


AI SEARCH REPORT / 01 / BUSINESS IMPACT

AI SEARCH. THE LATEST FINDINGS.

For B2B founders and small marketing teams. Pages 1-4: executive brief. Pages 5-8: evidence notes, annotated posts, and a worksheet.

The decision: Keep improving the pages buyers need to evaluate your product. Compare AI visibility with qualified interest before moving search budget.



Rate compares ChatGPT with Google organic. Volume compares **all LLMs** with Google organic. The volume chart is an index, not raw conversion counts. [1]

The wider sample: 100 websites and more than 145 million sessions over 12 months. ChatGPT sessions were 0.6% of Google organic sessions; all LLMs were 0.7%. These are ratios to Google organic, not shares of all website traffic. Industry mix, exact reporting dates, and conversion definitions were not specified.

AI research, then a website decision. Kieran Flanagan writes: "The buyer journey is getting condensed into AI Search (research) -> Vendor website (closing)." [SVP, Agentic GTM & Systems, HubSpot](#). [5] **JAC takeaway:** Answer the buyer's remaining objections. This is his August 2025 interpretation, not proof of a universal journey.

[1] [Marie Haynes: traffic and conversion analysis](#)




AI SEARCH REPORT / 02 / DISCOVERY AND CITATION

GETTING FOUND. GETTING CITED.

A source can enter an AI platform's candidate pool without appearing in the answer.

POOLED ENTRIES NEVER CITED



Steve Toth: "Getting retrieved is only the first step. The page also has to answer the exact question the model is writing about." [2] The comparison covers 15 three-turn conversations per platform, or 45 turns each, within a 277-conversation study.

Keep the limit visible. Captured in one day across three B2B SaaS categories. Pooled entries may repeat and do not prove that a complete page was read. Platform telemetry differs, so this is a directional comparison.

Source choice also changes with the question

82%**Earned sources**

For customer-review questions

50%**Owned sources**

For functionality questions

Omniscient Digital: 240 branded prompts across five systems and four industries; 23,387 unique sources. These percentages describe different intent groups, not a combined channel mix. [4]

Use this: Inspect both accurate product information on your site and the credible outside sources a buyer would consult. Record retrieval as unknown if the tool does not expose its source pool.

[2] Steve Toth: [How AI Search Works](#)

[4] Omniscient Digital: [sources cited for branded intent](#)




AI SEARCH REPORT / 03 / CONTENT PRIORITIES

GIVE EACH PAGE A CLEAR JOB.

Relevant information matters. Current research does not establish a universal page format or winning length.

CITATION RATE AMONG STRONG PRIMARY-QUERY MATCHES



AirOps + Kevin Indig: 16,851 ChatGPT queries, each run three times; 353,799 scraped pages and 815,484 scoring rows. Primary similarity ≥ 0.8 . [3]

What this supports: This comparison is observational and uses heading embeddings; the bands are not exhaustive. Shortening a page is not a proven citation tactic. [3]

Schema as a shortcut: Ahrefs followed 1,885 pages adding JSON-LD against 4,000 controls. Adjusted changes: -4.6% AI Overviews, +2.4% AI Mode, +2.2% ChatGPT. No clear uplift. Already-cited pages, pooled schema types, a 30-day follow-up, and possible coinciding changes limit the conclusion. [12]

Platform basics: Google emphasizes distinctive content and SEO foundations; extra AI files and schema are not requirements. Check Search Console inclusion. [6, 21] OpenAI separates search access (OAI-SearchBot) from training (GPTBot). Check host/CDN access; placement is not guaranteed. [19, 20]

Demand estimates: Soulo challenges exact prompt counts. Semrush models topic demand. Ask what is observed versus estimated; tracking can still provide a baseline. [13, 14, 15]

[3] AirOps + Kevin Indig: The Fan-Out Effect


[6] Google: Generative AI optimization guide



AI SEARCH REPORT / 04 / ATTENTION AND ACTION

WHAT MARKETERS ARE LOOKING FOR.

Search interest describes attention. It does not establish which tactic works.

GOOGLE TRENDS / US WEB SEARCH / SHARED 0-100 SCALE

Full phrases, search terms rather than Topics; complete-week starts 24 Oct 2021 to 11 Oct 2026. Partial boundary weeks excluded. Relative interest, not counts or market share; zeros can reflect insufficient volume. Each phrase had lower average interest in the latest 13 complete weeks than in the preceding 13. [16]

Interest is ahead of execution. In Semrush's 481-person survey, **77%** called AI search an extension of SEO; **22%** reported fully integrated work and reporting. These are different questions, not complementary shares. Self-reported vendor research, not a causal result. [17]

Buyer language matters. Fractl's survey of 343 US marketing decision-makers found 81% called their AI visibility strategy SEO; 36% flagged unexplained buzzwords. Self-reported vendor research. [18] The same-team preprints cover 158,099 answers / 468,037 citations (English US API sample) and a 31,200-answer location audit, not independent confirmations. [7, 11]

Your next check: Save one buying question and its answers. Verify the product facts and cited sources, fix the information gap, and repeat under comparable conditions. Review qualified interest separately from mentions. Keep specific buying prompts alongside branded ones. [10]

Read the originals

[1] Traffic and conversion study

[4] Branded-intent source study

[7] Citation preprint

[18] Fractl: buyer vocabulary

[12] Ahrefs: schema study

[15] Soulo: tracking value

[19] OpenAI: crawler controls

[2] Retrieval and citation telemetry

[5] Flanagan: buyer-journey interpretation

[11] Location preprint

[10] Marcus Tober: Page 2 Podcast

[13] Soulo: prompt volumes

[16] Google Trends comparison

[20] OpenAI: search eligibility

[3] Fan-out and citation research

[6] Google optimization guidance

[8] Google performance reporting

JAC: implementation guide

[14] Semrush: volume method

[17] Semrush: marketer survey

[21] Google: inclusion control

AI SEARCH REPORT / 05 / EVIDENCE NOTES

WHAT THE DATA CAN SUPPORT.

Keep each result attached to its population. These are separate studies, not one pooled experiment.

SOURCE / SAMPLE	OBSERVATION	DECISION / LIMIT
Haynes [1] 100 sites; 55 tracking conversions	ChatGPT 13.8% vs Google organic 3.0% conversion rate.	Use your own qualified-lead rate. Definitions and industry mix were not specified.
Toth [2] 45 turns per platform in this comparison	52-75% of pooled entries were never cited across four platforms.	Separate retrieval and citation. One-day B2B SaaS sample; entries can repeat.
AirOps / Indig [3] 16,851 queries; 353,799 pages	35.5%, 38.2%, 34.0% citation rates across selected coverage bands; strong primary match.	Check page relevance first. Observational; no universal length or shortening rule.
Ahrefs [12] 1,885 adding JSON-LD; 4,000 matched controls	-4.6% AIO; +2.4% AI Mode; +2.2% ChatGPT adjusted changes after 30 days.	No clear uplift in this sample. Pooled schema; already-cited pages; not randomized.
Omniscient [4] 240 prompts; 5 systems; 4 industries	Review intent: 82% earned. Functionality intent: 50% owned sources.	Check sources by intent. Different question groups; not one channel mix.

LinkedIn: whose content gets cited?



[Meltwater + LinkedIn, report p17](#). 9.5M citations, six systems, 16 B2B categories. Shares are within LinkedIn citations; the LinkedIn-only sample size is not stated. Vendor research, not a publishing experiment.

Format advice varies: [Ward's earlier Mentions analysis](#) reports 8x more citations for articles than posts. Meltwater reports 72% text posts / 12% articles; [LinkedIn's separate guide](#) describes about 60% articles / 40% posts. Samples and definitions differ. None is a citation probability adjusted for how much of each format was published.



THE POSTS.

THE PRACTICAL POINT.

Excerpt recreations from public LinkedIn posts. Counts come from dated retrievals; see page 8. cached pages may lag. Attention does not establish that the advice works.

**Tim Soulo**

CMO & Product Advisor, Ahrefs

2025-12-17

P01 / EXCERPT

"The biggest mistake I see people make in AEO/GEO is measuring their AI visibility based on a tiny list of hyper-relevant prompts."

[152 reactions / comments not returned](#)[Read original post](#)

Soulo warns that a narrow set of product-focused questions can miss discovery that starts with the buyer's problem.

Try this: Include a real buyer problem alongside a product comparison; keep the set consistent between runs.

**Lily Ray**

Founder, Algorythmic

2026-09-30

P02 / EXCERPT

"Better coverage can lead to better visibility"

[806 reactions / 193 comments](#)[Read original post](#)

A short question connects a new search behavior with established SEO work. It attracted substantial discussion, but does not quantify how much extra coverage changes visibility.

Try this: Map a real buyer question to the follow-up information your pages lack.

**Steve Toth**

Founder, Notebook Agency

2026-10-14

P03 / EXCERPT

"This is one sample of 277 conversations, so treat the skip rates as directional."

[74 reactions / 20 comments](#)[Read original post](#)

Original telemetry gives readers something concrete to inspect. Toth emphasizes citations as the win; Reynolds and Fishkin argue citation totals miss other influence.

Try this: Keep retrieved, cited, and recommended as separate observations.



WHERE EXPERTS AGREE AND DIFFER.

Reynolds questions citation totals as a standalone KPI. Ward and Solís connect useful content with information beyond your own site. These are attributed views.

**Wil Reynolds**

Founder, Seer Interactive

2026-09-26

P04 / EXCERPT

"Citations are a bad KPI for Generative engine optimization."[323 reactions / 67 comments](#)[Read original post](#)

The post challenges citation trend lines as a standalone KPI. Its useful business test is whether customers still use the platform even when citation share drops.

Try this: Compare the answer with qualified visits, buyer feedback, and factual accuracy.

**Jake Ward**

Cofounder, Known

2026-09-24

P09 / EXCERPT

**"Pick a small number of topics you want to own
Build depth around them with posts + articles
... Reinforce the same ideas beyond LinkedIn"**[506 reactions / 191 comments](#)[Read original post](#)

Ward connects a focused subject with useful depth and consistent information beyond LinkedIn.

Try this: Develop one substantiated topic across LinkedIn and your website.

**Aleyda Solís**

SEO & AI search consultant

2026-09-23

P06 / EXCERPT

"AI Search is an off-site corroboration problem, with an on-site quality floor"[175 reactions / 18 comments](#)[Read original post](#)

A practical framework links good owned information with credible outside corroboration. It points readers to worksheets, giving the post value beyond its headline.

Try this: Fix one product fact on your site, then check one credible independent source.

AI SEARCH REPORT / 08 / YOUR NEXT CHECK

MAKE ONE TESTABLE CHANGE.

JAC worksheet: choose a real buying question and record what you can actually observe. A useful improvement is not necessarily a measurable citation lift.

01 / BASELINE

Buyer question + target market: _____
 System / model / date / search mode: _____

02 / OBSERVE

Brand recommended? Facts correct? Cited pages? _____
 Exposed queries / retrieved sources (or unknown): _____

03 / CHANGE

One information gap + page / outside source to check: _____
 Owner / change date / repeat conditions: _____

Marcus Tober: "I wouldn't generalize that the homepage is always the most cited page. ... often most cited because the setup has prompts that would surface the homepage." [SVP](#), [Enterprise Solutions](#), [Semrush](#). [10] **JAC takeaway:** Repeat a consistent set of buyer questions before interpreting the page mix.

The complete selected post list

Nine accessible posts, seven authors, including older measurement guidance added October 27. Each had at least 100 reactions or 20 comments in the returned extraction. Relevance and variety also guided selection. This is a curated collection, not an exhaustive ranking.

AUTHOR / SUBJECT (LINKED)	REACTIONS	COMMENTS
P01 / Tim Soulo: Choose a useful prompt set	152	n/a
P02 / Lily Ray: SEO inside the fan-out	806	193
P03 / S. Toth: Retrieval and attribution	74	20
P04 / W. Reynolds: Citations in context	323	67
P05 / R. Fishkin: What happens before the citation	344	46
P06 / A. Solís: Quality and corroboration	175	18
P07 / Tim Soulo: Exact prompt-volume claims	161	34
P08 / Tim Soulo: Ahrefs research roundup	1,749	n/a
P09 / Jake Ward: LinkedIn as a search asset	506	191

Counts retrieved via Exa October 22, except P01 checked October 27; not audited live counters. No impressions, paid-distribution, or lead data. P01 and P08 comments were not returned. P07 critiques exact prompt volume; P08 is a vendor research roundup with different study samples. No paid status is inferred.

Portrait credits: [Tim Soulo](#) / [Lily Ray](#) / [Steve Toth](#) / [Wil Reynolds](#) / [Jake Ward](#) / [Aleyda Solís](#)