FIELD GUIDE • 17 AUGUST 2026

AI Search Citation Analysis

A portable method for deciding what an AI-search citation actually proves - and what it does not.

> **Core rule**
> Do not optimize against a citation count until you can connect **response → citation → exact URL → answer claim → supporting passage → source role → repeatability → permitted conclusion**.

Use this guide when

- You see a competitor or third-party source cited.
- Your own page is cited but the claim looks stretched.
- You need to compare source inclusion over prompts or time.
- You need a handoff another reviewer can reproduce.

Do not use it as

- A ranking-factor list.
- A synthetic AI authority score.
- A promise that one content change will create citations.
- Proof that citation equals endorsement or durable visibility.

> The goal is not “more citations.” The goal is a smaller, better-supported decision.

01 • CANONICAL METHOD

SCOPE → CAPTURE → TRACE → CLASSIFY → COMPARE → DECIDE → REOPEN

Stage	Question	Output
SCOPE	Which prompt family, platform, locale, window and decision?	Governed sample.
CAPTURE	What answer and visible source state did we actually observe?	Response locator + citation state.
TRACE	Which answer claim does the cited source appear to support, and where?	Claim + passage locator.
CLASSIFY	What job is the source performing and what is the passage-support state?	Role + support code.
COMPARE	Does the finding repeat across material variants?	Within-sample stability.
DECIDE	What is the smallest justified action and owner?	Action route.
REOPEN	What change would make the old conclusion unsafe?	Reopening trigger.

Stop rule: when additional observations are unlikely to change the source-role diagnosis, passage-support conclusion, owner or intervention, the sample has reached operational decision saturation.

The minimum record that can survive handoff

Scope and capture

1. Observation ID
2. Prompt / query
3. Platform / surface
4. Locale
5. Observation date/time
6. Response locator
7. Citation state
8. Exact cited URL

Meaning and review

1. Answer claim
2. Passage locator
3. Source role
4. Passage-support code
5. Brand-mention state
6. Reproduction state
7. Owner / reviewer
8. Limitation + reopen trigger

Source roles

Official / first-party authority
Primary evidence
Specialist interpretation
Publisher / reference
Commercial source
Community / witness
Directory / marketplace
Unresolved

A cited source may not be your true content competitor. First identify the source job; only then decide whether the opportunity is content, evidence, authority, PR, technical eligibility, or no action. **UNRESOLVED is not a decision-ready source role:** preserve the uncertainty until the role can be classified.

Do not confuse topic similarity with entailment

Code	Reviewer meaning	Permitted conclusion
ENTAILS	Passage supports the claim with compatible scope and qualifiers.	Verified support in this observation.
PARTIAL	Only part of the claim is supported.	Narrow the proposition.
CONTEXT	Source is relevant but does not substantiate the claim.	Context, not proof.
CONTRADICTS	Source conflicts with the generated statement.	Support defect; preserve contradiction.
STALE / JURISDICTION MISMATCH	Wrong period, market, product version, law or geography.	Qualify or reject support.
ORPHAN / CIRCULAR	No stable upstream locator, or secondary claims repeat each other.	Provenance insufficient.
NOT CHECKED	Passage review not performed.	Citation observation only.

Refusal condition: if the URL, claim, passage or scope is missing, do not upgrade the record to verified support.

Keep captured answers and platform reporting as separate evidence layers

Sample design

Dimension	Control
Prompt family	Define the user task: define, compare, troubleshoot, select, verify or buy.
Variants	Change wording while preserving the underlying task.
Platform / mode	Do not merge different answer surfaces by default.
Locale	Keep language/market differences explicit.
Window	Preserve when the observation occurred.

Measurement boundaries

- **Google:** generative-AI Search Console reporting can expose site visibility dimensions where available; it is not a universal claim-level citation ledger.
- **Bing:** AI Performance exposes citation activity and grounding-query samples; Microsoft explicitly separates these from ranking, authority, importance and answer placement.
- **OpenAI:** documented referral parameters can measure visits from ChatGPT search; referral traffic is not complete exposure or citation measurement.
- **Third-party monitoring:** useful for prompt-scale discovery and competitor/source comparison, but not platform-owned causal explanation.

> **Research lens:** recent 2026 work distinguishes **citation selection** from **citation absorption**. A URL can be exposed as a source without contributing equally to the generated answer. Treat this as observational research, not a universal platform rule.

Find the missing source job before changing a page

Symptom	Likely question	Next move
Official source cited	Does the answer need an official requirement?	Interpret accurately; do not try to replace the authority.
Competitor research cited repeatedly	Is original evidence the missing job?	Decide whether legitimate first-party evidence is worth producing.
Third-party article carries your brand	Is independent validation doing the work?	Separate third-party mention strategy from own-domain citation strategy.
Your page is cited but PARTIAL	Is a qualifier/evidence condition missing?	Strengthen explicit support only where user value justifies it.
Your page is not cited and technical eligibility fails	Is technical eligibility blocking the test?	Fix the first technical prerequisite failure.
Your page is not cited and technical eligibility passes	Is the gap passage contribution, source role, authority or entity association?	Hand off to extraction-ready content, competitor/source-set, authority or entity analysis instead of continuing generic technical work.

> A citation gap is not automatically a content gap. It can be an evidence gap, source-role gap, external-validation gap, technical gate, or simply a task your page should not own.

06 • WORKED EXAMPLE

A competitor is cited - but the real gap is primary evidence

A team samples six variants about enterprise SEO-platform migration risk. An official migration guide appears five times, a specialist implementation article four times, and a vendor comparison page twice. The team's product page is not cited.

Source	Role	Support	Decision
Official migration guide	Official authority	ENTAILS continuity/ export requirements.	Use as authority; do not imitate.
Specialist article	Specialist interpretation	ENTAILS governance/ team dependency.	Possible opportunity for stronger implementation guidance if expertise exists.
Vendor comparison	Commercial source	PARTIAL for product capability.	Useful for facts; too narrow for neutral migration-risk conclusions.
Own product page	Commercial source	Not observed.	No evidence that “write more SEO copy” is the correct fix.

Permitted action: decide whether the business can create a distinct evidence object - original migration data, a reproducible implementation guide, or a genuinely expert explanation. If not, stop rather than manufacture reference-looking content.

What to share, what to refuse, and when to reopen

Decision Brief

1. Target page/domain + platform + locale + observation window
2. Prompt family / user task
3. Primary finding
4. Strongest evidence
5. Source-role pattern
6. Passage-support conclusion
7. Key limitation / uncertainty
8. Prohibited inference
9. Smallest next action
10. Owner
11. Reopen trigger + decision date

Three sentences that keep the analysis honest

> **Safe:** “This source was cited and its passage entailed the claim in the captured observation.”
> **Stronger but still bounded:** “Verified support repeated across the defined sample.”
> **Important negative:** “The passage entailed the claim, but the citation did not reproduce in the controlled retest.”
> **Unsafe:** “This source has durable AI authority because it has a high citation percentage.”

Primary references

Google Search Central: AI features and Search Generative AI performance reports • OpenAI Publishers and Developers FAQ • Bing Webmaster Tools AI Performance public preview • W3C PROV-O. Independent research: Zhang et al. (2026) on citation selection vs absorption; Allaham & Diakopoulos (2026) on synthetic-source citations. Market-method benchmarks: Ahrefs Brand Radar, Semrush, OtterlyAI, Peec AI, Scrunch and AthenaHQ.

Evidence limit: The method does not claim access to proprietary ranking systems. It keeps platform documentation, captured observations, third-party monitoring and SERPIFAI synthesis separate.