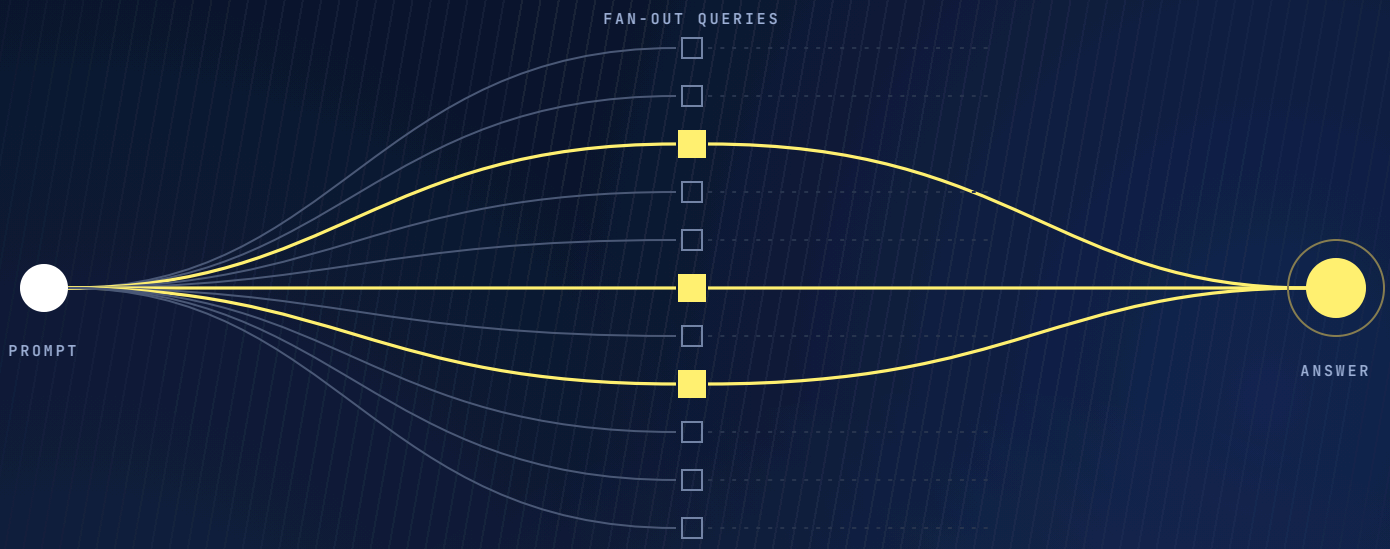
CHATGPT / CLAUDE / GEMINI

AI Algorithm Ranking Factors

How answer engines decide which brands to name, **what the evidence supports**, and where to spend effort first.



84%
of AI citations come from earned media

0.664
correlation of brand web mentions with AI Overview visibility

44.2%
of ChatGPT citations come from the first 30% of a page

<1%
chance the same prompt returns the same brand list twice

Contents

01 **Executive Summary** 03
The findings in six numbers

02 **Selection Architecture** 04
Two layers: model memory and live retrieval

03 **Layer 1: Parametric Memory** 05
What the model already believes about your brand

04 **Layer 2: Live Retrieval** 06
Query fan-out, retrieval rank, synthesis

05 **Platform Profiles** 07
ChatGPT, Claude, Gemini and Google AI

06 **Platform Comparison** 10
Where the three agree and where they split

07 **Ranking Factor Scorecard** 11
Eight factors graded by evidence

08 **The Eight Factors** 12
Data, mechanism and actions for each

09 **Tactics Without Evidence** 16
What the data does not support

10 **Crawler Access Reference** 17
User agents, robots.txt and rendering

11 **Measurement** 18
Tracking visibility when answers never repeat

12 **90-Day Action Plan** 19
Owned and off-site work in sequence

13 **Sources and Method** 20
Every study cited, with links

ABOUT THIS REPORT

No AI vendor publishes a ranking algorithm. What exists is evidence.

ChatGPT, Claude and Gemini do not rank pages the way a classic search engine does. A brand gets named when it survives two filters: what the model learned in training, and what it retrieves from the web at answer time.

"Ranking factors" in this report means the inputs that measurably change whether a brand is named or cited. Each one is tied to a published study or to vendor documentation, with the source numbered and linked.

EVIDENCE LABELS

EXPERIMENT	Controlled test that changed one thing and measured the result
AT SCALE	Observed across millions of answers or citations
CORRELATION	Statistical association. Not proof of cause
OFFICIAL DOCS	Stated by OpenAI, Anthropic or Google

READING THE NUMBERS

Superscript numbers link to the source list on page 20. Many studies come from vendors that sell AI visibility tools, and results shift as models update. Where two studies disagree, both are shown with the setup that explains the gap.

29	8	3	2
SOURCES	FACTORS	PLATFORMS	LAYERS

Brands are named by consensus, then checked against the source.

Third-party coverage decides whether an answer engine considers your brand. Your own site decides whether it can confirm the details. Most of what is sold as AI optimization sits outside both.

84%

of citations across ChatGPT, Claude and Gemini point to earned media. Paid and advertorial content: 0.3%.[7]

MUCK RACK, 25M+ LINKS, MAY 2026

3x

Brand web mentions correlate with AI Overview visibility at 0.664. Backlinks: 0.218.[2]

AHREFS, 75,000 BRANDS

95%

of the hidden sub-queries AI runs have zero measurable search volume.[10,16]

AIROPS (CHATGPT), SEER (GEMINI 3)

58%

citation odds for the top retrieved result in ChatGPT. By position ten: 14%.[11]

AIROPS, 16,851 QUERIES

44.2%

of ChatGPT citations are pulled from the first 30% of a page's text.[9]

GROWTH MEMO, 18,012 CITATIONS

<1%

chance that the same prompt returns the same list of brands twice.[8]

SPARKTORO AND GUMSHOE, 2,961 RUNS

FIVE FINDINGS

01 **Selection happens in two layers.** The model first draws on what it learned in training, then decides whether to search. Each layer responds to different work.

02 **Each platform searches a different index.** Claude leans on Brave, Gemini on Google, ChatGPT on its own index plus outside providers. Sources diverge sharply. The brands named overlap far more.[18,20]

03 **Third parties win discovery. Your site wins verification.** On factual questions, 84.4% of ChatGPT answers cite at least one official source.[3]

04 **Structure and specifics beat tone.** Statistics, cited sources and quotations lifted visibility 30 to 40% in controlled tests. Keyword stuffing did not.[1]

05 **Several popular tactics show no measured effect.** llms.txt files and schema markup added only to win citations both came back null in testing.[6,28]

THE SHORT VERSION

Be the brand other sites describe the same way, and be the site that confirms it in plain HTML.

- Earn mentions and list placements on sites AI already cites
- Publish exact pricing, specs and comparisons, answer first
- Rank in the index each platform reads
- Serve content without client-side JavaScript
- Measure appearance rate, never position

ChatGPT

PAGE 07

Own index plus outside providers. Verifies against official sites. Strongest pull toward recent pages.

Claude

PAGE 08

Tracks Brave Search. Searches less, so model memory decides more. Leans on reviews.

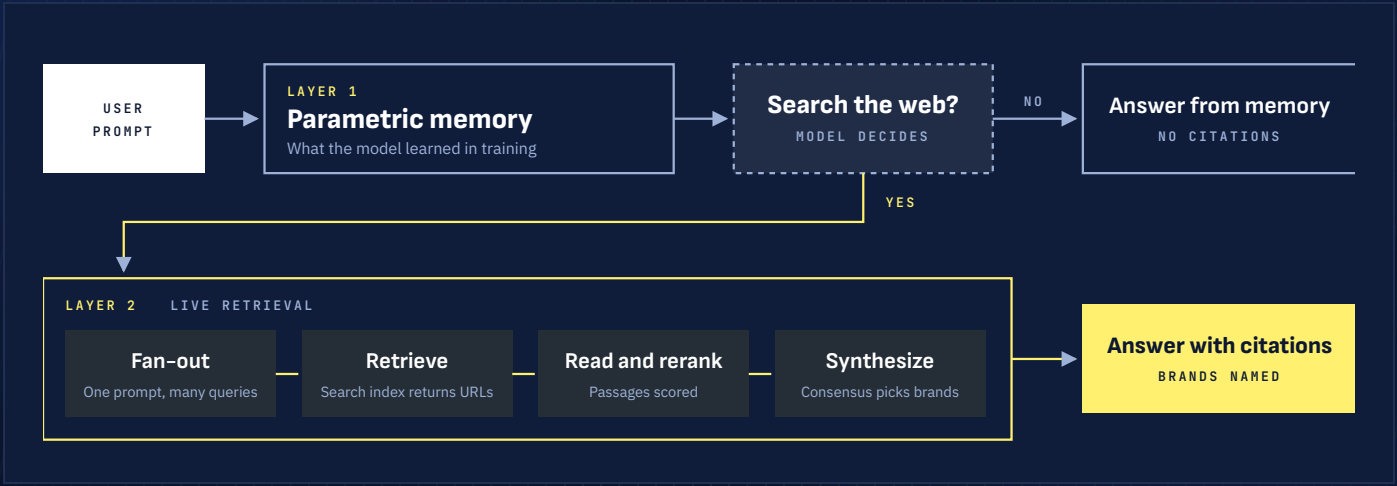
Gemini

PAGE 09

Grounded in Google's index. Widest fan-out. Favors first-party sites and listings.

Two layers decide every recommendation

There is no single ranking algorithm. A brand is named when it clears model memory, live retrieval, or both.



LAYER 1: PARAMETRIC MEMORY

Before any search, the model already holds an impression of your category: which brands belong to it, what each is known for, and who they compete with. That impression comes from how often and how consistently a brand appears next to category terms across the text the model trained on.

Memory also shapes Layer 2. The model writes its own search queries, and those queries often name brands directly. In one study, 26.4% of Gemini 3 fan-out queries did.[16]

LAYER 2: LIVE RETRIEVAL

When the model searches, it breaks the prompt into sub-queries, pulls candidate pages from a search index, reads passages, and writes an answer. Brands that several independent pages agree on are the ones that reach the final list.

Three things gate this layer: whether your page is in the index, where it ranks for the sub-query, and whether the passage answers it directly.

HOW OFTEN DOES LAYER 2 RUN?

14.4%

of free, logged-out ChatGPT responses contained fan-out queries. Six in seven were answered without them.[3]

AHREFS, 66.6M RESPONSES, AUG TO SEP 2026

36.6%

search-trigger rate for Claude in a separate test, against roughly 90 to 95% for ChatGPT.[18]

PROFOUND, JUNE 2026, VIA XPONENT21

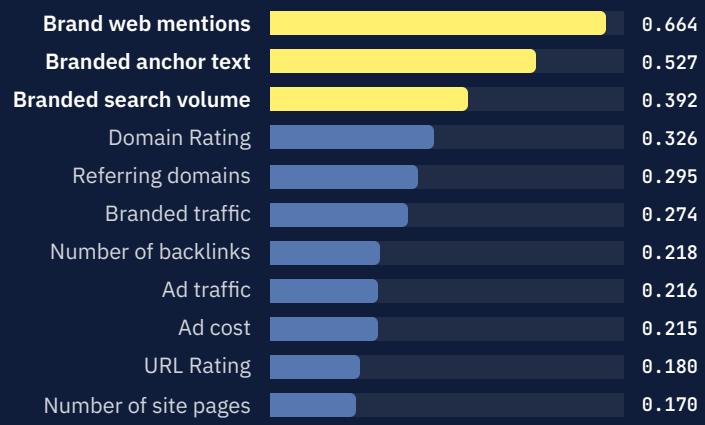
Why the gap. Search frequency depends on the prompt and the product tier. Broad everyday prompts rarely trigger search. Commercial "best X" prompts, paid plans and reasoning models trigger it far more.[3] Plan for both layers.

What the model already believes

In the largest public study of the question, the best predictor of AI brand visibility was how widely the brand is mentioned across the web.

Off-site signals predict AI visibility better than links

Spearman correlation with brand mentions in Google AI Overviews. 75,000 brands.[2]



Yellow bars are the three strongest factors. All three are off-site. Correlation is not causation, and the study measured AI Overviews, not ChatGPT or Claude directly.

10x

Brands in the top quartile for web mentions earn up to ten times more AI Overview mentions than the next quartile.[2]

26%

of the 75,000 brands studied had zero mentions in AI Overviews.[2]

READING IT CORRECTLY

This finding is often quoted as "unlinked mentions beat backlinks." The study counted mentions whether linked or not. The point is reach across many sites, not the absence of a link.

WHY MENTIONS MATTER

A language model learns associations from text. If "Acme" appears next to "payroll for restaurants" on hundreds of independent pages, the model links the two. A backlink profile carries no such meaning to a model that never sees link graphs.

WHAT BUILDS MEMORY

- Press and trade coverage that names the brand with its category
- Podcast and video appearances with published transcripts
- Reviews and community discussion
- The same one-line description everywhere: site, LinkedIn, Crunchbase, Wikidata, review profiles

TRAINING CRAWLERS

These collect content that may be used to train future models. Blocking them is a legitimate choice, with a likely cost to Layer 1 over time.[23,24]

GPTBOT · CLAUDEBOT · GOOGLE-EXTENDED

Memory is slow to change. A model's view of a brand updates only when a new model ships.

That makes Layer 1 the long game and Layer 2 the short one. Mentions earned this quarter can show up in retrieved answers within weeks. They reach model memory only after the next training cycle. Consistency compounds: conflicting descriptions of what a company does split the association across several weak signals.

One prompt becomes many searches

The queries that decide your visibility are written by the model, not the user. Keyword tools cannot see most of them.

95%
of fan-out queries show zero monthly search volume, for ChatGPT and for Gemini 3.[10,16]

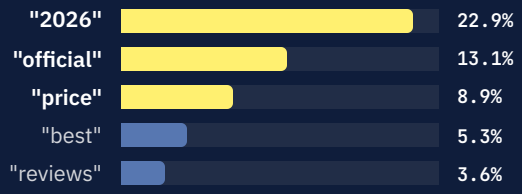
10.7
average fan-out queries per prompt for Gemini 3 with grounding forced on.[16]

2.1
average for ChatGPT. Perplexity runs 1.4, Grok 6.8.[13]

89.6%
of ChatGPT searches on 15,000 prompts generated two or more fan-out queries.[10]

Words ChatGPT adds to its own searches

Share of English fan-out queries containing each term, September 2026.[3]



"Official" rose 3.0 points and "price" 2.8 points in a single month. The model is steering itself toward current, first-party facts.

Retrieval rank sets citation odds

Chance a retrieved page is cited in ChatGPT, by its position in the results the model received.[11]



A four-fold gap from rank alone. The same study found pages whose headings closely match the query were cited 41% of the time, against 29% for weak matches.[11]

SETUP CHANGES THE NUMBERS

STUDY	SETUP	FAN-OUTS PER PROMPT	SITE: SHARE
Ahrefs	Free, logged out, 66.6M responses	About 1.6	6.8%
Peec AI	Tracked prompts, size undisclosed	56.5% had 2 or more	About 23%
Nectiv	API, about 4,000 prompts	7.61	64%
Wil Reynolds	Small tracked subset	Not reported	76%

As compiled by Ahrefs.[3] Broad consumer prompts trigger few brand lookups. Commercial "best X" prompts and paid reasoning models trigger many. Any single figure quoted for fan-out behavior describes one setup.

SITE LOOKUPS GO TO BRANDS

When ChatGPT restricts a search to one domain, 75.9% of those lookups among the most-searched sites target the brand's own website. Retailers take 11.9%. Reddit and Facebook together take 1.9%.[3]

WHAT THIS MEANS

- Write for the sub-question, not the head keyword
- Put the current year, exact prices and specs in the text
- Give each page one job. Pages covering 26 to 50% of a prompt's fan-outs were cited more than pages covering all of them[11]
- Rank in the index the platform reads. See pages 07 to 09

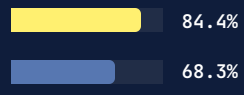
ChatGPT

Runs on its own search index plus outside providers. Checks facts against official sources and strongly prefers recent pages.

SEARCH INDEX	OpenAI's own index, supplemented by external providers. Peec AI observed ChatGPT labeling results as coming from an internal index or from third-party scraping services during a two-month window in 2026.[14] The Information reported in August 2025 that OpenAI had used Google results scraped through SerpApi for timely topics.[27]
CRAWLERS	`OAI-SearchBot` builds search visibility. `ChatGPT-User` fetches pages on a user's behalf. `GPTBot` collects training data. Sites that block OAI-SearchBot are not shown in ChatGPT search answers.[23]
QUERY STYLE	Rewrites the prompt before searching. Adds the current year, "official", "price" and "best". Runs `site:` lookups that mostly target brand websites.[3]
SOURCE LEAN	Official sites for facts, comparison lists for discovery. "Best X" lists made up 43.8% of cited source pages across 750 top-of-funnel prompts.[5]
RECENT SHIFT	Between March and September 2026, Reddit citations fell 88% and Wikipedia 80%. PubMed rose from 64th to 4th most-cited domain. Gains went to regulators, professional bodies, brand and support sites, and hands-on testers.[3]

Answers that cite an official source

Share of ChatGPT answers citing at least one official website. 20,000 US prompts.[3]



Official sites received 63.4% of all citations on factual questions and 46.4% on comparisons.

393 days

ChatGPT's in-text citations are this much newer, on average, than the pages ranking in Google for the same query. The strongest freshness preference of any assistant studied.[4]

88%

drop in Reddit citations between March and September 2026[3]

43.8%

of cited source pages on discovery prompts were "best X" lists[5]

6.8%

of fan-out queries use a `site:` lookup, at scale[3]

~24 hrs

for a robots.txt change to register with OpenAI[23]

PRIORITIES FOR CHATGPT

- **Allow OAI-SearchBot** in robots.txt and at the CDN. Changes take about 24 hours to register[23]
- **Keep official fact pages current:** pricing, specifications, integrations, support. Show a real last-updated date
- **Earn placement on third-party comparison lists** the model already cites for your category
- **Hold rankings in Google and Bing.** Both have fed ChatGPT retrieval

CAUTION

OpenAI changes retrieval often. Peec AI recorded at least five shopping experiments running at once, including one that preferred the internal index over outside results for 8% of chats.[14] Treat any snapshot of ChatGPT's sources as provisional, and re-check quarterly.

ChatGPT-User acts on a user's request, and OpenAI states that robots.txt rules may not apply to it.[23]

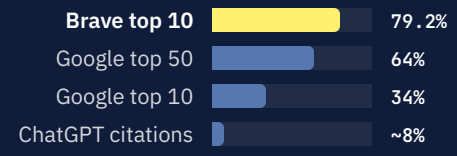
Claude

Searches less often than ChatGPT, and when it does, its citations track Brave Search far more closely than Google.

SEARCH INDEX	Brave Search, listed by Anthropic as a subprocessor since March 19, 2025. A second web-search subprocessor, turbopuffer, was added in May 2026. Its role in citations is not public.[18]
CRAWLERS	Anthropic documents three: `ClaudeBot` (training), `Claude-User` (fetches pages when a user asks) and `Claude-SearchBot` (improves search result quality). All honor robots.txt.[24]
BRAVE ACCESS	Brave's crawler has no separate user agent to allow. Per Brave's documentation, if Googlebot cannot crawl a page, Brave will not either. Googlebot access is the prerequisite.[18]
SEARCH RATE	Triggered a search on 36.6% of prompts in one June 2026 test, against roughly 90 to 95% for ChatGPT.[18] Model memory carries more of the answer here.
SOURCE LEAN	Cites reviews and social content at two to four times the rate of other models, and cites journalism less.[18,19]

Where Claude's cited URLs also appear

Share of Claude citations found in each result set for the same query.[18]



ABOUT THE 86.7% FIGURE

The widely quoted "87% overlap with Brave" comes from 13 of 15 results across three queries, measured the day after launch in March 2025.[17] A June 2026 re-run on about 400 queries found 79.2%.[18] The direction holds. The precision never existed.

79.2%
of cited URLs appear in Brave's top ten[18]

36.6%
of prompts triggered a web search in one test[18]

2 to 4x
the review and social citation rate of other models[19]

3
documented crawlers, each with its own robots.txt token[24]

PRIORITIES FOR CLAUDE

- **Check your rankings in Brave Search** for category and comparison queries. Few teams have looked
- **Keep Googlebot unblocked**, then allow Claude-SearchBot and Claude-User
- **Invest in review platforms and community presence.** Claude draws on them more than its peers
- **Build Layer 1.** With fewer searches, what the model already knows decides more answers

BY INDUSTRY

In food and beverage, 24.35% of Claude's citations came from reviews and social sources, nearly ten times Gemini's rate.[19] A Spring 2026 update put reviews at about 15% of Claude's citations overall, against 3 to 7% for other models.[19]

A brand with strong owned content and thin review coverage can be visible in Gemini and close to absent in Claude.

Gemini and Google AI

Grounded in the Google Search index. Fans out further than any other platform and leans hardest on first-party websites.

SEARCH INDEX	Google Search. AI Overviews and AI Mode draw on the same index and ranking systems as regular results. A page must be indexed and eligible to show a snippet.[25,26]
CRAWLERS	`Googlebot` for Search, including its AI features. `Google-Extended` is a separate control that limits AI training and grounding in some of Google's other systems. It does not remove a site from Search.[25]
QUERY STYLE	Gemini 3 averaged 10.7 fan-out queries per prompt at 6.7 words each. 26.4% named a brand and 21.3% contained a year.[16]
SOURCE LEAN	First-party websites and business listings. Google's own guide tells site owners to keep local business and ecommerce details accurate.[19,26]
FRESHNESS	Lower sensitivity than other assistants. AI Overviews cite pages about as old as those in regular organic results.[4]

Gemini cites what brands control

Share of Gemini citations by source type. Yext, 17.2M citations, Q4 2025.[19]



Yext's prompts carry location and brand context, which favors owned sources. Compare with the earned-media finding on page 12.

0.664

Correlation between brand web mentions and AI Overview visibility. The mentions finding on page 05 was measured on Google's own AI surface.[2]

10.7

fan-out queries per prompt for Gemini 3[16]

95%

of those queries have zero search volume[16]

93%

of citations come from sources organizations manage[19]

0

special files or markup required, per Google[26]

PRIORITIES FOR GEMINI AND GOOGLE AI

- **Do the SEO.** Indexing, internal links, page experience and helpful content are the entry ticket
- **Cover the sub-questions** a buyer would ask, in the page that owns the topic
- **Maintain Google Business Profile and Merchant Center** data where they apply
- **Build first-party depth.** Product, pricing and support pages carry more weight here than anywhere else

GOOGLE'S STATED POSITION

Working on AI search visibility is "optimizing for the search experience, and thus still SEO."[26]

Google's May 2026 guide lists what it says is unnecessary for its AI features: llms.txt files, special markup, content chunking, rewriting for AI, and seeking inauthentic mentions. It also warns against a separate page for every fan-out variation.[26]

Different sources, similar brands

The three platforms read different parts of the web and still land on a similar short list of names.

DIMENSION	CHATGPT	CLAUDE	GEMINI AND GOOGLE AI
Search index	Own index plus external providers. Google and Bing results have both been observed[14,27]	Brave Search. Second provider listed, role unknown[18]	Google Search index[25]
Search crawler	`OAI-SearchBot`	`Claude-SearchBot`, plus Googlebot access for Brave	`Googlebot`
User fetcher	`ChatGPT-User`	`Claude-User`	Google's user-triggered fetchers
Training crawler	`GPTBot`	`ClaudeBot`	`Google-Extended` (control token)
Fan-out	About 2 queries per prompt on average. Far more on paid reasoning models[3,13]	Fewer searches overall[18]	10.7 queries per prompt for Gemini 3[16]
Source lean	Official sites for facts, comparison lists for discovery[3,5]	Reviews and social at 2 to 4x the rate of peers[19]	First-party sites and listings[19]
Freshness	Strongest. Cites pages 393 days newer than organic[4]	Not separately measured	Weakest. Close to organic results[4]
Renders JavaScript	No[22]	No[22]	Yes, through Google's rendering[22]
First lever	Official fact pages and third-party lists	Brave rankings and reviews	Classic SEO and first-party depth

Engines agree on brands, most of all in retail and travel

Average overlap in named brands between pairs of AI engines, by category. Five engines.[20]



BrightEdge, May 2026: ChatGPT, Perplexity, Gemini, AI Mode, AI Overviews. Across all categories combined, top-30 overlap was 36 to 55%. Sample size not published.

~8%

of the URLs Claude cites are also cited by ChatGPT for the same query.[18]

IMPLICATION

Chasing each platform's favorite sources is a treadmill. The durable target is the outcome they share: a brand that many independent sources describe the same way.

Platform-specific work is then narrow: access, and the index each one reads.

TWO STUDIES, TWO HEADLINES

Muck Rack reports 84% of AI citations go to earned media.[7] Yext reports that 86 to 90% go to sources brands control.[19] Both hold. Muck Rack sampled broad prompts and counts journalism, research and third-party sites as earned. Yext sampled prompts with brand and location context, where websites and listings dominate.

THE SPLIT IS THE STRATEGY

Open-ended discovery prompts lean on third parties. Specific, factual prompts lean on the brand's own pages.

Eight factors, graded by evidence

Each factor below changed whether brands were named or cited in at least one large study. The grade reflects how solid that evidence is, not how hard the work is.

NO.	FACTOR	LAYER	EVIDENCE TYPE	STRENGTH	MATTERS MOST ON
01	Third-party consensus	Both	AT SCALE	▬▬▬▬▬	All three
02	Brand mention volume	Memory	CORRELATION	▬▬▬▬▱	All three. Claude most, since it searches least
03	Rank in the underlying index	Retrieval	AT SCALE	▬▬▬▬▱	Each platform, in its own index
04	Semantic alignment and structure	Retrieval	AT SCALE, EXPERIMENT	▬▬▬▬▱	ChatGPT, Gemini
05	Fact density	Retrieval	EXPERIMENT	▬▬▬▬▱	All three
06	Official source depth	Retrieval	AT SCALE	▬▬▬▬▱	ChatGPT, Gemini
07	Freshness	Retrieval	AT SCALE	▬▬▬▱▱	ChatGPT. Least on Google AI Overviews
08	Crawler access and rendering	Both	AT SCALE, OFFICIAL DOCS	▬▬▬▬▬	ChatGPT, Claude

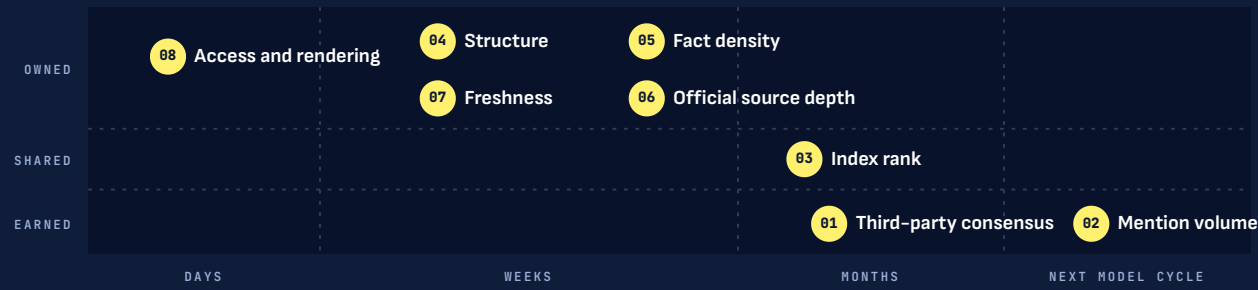
Strength is an editorial assessment of the evidence base: number of independent studies, sample size, and whether the result was observed or tested. It is not a published weighting from any vendor.

WHERE EACH FACTOR ACTS

Training	Fan-out	Retrieve	Read	Synthesize
Model learns brand and category associations	Model writes its own search queries	Index returns ranked candidate pages	Passages are extracted and scored	Agreement across sources picks the brands
01 02 08	02 06	03 07 08	04 05 08	01 06

Control against time to effect

Who controls the outcome, and how soon a change can show up in answers. Work left to right: access first, owned pages second, earned work continuously. Editorial placement.



Data, mechanism and work for each factor

01 Third-Party Consensus

AT SCALE EVIDENCE STRENGTH

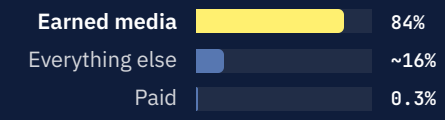
Answer engines name the brands that independent sources agree on. A claim repeated across a review site, a trade publication and a comparison list outweighs the same claim on your homepage.

THE DATA

- Earned media holds **84%** of citations across ChatGPT, Claude and Gemini, and has stayed between 82 and 89% over three editions. Paid and advertorial: 0.3%[7]
- "Best X" lists were **43.8%** of the source pages ChatGPT cited on 750 top-of-funnel prompts. Brands placed higher on those lists were recommended more often[5]
- Listicles earn 21.9% of AI citations overall and about **40%** on commercial queries[21]
- Ranking first in a frequently cited third-party list moved brands 0.8 to 1.8 positions earlier in AI answers[15]

Who AI cites

Share of 25M+ cited links, May 2026.[7]



DO THIS

- Run your category prompts and list the third-party pages that get cited
- Earn a top-three placement on those specific pages
- Keep review-platform profiles complete and current
- Pitch trade press with data, not announcements

02 Brand Mention Volume

CORRELATION EVIDENCE STRENGTH

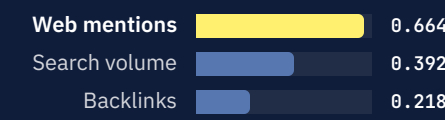
Models learn which brands belong to a category from how often the name appears beside category language. Reach across many sites matters more than authority concentrated on a few.

THE DATA

- Brand web mentions correlate with AI Overview visibility at **0.664**, the strongest of eleven factors tested. Backlink count: 0.218[2]
- Branded anchor text (0.527) and branded search volume (0.392) rank second and third. All three are off-site[2]
- Top-quartile brands for mentions earn up to **10x** more AI Overview mentions than the next quartile[2]
- Google tells site owners not to chase inauthentic mentions. Volume has to be earned[26]

Mentions against links

Correlation with AI Overview brand visibility.[2]



DO THIS

- Pair the brand name with one fixed category descriptor everywhere
- Prioritize placements that produce text: interviews, transcripts, guest articles
- Align facts across Wikidata, Crunchbase, LinkedIn and review profiles
- Track branded search volume as a leading indicator

03 Rank in the Underlying Index

AT SCALE EVIDENCE STRENGTH

A model can only cite what its search backend returns. Position in that backend, for the query the model actually wrote, is the gate to everything downstream.

THE DATA

- The top retrieved result is cited **58%** of the time in ChatGPT. By position ten: 14%[11]
- **79.2%** of Claude's cited URLs sit in Brave's top ten for the same query[18]
- Google's AI features draw only on pages that are indexed and eligible for a snippet[25]
- Domain authority and backlinks showed no positive correlation with citation once a page was retrieved[11]
- Brave follows Googlebot's access rules, so a page blocked to Googlebot is also out of Claude's reach[18]
- Fan-out queries are longer and more specific than typed searches: 6.7 words on average for Gemini 3[16]

Each platform, its own index

ChatGPT	OpenAI index, Google, Bing
Claude	Brave Search
Gemini	Google Search

DO THIS

- Check rankings in Google, Bing and Brave for comparison and "best" queries
- Target the phrasing models add: year, "official", "price", "reviews"
- Submit sitemaps to Google Search Console and Bing Webmaster Tools
- Fix indexing gaps before writing anything new

04 Semantic Alignment and Structure

AT SCALE EXPERIMENT EVIDENCE STRENGTH

Retrieved pages are read as passages. The passage that answers the sub-question directly, near the top and under a matching heading, is the one that gets quoted.

THE DATA

- **62.2%** of citation failures among 949 matched page pairs traced to semantic alignment. Content quality: 27.1%. Technical faults: 10.1%[29]
- **44.2%** of ChatGPT citations come from the first 30% of a page[9]
- Headings that closely match the query: cited 41% of the time, against 29% for weak matches[11]
- 68.7% of ChatGPT-cited pages use a sequential heading hierarchy, and 87% use a single H1[12]
- Pages of 500 to 2,000 words were cited most. Pages over 5,000 words were cited less than pages under 500[11]

Where on the page citations come from

Share of ChatGPT citations by position in the text.[9]



DO THIS

- Open with the answer in two or three sentences, then support it
- Phrase H2s as the questions buyers ask
- One page, one job. Split sprawling guides
- Use tables for comparisons and specifications

05 Fact Density

EXPERIMENT EVIDENCE STRENGTH

Specific, verifiable statements are easier for a model to lift and safer for it to repeat. This is the one factor with controlled experimental support.

THE DATA

- Adding statistics, citing sources and adding quotations each lifted visibility by **30 to 40%** across a 10,000-query benchmark[1]
- For sites ranked fifth, citing sources raised visibility by **115.1%**. Lower-ranked pages gain most[1]
- Keyword stuffing performed poorly[1]
- Cited passages are about twice as likely to contain a question mark, and they favor definitive wording and named entities[9]
- The effect varies by subject area, so test within your own category[1]

+40%

Upper bound of visibility lift from the best methods[1]

+115%

Lift from citing sources, for a site ranked fifth[1]

DO THIS

- Replace adjectives with numbers: prices, limits, timings, counts
- Link claims to primary sources
- Quote named people with titles
- State what the product does not do

06 Official Source Depth

AT SCALE EVIDENCE STRENGTH

After discovery, models verify. They go to the brand's own site for pricing, specifications and policy, and they are adding "official" to their searches more each month.

THE DATA

- **84.4%** of ChatGPT answers to factual questions cite at least one official source. For comparison questions: 68.3%[3]
- Official sites take 63.4% of citations on factual questions and 46.4% on comparisons[3]
- "Official" appeared in 13.1% and "price" in 8.9% of ChatGPT's English fan-out queries in September 2026[3]
- 75.9% of `site:` lookups among the most-searched sites target brand websites[3]
- 93% of Gemini citations come from sources organizations manage[19]

Pages models look for

PRICING · SPECIFICATIONS · INTEGRATIONS · SECURITY AND COMPLIANCE · SUPPORT DOCS · A VS B COMPARISONS · WHO IT IS FOR · WHO IT IS NOT FOR

DO THIS

- Publish real prices in text, not behind a form or in an image
- Give each use case, integration and compliance standard its own plain page
- Write honest comparison pages with concrete numbers
- Keep product facts identical across site, docs and listings

07 Freshness

AT SCALE EVIDENCE STRENGTH

Assistants cite newer pages than search engines rank, and they add the current year to their own queries. The size of the effect depends heavily on the platform.

THE DATA

- AI-cited URLs average 1,064 days old, against 1,432 for organic results: **25.7%** fresher[4]
- ChatGPT shows the strongest preference. Google AI Overviews show almost none[4]
- "2026" appeared in **22.9%** of ChatGPT fan-out queries, the most common added term[3]
- Pages not updated in a quarter were three times more likely to lose citations. Over 70% of cited pages were refreshed within a year[12]
- 79.1% of the "best" lists ChatGPT cited had been updated in 2025, the year of the study[5]

Average age of cited pages

Days since publication. Shorter is fresher.[4]



ChatGPT value derived from "393 days newer than organic."

DO THIS

- Review revenue pages every quarter and change what is out of date
- Show a visible last-updated date and matching `dateModified`
- Use the year in titles only where the content is truly current
- Never change a date without changing the content

08 Crawler Access and Rendering

AT SCALE OFFICIAL DOCS EVIDENCE STRENGTH

A page that cannot be fetched, or that arrives as an empty shell, is invisible however good it is. This factor is binary, which makes it the first thing to verify.

THE DATA

- None of the major AI crawlers tested rendered JavaScript: OAI-SearchBot, ChatGPT-User, GPTBot, ClaudeBot, PerplexityBot and others[22]
- GPTBot requested JavaScript files in 11.5% of fetches and ClaudeBot in 23.84%. Neither executed them[22]
- Gemini relies on Google's infrastructure, which does render[22]
- Sites that block OAI-SearchBot do not appear in ChatGPT search answers[23]
- Technical faults explain about **10%** of citation failures. Small, but total when it applies[29]

Two-minute test

```
curl -sL https://example.com/pricing/ \
  | grep -i "your price or key fact"
```

No match means AI crawlers cannot see it either.

DO THIS

- Server-render or statically generate all content that matters
- Review CDN and firewall bot rules, not just robots.txt
- Keep key facts out of modals, tabs loaded on click, and images
- See the crawler reference on page 17

What the data does not support

A good share of AI optimization advice has been tested and came back empty, or traces to a number that was never as firm as it sounded.

TACTIC OR CLAIM	WHAT THE EVIDENCE SHOWS	VERDICT
Publishing an llms.txt file	Three separate analyses, including one of 300,000 domains, found no meaningful link between having the file and being cited.[28] Google states it is not needed for its AI features.[26] Harmless to add. Not a lever.	NO MEASURED EFFECT
Adding schema to win citations	A matched test of 1,885 pages that added JSON-LD found no citation lift in ChatGPT or AI Mode, and a slight decline in AI Overviews.[6] Cited pages do carry schema more often (38.5% against 32.0% cited), which reflects well-built sites rather than cause.[11]	KEEP IT, EXPECT NOTHING
Keyword stuffing	Performed poorly in the one controlled experiment on generative engines.[1]	COUNTERPRODUCTIVE
Chunking or rewriting content "for AI"	Google says neither is required, and that its systems understand synonyms and multi-topic pages.[26] Clear structure helps readers and models alike. A separate AI voice does not.	UNNECESSARY
Very long comprehensive guides	Pages over 5,000 words were cited less often than pages under 500. The most-cited range was 500 to 2,000 words.[11]	COUNTERPRODUCTIVE
Building backlinks for AI visibility	Backlink count correlated at 0.218, a third the strength of brand mentions.[2] Once a page is retrieved, authority showed no positive relationship with being cited.[11] Links still matter for ranking in the index.	INDIRECT ONLY
Ranking yourself first on your own "best" list	These lists do appear in ChatGPT's sources today.[5] Google's guide warns against inauthentic signals, and a platform can discount self-ranked lists at any time.[26]	WORKS NOW, FRAGILE
Tracking your "AI rank"	The same prompt returns the same brand list less than 1% of the time, and order is close to random.[8] Position is noise. Appearance rate is signal.	MEANINGLESS METRIC
Optimizing for Google alone	Only 34% of the URLs Claude cites rank in Google's top ten for the same query.[18] Each platform reads a different index, and ChatGPT's source mix moved sharply within six months.[3]	INCOMPLETE
Allowing "Bravebot" or "Claude-Web"	Brave's crawler has no distinct user agent to allow.[18] Anthropic's documentation lists ClaudeBot, Claude-User and Claude-SearchBot. Claude-Web is not among them.[24]	OUTDATED

WHY SCHEMA STILL EARNS ITS PLACE

Structured data tells Google what an entity is. It feeds rich results and the Knowledge Graph, and it keeps product facts unambiguous. That supports the consistency behind Factors 02 and 06. It does not, by any controlled measure so far, cause a model to cite the page.

A TEST FOR ANY NEW TACTIC

- Does a platform's own documentation mention it?
- Has anyone measured citations before and after, with a control?
- How many data points sit behind the headline number?
- Who is selling the tool that measures it?

User agents, robots.txt and rendering

Search access and training access are separate switches on every platform. Decide each one on purpose.

TOKEN	OPERATOR	PURPOSE	ROBOTS.TXT	IF BLOCKED
`OAI-SearchBot`	OpenAI	Surfaces sites in ChatGPT search	Honored	Not shown in ChatGPT search answers. May still appear as a navigational link[23]
`ChatGPT-User`	OpenAI	User-initiated actions in ChatGPT and Custom GPTs	May not apply	Does not control Search inclusion[23]
`GPTBot`	OpenAI	Content for model training	Honored	Signals content should not be used for training[23]
`Claude-SearchBot`	Anthropic	Improves search result quality	Honored	May reduce visibility and accuracy in search results[24]
`Claude-User`	Anthropic	Fetches pages when a user asks	Honored	May reduce visibility for user-directed search[24]
`ClaudeBot`	Anthropic	Content for model training	Honored, with Crawl-delay	Future content excluded from training sets[24]
`Googlebot`	Google	Search index, including AI Overviews and AI Mode	Honored	Cannot be crawled for Search or its AI features[25]
`Google-Extended`	Google	Control for AI training and grounding in some other Google systems	Honored	No effect on Google Search[25]
Brave Search	Brave	Index behind Claude's web search	No distinct token	Follows Googlebot rules: blocked there, blocked here[18]

EXAMPLE: VISIBLE IN SEARCH, OPTED OUT OF TRAINING

```
# Search and user-directed access
User-agent: OAI-SearchBot
Allow: /

User-agent: Claude-SearchBot
Allow: /

User-agent: Claude-User
Allow: /

# Training: a business decision
User-agent: GPTBot
Disallow: /

User-agent: ClaudeBot
Disallow: /
```

One valid configuration, not a recommendation. Blocking training crawlers keeps content out of future models, which works against Factor 02. To maximize visibility, allow all of them.

BEYOND ROBOTS.TXT

- **Check the CDN and firewall.** Bot-management rules can block AI crawlers that robots.txt allows. Review the bot settings in Cloudflare or your host, and confirm with server logs
- **Verify by IP, not user agent.** OpenAI and Anthropic publish the IP ranges their crawlers use[23,24]
- **Allow time.** OpenAI says robots.txt changes take about 24 hours to register[23]
- **Repeat per subdomain.** Each needs its own robots.txt[24]

RENDERING CHECKLIST

- [ ] Key content present in the raw HTML response
- [ ] No facts that load only on click, scroll or after a consent modal
- [ ] Prices and specs as text, not images
- [ ] One H1, sequential headings, descriptive title
- [ ] Canonical URL returns 200 without a redirect chain

Tracking visibility when answers never repeat

Answers are drawn from a probability distribution. Measure how often a brand appears across many runs, not where it placed in one.

Visibility is a rate

Illustration: one prompt run 100 times. Each dot is one answer.



62%

Appearance rate: the brand was named in 62 of 100 answers. Position within each answer is not recorded, because it does not hold still.

<1%

chance that ChatGPT or Google's AI returns the same brand list twice for the same prompt.[8]

2,961

prompt runs by about 600 volunteers across ChatGPT, Claude and Google AI, each prompt run 60 to 100 times.[8]

WHAT THE STUDY FOUND

- Nearly every response was unique in three ways: which brands, in what order, and how many[8]
- Ranking position is unstable enough to be meaningless[8]
- The set of brands that shows up across many runs is far more consistent, even when people phrase the prompt differently[8]

That last point is what makes measurement possible. Presence across runs is stable enough to track. A single screenshot is not.

METRICS WORTH TRACKING

METRIC	WHAT IT TELLS YOU	HOW TO COLLECT IT
Appearance rate	Share of runs in which the brand is named, per prompt and per platform	Fixed prompt set, repeated runs, same cadence each month
Share of voice	Your appearances against named competitors on the same prompts	Same runs, counted by brand
Cited but not chosen	Prompts where your page is a source and the brand is still left off the list	Compare citation URLs with the brands named
Cited sources	Which third-party pages the model leans on for your category	Log every cited URL. Rank by frequency
AI referral sessions	Visits arriving from chatgpt.com, claude.ai, gemini.google.com and similar	GA4 referral source, grouped into one channel
Crawler activity	Whether search and user fetchers reach your key pages	Server logs filtered by the tokens on page 17
Branded search volume	Leading indicator for Layer 1 (correlation 0.392)[2]	Google Search Console, brand queries

BUILD THE PROMPT SET

Twenty-five to fifty prompts written the way buyers ask: by use case, company size, budget and integration.

RUN IT REPEATEDLY

Several runs per prompt, per platform, on a fixed schedule. Report rates with the number of runs behind them.

KEEP PLATFORMS SEPARATE

A blended score hides the pattern. A brand can lead on Gemini and be missing on Claude for reasons page 08 explains.

Access, then owned, then earned

The sequence follows the evidence: remove blockers, make your own pages the best source of fact, then build the outside consensus that gets you considered.

DAYS 1 TO 30

Access and Baseline

FACTOR 08 · FACTOR 03

- [ ] Fetch key pages without JavaScript and confirm the content is there
- [ ] Audit robots.txt and CDN bot rules against page 17
- [ ] Confirm indexing in Google, Bing and Brave
- [ ] Write the prompt set and record baseline appearance rates on all three platforms
- [ ] Log every third-party URL the models cite for your category
- [ ] Set up an AI referral channel in GA4

DAYS 31 TO 60

Owned Pages

FACTORS 04 TO 07

- [ ] Rebuild pricing, product and comparison pages answer-first
- [ ] Put exact prices, limits, integrations and compliance facts in plain text
- [ ] Add a use-case page for each segment you serve, and say who it is not for
- [ ] Rewrite headings as buyer questions
- [ ] Add sourced statistics and named quotes to top articles
- [ ] Show last-updated dates and set a quarterly review

DAYS 61 TO 90

Earned Consensus

FACTORS 01 AND 02

- [ ] Pursue placement on the cited lists logged in month one
- [ ] Complete and refresh review-platform profiles
- [ ] Pitch two or three data-led stories to trade press
- [ ] Book podcast or video appearances that publish transcripts
- [ ] Align the brand description across every profile and directory
- [ ] Re-run the prompt set and compare with baseline

OFF-SITE WINS DISCOVERY

Open-ended prompts such as "best payroll tool for restaurants" are answered from third-party lists, reviews and press. If the brand is absent there, it is rarely considered, however good the website.

COMPARISON LISTS · REVIEW PLATFORMS · TRADE PRESS · TRANSCRIPTS

OWNED WINS VERIFICATION

Once a brand is in the running, the model checks price, features and fit against the official site. Missing or vague facts hand the detail, and often the recommendation, to a competitor.

PRICING · SPECS · COMPARISONS · DOCS

AFTER DAY 90

Repeat the measurement monthly, the page review quarterly, and the earned work continuously.

CADENCE	TASK
Monthly	Re-run the prompt set. Review cited sources and crawler logs
Quarterly	Refresh revenue pages. Re-check each platform's index and source mix, which shift with model releases
Ongoing	List placements, reviews, press and transcripts

Sources

01 **Aggarwal et al.** GEO: Generative Engine Optimization. KDD 2024. arxiv.org/abs/2311.09735

02 **Ahrefs.** An Analysis of AI Overview Brand Visibility Factors (75K Brands Studied). ahrefs.com/blog/ai-overview-brand-correlation

03 **Ahrefs.** How ChatGPT Does Query Fan-Out (15 Million Queries Studied). October 2026. ahrefs.com/blog/query-fan-out-chatgpt

04 **Ahrefs.** AI Assistants Prefer to Cite "Fresher" Content (17 Million Citations Analyzed). ahrefs.com/blog/do-ai-assistants-prefer-to-cite-fresh-content

05 **Ahrefs.** Do Self-Promotional "Best" Lists Boost ChatGPT Visibility? Study of 26,283 Source URLs. ahrefs.com/blog/best-lists-research

06 **CMSWire.** Does Structured Data Actually Improve Answer Engine Citations? Reports the Ahrefs matched study of 1,885 pages. September 2026. cmswire.com

07 **Muck Rack, Generative Pulse.** What Is AI Reading? May 2026 edition. generativepulse.ai/whatisaireading

08 **SparkToro and Gumshoe.** AIs are highly inconsistent when recommending brands or products. January 2026. sparktoro.com/blog

09 **Kevin Indig, Growth Memo.** The science of how AI pays attention. February 2026. searchenginejournal.com

10 **AirOps.** The Influence of Retrieval, Fan-Out and Google SERPs on ChatGPT Citations. airops.com/report

11 **AirOps.** The Fan-Out Effect. 16,851 queries, 353,799 pages. airops.com/lp/the-fan-out-effect-report

12 **AirOps.** The 2026 State of AI Search, and Structuring Content for LLMs. airops.com/report/the-2026-state-of-ai-search

13 **Peec AI.** Fan-outs per prompt by engine. 5 million fan-outs, April 2026. docs.peec.ai/research

14 **Peec AI.** ChatGPT built its own search index. 2026. peec.ai/blog/chatgpt-built-its-own-search-index

15 **Peec AI.** The Listicle Rank Effect. peec.ai/blog

16 **Seer Interactive.** Initial Research: Gemini 3 Query Fan-Outs. seerinteractive.com/insights

17 **Profound.** What is Claude Web Search? March 2025. tryprofound.com/blog

18 **Xponent21.** Anthropic Lists Two Web-Search Subprocessors Behind Claude. 2026. xponent21.com/insights

19 **Yext Research.** AI Citation Behavior Across Models: Evidence from 17.2 Million Citations, with Spring 2026 update. yext.com/research

20 **BrightEdge.** Where AI Engines Agree on Brands. May 2026. brightedge.com/resources

21 **Wix Studio AI Search Lab.** Content types most cited by LLMs. 75,000 answers. wix.com/studio/ai-search-lab

22 **Vercel and MERJ.** The rise of the AI crawler. December 2024. vercel.com/blog/the-rise-of-the-ai-crawler

23 **OpenAI.** Overview of OpenAI Crawlers. developers.openai.com/api/docs/bots

24 **Anthropic.** Does Anthropic crawl data from the web, and how can site owners block the crawler? support.claude.com/en/articles/8896518

25 **Google Search Central.** AI features and your website. developers.google.com/search/docs/appearance/ai-features

26 **Google Search Central.** Optimizing your website for generative AI features on Google Search. May 2026. developers.google.com/search/docs/fundamentals/ai-optimization-guide

27 **Search Engine Land.** ChatGPT's answers came from Google Search after all: Report. August 2025. searchengineland.com

28 **SE Ranking.** llms.txt: What It Is, What the Data Says, and Whether You Need One. visible.seranking.com/blog/llms-txt

29 **Tian et al.** Diagnosing and Repairing Citation Failures in Generative Engine Optimization. March 2026. arxiv.org/abs/2603.09296

METHOD AND LIMITS

Research compiled in November 2026 from published studies and vendor documentation. Figures are reported as published, with sample and setup noted where they change the reading. Most studies are observational and many come from companies that sell AI visibility software. Sources 1 and 29 tested simulated engines, not commercial products. Sources 6 and 7 are cited through secondary reports or press releases, and the failure percentages attributed to source 29 come from secondary coverage of that paper. Platforms change retrieval behavior frequently. Verify anything decision-critical against the linked source before acting on it.

ABOUT THE AUTHOR

Kevin Leary

Fractional WordPress, AEO and AI consultant and hands-on developer in Boston, working directly with growth-focused companies since 2008.

{ AEO Consulting }

{ Work With Me }