GEO FOUNDATIONS

Generative Engine Optimization: getting retrieved, cited, and recommended inside AI answers from ChatGPT, Perplexity, Claude, Gemini, and Google's AI surfaces. The mechanics, the evidence, and what nobody actually knows yet.

SELWYN UY · SELWYNUY.DEV

EDITION 2026.07

CONTENTS

IN THIS GUIDE

01 — THE SHIFT

THE NEW CONSIDERATION SET

IN PLAIN WORDS

When someone asks an AI for the best option, it hands back a short list of names. If your client's brand is not on that list, the buyer never even hears about them, even when they would be a perfect fit. This section is about why getting your brand onto that list is the whole game now.

Buyers increasingly ask an assistant "what's the best X" and receive a synthesized shortlist. If a brand is not in the model's consideration set or the retrieval layer's citation pool, it is invisible at the moment of highest intent. The traffic side is already measurable: Semrush's clickstream analysis found ChatGPT's outbound referral traffic grew 206% during 2025 even as its own visits plateaued near a billion per month, and 30% of all ChatGPT referral traffic flows to just ten domains. Pew Research found users click a link inside an AI summary on just 1% of visits: the mention itself is most of the value.

0.664

Spearman correlation of branded web mentions with AI Overview visibility across 75K brands, vs 0.218 for backlinks (Ahrefs, 2025).

+206%

Growth in ChatGPT outbound referral traffic during 2025 (Semrush clickstream, 2026).

55-77%

Stable inclusion rate of top brands across repeated AI runs, while exact lists repeat less than 1 in 100 times (SparkToro/Gumshoe, 2026).

GEO matters not because it replaces SEO but because a growing slice of demand is mediated by generative engines whose selection mechanics differ from ranking algorithms in documented, exploitable ways. This guide covers those mechanics, the evidence for what moves them, and the discipline needed to not fool yourself.

02 — THE MAP

ONE FOUNDATION, THREE LAYERS

IN PLAIN WORDS

Getting picked by an AI happens in three stacked steps. First your page has to be findable at all, then a specific piece of it has to be the exact part the AI quotes, and finally your brand name has to be one it recommends out loud. Each step sits on top of the one below it, so you cannot skip ahead.

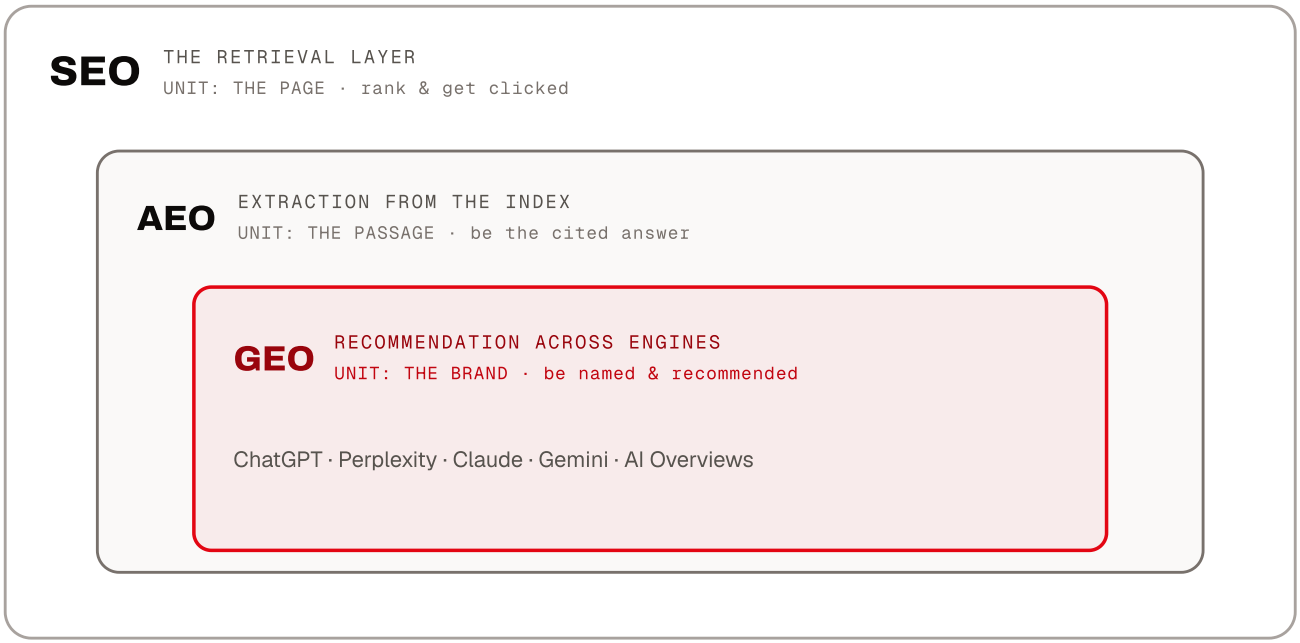

GEO is the outermost layer: it needs SEO's retrieval and AEO's citable passages, then adds the brand and training-data layer neither touches.

SEO, AEO, and GEO are concentric layers on one shared foundation, not rival disciplines. **SEO** is the retrieval layer everything else sits on: every major answer and generative engine retrieves from a conventional crawler-built index. Google's index feeds AI Overviews, AI Mode, and Gemini; Bing's feeds ChatGPT search and Copilot; Perplexity runs its own crawler. **AEO** sits on top and changes the unit of optimization from the page to the passage: being the specific chunk an answer surface extracts and attributes. **GEO** sits on top of both and changes the unit again, from the passage to the brand: what generative engines already know about you, plus earned presence in the third-party sources they retrieve and cite.

What is genuinely new in GEO: a second battlefield (training-time presence, which classic SEO never touched) and a new measurement regime (non-deterministic answers require sampled visibility rates instead of rank tracking). What is not new: the prerequisites. Google's May 2026 guidance calls all of this "still SEO", and mechanically it is right.

READ THIS GUIDE IF

You need brands and LLMs to recommend you. If you first need to be retrievable and ranked, read **SEO Foundations**. If you rank but are not the extracted answer, read **AEO Foundations**. All three at selwynuy.dev/guides.

03 — THE MACHINE

THE CITATION PIPELINE

IN PLAIN WORDS

This walks through the steps an AI takes between reading a question and writing its answer. Sometimes it answers from memory, sometimes it searches the web first, and it usually breaks one question into several smaller searches. A citation is just the AI naming its source, and knowing these steps shows you exactly where your brand can slip in.

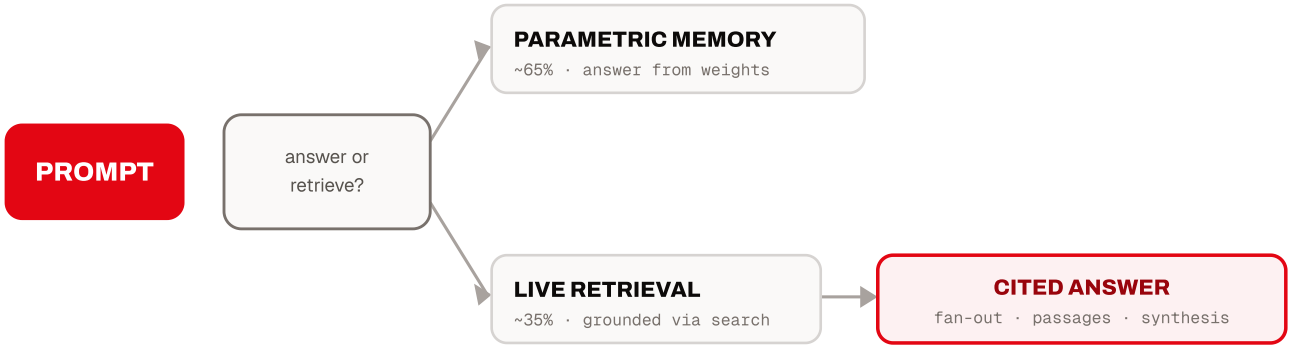

Every answer forks first: reply from training memory (about 65% of ChatGPT queries) or retrieve live and cite. GEO fights on both paths.

Every major AI answer engine is a retrieval-augmented generation (RAG) system layered on a conventional search index. GEO is impossible to practice well without the four stages.

Stage 1: Answer or retrieve

The model first decides whether to answer from parametric memory (knowledge baked into weights at training time) or ground the answer with live retrieval. Semrush's clickstream data shows ChatGPT enabled web search on only 34.5% of queries as of February 2026, down from 46% in late 2024. The split matters enormously: for memory-led answers only your training-data footprint counts; for grounded answers your retrievability counts.

Stage 2: Query fan-out

The engine rewrites the user's prompt into multiple synthetic sub-queries. Google's documentation confirms AI Overviews and AI Mode "may use a query fan-out technique, issuing multiple related searches across subtopics and data sources". Search Engine Land's reverse-engineering of ChatGPT's internal web.run tool found newer models run five to ten-plus rounds of search per response, refining queries from earlier results, and formulating queries around sources the model already knows. Their line worth memorizing: a brand absent from parametric memory won't even be considered. You are optimizing for a cloud of machine-written queries, not one keyword.

Stage 3: Index dependency

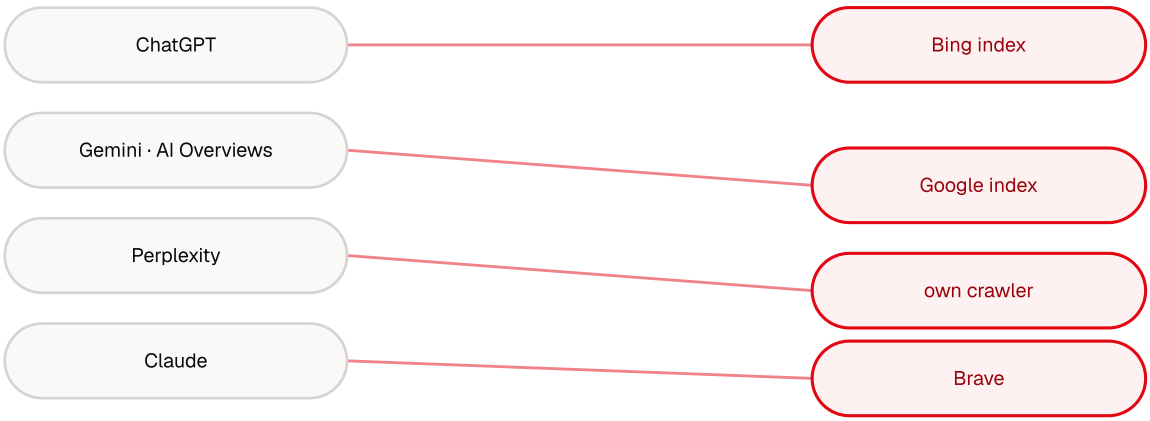

Each engine rides a specific index. Absent from Bing means absent from ChatGPT, whatever your Google rank. Verify in Bing Webmaster Tools and adopt IndexNow.

Each engine rides a specific index: ChatGPT retrieves via Bing-derived and third-party search APIs, then fetches selected pages with ChatGPT-User; Gemini, AI Overviews, and AI Mode use Google's index; Perplexity runs PerplexityBot; Claude's web search has used Brave. Consequence: a site absent from Bing effectively does not exist for ChatGPT search, whatever its Google rankings. Bing Webmaster Tools and IndexNow are first-class GEO infrastructure. The foundations each index requires are the SEO guide's territory.

Stage 4: Passage selection and synthesis

Retrieved pages are chunked into passages; the model selects, compresses, and attributes. Citation slots are scarce and shrinking: after a ChatGPT model update, unique domains cited per response fell from 19 to 15 (Search Engine Land). The unit of competition is the passage, not the page. A self-contained, quotable block beats a great page whose answer is smeared across two thousand words.

04 — THE EVIDENCE

THE GEO PAPER AND ITS LIMITS

IN PLAIN WORDS

There is one solid research study behind all of this, and it found that adding real quotes, real numbers, and links to sources makes an AI more likely to feature your content. Cramming in repeated keywords actually made things worse. Just remember the study was run in a lab, so treat its findings as a strong hint, not a guarantee.

The discipline's name comes from the one piece of peer-reviewed evidence: Aggarwal et al., "GEO: Generative Engine Optimization", ACM KDD 2024. It tested nine content modifications across GEO-bench (10,000 queries) and found visibility gains of up to 40%. Tactic-level results, measured as position-adjusted word count: quotation addition +41% over baseline, statistics addition +33%, citing sources +28%. Keyword stuffing dropped visibility below baseline, the single most quotable anti-habit finding. And the equity result: citing sources lifted rank-5 sites' visibility 115.1% while rank-1 sites lost 30.3%. GEO is highest-leverage for challengers.

THE LIMITATION, STATED EVERYWHERE THE PAPER IS CITED

The experiments ran on 2023-era models in a sandboxed RAG setup, not live commercial engines; visibility measured answer share, not traffic or revenue; and tactics were applied to sources already retrieved, so the paper says nothing about winning retrieval itself. Treat it as proof that content features causally affect citation in RAG systems, not as a guaranteed percentage.

The evidence hierarchy

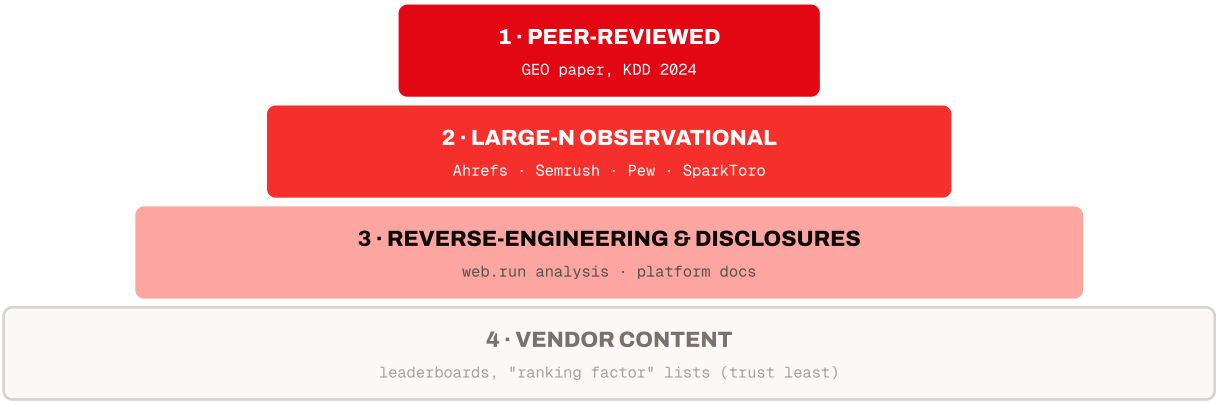

Weigh every claim by its tier: strongest and rarest at the top, weakest and most common at the base. Most published GEO content is tier 4 quoting tier 4.

Apply it ruthlessly: every tactical claim gets tagged to a tier, and tier-4 numbers are quoted only with named methodology. Most published GEO content is vendors quoting each other.

05 — THE CRAFT

WHAT MAKES A PASSAGE CITABLE

IN PLAIN WORDS

A passage is just a chunk of your writing that an AI might lift out and use in its answer. To make it easy to lift, open each section with a short answer that stands on its own, back it with real numbers and named sources, and say your brand name plainly. If a chunk only makes sense when it is glued to the rest of the page, the AI will pass it over.

The on-page directives, translated from the validated evidence (the full question-first architecture is the AEO guide's territory):

- A direct, self-contained answer of roughly 40 to 60 words leading every section, written to survive being lifted out of context: subjects restated, no pronoun dependence.
- Specific, sourced numbers ("converts at 7.1%" beats "converts well"), named-expert quotations, and outbound citations to primary sources. These are the three empirically validated boosts.

- Structure for retrieval: descriptive H2s and H3s phrased as questions, tables for comparisons, tight lists, one idea per paragraph. Formatting is retrieval engineering, not decoration.
- Entities by name: your brand, product, and category in the citable passages themselves. Models attribute to named entities, not to "we".
- Freshness: dateline content and update it. Retrieval favors recent documents for time-sensitive fan-out queries.

One honest complication: Google's May 2026 guide says publishers "can ignore tactics like chunking content" because its systems understand multi-topic pages. Keep the self-contained sections anyway, framed correctly: it is clear writing, and non-Google pipelines still retrieve chunks. What the citation-share data adds: Reddit, Wikipedia, YouTube, and LinkedIn dominate cited domains across engines, and 30% of ChatGPT referral traffic flows to ten domains. Being cited *about* on high-trust third-party surfaces is often easier than getting your own domain cited, which is the next section's subject.

06 — THE TWO LAYERS

TRAINING DATA VS RETRIEVAL

IN PLAIN WORDS

An AI can know your brand in two ways. Either it learned about you when it was first built, or it looks you up fresh the moment someone asks. You help the first by keeping your facts the same everywhere online, and you help the second by earning other trusted websites to mention you by name. This section shows you how to work both.

Layer 1: Parametric presence

Models answer roughly 65% of ChatGPT queries without live search (Semrush). What a model knows about a brand comes from its training corpus: Wikipedia, news archives, Reddit and forums, review sites, GitHub, the general crawl. You cannot inject yourself into a frozen model, but you can influence the next snapshot and, more immediately, the model's query-writing behavior, since engines formulate retrieval queries around entities they already know. The levers: a factually consistent Wikipedia and Wikidata presence where notability genuinely supports it; identical entity facts everywhere (same name, category description, and positioning across your site, LinkedIn, Crunchbase, G2, directories); earned coverage in publications and communities that are demonstrably in training corpora.

Layer 2: Retrieval presence

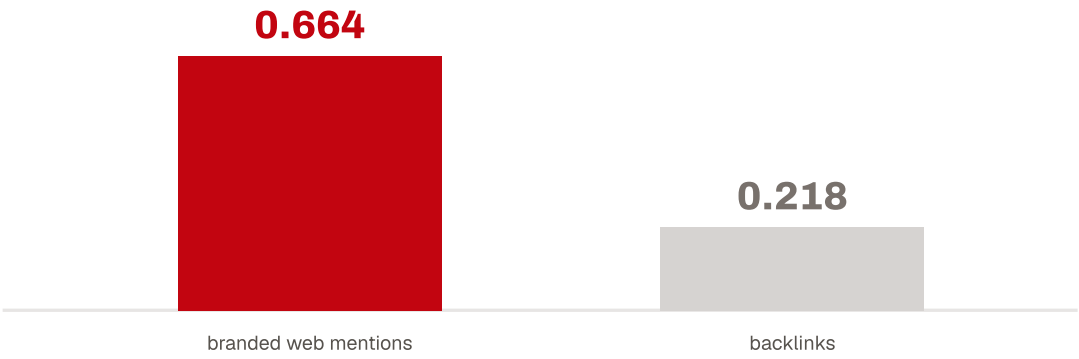

How strongly each signal correlates with AI Overview visibility across 75,000 brands (Spearman). Mentions beat backlinks roughly three to one (Ahrefs, 2025).

For grounded answers the battle is being in the index, matching fan-out queries, and being quotable (previous section). The anchor dataset for where to spend: Ahrefs' 75,000-brand study found branded web mentions correlate with AI Overview visibility at 0.664 versus 0.218 for backlinks, and top-quartile brands for mentions averaged 169 AI mentions versus 14 for the next quartile.

- Reallocate digital PR from link acquisition to named mentions in credible, crawlable publications. A mention with your brand next to category language is the unit of value.
- Be genuinely present on Reddit, industry forums, and YouTube. Astroturfing risk is real: platforms police it, models cross-check, and Google's May 2026 guide now explicitly warns against "pursuing inauthentic mentions".
- Seed best-of and comparison coverage on third-party sites; engines prefer citing neutral-seeming lists for recommendation queries.
- Maintain a machine-consistent fact base: a plain-language About page, sameAs schema linking official profiles, identical descriptions across profiles. One coherent entity for both training and retrieval to encounter.

The mention system

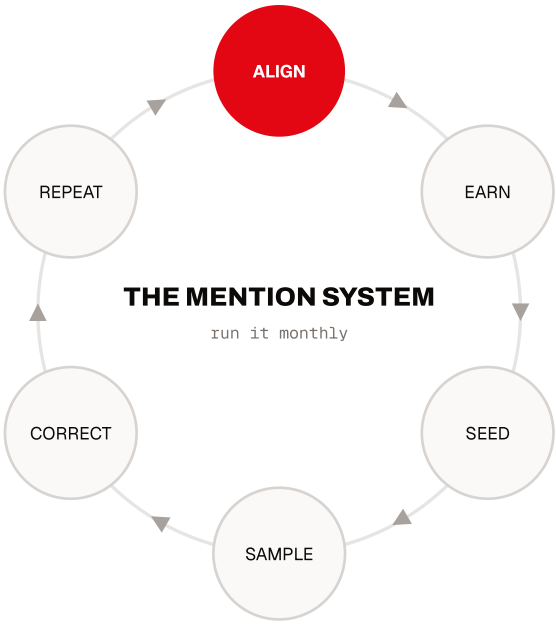

A monthly cycle: align your entity facts everywhere, earn a named mention in a crawlable publication, seed comparison coverage, sample AI visibility, correct what the models get wrong, and repeat. Brand presence is built, not shipped once.

GEO's off-site work only compounds if it recurs. Each pass is a handful of concrete tasks: reconcile any drifted entity fact, land one credible mention next to your category language, add or refresh a third-party comparison, run the prompt panel, and publish a correction for any hallucinated fact. None of it is a campaign; it is a standing routine a VA can own.

The monthly GEO checklist

MONTHLY TASK	WHY IT MATTERS
Ask each engine "what is [brand]" and log the answer	Catches drift and hallucinated facts early
Reconcile any fact that differs across your profiles	One coherent entity feeds both training and retrieval
Earn one named mention in a crawlable publication	Mentions correlate with AI visibility 3x more than links
Add or refresh one third-party comparison / listicle	Engines prefer citing neutral lists for "best X"
Run the fixed prompt panel across engines	Repeated sampling turns noise into a trend
Publish a correction for any wrong fact an engine stated	Authoritative pages are how you fix the record
Confirm your quotable pages are in Bing's index	Absent from Bing means absent from ChatGPT search

07 — PREREQUISITES

TECHNICAL READINESS

IN PLAIN WORDS

AI companies send out little reader programs, called crawlers, that scan the web, and some of them are the ones that let an AI find and quote you. This section makes sure you have not accidentally shut the door on the readers you actually want, and that your pages can be read without any fancy loading tricks. Small settings here can quietly make you invisible.

The crawler taxonomy (canonical for the set)

CLASS	AGENTS	ROBOTS DECISION
Training	GPTBot, Google-Extended, ClaudeBot	Blocking keeps you out of future corpora, not out of search or answers
Search index	OAI-SearchBot, PerplexityBot, Bingbot, Googlebot	Blocking removes you from that engine's answers entirely
User-triggered fetch	ChatGPT-User, Perplexity-User, Claude-User	Blocking prevents live citation at answer time

The two facts every robots.txt audit needs: Google-Extended controls Gemini training only, NOT AI Overviews (those use normal Googlebot, and cannot be opted out of separately except via snippet controls); and blocking GPTBot does not remove you from ChatGPT search, that is OAI-SearchBot and ChatGPT-User. Many sites blocked AI bots wholesale in 2023-2024 and are still invisible in AI search without knowing it.

The rest of the checklist

- **Bing coverage.** Verify in Bing Webmaster Tools, submit sitemaps, adopt IndexNow. Absent from Bing means absent from ChatGPT search.
- **Rendering.** Most AI fetchers execute little or no JavaScript. Server-render anything you want quoted; fetch your pages with JS disabled and check server logs for AI user agents.
- **Structured data.** Google says no special schema is required for AI features. Keep Organization, Article, and sameAs for entity disambiguation; do not sell schema as a citation hack.
- **llms.txt.** No confirmed consumer; Mueller: "no AI system currently uses llms.txt". Optional, low priority. The AEO guide carries the full debate.
- **Performance.** Real-time fetchers work under tight time budgets; slow pages get skipped at answer time. Fast TTFB is retrieval insurance.

08 — SCOREBOARD

MEASUREMENT WITHOUT SELF-DECEPTION

IN PLAIN WORDS

An AI gives a slightly different answer every time you ask, so checking your brand just once tells you almost nothing. You have to ask the same questions many times and track how often your brand shows up, the way you judge weather over a week instead of a single day. Any tool that claims you are ranked number three in an AI is making it up.

SparkToro and Gumshoe.ai ran 2,961 tests across ChatGPT, Claude, and Google AI with 600 volunteers: under a 1-in-100 chance of the same brand list twice for the same prompt, roughly 1-in-1,000 for the same order. Fishkin's conclusion: any tool that gives a ranking position in AI is full of baloney. But visibility percentage across many runs is stable (top brands appeared in 55-77% of responses), so the correct metric is share of responses mentioning the brand, measured across dozens of prompts, each run many times, per platform, as a trend.

The metric stack

- Mention rate and share of voice across a fixed prompt set, weekly, per engine.
- Citation rate: how often your domain is a linked source. A different number from mentions; tools report them differently.
- Sentiment and accuracy of what engines say about you. Hallucinated pricing and stale facts are common and fixable with authoritative pages.
- AI referral traffic, segmented by referrer (ChatGPT appends utm_source=chatgpt.com on many links). Track conversion separately and distrust vendor conversion claims.
- Bot activity in server logs: ChatGPT-User and Perplexity-User hits reveal which pages are being pulled into answers right now.

Tools, and how to buy them

Dedicated: Profound (enterprise leader), Peec AI (mid-market), Otterly.AI (entry tier), Scrunch AI. Incumbent add-ons: Semrush AI Toolkit, Ahrefs Brand Radar. Free: your own prompt panels via API scripts, plus Search Console (knowing AI Mode clicks are folded into web totals) and log analysis. The buyer's checklist: disclosed sample sizes and re-run counts, per-engine breakdowns, a citation-vs-mention distinction, API access to raw runs, and no "AI rank #3" theater. A vendor that cannot explain sampling error is selling decoration. The full tool stack, and how to wire it together, lives in The Stack companion.

09 — THE PLAYBOOK

DO THIS, IN ORDER

IN PLAIN WORDS

This is the whole guide turned into a to-do list you can work through from top to bottom. It starts with the basics and builds up, and every step names the tool to use and tells you how to know it is done. If you act on only one section of this guide, make it this one.

The whole discipline as an ordered plan. Follow it top to bottom if you are starting out; jump to your level if you are not. Every step names the tool and how you know it worked.

LEVEL 1 STARTER

become retrievable

1

Get into Bing's index.

`bing.com/webmasters` — verify your site, submit your sitemap, turn on IndexNow. This is what ChatGPT search retrieves from.

Done when: Bing Webmaster Tools shows your pages indexed.

2

Audit your AI crawler access.

In `robots.txt` and your CDN/WAF, decide per bot class: training (GPTBot, Google-Extended), search-index (OAI-SearchBot, PerplexityBot, Bingbot), user-fetch (ChatGPT-User, Perplexity-User).

Done when: the search-index and user-fetch bots you want citations from are allowed (verify with a fetch or your logs).

3

Confirm your answers survive without JavaScript.

Load a key page with JS disabled (or `curl` it) and check the passage you want quoted is present.

Done when: your quotable content is in the server HTML.

4

Fix your entity fact base.

Make your name, category, and one-line positioning identical across your About page, LinkedIn, Crunchbase, and G2; add `sameAs` schema.

Done when: an AI asked "what is [brand]" describes you the way your About page does.

LEVEL 2 PRACTITIONER

become citable

5

Lead every section with a quotable answer.

Rewrite section openings as a self-contained 40-60 word answer with the subject restated (no pronouns).

Done when: any single section, copied alone, still makes sense.

6

Embed the three validated boosts.

Add specific sourced numbers, named-expert quotations, and outbound citations to primary sources in the passages themselves.

Done when: each key passage carries a stat, a quote, or a citation.

7

Name the entities in the passage.

Put your brand, product, and category in the citable text, not "we".

Done when: a passage read in isolation attributes to your brand by name.

8

Reallocate PR from links to mentions.

Aim outreach at earning your brand named in credible, crawlable publications and genuine presence on `Reddit`, `YouTube`, not backlinks.

Done when: you are earning mentions next to category language, not just links.

LEVEL 3 ADVANCED

measure without self-deception

9

Build a visibility-rate panel.

Fix a set of prompts, run each 30-100 times per engine on a schedule, and measure the share of responses that mention or cite you.

Done when: you have a mention-rate trend per engine, and you have stopped tracking "AI rank position".

10

Separate mention from citation.

Track how often you are named versus how often your domain is a linked source; they move differently.

Done when: your dashboard reports both, distinctly.

11

Watch AI referrals and bot hits.

Segment referral traffic by `chatgpt.com`, `perplexity.ai`, `gemini.google.com` in GA4, and watch ChatGPT-User and Perplexity-User hits in your server logs.

Done when: you can see which pages are pulled into answers in near-real time.

12

Correct what the models get wrong.

When an engine states outdated pricing or a wrong fact about you, publish an authoritative page stating it plainly, then re-sample.

Done when: the corrected fact shows up in later answers.

10 — ANTI-PATTERNS

THE MISTAKES THAT COST THE MOST

IN PLAIN WORDS

This is a list of the traps that waste the most time and money, so you can steer around them. Things like checking your results only once, forgetting about Bing, or spending on the wrong kind of links. Knowing what not to do can save you more than any single clever tactic.

- Treating GEO as keyword SEO with a new name. Engines fan out into five to ten-plus machine-written sub-queries, and keyword stuffing measurably reduced visibility in the KDD study.
- Tracking "AI rank position" from single prompt runs. Single-sample dashboards measure noise.
- Ignoring Bing. A site unverified in Bing Webmaster Tools can rank #1 on Google and be absent from ChatGPT answers.
- Blocking AI crawlers wholesale, or blocking OAI-SearchBot and ChatGPT-User while meaning to block only the training bot GPTBot.
- Shipping client-side JavaScript content that AI fetchers cannot render, so the passages meant to be quoted are invisible.
- Spending the off-site budget on backlinks (0.218 correlation) instead of earned branded mentions (0.664) and presence where citations concentrate: Reddit, YouTube, Wikipedia.
- Burying the answer: no self-contained, stat-and-source-rich lead in a section means nothing survives passage-level chunking.
- Presenting llms.txt, schema hacks, or single-anecdote case studies as proven tactics instead of labeled experiments. Informed readers notice, and credibility is the product.

11 — THE HEAD GAME

OPTIMIZING WHAT YOU CANNOT CONTROL

IN PLAIN WORDS

This part is about staying patient, because the results here are slow and jumpy by nature. Your brand might show up one day and vanish the next, and it can take a couple of months before mentions turn steady. That wobble is completely normal, so do not panic and do not quit during the quiet stretch.

GEO is the hardest surface to stay sane on, because the output is probabilistic: the same prompt gives different answers. Volatility is the baseline, not a sign you are failing.

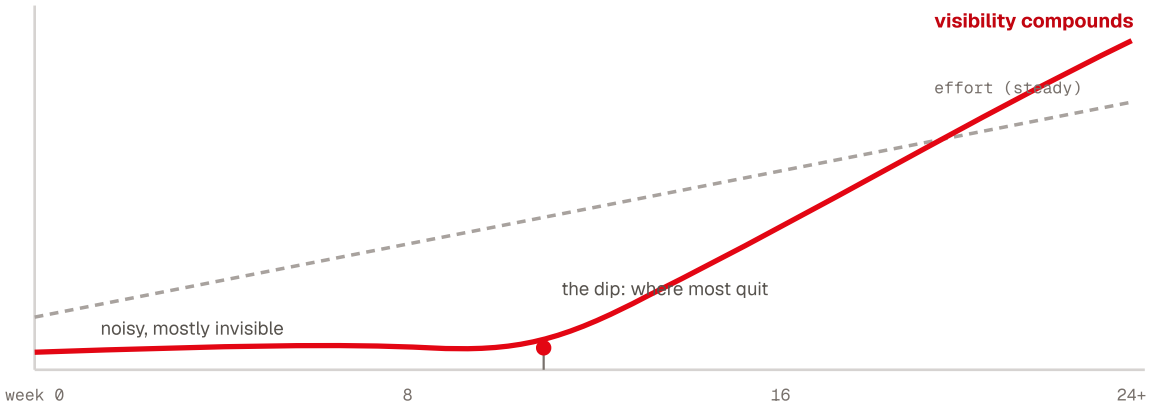

GEO's payoff is the slowest and noisiest: eight to sixteen weeks to consistent citations. In one 2026 study only ~30% of brands stayed visible between two consecutive answers (AirOps). Expect the wobble.

What to expect

Not a rank, a probability. The same prompt varies across sessions, users, and geographies. Consistent citations take eight to sixteen weeks; stable multi-engine visibility four to six months. In that 2026 study only 30% of brands stayed visible between two consecutive answers and 20% across five runs. A single good answer this morning tells you almost nothing.

Traps to avoid

1 **Measuring noise as signal**

Single-session spot checks are noise. Nothing you conclude from one prompt run is real.

2 **Hype in both directions**

Agencies stoke FOMO, yet most expensively marketed "GEO measures" are just good classic SEO. Cynical dismissal during the trough is the same error. Demand evidence either way.

3 **Chasing GEO "hacks"**

Most are unproven or repackaged SEO. The durable levers are retrievability, quotable content, and earned brand presence.

INSTRUMENT FOR NOISE

Fix a prompt library (30 to 50 across brand, category, and comparison queries), run it weekly across engines, report monthly, and smooth with rolling averages. Track share of voice, not rank. Rolling averages are what turn a noisy surface into a trend you can act on.

12 — HONEST UNCERTAINTY

WHAT NOBODY KNOWS YET

IN PLAIN WORDS

This is an honest list of the things even the experts have not figured out yet. Because AI answers change so much, it is genuinely hard to prove what works, and the rules keep shifting as the tools get updated. Knowing these limits keeps you from trusting anyone who claims to have it all solved.

- **Non-determinism makes attribution hard.** Identical prompts produce different brand lists over 99% of the time. Any before/after case study on small samples is indistinguishable from noise, and almost no published case study meets the bar.
- **The flagship research may not generalize.** The paper's 30-40% gains came from a controlled sandbox on 2023-era models. Whether quotation addition still moves citations in today's multi-round fan-out pipelines is publicly untested.
- **Correlation is not causation for mentions.** The 0.664 figure could partly reflect that big, mentioned brands are better at everything. No public controlled experiment isolates "earn 100 mentions, gain X% visibility".
- **The mechanics churn under your feet.** Citation counts per response, search-trigger rates, and fan-out behavior all shifted materially within a year. Principles (retrievability, quotability, entity consistency) age better than tricks.
- **Personalization is unmeasurable from outside.** Memory, account context, and location shape answers per user; tools query from clean accounts and may not see what real prospects see.
- **There is no ChatGPT Search Console.** All measurement is outside-in sampling plus referral crumbs. Prompt-volume estimates sold by tools are modeled, not observed.
- **Contested tactics:** llms.txt (no consumer), schema as a citation lever (weak correlational evidence; the AEO guide's schema section carries the detail), tone hacks for "writing for AI", and blocking crawlers as negotiating leverage (publishers who blocked lost visibility without compensation).

Sources

- Aggarwal et al.: GEO, ACM KDD 2024 arxiv.org/abs/2311.09735
- Ahrefs: AI Overview brand visibility factors, 75K brands ahrefs.com/blog/ai-overview-brand-correlation/
- Semrush: ChatGPT clickstream study, 17 months semrush.com/blog/chatgpt-search-insights/
- Pew Research Center: clicks with AI summaries pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/
- SparkToro: AI brand-recommendation inconsistency study sparktoro.com/blog/new-research-ais-are-highly-inconsistent-when-recommending-brands-or-products-marketers-should-take-care-when-tracking-ai-visibility/
- Search Engine Land: inside ChatGPT search, web.run and fan-out searchengineland.com/inside-chatgpt-search-web-run-fan-out-queries-ai-visibility-477339
- Google Search Central: AI features and your website developers.google.com/search/docs/appearance/ai-features
- Search Engine Journal: llms.txt has no current implementation searchenginejournal.com/google-says-llms-txt-is-purely-speculative-for-now/577576/
- Ahrefs: RAG explained for search marketers ahrefs.com/blog/retrieval-augmented-generation/
- iPullRank: how AI platforms expand queries with fan-out ipullrank.com/expanding-queries-with-fanout
- Digiday: WTF are GEO and AEO digiday.com/media/wtf-are-geo-and-aeo-and-how-they-differ-from-seo/
- Zapier: AI visibility tools, 2026 landscape zapier.com/blog/best-ai-visibility-tool/