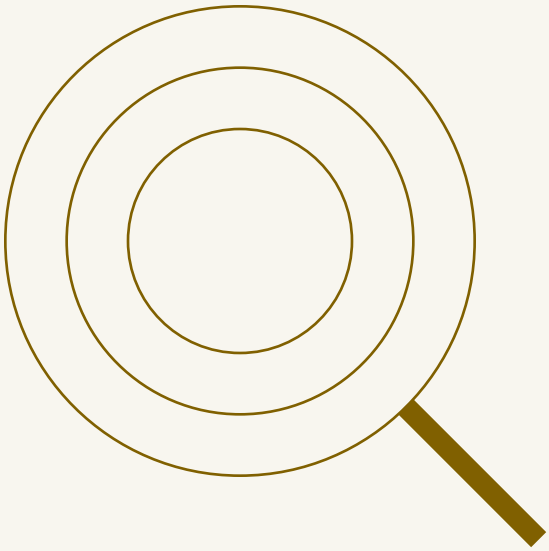
EVIDENCE FOR EVERYDAY DECISIONS

Named. Cited. Chosen?

Read AI visibility without mistaking attention for a recommendation.

Named.
Cited.
Chosen?



Separate names, citations, recommendations and enquiries. Record where and when an answer appeared, then improve the facts a buyer needs.

01 / PRACTICAL FIELDNOTE

The short version

A mention tells you less than it appears to.

ASK
A useful buyer question

READ
The complete answer

CHECK
The source and context

ACT
One evidence-led decision

An AI answer names a furniture workshop. Is that good news? Perhaps. It might be a recommendation, a comparison, a warning or a passing reference. The useful evidence is the sentence around the name and the question that produced it.

This issue gives a small business a proportionate way to inspect that evidence. Use it after checking the accuracy of your own website and listings. Start with the questions customers actually ask and the business facts you can maintain.

A periodic check may be enough when there is little evidence that AI discovery matters to your customers. Increase attention only when the information could change a business decision. Treat a dramatic screenshot as one observation and ask whether it changes a practical decision.

Read the answer in context before turning it into a metric.

Fictional workshop. Documented data and source scope are identified in the text.

02 / PRACTICAL FIELDNOTE

Five different observations

Keep each label attached to its evidence.

MENTION	The name appears
CITATION	A source is referenced
RECOMMENDATION	The wording expresses fit
FIRST-NAMED	Position in the text
ENQUIRY	A separate business event

Separate observations. No automatic progression or attribution is implied.

A mention means the brand is named. A citation means a source is linked or otherwise identified. A recommendation requires wording that expresses a preference or suitable choice. First-named order records position in the text; it does not establish endorsement. An enquiry is a separate event in the business's own records.

These observations can coexist, but none guarantees the next. An answer can cite your website while recommending another supplier. It can recommend your service without linking your domain. A prospective buyer can read an answer and contact you later through an untracked route.

Keep a copy of the complete answer and its sources. If a dashboard supplies only a score, ask how that score is calculated and inspect examples. A number without a definition is mostly typography.

Mention, citation, recommendation, order and enquiry are separate fields.

Fictional workshop. Documented data and source scope are identified in the text.

03 / PRACTICAL FIELDNOTE

An answer under the magnifier

Invented example, not a captured AI result.

QUESTION

What should I check before ordering?

ANSWER

Workshop A offers a compact model.

READING

Named here; no preference established

Question: “What should I check before ordering a desk for a small room?” Invented answer: “Measure the usable space and access route. Workshop A offers a compact model. Compare its dimensions and assembly terms with other suppliers before ordering.”

Workshop A is named. The wording does not establish that it is the best option. If a link points to Workshop A’s specification page, that would be an owned-site citation, but the answer still asks the reader to compare.

Now change the sentence to “For the stated width limit, Workshop A is the best fit among these options.” That expresses a preference within stated conditions. In a real AI answer, that wording would express a recommendation to verify, not proof of product quality or a customer purchase. The words around the brand do the work.

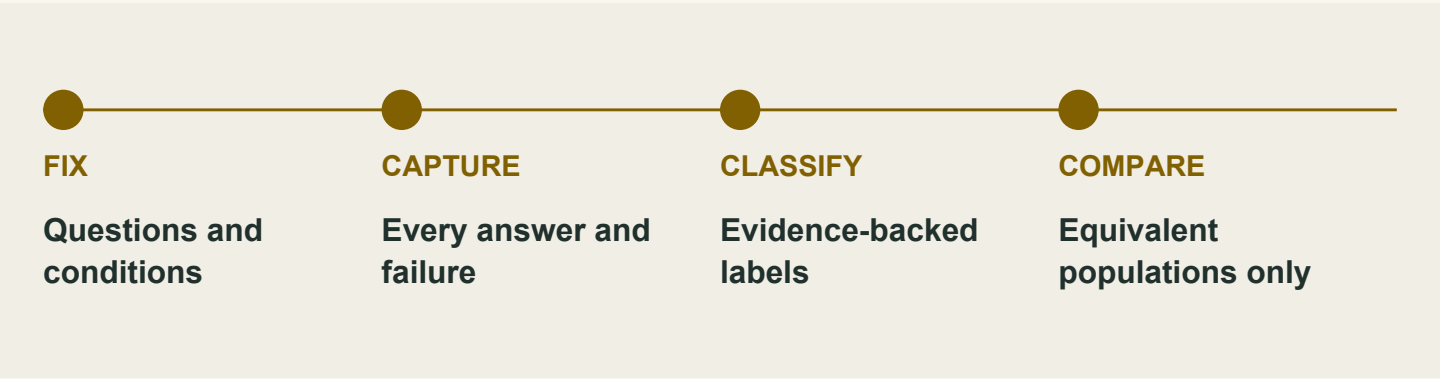
Classify what the answer actually says, including its conditions.

Fictional workshop. Documented data and source scope are identified in the text.

04 / PRACTICAL FIELDNOTE

Choose questions before collecting answers

Avoid building a test around the result you want.



Write a small panel of genuine buyer questions: suitability, alternatives, comparisons and practical requirements. Separate questions that name the brand from those that do not. A prompt containing your own name naturally changes what an answer might discuss.

Record the language, market, surface and settings. Consumer chat and API outputs should not be treated as interchangeable observations. Note whether browsing was enabled and whether the session carried earlier context.

Repeat questions consistently if repeatability matters, while retaining every attempt and failure. Do not keep asking until a pleasing answer appears and then call that the baseline. A practical panel samples specific questions. Its coverage of customer demand and possible answers remains limited.

Freeze the question panel and collection conditions before comparison.

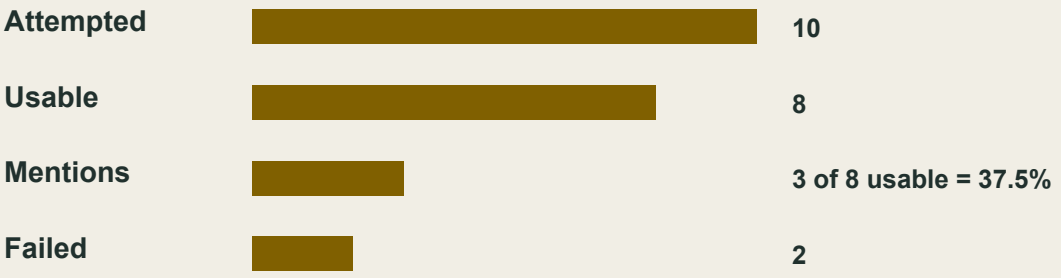
Fictional workshop. Documented data and source scope are identified in the text.

05 / PRACTICAL FIELDNOTE

The denominator deserves a label

A failed capture is not a negative answer.

ILLUSTRATIVE NUMBERS - NOT CLIENT RESULTS



Bars use a shared zero and scale within this figure.

Invented collection: 10 attempted captures; 2 fail; 8 return usable answers; 3 of those mention the workshop. The observed mention fraction is 3 of 8 usable answers, or 37.5%. Report the 2 failures alongside it. Do not quietly call the result 3 of 10 unless you explicitly define a different attempted-capture metric.

Small samples are volatile. One extra mention would move this example to 4 of 8. That is a large percentage change with very little underlying evidence. Repeated answers from one setup may also be dependent, so simple statistical assumptions need care.

Use counts and descriptions before ranking brands by small differences. The purpose is to find inspectable patterns and questions worth investigating, not to manufacture a precise-looking league table.

Show attempted, failed, usable and positive counts separately.

Fictional workshop. Documented data and source scope are identified in the text.

06 / PRACTICAL FIELDNOTE

Reading the October study’s scope

A worked example of checking the population behind a chart.

44 questions

Five runs each: 220 answers

40 questions

Non-branded explorer subset

6 Oct 2026

Historical collection date

Recallant’s published helpdesk study records 44 questions repeated five times, producing 220 answers collected on 6 October 2026. The public explorer uses a 40-question non-branded subset. Those populations must not be mixed.

Automated first-named order in that record is a positional measure, not a verified recommendation. The study also reports a separate ten-answer manual ChatGPT check; its different collection method makes it a separate observation. For a concrete scope check, 44 minus 40 leaves four questions outside that explorer subset. A chart of the subset cannot answer a question about all 44 without examining the exclusions.

This magazine does not rerun that research or claim the results describe today’s answers. It uses the study as an example of why dates, question sets and definitions belong beside a chart. It is not evidence of a client outcome or of an improvement caused by Recallant.

Historical observations stay dated, scoped and separate from current claims.

Fictional workshop. Documented data and source scope are identified in the text.

07 / PRACTICAL FIELDNOTE

Google reporting, checked November 2026

A dedicated view with a defined measurement scope.

IMPRESSION	**A reported visibility observation**
CLICK	**A different interaction**
ENQUIRY	**A business record, not an impression**

Separate observations. No automatic progression or attribution is implied.

Google now documents a dedicated Generative AI performance report for Search, covering impressions in AI Overviews and AI Mode. Its help page states worldwide rollout as of 31 August 2026. A missing report can reflect insufficient impressions.

The report's documented metric is impressions, not a separate AI click or revenue measure. Its data is included in the broader Web performance reporting: do not add it to the overall total as extra traffic. Check date, country, device and page scope when comparing views.

This section was checked on 19 November 2026 against the current report reference, including its "Not seeing the report?" and "More about the data" sections. Recheck that reference before building a recurring report around a tutorial or screenshot.

A new reporting view does not create a new revenue attribution method.

Fictional workshop. Documented data and source scope are identified in the text.

08 / PRACTICAL FIELDNOTE

Improve the page a buyer lands on

Answer the practical questions you can control.

FIT	TERMS	PROOF	NEXT STEP
Dimensions and materials	**Delivery and assembly**	**Relevant product photographs**	**How to confirm an order**

For our invented workshop, the useful page explains dimensions, materials, delivery coverage, assembly and how to confirm a custom order. It shows the actual item and links the applicable terms. These details help a buyer assess fit even when no AI system chooses to cite the page.

Google’s current guidance continues to connect AI-search visibility with SEO fundamentals, useful original content and relevant images or video. It cautions against mass-producing query variations primarily to manipulate results.

A small business can therefore start with one important page and its genuine information gaps. Collect the facts from the people doing the work. Add photographs that show the feature being discussed, rather than a decorative image that merely shares the mood. No page format guarantees inclusion.

Make the page useful on its own; discovery is not guaranteed.

Fictional workshop. Documented data and source scope are identified in the text.

09 / PRACTICAL FIELDNOTE

The visibility record

A practical worksheet for a bounded check.

Question / surface	Capture / date	Label + passage	Next action
Desk suitability	Saved answer / date	Mention / passage	Check page

First row is illustrative. Print and complete the blank rows; protect customer information.

Record the buyer question, reason for including it, exact wording, language, market, product surface, collection date and relevant settings. Save the complete answer and cited URLs where permitted. Mark failed or unavailable captures explicitly.

Classify mention, owned-domain citation, recommendation wording and first-named position separately. Add the sentence supporting each classification and note ambiguous cases. Let another reader challenge the labels before aggregating them.

Finish with the business implication. Does the answer misstate an important specification? Is the correct source absent from your own page? Or is the observation interesting but not actionable? Keep this decision alongside the evidence so another round of collection has a purpose.

Every metric should lead back to an answer and a defined question.

Fictional workshop. Documented data and source scope are identified in the text.

10 / PRACTICAL FIELDNOTE

A sensible next move

Fix a fact before buying a bigger number.

FIND

One information gap

FIX

The source you control

RECORD

The change and its date

Choose one recurring information gap and correct the source you control. For the workshop, that might be a delivery page that never says which areas it serves. Record the change date and check that the page works for a person on a phone.

If you later repeat an answer check, retain the original panel and note what else changed. Different answers after a page edit do not prove the edit caused them. Product updates, source changes and ordinary variation remain plausible explanations.

For the next question—whether visibility leads to useful conversations—turn to The Enquiry Ledger. It uses the business's own records rather than repeating the answer sample. Keep those two investigations distinct.

Choose one useful change, then observe without inventing causality.

Fictional workshop. Documented data and source scope are identified in the text.

The evidence behind this edition.

Google: optimising for generative AI search

https://developers.google.com/search/docs/fundamentals/ai-optimization-guide

Google: Generative AI performance report

https://support.google.com/webmasters/answer/16984139?hl=en

Recallant: dated helpdesk study and method

https://recallant.com/helpdesk-study.html

Source check: 19 November 2026. References support the specific factual passages marked in the online edition. They do not endorse Recallant or establish customer outcomes.

How this edition was made

Original editorial writing and diagrams, with AI-assisted drafting and independent critique. Arithmetic, claims and final layouts were checked before release. This is a practical guide, not a controlled evaluation, professional advice or a promise of commercial results. The accompanying HTML provides structured headings, readable text and source links; this visual PDF is not claimed to be PDF/UA certified.

Read the series / report a correction

recallant.com/fieldnotes.html
hello@recallant.com