


The Complete Guide to Generative Engine Optimization (GEO)

Aug 2, 2026



Rempa AI · The Complete Guide to GEO · October 2026 edition

*A practical playbook by **Rempa AI** — be the brand AI recommends.*

Your buyers now ask ChatGPT, Gemini, Perplexity, Copilot and Google's AI Mode the questions they used to type into a search bar. Each answer names only a handful of brands. Generative Engine Optimization (GEO) is the discipline of becoming one of them — and this guide shows you how, step by step.

Who this guide is for

- **Marketing leaders** who need to know whether AI assistants recommend their brand or a competitor.
- **SEO and content teams** adapting their playbook from ranked links to written answers.
- **PR and communications teams** whose earned media now feeds AI answers directly.

- **Founders** of challenger brands, for whom AI search is a rare chance to leapfrog bigger incumbents.

What's inside

1. Search has changed — the numbers behind the shift to AI answers
2. What GEO is, and how it differs from SEO and AEO
3. How generative engines choose the sources they cite
4. The five pillars of GEO
5. Content that gets cited
6. Technical GEO
7. Off-site authority: getting mentioned where AI looks
8. The platform playbook: ChatGPT, Google, Perplexity, Gemini, Copilot, Claude, DeepSeek
9. Measuring GEO
10. Your 90-day roadmap
11. Myths and mistakes to avoid
12. The GEO checklist, glossary and how Rempa AI can help

How to read it. Short on time? Read the Executive summary, then jump to the 90-day roadmap and the checklist. Building a GEO program from scratch? Read it front to back.

Executive summary

AI answers are now a mainstream discovery channel, and they reward a different kind of visibility than a ranked list of links. Eight facts frame everything in this guide.

1. **The audience is already there.** ChatGPT reaches 1.2 billion weekly users (OpenAI DevDay, 29 Sep 2026). Google's AI Overviews reach more than 2.5 billion monthly users and AI Mode more than 1 billion.
2. **Clicks are shrinking.** 68% of US Google searches now end without a click (SparkToro, 2026). When an AI Overview appears, the top result loses about 58% of its clicks (Ahrefs).
3. **Being cited pays.** On AI Overview queries, brands cited inside the answer earn about 120% more clicks per impression than brands that are not (Seer Interactive, 2026).

4. **Other people's pages decide.** Across 102,025 AI answers, 75.2% of citations pointed at third-party pages and only 2.9% at the brand's own site. "Best X" listicles alone took 21%.
5. **Facts beat formatting.** In a 252,000-trial experiment, topic fit, a stated price, a recent date and list position decided the first citation. Formatting moved nothing.
6. **Recognition wins ties.** When ten products had identical specs, models picked the brand they recognised in 100% of trials. Brand building is GEO.
7. **Every engine is different.** Google's AI Overviews and AI Mode cite the same URL only 13.7% of the time, and AI Mode replaces 56% of its cited domains every week.
8. **One-off checks mislead.** Answers vary between runs, by language and by country. GEO is measured as a trend, across engines and markets.


The five moves that matter most: map the prompts your buyers ask; publish clear, quotable facts on your own site; get present on the third-party sources AI already cites; let AI search crawlers in; and measure every week, engine by engine.

Chapter 1 — Search has changed

In under three years, AI assistants went from novelty to the first place millions of buyers ask for a recommendation. ChatGPT tripled its weekly audience in 19 months, and Google turned its own results page into an AI answer.

ChatGPT weekly users tripled to 1.2 billion in 19 months

Weekly active users as announced by OpenAI, millions



ChatGPT weekly active users as announced by OpenAI · Feb 2025 to Sep 2026

Each point is a figure OpenAI announced. The weekly audience grew by 300 million in the last seven months alone.

Google turned search into an answer

Google no longer just lists pages; on a growing share of queries it writes the answer itself. At I/O in May 2026, Google merged AI Overviews and AI Mode into one experience and put AI Overviews at more than 2.5 billion monthly users.

- **Coverage keeps rising.** BrightEdge measured AI Overviews on about 48% of tracked queries in February 2026: 88% in healthcare, 83% in education and 82% in B2B tech.
- **Desktop too.** Comscore's panel saw AI Overviews on 39.4% of US desktop Google searches in June 2026, up from 25.8% a year earlier.
- **Brand searches are next.** In late September 2026, AI Overviews appeared on 80% to 90% of the branded queries DemandSphere tracks, up from 26% at the start of the month.
- **The "People also ask" box is now AI.** AlsoAsked counted 97% of those answers as AI Overviews in early September 2026.

Fewer clicks, higher stakes

When the answer sits on the page, fewer people visit the sources behind it. Pew found users clicked a traditional result on 8% of visits with an AI summary, against 15% without one. The first randomized field experiment put the loss at 39.8% of outbound clicks.

The flip side matters more for marketers. Clicks that survive arrive later in the journey, after the AI has already compared options. And brands cited inside the answer recover a large share of the traffic, as Chapter 9 shows.

From ranking to being named

A results page offered ten slots; an AI answer typically names three to five brands, with a reason for each. That turns visibility into a shortlist game: you are either in the sentence the buyer reads, or invisible at the moment of decision.

What this means for you. The question is no longer "what position do we rank?" but "when a buyer asks AI about our category, are we named, in what order, in what words, and on the strength of which sources?" That is the question GEO answers.

Chapter 2 — What is GEO?

Generative Engine Optimization (GEO) is the practice of increasing how often, how prominently and how favourably a brand appears in answers written by AI engines such as ChatGPT, Gemini, Perplexity, Copilot, Claude, DeepSeek, Google AI Mode and AI Overviews.

The term comes from research. Pranjal Aggarwal and colleagues at Princeton, Georgia Tech, the Allen Institute and IIT Delhi coined it in a 2023 paper published at KDD 2024. On their own benchmark, content changes raised visibility in generative answers by up to 40%. That is a lab result, not a traffic promise, but it proved that what you publish changes what AI says.

AI search replaces ten links with one answer that names a shortlist




Illustration · fictional brands from Rempa's sample data

In classic search the buyer does the comparing. In AI search the engine compares for them, and only the brands in its sentence make the shortlist.

GEO, SEO and AEO compared

The three disciplines share a foundation but chase different outcomes. Many teams use GEO and AEO interchangeably; what matters is measuring the result in the answers themselves.

Discipline	Goal	Where it plays	How success is measured	Main levers
SEO	Rank pages high in a list of links	Google and Bing results pages	Position, impressions, clicks	Keywords, backlinks, technical health
AEO	Become the single answer to a question	Featured snippets, voice assistants, "People also ask"	Owning the answer box	Concise Q&A, structured data

Discipline	Goal	Where it plays	How success is measured	Main levers
GEO	Be named, recommended and cited in AI-written answers	ChatGPT, Gemini, Perplexity, Copilot, Claude, DeepSeek, AI Mode, AI Overviews	Mention rate, position in the answer, sentiment, cited sources	Quotable facts, third-party presence, brand recognition, crawler access

What stays the same

GEO is built on SEO, not instead of it. AI search engines retrieve pages from web indexes, so a page that cannot be crawled, indexed or trusted will not be cited. Fast, accessible pages, accurate information and a credible brand still do the heavy lifting.

What changes


- **The prompt replaces the keyword.** Buyers ask full questions with context: "best CRM for a 20-person team that needs GDPR compliance". You optimize for intents and personas, not single terms.
- **One answer blends many sources.** The engine synthesizes several pages, so your story must be consistent everywhere it appears.
- **Mentions count, even without links.** A brand named on a trusted review site can be recommended without a single backlink.
- **Results move constantly.** The same prompt can return different brands from one run to the next, so visibility is a rate measured over time, not a fixed rank.

Chapter 3 — How generative engines choose their sources

An AI answer draws on two kinds of knowledge: what the model learned during training, and what it retrieves from the web at the moment of the question. GEO has to work on both paths.

Every AI answer blends model memory with pages retrieved at answer time

PATH 1 · MODEL MEMORY · slow to change



How generative engines build an answer · simplified

Memory supplies what the model already believes about your brand; retrieval adds what it finds on the web today. The answer blends both.

Path 1: the model's memory

Large language models are trained on huge snapshots of the web, books and other text. What they absorbed about your brand, your category and your competitors becomes their default view. That memory is slow to change: it updates only when a new model is trained, and it reflects how often and how consistently your brand was described across the web.

Memory matters more than most teams assume. Similarweb measured ChatGPT citing the web in only 6.8% of answers in May 2026, and a retrieval study found free instant answers opened no page at all 93% of the time. Many recommendations come straight from what the model already believes.

Path 2: live retrieval

For fresh, specific or commercial questions, engines search the web before answering, a process called grounding or retrieval-augmented generation. The steps are broadly the same everywhere:

1. **Understand the prompt.** The engine reads the question and the conversation around it. Removing earlier messages changed the answer in 44.7% of tested conversations.
2. **Fan out.** The prompt is split into several sub-queries (Google calls this query fan-out in AI Mode).

3. **Retrieve.** Candidate pages come from a search index, a page cache or a live visit. ChatGPT combines its own index, a shared cache and live page opens.
4. **Select.** Passages are scored for relevance, freshness and trust; only a few survive.
5. **Synthesize and cite.** The model writes one answer, blending retrieved passages with its memory, and attaches some sources.

What the evidence says about selection


- **Page one is not the gate.** In a 55,393-query audit, nearly 30% of domains cited in AI Overviews were absent from the first page of results. A peer-reviewed study put the overlap between AI Overview sources and first-page results at 18%.
- **Specific facts win.** In 252,000 controlled trials, the first citation went to content with topic fit, a stated price, a recent date and a place on a list. Formatting alone moved nothing.
- **Language picks the market.** Asking in the local language surfaced local suppliers nearly always; the same question in English rarely did.
- **Sources can shape without being shown.** In lodging answers, Tripadvisor was a source in 61% of responses but a visible citation in only 21% (Comscore). Influence is wider than the links you see.

Takeaway. Win memory with consistent, widespread brand descriptions over time. Win retrieval with crawlable pages that state precise, current facts. Do both, because every engine weighs them differently.

Chapter 4 — The five pillars of GEO

Every GEO program rests on five pillars: be reachable, be clear, be vouched for, be recognised, and keep measuring. Weakness in any one caps the results of the other four.

Five pillars hold up the goal of being the brand AI recommends



The five pillars of GEO · Rempa AI framework

1. **Discoverability.** AI search crawlers can reach, render and index your pages. Without it, nothing else counts for live retrieval (Chapter 6).
2. **Clarity.** Your site states the facts AI needs in a form it can quote: what you do, for whom, at what price, compared with whom (Chapter 5).
3. **Credibility.** Independent sources that AI trusts — comparison sites, reviews, forums, media — mention and recommend you (Chapter 7).
4. **Recognition.** Your brand is a well-defined entity, described the same way everywhere, so models know who you are and what you are best at.
5. **Measurement.** You track a stable set of prompts across engines and markets, and tie changes in the answers to the actions you took (Chapter 9).

Diagnose your weakest pillar

The fastest gains usually come from fixing the weakest pillar first. Use the symptoms you observe in AI answers to find it.

What you see in AI answers	Likely weak pillar	First fix
You are never mentioned, even for your own brand name	Discoverability or Recognition	Check robots.txt for search bots; publish a clear "About" and entity profile
Competitors are named for prompts you should win	Credibility	Get listed on the comparison and review sites cited for those prompts

What you see in AI answers	Likely weak pillar	First fix
You are mentioned, but with outdated prices or features	Clarity	Publish current pricing and feature pages with visible update dates
You are named, but last, or with a weak reason	Clarity and Credibility	Create comparison pages; earn reviews that state your strengths
Tone about you is negative or wrong	Credibility	Find the sources driving it; correct facts and respond on those platforms
Visibility swings and you cannot say why	Measurement	Track a fixed prompt set daily and read weekly trends, not single answers

Chapter 5 — Content that gets cited

AI engines reuse what a page states clearly. The winning content answers a buyer's question in plain sentences, with specific, current and verifiable facts the model can lift into its answer.

The pages AI looks for

When buyers ask AI to compare, shortlist or price, engines go looking for pages that answer exactly that. Most brand sites lack several of them.

- **Pricing page** with real numbers per plan, seat or unit, and what each plan includes.
- **Comparison pages** ("You vs Competitor") with an honest feature table.
- **Alternatives pages** ("Best alternatives to Competitor") that place you on a list.
- **Use-case and industry pages** stating who you are best for and why.
- **FAQ pages** written in the buyer's own words, one question per heading.
- **Proof pages:** case studies with numbers, security and compliance facts, integrations lists.

A page AI can quote states facts, not adjectives

Lumen CRM costs \$19 per seat per month
Billed annually · 14-day free trial
Updated 1 October 2026

Starter	Growth	Scale
\$19 per seat / month	\$29 per seat / month	\$49 per seat / month

Best for sales teams of 5 to 50 people in the EU

Lumen CRM vs Orbita	Lumen	Orbita
Hosted in France	✓	✗
Setup in a day	✓	✓
Price per seat	\$19	\$35

Is Lumen CRM GDPR compliant?
Yes. Customer data is hosted in France.

- 1 Answer-first headline**
The price in the very first sentence
- 2 A visible update date**
Recency helps win the first citation
- 3 Real numbers for every plan**
A stated price is a citation gatekeeper
- 4 Who you are best for**
Gives AI a reason to recommend you
- 5 An honest comparison**
Answers A-vs-B prompts directly
- 6 FAQs in buyer language**
One question per heading, one answer

Illustration · fictional brand and prices

One pricing page can answer half a dozen buyer prompts at once, if every fact on it is stated plainly.

Eight writing rules backed by evidence

Rule	Why it works	Example
Lead with the answer	Engines extract short passages; the first sentence under a heading is the most quotable	"Lumen CRM costs \$19 per seat per month, billed annually."
State numbers, not adjectives	A stated price was one of four gatekeepers of the first citation in 252,000 trials	"Setup takes one day" instead of "fast setup"
Add statistics, quotes and sources	Citations, quotations and statistics were the strongest levers in the original GEO paper	Cite the study, quote the customer, link the source
Say who you are best for	Answers recommend with a reason; give the model that reason	"Best for EU teams that need data hosted in France"

Rule	Why it works	Example
Show a recent date and keep it true	A recent date was another gatekeeper of the first citation	"Updated 1 October 2026" on pricing and comparison pages
Get onto lists	List position was a gatekeeper; "best X" listicles took 21% of brand citations	Publish your own honest comparison, and earn places on others'
Use the buyer's words and language	The language of the question selects the market and the suppliers named	Localized pages for each market you sell in
Say the same thing everywhere	Answers blend many sources; contradictions weaken your description	One positioning line reused on site, profiles and press kit

From vague to quotable

Hard for AI to quote	Easy for AI to quote
"Our innovative platform empowers teams of all sizes."	"Lumen CRM is a CRM for sales teams of 5 to 50 people."
"Competitive pricing, contact us."	"Plans start at \$19 per seat per month; a 14-day trial is free."
"We take security seriously."	"Data is hosted in France and the service is GDPR compliant."
"Loved by customers worldwide."	"Rated 4.7 out of 5 from 1,200 reviews on two review sites."

Examples use the fictional brand Lumen CRM.

Formatting is not a shortcut. Tables, FAQs and clean headings help people and machines read a page, but in controlled tests formatting alone did not win citations. The facts inside the format do. And keep every claim true: invented authority claims can sway a model in a lab, but they expose you to reputational and legal risk when a buyer checks.

Chapter 6 — Technical GEO

An engine can only cite what its crawler can fetch and read. For most brands, technical GEO comes down to four checks: let the right bots in, serve content in raw HTML, keep facts consistent in structured data, and signal freshness.

Know the three kinds of AI bot

Not all AI bots do the same job, and blocking the wrong one costs reach for nothing. A **training crawler** collects pages to train future models. A **search crawler** indexes pages so an assistant can cite them. A **user fetcher** opens a page because someone asked about it.

Token	Operator	Kind	What it feeds	Recommendation for brands
OAI-SearchBot	OpenAI	Search	ChatGPT search answers; blocked sites are not shown there	Allow
GPTBot	OpenAI	Training	Future OpenAI models	Allow, unless you have a licensing reason to block
ChatGPT-User	OpenAI	User fetcher	Pages opened during a ChatGPT conversation	Allow
Claude-SearchBot	Anthropic	Search	Claude's search results	Allow
ClaudeBot	Anthropic	Training	Future Claude models	Allow, unless you have a licensing reason to block
PerplexityBot	Perplexity	Search	Perplexity answers; not used for model training	Allow
Googlebot	Google	Search	Google Search, including AI Overviews and AI Mode	Allow — it is the only crawler-level gate to Google's AI answers
Google-Extended	Google	Control token	Gemini training and grounding in Gemini apps	Allow; one study links blocking it to fewer AI Overview citations

Token	Operator	Kind	What it feeds	Recommendation for brands
Bingbot	Microsoft	Search	Bing and Copilot; AI use is controlled by NOCACHE and NOARCHIVE meta tags	Allow, and avoid NOARCHIVE on pages you want cited

Publishers and brands differ. A news publisher may block training bots to protect licensing value. A brand that wants to be recommended usually gains from being known by every model, so the default for brands is to allow.

A robots.txt that welcomes AI search while keeping your private areas closed can be as simple as this:

```
# Welcome AI search and assistant crawlers
User-agent: OAI-SearchBot
User-agent: ChatGPT-User
User-agent: Claude-SearchBot
User-agent: PerplexityBot
User-agent: Googlebot
User-agent: Bingbot
Allow: /
Disallow: /account/

Sitemap: https://www.example.com/sitemap.xml
```

Then check your CDN or firewall: some bot-protection settings block AI crawlers by default, silently overriding robots.txt.

Serve the content in raw HTML

Google renders JavaScript, but the crawlers behind ChatGPT, Claude and Perplexity do not. Vercel and MERJ's crawl-log study found that none of the major AI crawlers rendered JavaScript: GPTBot and ClaudeBot downloaded script files but never ran them. A page built in the browser can look perfect to you and empty to AI.

AI crawlers read raw HTML: content built by JavaScript stays invisible

Same page, three views. Google renders JavaScript; GPTBot, ClaudeBot and PerplexityBot do not.




Illustration · fictional brand

- Use server-side rendering or pre-rendering for every page you want cited.
- Test it: fetch a page with a bot user agent (`curl -A "GPTBot"` `https://yoursite.com/pricing`) and check that your key facts appear in the raw response.
- Never hide prices, specs or FAQs behind tabs, accordions or widgets that load on click.

Structured data: reinforce, don't replace

Schema.org markup (Organization, Product, Offer, FAQPage, Review) helps machines connect your facts to your brand. Keep it identical to the visible text, and server-render it: JSON-LD injected by JavaScript is invisible to non-rendering crawlers. Google's own guidance says no special markup is required for its AI features, so treat schema as consistency, not a ranking switch.

What about llms.txt?

The llms.txt file is a proposed summary of a site for AI tools. The evidence so far is flat: Google says it does not use it for AI Overviews or AI Mode, and independent studies found no measurable link between the file and AI citations, with 97% of published files receiving no requests in a month. It may help coding agents read documentation; it will not make you visible.


Signal freshness

Recency helps win citations, so make updates visible. Keep `lastmod` dates in your sitemap accurate, show an "updated on" date on key pages, and submit changes to Bing through IndexNow so Copilot sees them quickly.

Chapter 7 — Off-site authority: get mentioned where AI looks

AI trusts what others say about you more than what you say about yourself. In a study of 102,025 brand-recommendation answers across five assistants, 75.2% of citations pointed at other companies' pages, and only 2.9% at the brand's own domain.

Three in four AI citations point to pages you do not own



Share of citations across 102,025 AI answers, June 2026

Your site still matters as the source of truth others repeat, but most of the evidence an AI weighs about you lives elsewhere.

Mentions are the new backlinks

Links still matter for classic SEO, but AI answers are driven by how often and how well your brand is talked about. One widely shared analysis found branded web mentions correlated with AI visibility about three times more strongly than backlinks (0.664 against 0.218). And recognition decides close calls: when ten products had identical specs, models picked the brand they knew in 100% of trials.

That favours incumbents. In more than 20,000 shopping prompts run for Lightspeed, AI answers chose the national chain over a smaller store up to 94% of the time. Challenger brands have to earn their place on the sources AI reads.

The source playbook

Source type	Why AI uses it	What to do
Comparison and "best X" lists	They answer shortlist prompts directly; listicles took 21% of brand citations	Identify the lists cited in your category; pitch inclusion with facts, pricing and a clear "best for"
Review platforms	They supply sentiment and the attributes attached to your name	Ask happy customers for detailed reviews; reply to every negative one with facts
Communities and forums	Real users' experiences; Reddit is among the most-cited domains	Take part openly as the brand, answer questions, never seed fake posts
Video	YouTube leads citations in several categories on Google's AI surfaces	Publish demos, comparisons and tutorials with accurate titles and transcripts
Reference sources	Wikipedia and Wikidata anchor entity facts; ChatGPT cites Wikipedia more than other engines	Keep your entity data accurate; follow each site's conflict-of-interest rules
Media and analysts	News and expert coverage build authority and memory	Pitch data-driven stories; publish original research others will quote
Marketplaces and directories	Product facts, prices and availability for shopping prompts	Complete every listing with consistent specs and pricing

Find the sources that matter for you

Every category has its own handful of influential sources. In Comscore's retail sample, YouTube appeared among the top citations in 23% of answers, Amazon in 14%, Reddit in 12% and Wikipedia in 11%. Your market will look different, so measure it: list the domains cited in answers to your prompts, mark where competitors appear and you do not, and start with the sources cited most often.

Example from Rempa: softwarecompare.io is cited in 41% of answers in the CRM sample, and Lumen CRM is absent from it. That single gap is a better PR target than ten generic press releases.

Play it straight. Fake reviews, undisclosed paid posts and seeded forum threads break platform rules and consumer-protection law, and communities detect them quickly. Durable AI visibility comes from genuine customers and real coverage.

Chapter 8 — The platform playbook

Each engine retrieves, selects and cites differently, so the same brand can lead on one and be absent on another. Use the shared foundations from Chapters 5 to 7 everywhere, then tune for each platform.

Rempa sample data · fictional brand, illustrative values

The same brand leads on Gemini in France and barely registers on DeepSeek in Japan. Strategy has to be set engine by engine and market by market.

Engine	Latest public reach	Where answers come from	Crawler or control to allow
ChatGPT	1.2 billion weekly users (Sep 2026)	Model memory, OpenAI's own index, page cache, live fetches	OAI-SearchBot, ChatGPT-User
Google AI Overviews	2.5 billion+ monthly users (May 2026)	Google's search index, Gemini models	Googlebot
Google AI Mode	1 billion+ monthly users (May 2026)	Google's index with query fan-out	Googlebot
Gemini app	Hundreds of millions of monthly users	Model memory, Google Search grounding	Google-Extended
Perplexity	Not comparable	Live search on every answer, many citations	PerplexityBot
Microsoft Copilot	Copilot Search on 17.3% of Bing desktop searches (Jun 2026)	Bing's index	Bingbot (meta tags control AI use)
Claude	Fast-growing referrer from a small base	Model memory, web search	Claude-SearchBot, Claude-User
DeepSeek	Strong in some markets	Mostly model memory; optional web search	No documented crawler token

ChatGPT

The biggest single audience, and the source of about 78% of AI referral traffic in March 2026 (Statcounter). Most answers still come from memory: ChatGPT cited the web in 6.8% of answers in May 2026. It also states a first-person product preference in 79% of recommending answers, far more than Google's surfaces.

- Build long-term memory through widespread, consistent mentions on trusted sites, including Wikipedia where you qualify.
- For search-enabled answers, keep OAI-SearchBot allowed and key facts in raw HTML; retrieval often works from short snippets, so put the answer in the first lines.

Google AI Overviews and AI Mode

Since May 2026, one experience: the Overview sits on the results page and AI Mode carries the follow-up conversation. Both run on Google's normal search index, so classic SEO health is the entry ticket.

- Expect different sources on each surface: they cite the same URL only 13.7% of the time. AI Overview sources barely change for months; AI Mode replaces 56% of its cited domains every week.
- Invest in YouTube, which Google's AI surfaces cite heavily in several categories.
- Use Search Console's Generative AI performance report, live worldwide since 31 August 2026, to see AI impressions by page, country and device.

Gemini

The Gemini app answers from model memory and Google Search grounding. Google-Extended controls whether your content trains Gemini and grounds Gemini app answers, so keep it allowed if you want to be recommended there.

Perplexity

Perplexity searches the web for nearly every answer and lists its sources prominently, which makes it the most "SEO-like" engine. It favours well-sourced, factual pages; in one health audit it leaned on academic sources more than any other engine.

- Keep PerplexityBot allowed and publish pages with clear facts, data and references.

Microsoft Copilot

Copilot builds on Bing's index, so Bing is the gate. Submit sitemaps in Bing Webmaster Tools, use IndexNow for fast updates, and check its AI Performance report, which Microsoft calls an early step toward GEO tooling. Avoid the NOARCHIVE meta tag on pages you want cited.

Claude

Claude's referrals are still small but growing fast: they more than doubled in Q2 2026 on DataDome's network. Claude app visits often arrive without a referrer header, so they can hide in your "Direct" traffic. Allow Claude-SearchBot and Claude-User.

DeepSeek

DeepSeek matters in specific markets and relies heavily on what its models learned in training. Broad, consistent coverage of your brand, in the local language, is the main lever. Track it per country rather than assuming it mirrors ChatGPT.

Rule of thumb. Never judge your AI visibility from one engine or one country. As the heatmap at the start of this chapter shows, the gaps between them can be larger than the gap between you and your competitors.

Chapter 9 — Measuring GEO

You cannot manage AI visibility from a handful of screenshots. Answers change from run to run, from engine to engine and from country to country, so GEO is measured like an audience: a fixed set of prompts, asked repeatedly, read as trends.

The four core metrics

Every AI answer can be read on four dimensions. Aggregated across your prompts, they become your GEO scorecard.

Answer analysed the way Rempa reads it · Rempa sample data, fictional brands

Metric	What it tells you	How it is calculated
Visibility	How often AI names you — your AI share of voice	Share of answers that mention your brand
Position	Whether you are the first recommendation or an afterthought	Average rank of your brand among the brands named
Sentiment	The tone and the attributes attached to your name	Score of the language used about you (for example 0 to 100), plus the recurring descriptors
Sources	Which pages shape the answer, and whether yours are among them	Domains and URLs cited in answers, with your presence on each

Add two business metrics on top: **AI referral traffic and conversions** from your analytics, and **AI impressions** from Google Search Console's Generative AI report and Bing Webmaster Tools' AI Performance report.

Build a prompt set that mirrors real buyers

1. **Start from personas and jobs to be done**, not keywords: "What's the best CRM for a 20-person company?" from a sales manager; "Is there a CRM hosted in France?" from an IT lead.

2. **Cover the funnel:** category discovery ("best X for Y"), comparison ("A vs B"), validation ("is A reliable?") and branded questions.
3. **Localize:** write prompts in each market's language, because the language of the question selects the suppliers named.
4. **Group prompts** by topic, persona and funnel stage so you can see where you win and where you lose.
5. **Keep the core set stable** for months; add new prompts in a separate group so trends stay comparable.

Measure trends, not single answers

Ask every prompt on a fixed schedule in clean sessions, with no history or personalization, so results are comparable over time. Real users bring context: removing earlier messages changed answers in 44.7% of tested conversations. A clean-session baseline is therefore a standardized benchmark, not a replay of every user's chat, which is exactly why volume and repetition matter.

The GEO improvement loop · Rempa AI method

Measurement is not the last step but the engine of the loop: each week's data shows which prompts to work on next.

Prove the impact

- **Use a control group.** Glasp optimized part of its site and left the rest untouched. Optimized pages' AI referrals grew 5.7 times, but untouched pages grew 3.5 times; the real effect of the work was about 1.8 to 2.3 times. Without the control, it would have claimed all the platform's growth.
- **Fix your analytics.** Create an "AI assistants" channel for referrers such as chatgpt.com, perplexity.ai, gemini.google.com, copilot.microsoft.com and claude.ai. Expect leakage: one site found 22% of its AI Overview traffic filed as Direct.
- **Value citations, not only clicks.** On AI Overview queries, cited brands earned about 120% more clicks per impression than uncited ones (Seer Interactive), and visitors who click after reading an AI answer arrive better informed.

Chapter 10 — Your 90-day GEO roadmap

A first GEO cycle fits in one quarter: a month to measure, a month to fix, a month to amplify and prove impact. Each phase ends with a gate that decides what the next one works on.

90-day GEO roadmap · phases of equal length, not to scale

Do not skip the first gate: without an agreed baseline, no later result can be credited to your work.

With Rempa, Phase 1 takes a morning. Define your brand, competitors and markets; Rempa suggests prompts from your keywords and personas, runs them on the engines you choose, and delivers your first baseline in under 24 hours.

Who does what

GEO is cross-functional. One owner, usually the head of SEO or content marketing, runs the program and the weekly review.

- **Content team:** pricing, comparison, alternatives, use-case and FAQ pages.
- **PR and communications:** presence on cited sources, media coverage, original research.
- **Web team:** crawler access, server-side rendering, structured data, sitemaps.
- **Customer and community teams:** reviews, forum participation, accurate public answers.
- **Analytics:** prompt tracking, AI referral channel, impact readouts.

The operating rhythm

Data arrives daily, but decisions should not. Review trends weekly, report monthly, and reset priorities each quarter. Answers change as models and competitors move, so the loop never really ends: after day 90, start the next cycle from the prompts you still lose.

Chapter 11 — Myths and mistakes to avoid

GEO is young, and much of the advice circulating about it is untested. Here is what the evidence says about the most common beliefs.

Myth	What the evidence shows
"GEO replaces SEO."	AI engines retrieve from search indexes. A page that is not crawlable, indexable and trusted will not be cited. SEO is the foundation of GEO.

Myth	What the evidence shows
"Ranking first on Google means AI will cite us."	Nearly 30% of domains cited in AI Overviews were absent from page one, and source overlap with page one was just 18%.
"An llms.txt file will get us cited."	No measurable effect on citations in independent studies; Google says it does not use the file.
"Blocking GPTBot takes us out of ChatGPT search."	GPTBot feeds training. OAI-SearchBot controls ChatGPT search; only blocking that removes you.
"Schema and formatting are the secret."	In 252,000 trials, formatting alone did not change who got cited. Facts did: fit, price, recency, list position.
"Our own website is enough."	75.2% of citations in brand recommendations point to third-party pages, and 2.9% to the brand's own site.
"What works on ChatGPT works everywhere."	Even Google's two AI surfaces cite the same URL only 13.7% of the time. Each engine needs its own measurement.
"One check shows where we stand."	Answers vary between runs, languages and conversations. Only repeated tracking reveals the real rate.
"AI traffic shows up clearly in analytics."	AI visits are often misfiled as Direct, and many AI answers influence buyers without any click.

Five costly mistakes

1. **Tracking only branded prompts.** You will look strong while competitors win every "best X for Y" question.
2. **Measuring in English only.** The language of the question selects the market; each country needs its own prompts.
3. **Letting stale facts live.** Old prices and discontinued features keep resurfacing in answers until you correct them at the source.
4. **Gaming communities and reviews.** Seeded posts and fake reviews get detected, removed and remembered — by people and by models.

5. **Reacting to daily noise.** Single-day swings are normal; act on weekly trends across many prompts.

The GEO checklist

Tick these off over your first quarter. Each line maps to a chapter of this guide.

Scope and measurement (Chapter 9)

- ☐ List your brand, three to five main competitors and every market you sell in
- ☐ Write 100 or more prompts across personas, topics and funnel stages, in each market's language
- ☐ Track them daily across ChatGPT, Gemini, Perplexity, Copilot, Google AI Mode, AI Overviews and DeepSeek
- ☐ Record your baseline: visibility, average position, sentiment and cited sources
- ☐ Create an "AI assistants" channel in your analytics and connect Search Console and Bing Webmaster Tools AI reports

Technical access (Chapter 6)

- ☐ Allow OAI-SearchBot, ChatGPT-User, Claude-SearchBot, PerplexityBot, Googlebot and Bingbot in robots.txt
- ☐ Decide deliberately on training crawlers (GPTBot, ClaudeBot, Google-Extended)
- ☐ Check that your CDN or firewall does not block AI crawlers by default
- ☐ Confirm key pages show their content in raw HTML, without JavaScript
- ☐ Keep sitemap `lastmod` dates accurate and use IndexNow for Bing

Content on your site (Chapter 5)

- ☐ Publish a pricing page with real numbers
- ☐ Publish honest comparison pages against your top competitors
- ☐ Publish an alternatives page and use-case pages that say who you are best for
- ☐ Rewrite key pages answer-first, with numbers, sources and a visible update date
- ☐ Align your one-line positioning across site, profiles and press kit
- ☐ Add structured data that matches the visible text, rendered on the server

Off-site presence (Chapter 7)

- ☐ List the 20 domains most cited in answers to your prompts
- ☐ Mark where competitors appear and you do not; prioritize the most-cited gaps

- ☐ Pitch inclusion on comparison and "best X" lists with facts and pricing
- ☐ Run a review campaign with existing customers and answer every negative review
- ☐ Take part openly in the communities and forums cited in your category
- ☐ Publish video demos and comparisons
- ☐ Correct outdated or wrong facts on the third-party pages that carry them

Review and iterate (Chapter 10)

- ☐ Hold a weekly trend review and a monthly report
- ☐ Compare optimized pages and prompts against an untouched control group
- ☐ Set alerts for visibility drops, new competitors and negative answers
- ☐ Re-plan the next quarter from the prompts you still lose

Glossary

Term	Meaning
AEO (Answer Engine Optimization)	Optimizing to become the direct answer to a question, from featured snippets to AI assistants; often used interchangeably with GEO
AI Mode	Google's conversational search experience, merged with AI Overviews into one flow since May 2026
AI Overviews	AI-written summaries shown at the top of Google results, with linked sources
AI share of voice	Your share of AI answers that mention your brand, compared with competitors
Citation	A source an AI answer links to or lists as the basis for its claims
Crawler (search vs training)	A bot that fetches pages; search crawlers index pages for citation, training crawlers collect them to train models
Entity	A uniquely identifiable thing — your company, product or founder — that engines connect facts to
GEO (Generative Engine Optimization)	Practices that increase a brand's presence, position and sentiment in AI-generated answers

Term	Meaning
Grounding	Anchoring an AI answer in documents retrieved at answer time, usually from a web search
llms.txt	A proposed text file summarising a site for AI tools; no measured effect on AI citations so far
Model memory	What a model learned during training, used when it answers without searching
Prompt	The full question a user asks an AI assistant; the unit GEO tracks, as keywords are for SEO
Prompt tracking	Asking a fixed set of prompts to AI engines on a schedule and analysing the answers
Query fan-out	Splitting one prompt into several sub-searches to gather sources, used by Google's AI Mode
RAG (retrieval-augmented generation)	The technique of retrieving documents and feeding them to a model before it writes an answer
Sentiment	The tone and descriptors an AI answer attaches to your brand
Zero-click search	A search that ends without a click to any website, because the answer was on the page

How Rempa AI can help

[Rempa AI](#) shows you your brand as AI tells it. It sends your prompts to the answer engines every day and reads each answer for you: who is mentioned, in what order, in what tone, and on the strength of which sources. Everything in Chapters 4 to 10 becomes a weekly routine instead of a research project.

Mock-up of the Rempa visibility view · sample data from rempa.ai

What Rempa brings to your brand

AI answers are becoming a brand's most influential shop window, yet most teams cannot see them. Rempa makes that channel visible, measurable and actionable.

- 1. **See what your buyers hear.** Rempa asks the questions your buyers ask and shows exactly what ChatGPT, Gemini, Perplexity, Copilot, DeepSeek and Google's AI surfaces say about you. No more guessing from a few manual screenshots that change on every refresh.
- 2. **Know where you lose before the sale starts.** Rempa lists the prompts where AI names a competitor and not you: the shortlists you miss at the very moment a buyer decides.
- 3. **Turn answers into a work plan.** Because Rempa shows the sources behind each answer, content teams know which pages to create and PR teams know which sites to win, ranked by how often AI cites them.
- 4. **Protect your reputation.** Sentiment tracking and alerts flag outdated prices, wrong features or negative descriptions as soon as they appear, so you can correct them at the source before buyers believe them.
- 5. **Prove the return on GEO.** Week-over-week trends by engine and market show which actions moved visibility, position and sentiment: the evidence leadership needs to keep investing.
- 6. **Win market by market.** With 52 countries and localized prompts, you see where you lead, where a local rival dominates, and where to focus first.

The loop in practice (sample data). Rempa's demo project shows the full cycle for the fictional brand Lumen CRM: 12 prompts where AI cited its rival Orbita and not Lumen, two pages to create (a comparison page and a pricing page), and visibility up 6.2 points, from 31.8% to 38%, over four weeks.

Who uses Rempa, and for what

Team	The question Rempa answers	What they do with it
CMO and brand leads	Does AI recommend us, or our competitors?	Set AI share-of-voice targets and report progress to leadership
SEO and content	Which prompts do we lose, and which pages would win them?	Create and refresh pricing, comparison and FAQ pages, then track the effect
PR and communications	Which sites does AI trust in our market, and are we on them?	Focus outreach on the most-cited sources where competitors appear
Product marketing	Which words does AI attach to our name?	Sharpen positioning and correct outdated claims

Team	The question Rempa answers	What they do with it
International teams	Is it the same in every country and every engine?	Localize prompts and content market by market
Founders of challenger brands	Where can we beat bigger rivals?	Target the prompts and markets where incumbents are weak

Without Rempa, with Rempa

Without Rempa	With Rempa
A few prompts checked by hand, in one engine	Your full prompt set, across the engines you choose, on a fixed schedule
Screenshots that change on every refresh	Trends that smooth out the noise
No idea which sites shape the answer	A ranked list of cited sources, with your gaps highlighted
Wrong descriptions discovered by customers	Email and Slack alerts as soon as an answer turns negative
GEO work that is hard to justify	Visibility, position and sentiment tied to the actions you took

How Rempa maps to this guide

GEO step in this guide	What Rempa gives you
Find the prompts you lose	Prompt tracking with suggestions from your keywords, grouped by topic, persona or funnel stage; competitor benchmarking that shows the prompts where rivals are named and you are not
Fix what AI reads about you	Sentiment and the attributes attached to your name, so you catch outdated or negative descriptions; the pages to create or improve
Get mentioned where AI looks	Source analysis: the domains and URLs AI cites most in your category, and where competitors appear without you

GEO step in this guide	What Rempa gives you
Measure the effect	Visibility, position, sentiment and sources per model and market, week after week, with email and Slack alerts

Coverage. ChatGPT, Gemini, Perplexity, Copilot, DeepSeek, Google AI Mode and Google AI Overviews, across 52 countries with localized prompts.

Method. Prompts run in clean sessions with no history or personalization, daily (every three days on the Starter plan), so trends are stable and comparable.

Integrations. CSV exports, Looker Studio and an API to bring GEO data into the reports you already use.

Trust. Built and hosted in the European Union, GDPR compliant, with a DPA available on request.

See what AI says about you by tomorrow. Start a 14-day free trial at rempa.ai — no credit card, first results in under 24 hours. Plans start at \$75 a month.

© 2026 Rempa AI. *Visibility analytics for AI search, built and hosted in Europe.*
hello@rempa.ai