



GENERATIVE ENGINE OPTIMIZATION

# The Ultimate Guide to Getting Mentioned by AI

Written for South African SMEs, practitioners, and multinationals who want to show up where buyers are actually looking in 2026 and 2027 — not just on page one, but inside the answer itself.

# What's inside

**Introduction** Something changed, and most brands missed it

**Chapter 1** What GEO is, and why your SEO playbook doesn't carry over

**Chapter 2** How AI models actually decide who to recommend

**Chapter 3** Find the prompts your buyers are already using

**Chapter 4** Run your own AI visibility audit

**Chapter 5** Fix the five things making AI skip you

**Chapter 6** Write content AI actually cites

**Chapter 7** Build citation presence beyond your own site

**Chapter 8** Track it, repeat it, don't stop

**Appendix A** The data behind GEO

**Appendix B** Your GEO checklist

**FAQ** Answers to the questions everyone asks first

# Something changed, and most brands missed it

A few years ago, a business owner looking for an accountant, a physiotherapist, or a consulting firm would Google it, scan ten blue links, and click around until something looked credible. That's not how it works anymore.

Today that same owner opens ChatGPT or Gemini and asks the question directly. They get a short answer with two or three names in it. They pick one and move on. No scrolling, no comparison shopping, no ten tabs open at once.

Here's the uncomfortable part: if you're not in that short answer, you were never in the running. The buyer never saw a list to choose from — they saw a recommendation, and it wasn't you.

Google hasn't gone away. But it's no longer the only front door, and the new doors run on a different lock. This guide is Ameliorate's own playbook for getting through them — built the same way we build every client system: as a process, not a trick.

## Why the numbers matter

**65%**

of consumers now start their search directly inside an AI tool, skipping the search engine entirely.

*Search Engine Land, Jan 2026*

**4.4×**

higher conversion rate from visitors who arrive via an AI recommendation, versus organic search traffic.

*Semrush, 2025*

**78%**

of marketing leaders are already building a dedicated AI-visibility strategy into next year's plan.

*Gartner, Jun 2025 · 413 leaders surveyed*

That middle number is the one most businesses underrate. A visitor who arrives because an AI model told them you're worth a look isn't a casual browser — an algorithm has already vouched for you. They arrive closer to a buying decision, which is exactly why a single mention in an AI answer is worth more than an equivalent ranking on a results page.

## Who this guide is for

---

- The owner who's noticed direct enquiries changing shape — fewer form-fills from Google, more "ChatGPT said I should call you" — and wants to understand what to do about it.
- The marketer who's heard the term GEO and needs a practical, step-by-step process to run without a six-month project plan.
- The practice or firm that wants this handled as a system, not a once-off campaign — which is where Ameliorate comes in.

You don't need to be technical to use this guide. You need to understand the four signals AI models actually weigh, then build the habits that push your brand into that decision. That's the rest of this document.

## CHAPTER 01

# What is GEO?

SEO earns you a ranking. GEO earns you a mention. The difference is bigger than it sounds.

GEO — Generative Engine Optimization — is the discipline of making your business the name an AI model reaches for when someone asks a question in your category. Google ranking rewarded volume: rank on page one alongside nine competitors, and the buyer still had to choose between all ten. An AI answer doesn't work that way. It names two or three options. Sometimes one. Being in that answer isn't a tie — it's the whole game.

## Why your SEO work doesn't automatically transfer

AI models don't read rankings. They don't know or care that you're first on Google. What they weigh is whether you've been cited, referenced, and discussed consistently enough across the web that the model has learned to associate your name with your category.

A business that's been written up in industry publications, reviewed on third-party platforms, and discussed in forums has effectively trained the model to know who it is. A business with strong rankings but a thin footprint elsewhere hasn't — and the model has nothing to work with. This doesn't mean SEO is wasted effort: one large-scale study across 10,000 prompts in finance and SaaS found a 0.65 correlation between organic Google rankings and how often a brand gets mentioned by an LLM. Good SEO is the foundation. It's not the whole building.

**0.65**

correlation between Google organic rankings and LLM brand mentions — versus 0.25 for domain rank and just 0.10 for backlink count.

*Large-scale LLM correlation study, finance & SaaS*

## SEO vs GEO, side by side

|                      | SEO                         | GEO                                                 |
|----------------------|-----------------------------|-----------------------------------------------------|
| Goal                 | Rank on Google for keywords | Get named in AI answers for buying prompts          |
| Primary signal       | Backlinks, on-page keywords | Entity clarity and citation signals                 |
| Where you work       | Your own website            | Your website, plus everywhere else you're mentioned |
| Measured by          | Rankings and clicks         | Mention rate and share of voice                     |
| Keyword stuffing     | Moderately effective        | Actively counterproductive                          |
| Third-party mentions | Helpful                     | Essential                                           |

## Three types of AI you're optimizing for

---

- Pure LLMs — ChatGPT without search, Claude, Gemini in chat. Answers come from training data. This is the long game: brand building.
- AI search engines — ChatGPT with search on, Google AI Overviews, Gemini. LLM plus a live web crawl. This is the fast lane: content and citations move the needle here within weeks.
- Embedded AI — Siri, Copilot, Notion AI. These pull from the same sources as the two above, so the work you do for either one covers this automatically.

**One thing unifies all three: clarity, credibility, and consistency. The model needs to be confident about who you are, what you do, who you serve, and why you're credible. Businesses that nail those three things across the web get mentioned. Everything from here is about building them.**

## CHAPTER 02

# How AI models decide

It isn't a black box. Researchers tested 10,000 queries, and the patterns are clear enough to act on.

The single most useful piece of research behind this guide comes from a Princeton University and IIT Delhi study — 10,000 real queries tested across multiple domains, measuring exactly which optimization methods move the needle on AI visibility. Presented at KDD '24 in Barcelona, it's the closest thing GEO has to a rulebook.

## The four signals AI weighs

---

- Entity clarity — how confidently the model "knows" your brand: your category, the problem you solve, who you serve. Inconsistent descriptions across your own channels create ambiguity, and ambiguous brands get skipped.
- Citation signals — are other credible sources writing about you, reviewing you, listing you? Adding credible citations to a page measurably lifted visibility in the study.
- Structured content — models pull clean, extractable chunks. Plain, well-organised writing consistently outperformed dense or jargon-heavy copy.
- Freshness — for AI search platforms specifically, content updated in the last 90 days was cited far more often than stale pages. Keep your core pages current.

## GEO methods, ranked by impact

---

Two moves produced the largest measurable lift in the study: adding real quotations from credible sources, and adding real statistics with a named source. Neither requires a developer. Fluency and readability improvements were close behind. Keyword stuffing measurably hurt visibility — models parse language semantically, not by frequency count.

## What doesn't work

---

- Keyword stuffing — explicitly tested as one of the worst-performing tactics, reducing visibility in some cases.
- Long, low-quality backlink lists — a thousand thin links to thin content won't earn you an AI mention. Destination quality outweighs link count.
- Vague authority claims — "industry-leading," "world-class team." Models recognise hollow marketing language the same way a sceptical buyer does. Specific, verifiable claims win in every version of this study.

**+115%**

visibility lift for a site ranked fifth in traditional search, from the single Cite Sources method. You don't need to be the market leader — you need to be the best-structured, most clearly described option in your tier.

*Princeton GEO study, KDD '24*

## CHAPTER 03

# Find your prompts

Not keywords — full, conversational questions. Build a list of 20 to 30. Everything else in this guide builds on it.

Before you optimise anything, you need to know what your buyers are actually typing into an AI tool. These are full, conversational questions — closer to how a client actually talks to you in a first call than to a search-bar keyword:

- “Who’s a good business consultant for a 15-person manufacturing firm in Gauteng?”
- “Which firm handles both lean process work and RFQ management for SMEs?”
- “Is [competitor] worth it, or is there someone better for a multinational subsidiary?”

## The buying journey, in three layers

---

- Awareness — someone learning what the category even is. “What does a business systems consultant actually do?”
- Comparison — someone choosing between named options. “Best lean consulting firms for SMEs in South Africa.”
- Decision — almost ready to buy, looking for confirmation. “Is [Firm] worth the retainer?”

## Five ways to build your prompt list

---

- Think like the buyer, not the marketer. Not “best consulting firm,” but “I need someone who can fix our procurement process without a six-month engagement.” Specific prompts are less competitive and easier to win.
- Watch autocomplete. Type your category into ChatGPT or Gemini and see what it suggests — that reflects real usage patterns, not guesswork.
- Mine your own sales calls. Whatever a prospect asks in the first fifteen minutes is the real buyer's language. Write it down verbatim.
- Run competitor comparisons. “[You] vs [Competitor]” is the highest-intent prompt in any category. Run it yourself and see who shows up — and where the gap is.
- List the category-defining questions. The foundational questions buyers ask before they know which firms exist. Getting into these answers builds awareness early, before the buyer is even comparing.

**Sort your finished list into the three buckets above, then prioritise comparison and decision prompts first — that's where the buying actually happens.**

## CHAPTER 04

# Run your own audit

Take your prompt list and run the experiment. This is your baseline — where you stand before you change anything.

Run each of your 20 to 30 prompts, one at a time, across every platform your buyers actually use. For each one, record whether you appear, where, the exact language used to describe you, and which sources the model cites.

- ChatGPT
- Claude
- Gemini
- Grok
- DeepSeek

## Build the audit sheet

| Prompt                           | Appears? | Position | How AI describes you          |
|----------------------------------|----------|----------|-------------------------------|
| "Best lean consultant for SMEs"  | Yes      | 2nd      | "process improvement firm"    |
| "RFQ management partner Gauteng" | No       | —        | —                             |
| "Ameliorate vs [Competitor]"     | Yes      | 1st      | "full-stack business systems" |

## What you'll likely find

- You appear inconsistently — visible on ChatGPT, invisible on Gemini, present for one prompt but absent from five nearly identical ones.
- Unexpected competitors win — smaller or newer firms that simply did more GEO-friendly work outrank you. This is fixable, not permanent.
- The language isn't yours — competitors get specific, differentiated descriptions; you get generic filler. The model hasn't found a clear differentiator to extract.
- Brand prompts work, category prompts don't — people find you by name but not by "best [category] firm." This is the most common gap, and the highest priority to close.

**Run every comparison and decision prompt five times, at different times of day, and average the results into a single share-of-voice number. That number becomes your baseline — and your new north star metric.**

## CHAPTER 05

# Fix what's costing you

Most businesses have the same five problems. Here's what each looks like, why it happens, and exactly what to do.

## 01 — The AI doesn't clearly know what you do

Perhaps you appear in broad searches but not specific ones, described in language vague enough to fit almost anyone.

Why: your homepage, About page, and social profiles each describe you slightly differently. The model can't synthesise one clean picture from three conflicting ones.

**The fix: write one positioning statement and use it everywhere — homepage, About, Google Business, LinkedIn, meta description. “[Business] is a [category] for [customer] that helps them [outcome]. Unlike [alternative], we [differentiator].” Write it for the model first, the human second.**

## 02 — You're not mentioned anywhere outside your own site

---

Perhaps zero or one AI mention per prompt, despite years in business and genuinely happy clients.

Why: it's all owned media. Models weight third-party mentions far more heavily — they're independent evidence that someone else found you credible.

**The fix, 30-day sprint: get listed and fully completed on every relevant review platform. Pitch ten articles that already list your competitors but not you. Get quoted once in an industry publication. Every mention compounds.**

## 03 — You have no structured FAQ content

---

Perhaps you appear in almost none of the informational or comparison answers, even where you're clearly relevant.

Why: models prefer question-and-answer structure because it's a clean, pre-extracted chunk. An FAQ page is already in the model's preferred shape.

**The fix: build FAQ pages that answer the prompts from your Chapter 3 list, in the buyer's exact phrasing, answered in the first one or two sentences, under 200 words, with FAQ schema markup added.**

## 04 — Your content doesn't directly answer buying questions

---

Perhaps you rank for Google keywords but appear in almost no AI comparison or decision prompts.

Why: your blog was written to rank for keywords, not to answer the specific question a buyer types into an AI tool.

**The fix: one page per high-priority prompt. Title mirrors the prompt. Answer it in the first paragraph. Support with two or three pieces of evidence. Add a short FAQ and schema. Aim for 600–900 words, and re-run the audit in 30 days.**

## 05 — You have no schema markup

---

Schema is structured code that tells an AI crawler exactly what a page contains, rather than leaving it to infer. Without it, every page is interpreted in isolation, which creates inconsistency. Add three types at minimum:

- Organization — who you are: name, category, location, verified profiles.
- FAQ — what you answer, feeding your Q&A directly into retrieval.
- Product/Service — what you sell, who it's for, and what it costs.

## CHAPTER 06

# Write content AI cites

There's a format models consistently prefer when pulling content into an answer. Once you see it, it changes how you write everything.

## 11The four qualities of citeable content

- One topic per page — a page covering five topics can't be confidently cited for any of them. One clear topic, one clear answer.
- Answer first, expand second — the inverted pyramid. State the answer in the first paragraph; the rest of the page is evidence and context.
- Real evidence, named — “Semrush found AI-search visitors convert at 4.4× the organic rate” beats “studies show conversion improves.” Name the source every time.
- Natural language for the actual question — use the buyer's real phrasing and its natural variations, because you're genuinely answering that question, not stuffing a keyword.

## The format that outperforms

|                     |                                                                      |
|---------------------|----------------------------------------------------------------------|
| <b>Title</b>        | The specific question you're answering                               |
| <b>Opening</b>      | Direct answer in two to four sentences — this is what gets extracted |
| <b>Section 1</b>    | Why the answer is correct: evidence, data, context                   |
| <b>Section 2</b>    | Qualifications and nuance — this is what builds trust                |
| <b>Section 3</b>    | How to act on the answer: practical next steps                       |
| <b>FAQ + schema</b> | Three to five related questions, marked up with FAQ + Article schema |

## What gets cited, and what doesn't

- Frequently cited: FAQ pages with clear Q&A, “A vs B” comparison pages, numbered how-to guides, pages with sourced statistics, case studies with real numbers.
- Rarely cited: long opinion pieces with no structure, product pages written as marketing copy, keyword-stuffed intros, pages trying to answer too many questions at once, set-and-forget pages nobody's touched in a year.

**Updating existing content? For your top ten pages, check each one: does it answer a specific question directly, include sourced statistics, have an FAQ section, read clearly without jargon, and was it updated in the last six months? Fix whichever answers are “no.”**

## CHAPTER 07

# Build citation presence

Your website is one input. AI models weight third-party sources heavily — in some cases, more heavily than your own content.

Ahrefs research found that 28% of the pages most frequently cited by ChatGPT carry zero Google organic visibility. They rank nowhere in traditional search but appear constantly in AI answers, because they live on domains the model has learned to trust. Earning a place on those domains is the off-site half of GEO work.

## The citation hierarchy

---

- Tier 1 — highest value: Wikipedia, major publications, review platforms (G2, Capterra, Trustpilot), formal research.
- Tier 2 — high value: industry newsletters, comparison sites, active subreddits, credible LinkedIn voices in your space.
- Tier 3 — supporting value: guest posts, podcast appearances, YouTube descriptions, forum threads.

## A 30-day citation sprint

---

- Week 1 — review platforms: claim and fully complete every profile. Request reviews from 5–10 recent clients. Respond to every existing review.
- Week 2 — outreach: find 5–10 recent articles that list competitors but not you. Offer the author a specific data point or case study, not a generic pitch.
- Week 3 — original data: publish one piece of original research. When you're the source of a statistic, everyone who cites it references you.
- Week 4 — media: identify two journalists who cover your space. Engage genuinely, then offer one useful perspective — no immediate ask.

### #1

most-cited single domain across major AI search platforms is Reddit — ahead of YouTube, Wikipedia, and LinkedIn. The lesson isn't to spam it; it's to participate as a genuine expert. Answer questions, help people, don't pitch. Over time, those contributions become part of the pool AI pulls from.

*AI search citation research, 2025–2026*

## CHAPTER 08

# Track it, repeat it

GEO isn't a project with an end date. Two hours a month keeps you moving forward.

## Your monthly GEO routine

- Re-run your prompt list (05 min) — same prompts, same platforms. Update the spreadsheet. Where did you gain or slip, and what changed?
- Check citation sources (15 min) — which pages are getting referenced, and which competitors are cited instead?
- Publish one new page (90 min) — there's always at least one prompt where you should appear but don't. Write the page that answers it.
- One citation outreach action (15 min) — one review request, one article pitch, or one journalist reply. Small actions compound over a year.

## The metrics that matter

| Metric                       | Priority   | What it tells you                                                                  |
|------------------------------|------------|------------------------------------------------------------------------------------|
| Share of voice               | Primary    | Out of ten relevant prompts, how many name you? Your single most important number. |
| Position in answer           | Secondary  | First mention carries far more weight than being buried third in a list.           |
| Consistency across platforms | Tertiary   | Strong on one platform but absent on others signals a content or citation gap.     |
| AI referral traffic          | Supporting | Small today for most businesses, growing fast, and converting well above organic.  |

**When to push harder:** more than 20% of category answers consistently go to one or two competitors; your category shifts from “who do I call” to “who does AI recommend”; or you start hearing “ChatGPT said I should call you” from inbound leads. First-mover advantage in GEO is real — moving now is cheaper than catching up later.

# The data behind GEO

This guide is built on published research, not house opinion. Here are the studies referenced throughout, for your own reading.

## Princeton University / IIT Delhi GEO Study — KDD '24

---

Aggarwal, Murahari, Rajpurohit, Kalyan, Narasimhan, Deshpande · ACM SIGKDD, Barcelona, Aug 2024

- GEO methods lifted visibility in AI answers by up to 40% in aggregate
- Credible-source citations: up to 42% increase in specific cases
- Relevant statistics: up to 35% increase · Fluency optimisation: 15–20%
- Keyword stuffing: measurable negative effect · Lower-ranked sites benefited most

## Large-scale LLM correlation study — finance & SaaS

---

- 0.65 correlation between Google organic rankings and LLM brand mentions
- Backlinks: weak-to-neutral correlation · Page-one ranking modestly lifts AI appearance
- Confirms SEO and GEO are complementary disciplines, not competing ones

## Other cited research

---

- AI search citation research (2025–2026): Reddit is consistently the most-cited single domain; Wikipedia ≈10% of citation links; over 50% of cited links are brand-owned.
- Semrush AI Search (2025): AI search converts at 4.4× organic; referral visits up 50% YoY; 78% of decision-makers building AI-search budget.
- Ahrefs AI vs Search (2025): 28% of top ChatGPT-cited pages have zero Google visibility; AI and search visibility are largely uncorrelated.

# Your GEO checklist

Use this to audit where you stand today and prioritise what to fix next.

## Entity clarity

---

- ☐ One clear positioning statement, consistent across every channel
- ☐ Homepage states what you do, for whom, and the outcome
- ☐ About page uses the same description in the same language
- ☐ Google Business & LinkedIn match your website positioning
- ☐ Consistent business name everywhere — no variations

## Content structure

---

- ☐ Priority pages answer the question directly in the first paragraph
- ☐ FAQ page or sections for your top buyer questions
- ☐ Comparison pages for your top competitor match-ups
- ☐ FAQ, Organization & Product schema implemented
- ☐ Core pages updated within the last 90 days
- ☐ Content includes statistics with named, cited sources

## Citation presence

---

- ☐ Profile claimed and completed on relevant review platforms
- ☐ Recent client reviews requested
- ☐ Mentioned in 3+ industry publications or roundups
- ☐ Listed in relevant directories for your category
- ☐ Active presence in one relevant community
- ☐ No inconsistencies in how third parties describe you

## Tracking

---

- ☐ Prompt list of 20–30 real buyer questions created
- ☐ Baseline audit run across ChatGPT, Claude, Gemini, Grok
- ☐ Monthly audit schedule established
- ☐ AI referral traffic tracked in analytics
- ☐ Competitive share-of-voice baseline set

# Questions everyone asks first

## How long does GEO take to show results?

---

For live AI-search platforms — ChatGPT with search, Gemini, Google AI Overviews — you can see change within four to eight weeks of improving content and earning citations. Pure LLMs depend on the next training update, which takes longer. Focus on AI search first; it's the pragmatic starting point.

## Do I need to start SEO over?

---

No. Good SEO helps GEO — research shows a 0.65 correlation between organic Google rankings and LLM brand mentions. Your existing content and ranking investment is a GEO asset. You're building on it, not replacing it.

## Is GEO just another buzzword?

---

The acronym is new; the underlying shift is real and measurable. Buyers are using AI to make purchase decisions. Businesses inside AI answers get considered; businesses that aren't, don't. Whatever you call it, optimising for it has a clear, quantifiable payoff.

## Which platform should I focus on first?

---

Start with ChatGPT (search enabled) and Gemini — the two sending the most engaged referral traffic today. Google AI Overviews matter most for local and consumer businesses; Gemini is growing fastest in B2B. Good content and citations tend to cover most platforms at once.

## Can a small business compete with big national brands?

---

Yes — often more effectively than in traditional SEO. The Princeton research found that lower-ranked sites benefited disproportionately from GEO work. Clarity, specificity, and structure matter more here than domain authority or ad budget.

## How much does this actually cost?

---

The core work — positioning, content structure, schema, citations — is mostly labour, not ad spend. Unlike paid traffic, the improvements compound over time. You're building an asset, not renting visibility.

## CLOSING NOTE

# You know the loop. Now run it.

Find your prompts → Audit where you stand → Fix what makes AI skip you → Write content AI can cite  
→ Earn citations → Repeat monthly

Businesses that make this a habit now build a compounding advantage that's genuinely hard for late movers to close. The work isn't complicated — it takes clarity and consistency more than budget or technical skill.

*Don't want to run this manually? That's what Ameliorate is for.*

## This guide was produced by Ameliorate

the business systems partner that helps South African SMEs get found where buyers are actually looking — across search, AI answers, and everything in between.

## SOURCES & REFERENCES

Aggarwal et al. "GEO: Generative Engine Optimization." KDD '24, ACM, Barcelona, 2024. · Large-scale LLM correlation study, finance & SaaS. · AI search citation research, 2025–2026. · Semrush AI Search Statistics 2025. · Ahrefs AI vs Search Traffic Analysis 2025. · Search Engine Land, Jan 2026. · Gartner Marketing Leader Survey, Jun 2025.