

8%

of ChatGPT's citations come from pages that rank in Google's **top 10**. Being found by AI has quietly become its own channel.

The AI Search *Visibility* Playbook

Why being found in ChatGPT is *a different game* than ranking on Google, and what to actually do about it.

Contents

00	The one-paragraph version	03
01	The finding: AI and Google visibility have split apart	05
02	How AI Search actually works (the plumbing)	08
03	What actually drives AI citations (the evidence)	11
04	The practical AEO / GEO playbook	14
05	How to measure AI visibility	17
06	The balanced verdict: is it "just SEO"?	20

Being everywhere in ChatGPT and *invisible* on Google is now normal.

A brand can be cited constantly in ChatGPT, Claude and Perplexity while showing zero presence in Google Search Console. That is not a glitch; it is how the ecosystem now works. Independent analyses find only about 8% of ChatGPT and Gemini citations come from pages ranking in Google's top 10 ([Ahrefs](#)). AI assistants build answers from different indexes (Bing, Brave, their own crawlers) using different signals, brand mentions and claim-level relevance rather than backlinks.

So the old reassurance, "if you rank in Google you are visible in AI," is no longer safe. But the opposite over-correction is also wrong: fresh-content AI citations decay in about one to two months, AI is still under 1% of referral traffic, and Google still owns about 90% of search. Treat AI visibility as a measured, deliberately optimized channel of its own that sits alongside SEO, not instead of it.

800M

ChatGPT weekly users, late
2025 (OpenAI)

+43%

AI search visits YoY, Q1'25 to
Q1'26 (Cloudflare)

37%

of consumers now start a search
with an AI tool

The five ways the game is different.

01

Different index. ChatGPT leans on Bing and its own crawler, Claude uses Brave, Perplexity runs its own. Only Google AI Overviews behave like the search results you know, so a Google ranking does not carry over.

02

Different signals. Branded mentions correlate with AI visibility far more strongly than backlinks. Reddit, Wikipedia and YouTube dominate citations, so your presence off your own site can matter more than your homepage.

03

Different format wins. Listicles, comparison tables and answer-first passages get lifted into answers far more than long prose. Models retrieve at the passage level, so each section has to stand on its own.

04

Different failure mode. Most AI crawlers do not run JavaScript, so a page that ranks #1 on Google can be completely invisible to them. Server-rendered HTML is what makes you readable.

05

Different scoreboard. Google Search Console cannot see any of it. AI visibility needs its own measurement, tracking citations, share of voice and sentiment.

01

PART ONE

The finding: AI and Google visibility have split apart

A 100% AI-generated site can flood ChatGPT with citations and register nothing in Google. What that proves, and what it does not.

The finding


What the Surferstack experiment showed

The experiment that prompted this guide built a 100% AI-generated listicle site ("Surferstack"), connected its CDN to an AI-traffic monitor (Promptwatch), and watched what happened. The result: high-volume citations from ChatGPT, Claude and others almost immediately, alongside zero presence in Google Search Console. Bing correlated with the AI citations far more closely than Google ever did.

“Being visible in ChatGPT does not require being visible in Google.”

That headline holds up well against independent data. But the experiment mixes one robust finding with a couple of effects that need honest caveats. This guide separates them.

How little AI citations overlap Google's top 10



Share of each engine's cited URLs that rank in Google's top 10 for the same prompt. Standalone assistants barely overlap with Google; only Google's own AI Overviews do.

Source: [Ahrefs, AI vs Google overlap](#)

What holds up

- **The channels are genuinely independent.** Only about 8% of ChatGPT and 8.6% of Gemini citations come from URLs ranking in Google's top 10; Perplexity is closest at about 28%. For the standalone assistants, a Google ranking is a weak predictor of being cited.
- **Bing matters more than Google for chat assistants.** ChatGPT's lineage runs through Bing and Copilot reads Bing's index. Caveat: Perplexity, Claude and Gemini do not use Bing, and ChatGPT now has a Google fallback, so do not overstate this as "the whole AI stack runs on Bing."
- **AI weights different signals.** Branded web mentions correlate with AI visibility about 3x more strongly than backlinks, and nofollow links correlate almost as well as follow links. Correlations, not proven causation.

The honest caveats

"Google just penalizes AI listicle spam."

Largely true. Google's March 2024 scaled content abuse policy targets mass-generated pages regardless of author, so a 100% AI listicle site being absent from Google is the designed outcome.

"The AI boost is fragile." Supported. One large study found a median citation half-life of about 4.5 weeks. Fresh articles get boosted,

then fade within a month or two unless reinforced.

"You still need something to index it."

Correct. A crawler (usually Bing's) has to find the page first. The independence is between Google and AI, not between AI and all indexing.

"Cited is not recommended." The most important caveat. A model reading or linking your page is different from recommending your brand to a buyer. More in Part 3.

02

PART TWO

How AI Search actually works

Each assistant reads the web through a different index, using different bots, and most of those bots never run JavaScript.

The plumbing

You cannot optimize a system you do not understand. The single biggest reason AI and Google visibility diverge is that each assistant reads the web through a different index, using different bots, and most of those bots do not render JavaScript.

Where each assistant gets its answers

ASSISTANT	WHERE IT GETS ANSWERS	CLOSEST TO
ChatGPT	Own crawler plus a Bing lineage; a Google fallback is emerging and OpenAI is building its own index	Bing SEO
Claude	Brave Search	Structure and Brave
Perplexity	Its own index	Closest to Google (~28%)
Google AI Overviews	Google Search	Traditional SEO
Copilot	Bing	Bing SEO

There is no single "AI index." Optimizing for Google AI Overviews is close to SEO; for Claude it is a Brave and structure problem; for ChatGPT it is largely a Bing-and-own-crawler problem.

The JavaScript rendering gap: the silent killer

A landmark [Vercel and MERJ study](#) (500M+ fetches) found the major AI crawlers fetch raw HTML and do not execute JavaScript. Googlebot and Gemini, by contrast, do render it.

Renders JavaScript

Googlebot
Google Gemini
AppleBot

Does NOT run JavaScript

GPTBot / OAI-SearchBot
ClaudeBot
PerplexityBot

“A page can rank #1 on Google and be completely blank to every AI crawler.”

The fix is server-side rendering (SSR, SSG or ISR) or prerendering, so meaningful content is in the initial HTML. [IndexNow](#) (supported by Bing and Yandex, not Google) then pushes new URLs so they become quotable faster.

03

PART THREE

What actually drives AI citations


Brand mentions beat backlinks, third-party sources beat your homepage, and structure beats prose. The evidence, charted.

The evidence

Brand mentions beat backlinks

The headline shift from SEO: across 75,000 brands, branded web mentions correlate with AI visibility about 3x more strongly than backlinks. **Semrush** found nofollow links correlate almost as well as follow links. Models process the fact that you are mentioned, not link equity.

What correlates with AI visibility (vs backlinks)



Spearman correlation with brand presence in AI answers across 75,000 brands. Branded mentions correlate roughly 3x more strongly than backlinks. Correlation, not causation.

Source: [Ahrefs, 75,000-brand study](#)

Third-party sources dominate, not your own site

Across roughly 30M cited sources, Reddit, Wikipedia, YouTube, LinkedIn and Forbes lead. Wikipedia is about 48% of ChatGPT's most-cited reference sources; Reddit is heavily favored by Perplexity. Your presence off your own domain can matter more than your homepage.

Structure wins

Most-cited content formats



Share of all AI citations by format across ~25,000 URLs. Ranked 'Top N' lists made up the majority of cited listicles.

Source: [Evertune \(via Search Engine Land\)](#)

Answer-first. Lead each section with a direct two-to-three-sentence answer, then expand. Models extract the first sentences after a heading.


Self-contained chunks. Retrieval happens at the passage level (100 to 300 words). Each section should stand alone.

Stats, quotes, sources. The peer-reviewed [Princeton GEO study](#) found statistics (~+41%), quotations and citations each lifted visibility by 30 to 40%.

Keyword stuffing backfires. The same study found it performed worse than baseline.

The FAQ myth. [Omniscient's analysis](#) of 23,387 cited sources found standalone FAQ pages were cited only **0.41%** of the time for branded queries, while reviews and social proof drove about **57%**. The Q&A format inside your content helps; dedicated FAQ pages as a citation strategy do not.

Cited is not crawled is not recommended



04

PART FOUR

The practical AEO / GEO playbook

Fix the foundation first. There is no point optimizing content a crawler cannot read.

The playbook

1 · Technical foundation (do these first)

- Serve meaningful content in the initial HTML (SSR, SSG or prerendering). Disable JavaScript to test: if content vanishes, AI crawlers cannot see it.
- Audit robots.txt: allow the search and live-fetch bots (OAI-SearchBot, PerplexityBot, ChatGPT-User, Claude-User) even if you block training bots.
- Put critical facts (prices, specs, dates, definitions) in plain text and schema, not only JavaScript widgets or images.
- Adopt IndexNow to push new and updated URLs to Bing instantly.
- Keep crawl efficiency high: fix 404s, clean sitemaps, fast responses.

2 · Content structure

- Lead every section with a direct, self-contained answer, then support it.
- Use listicles, ranked "Top N" formats, comparison tables and a clear H2/H3 hierarchy.
- Write passage-level: each section should make sense lifted out on its own.
- Add original statistics, direct quotes and cited sources.
- Name entities explicitly and consistently so models attribute correctly.
- Build genuine comparison and "best X for Y" content.

3 · Off-site and distribution (often higher-leverage)

- Earn brand mentions, linked or unlinked, across credible sites.
- Show up where AI actually cites: Reddit, YouTube, LinkedIn, Wikipedia-eligible references, reputable roundups.
- Invest in digital PR and original data others repeat; repetition across trusted sources makes citations durable.

4 · Sustain it

- Refresh cornerstone content regularly; individual citations decay in one to two months.
- Do not rely on disposable AI-generated pages; they risk Google's scaled-content-abuse enforcement.
- Reinforce the same claims across multiple owned and third-party surfaces.

05

PART FIVE

How to measure AI visibility

Google Search Console cannot see any of it. You need a separate stack, tracking three layers.

Measurement

Google Search Console cannot see ChatGPT, Claude, Perplexity or Copilot. Even Google's own [Generative AI report](#) (launched June 2026) shows impressions only, with no clicks, CTR or query data.

What to measure: three layers

- **Crawl.** Which AI bots fetch your pages and whether they hit errors. Source: CDN and server logs.
- **Citation.** How often you are named or linked as a source, and for which prompts.
- **Recommendation and sentiment.** What models actually say about you. This is the layer that maps to revenue.

Citation half-life by engine (weeks)



Median time before half of a page's citations disappear. Overall median is 4.5 weeks. One large vendor study (3.5M citation events); directional.

Source: [Scrunch / Stacker](#)

Tools that track it

TOOL	WHAT IT TRACKS
Promptwatch	Mention rate, share of voice, citation sources, sentiment, plus crawl analytics
Profound	Enterprise answer share-of-voice, multi-country, competitor tracking
Peec AI / Otterly / Scrunch	Prompt and mention monitoring, share-of-answer
Semrush / Ahrefs	AI share of voice layered onto existing SEO suites

Core dna: a CMS built for AI discovery. Most AI crawlers never run JavaScript, so pages that depend on it are invisible to them. **Core dna** serves fully server-rendered HTML with structured content and clean markup, so your pages are readable and quotable by every AI engine by default. Built-in IndexNow support pushes new and updated URLs the moment they publish.

Do citations convert?

Early vendor data claims a premium over organic (figures range widely, 4x to 23x), but absolute volume is still around 1% of traffic and only about 14% of marketers track it separately. The near-term payoff is brand inclusion and recommendation. Measure share-of-voice and sentiment now; measure pipeline as volume grows.

06

PART SIX

The balanced verdict

It is not a replacement for SEO, and it is not 'just SEO.' It is a distinct, measurable channel on familiar foundations.

Is it "just SEO"?

Where the skeptics are right

- If a page cannot be crawled, indexed and read, it cannot be cited.
- Much of "AEO" is Bing SEO plus content hygiene good SEOs have done for years.
- AI is still about 1% of referral traffic and Google holds about 90% of search. Abandoning SEO is premature.

Where it is genuinely new

- A different index and bots (Bing, Brave, no JS rendering) mean Google success does not transfer.
- Off-site signals outrank backlinks as predictors.
- New metrics: citation share, share of voice, sentiment, none visible in GSC.
- New behaviors: query fan-out, passage-level retrieval, agentic commerce.

“Keep the SEO fundamentals that make you readable. Layer on the presence, structure and measurement that make you recommendable.”

Verdict. AI visibility is a distinct, measurable channel built on familiar foundations but with its own indexes, signals, formats and scoreboard. The winning posture is to do both. Treat it as its own line item, not a footnote under SEO.

core dna