations, or misplaced confidence ([Nature](#)).

This system is built for founders, business owners, creators, consultants, employees, and everyday people who want to use AI without getting fooled by it.

What this guide helps you do

Use this guide when you need to:

- Learn a new topic without drowning in random articles.
- Compare tools, vendors, agencies, markets, or strategies.
- Prepare for a sales call, client meeting, investor conversation, or internal decision.
- Research content ideas before filming or publishing.
- Validate a claim before repeating it.
- Turn messy information into a usable recommendation.
- Know when AI is useful and when human expertise is required.

Do not use this system as a replacement for expert review in legal, medical, tax, financial, safety, compliance, or high-stakes employment decisions.

The core idea

Bad AI research usually starts like this:

Tell me everything about [topic].

That prompt feels productive, but it usually creates a long answer that is hard to verify. It gives you paragraphs before it gives you judgment.

Good AI research starts like this:

I need to make a decision about [decision].
Help me define the research question, identify the source types I should trust most, list what evidence would change the recommendation, and build a research plan before answering.

That prompt changes the job. You are no longer asking AI to be the authority. You are asking AI to help you run a research process.

The AI Research System

Use this loop whenever the answer matters:

1. Define the decision.
2. Turn the decision into a research question.
3. Build a source ladder.
4. Ask AI for a research plan before the answer.
5. Collect claims, not just paragraphs.
6. Verify important claims with better sources.
7. Separate facts, interpretations, and recommendations.
8. Write uncertainty directly.
9. Turn the research into an action brief.
10. Save reusable notes for the next decision.

The point is not to make research slower. The point is to make it less sloppy.

Step 1: Start with the decision

Do not begin with a topic. Begin with the decision the research is supposed to improve.

Weak:

Research AI tools for small businesses.

Stronger:

I need to decide which AI tool stack a two-person service business should use for email, meeting notes, client follow-up, and content creation. The decision needs to balance cost, ease of use, reliability, and setup time.

The second version is better because it gives the research a job. It tells the AI what trade-offs matter.

Decision-first prompt

Copy this:

I need to make this decision:
[write the decision]

My context:

- Who this is for:
- Current situation:
- Constraints:
- Budget:
- Time available:
- Risk tolerance:
- What would make this successful:

Before answering, help me turn this into:

1. The main research question
2. 5 supporting questions
3. The source types we should trust most
4. The claims that must be verified before acting
5. The decision criteria we should use

Saveable tool: Decision frame

Field	Your answer
Decision	
Who it affects	
Why now	
Constraints	
Risk if wrong	
Evidence needed	
Final decision deadline	

Step 2: Build a source ladder

Not all sources deserve equal trust. A random thread, a vendor landing page, and official documentation can all be useful, but they should not carry the same weight.

Use a source ladder before you start researching.

Source tier	Examples	Best for	Watch out for
Primary source	Official docs, product pages, legal filings, government data, company announcements	Verifying what something is, costs, terms, limits, dates	Marketing spin, missing context
Direct evidence	Screenshots, demos, benchmarks, recorded examples, first-hand usage	Seeing whether something works in practice	Small samples, cherry-picking
Expert source	Researchers, experienced operators, technical specialists, respected analysts	Understanding trade-offs and interpretation	Bias, incentives, outdated expertise
Independent coverage	Reputable journalism, neutral comparisons, third-party reviews	Broader context and external validation	Surface-level summaries
Community evidence	Reddit, Discord, X, YouTube comments, forums, reviews	Pain points, complaints, real user language	Anecdotes, exaggeration, brigading
AI synthesis	ChatGPT, Claude, Perplexity, Gemini, deep research tools	Planning, summarizing, comparing, generating questions	Hallucinations, weak citations, false confidence

The University of Chicago’s SIFT method is useful here because it teaches you to stop, investigate the source, find better coverage, and trace claims back to the original context instead of relying on the first thing you see ([University of Chicago Library](#)).

Source ladder prompt

For this research question:
[question]

Build a source ladder. Rank the source types we should trust most and least. For each source type, explain:

- What it is useful for
- What it is bad at
- What bias it may have
- What kind of claim it can verify
- What kind of claim it cannot verify

Step 3: Ask for the research plan before the research

Deep-research tools are strongest when you give them a specific outcome, source preferences, and boundaries. OpenAI's deep research workflow asks the user to describe the outcome, review or modify the proposed research plan, and then review a structured report with citations and source links ([OpenAI Help Center](#)).

That is the habit to copy, even when you are not using a formal deep-research mode.

Research-plan prompt

You are helping me research this decision:
[decision]

Do not answer yet. First create the research plan.

Include:

1. The main research question
2. 5-8 supporting questions
3. The source ladder
4. The search terms or source types to use
5. What claims would need verification
6. What would count as strong evidence
7. What would count as weak evidence
8. What uncertainty may remain after the research

After I approve the plan, then complete the research.

This prevents the AI from sprinting into a polished answer before you know whether it is researching the right thing.

Step 4: Collect claims, not paragraphs

A polished paragraph can hide weak evidence. A claim table makes the research easier to check.

Use this format:

Claim	Source	Evidence type	Date	Confidence	Notes
		Primary / expert / community / AI synthesis		High / medium / low	

Claim-extraction prompt

Take the research below and extract the claims into a table.

For each claim, include:

- The exact claim
- The source
- The type of evidence
- The date of the source, if available
- Whether the claim is factual, interpretive, or opinion
- Confidence level: high, medium, or low
- What would be needed to verify it

Research:

[paste notes, links, transcript, article, or report]

This is where AI becomes extremely useful. It can turn messy research into a structured evidence table. You still have to verify the important parts.

Step 5: Verify anything consequential

You do not need to verify every sentence. You do need to verify anything that could change the decision.

Verify:

- Prices, fees, limits, terms, refund policies, and contract requirements.
- Product capabilities and restrictions.
- Legal, tax, compliance, safety, financial, or medical claims.
- Statistics you plan to publish.
- Claims that make one option look obviously better than another.
- Anything the AI says with confidence but no strong source.

NIST's Generative AI Profile warns that generative systems can produce confidently stated false content and can even produce citations or reasoning that appear to justify incorrect answers (NIST). That is why source checking is not optional when the answer matters.

Verification prompt

Review this research as a skeptical verifier.

Find:

1. Claims that are not supported strongly enough
2. Claims that rely on weak or biased sources
3. Claims that need a primary source
4. Claims that may be outdated
5. Claims that are opinion but written like fact
6. Claims that would be risky to act on without expert review

Return a table with:

- Claim
- Problem
- Why it matters
- Best next source to check
- Confidence after review

Step 6: Use lateral reading

Do not evaluate a source only by reading its own page. Open another tab. Search the author, organization, product, claim, and competing coverage.

This is called lateral reading. It is one of the most practical ways to avoid being fooled by a polished page.

Ask:

- Who created this?
- What do they sell?
- What incentive do they have?
- Are other credible sources saying the same thing?
- Is the claim being repeated from one original source?
- Is the original source being quoted fairly?
- Is the claim current?

Lateral-reading prompt

Help me evaluate this source:
[paste source or summary]

Use a lateral-reading approach.

Check:

1. Who created it
2. What incentives they may have
3. Whether better sources cover the same claim
4. Whether the claim traces back to an original source
5. What context may be missing
6. Whether this source should be trusted for this decision

Step 7: Separate facts from interpretation

A lot of bad AI research fails because it blends facts, assumptions, and recommendations together.

Use three separate buckets:

Bucket	Meaning	Example
Fact	A claim that can be checked	"This tool costs \$29 per month."
Interpretation	A judgment based on facts	"This is affordable for a solo founder."
Recommendation	The action to take	"Start with this tool for the first 30 days."

Separation prompt

Separate this research into three sections:

1. Facts that can be verified
2. Interpretations or judgments
3. Recommendations

For each recommendation, show which facts support it and what assumptions it depends on.

This makes the final answer cleaner and more honest.

Step 8: Write uncertainty directly

Good research does not pretend to know everything. It tells you what is known, what is likely, what is unclear, and what would change the answer.

Use this confidence scale:

Confidence	Meaning	Use when
High	Strong sources agree and the claim is current	Primary source or multiple credible sources support it
Medium	Evidence points one way, but context may be incomplete	Good sources exist, but there are trade-offs or gaps
Low	Claim is plausible but not proven	Source quality is weak, old, biased, or anecdotal
Do not act yet	The claim could create serious risk if wrong	Legal, medical, tax, compliance, safety, or major financial stakes

Uncertainty prompt

Rewrite this research with explicit uncertainty.

For each major conclusion, include:

- Confidence level
- Why that confidence level is appropriate
- What evidence supports it
- What evidence is missing
- What would change the recommendation
- Whether expert review is needed before acting

This is one of the easiest ways to make AI output more useful than a generic answer.

Step 9: Turn research into a decision brief

Research is not done when you have information. Research is done when you can make a better decision.

Use this one-page decision brief:

Decision:

Recommendation:

Why this is the recommendation:

Best evidence:

- 1.
- 2.
- 3.

Biggest risks:

- 1.
- 2.
- 3.

What is uncertain:

What would change the recommendation:

What to do next:

- 1.
- 2.
- 3.

Decision-brief prompt

Turn this research into a one-page decision brief.

Use this structure:

- Decision
- Recommendation
- Best evidence
- Alternatives considered
- Risks
- Uncertainty
- What would change the recommendation
- Next actions

Keep it practical. Do not overstate certainty. Flag anything that needs expert review.

Step 10: Save reusable research notes

If you research the same category again, do not restart from zero.

Save:

- The decision frame.
- The source ladder.
- The claim table.
- The best sources.
- The rejected sources.
- The final decision brief.
- What turned out to be wrong later.
- What you would check first next time.

This is how a founder, creator, consultant, or operator builds a personal research advantage. Your notes become a reusable judgment system.

The 30-minute AI research workflow

Use this when you need a practical answer fast.

Minute 0-5: Frame the decision

Write the decision, constraints, and what would make the research useful.

Prompt:

Help me turn this decision into a research question and source ladder.

Minute 5-15: Run research

Ask AI to gather sources, summarize competing viewpoints, and list claims.

Prompt:

Research this question using the source ladder. Return claims in a table, not just paragraphs.

Minute 15-25: Verify

Pick the 5 claims that matter most. Check them against primary sources or better coverage.

Prompt:

Identify the claims that must be verified before acting and tell me where to check them.

Minute 25-30: Decide

Turn the research into a decision brief.

Prompt:

Write the one-page decision brief with confidence levels and next actions.

The “do not trust this yet” checklist

Slow down if the AI answer:

- Gives a specific number with no source.
- Cites sources that do not actually support the claim.
- Treats a vendor claim as neutral fact.
- Uses old information for a fast-changing topic.
- Sounds certain about legal, medical, tax, safety, or financial advice.
- Ignores trade-offs.
- Recommends a tool without checking pricing, limits, or terms.
- Summarizes a community opinion as if everyone agrees.
- Gives a conclusion before showing the evidence.
- Cannot explain what would change the recommendation.

If two or more of these show up, run a verification pass.

The best everyday use cases

Buying software

Use AI to compare pricing, features, limits, onboarding time, customer complaints, and alternatives. Verify prices and terms on the official product site before acting.

Learning a new skill

Use AI to build a learning path, explain concepts, find beginner mistakes, and create practice drills. Verify technical claims with official documentation or current expert sources.

Researching a market

Use AI to map competitors, customer pains, pricing patterns, and content angles. Treat the result as a starting point until you verify with real customers and current sources.

Preparing content

Use AI to find questions people ask, summarize competing viewpoints, and identify gaps. Do not publish statistics or strong claims without checking the original source.

Making a business decision

Use AI to structure options, risks, assumptions, and next steps. Do not let AI make the decision for you. Use it to improve your judgment.

What AI research is good at

AI is useful for:

- Turning vague questions into better research questions.
- Creating search terms and source plans.
- Summarizing long documents.
- Comparing multiple sources.
- Extracting claims from messy notes.
- Finding gaps and contradictions.
- Creating checklists, briefs, tables, and next actions.
- Helping you explain a topic in plain English.

What AI research is bad at

AI is risky when:

- The source quality is weak.
- The topic changes quickly.
- The stakes are high.
- The answer requires licensed expertise.
- The model is guessing beyond the evidence.
- The citations are not checked.
- You want certainty where only probability exists.

AI should speed up research. It should not remove your responsibility to think.

One-page version

Copy this into your notes.

AI Research System

1. Define the decision.
What am I trying to decide?
2. Write the research question.
What question would improve the decision?
3. Build the source ladder.
Which sources should be trusted most?
4. Ask for the research plan first.
Do not let AI answer before it plans.
5. Collect claims.
Turn paragraphs into claim/source/evidence tables.
6. Verify consequential claims.
Check numbers, prices, dates, risks, and strong recommendations.
7. Separate facts, interpretation, and recommendations.
Do not let opinion hide inside "facts."
8. Write uncertainty.
High, medium, low, or do not act yet.
9. Create a decision brief.
Recommendation, evidence, risks, uncertainty, next steps.
10. Save the system.
Keep the sources, notes, and what you would check next time.

Final reminder

The goal is not to become a professional researcher. The goal is to stop being a passive consumer of confident answers.

Use AI to move faster. Use sources to stay grounded. Use uncertainty to stay honest. Use your judgment to decide.