

The GEO Playbook

The GEO Playbook

The 26 things you change on your website this week so AI engines actually cite you.

By Artem Novitckii

Read this first

Google just rolled out AI Mode. It's the first real overhaul of Search in 25 years and it's already eating your traffic.

The page-one blue links you've been optimizing for aren't where the answers live anymore. The answer is the AI summary at the top. And the AI summary cites 3-5 sources, not 10. If your site doesn't get cited, you don't get the click. If you don't get the click, the rest of your funnel doesn't fire.

This is true whether you're a SaaS company, a creator, an agency, a consultant, a blog, or an e-commerce brand. SEO got you ranked. **GEO gets you cited.**

I built a Claude skill (`geo-audit`) that scans any URL and produces a 100-point score plus a 15-point entity score with the exact fixes ranked by impact. Get it at the end of this playbook. But everything below works whether you install the skill or not. These are the 26 moves. Most sites are missing 18-22 of them.

Don't try to do all 26 at once. Pick the top 3 by impact, ship those this week, re-run the audit in 4 weeks. Repeat.

The shift in one paragraph

SEO was: write a page → rank #3 → get clicks. **GEO is:** write a page → get *cited* in an AI summary → get clicks.

You don't need to rank #1 to get cited. A page that's *extractable* and *attributable* can get cited over a page that ranks higher but reads like a corporate brochure. Two kinds of extractable:

- **Content extractable** — AI can pull a clean, self-contained answer from your page.
- **Entity extractable** — AI knows *who or what wrote it* and can connect it to your other platforms and offers.

Most sites fail both. Great content buried under fluff intros. No entity schema. No `sameAs` links. A "Subscribe" button with no name on what's behind it. AI summarizes what's there, but doesn't credit you, and can't surface your offer when someone asks for it.

Part 1 — Base GEO Score (100 pts)

The first 20 moves cover what Google's official AI optimization guide actually measures. These apply to every site type — SaaS, creator, agency, e-commerce, blog, doesn't matter.

Category 1 — Google AI Eligibility (25 pts)

These are the hard gates. Get them wrong and nothing else matters. Your content can be the best on the internet and you'll still be invisible.

Move 1 — Don't block Googlebot

Open your `robots.txt` and look for `User-agent: Googlebot` followed by `Disallow: /`. If it's there, delete it. This is the single biggest mistake sites make. A `Disallow: /` under Googlebot blocks the entire site from AI Overviews, AI Mode, and Search itself.

Bad:

```
User-agent: Googlebot
Disallow: /
```

Good:

```
User-agent: *
Allow: /
Sitemap: https://yoursite.com/sitemap.xml
```

Move 2 — Don't `noindex` pages you want cited

Check the `<meta name="robots">` tag in your page source AND the `X-Robots-Tag` HTTP response header. If either says `noindex`, the page is invisible to all Google AI features even if your content is perfect. This gets set accidentally all the time —

staging environments, CMS defaults, copy-pasted templates.

Bad: `<meta name="robots" content="noindex, nofollow">`

Good: `<meta name="robots" content="index, follow">` (or no robots meta at all — same effect)

Move 3 — Don't `nosnippet` your core content

`nosnippet`, `max-snippet:0`, and `data-nosnippet` wrappers around your main content block AI engines from quoting you. If they can't quote you, they can't cite you. The whole game is being quotable.

Bad: `<div data-nosnippet>This is your main content...</div>`

Good: Drop the `data-nosnippet` attribute. Or only use it on legal disclaimers, ads, and footer copy you genuinely don't want excerpted.

Move 4 — Set a self-referencing canonical

`<link rel="canonical" href="...">` tells Google "this URL is the canonical version of this content." Without it, duplicate-content signals can split your authority across variants (with/without trailing slash, with tracking params, etc.).

Bad: No canonical link at all, or one pointing to the homepage on every page.

Good: `<link rel="canonical" href="https://yoursite.com/this-exact-page/">` on every page.

Category 2 — Helpful Non-Commodity Content (25 pts) — the highest-leverage category

Google says this out loud in their guide: *"Success often requires no overt SEO at all."* What they mean: AI engines preferentially cite content that couldn't have been written by another AI summarizing the web. This is the hardest category to fix because no plugin solves it. You have to actually have something to say.

Move 5 — First-hand expertise visible in the first 200 words

Specific numbers from your work. Named tools you actually use. Opinions with reasoning. If the first paragraph could've been written by ChatGPT from the headline alone, you're commodity content and AI will summarize the commodity version instead of citing yours.

Bad:

"Email marketing is one of the most effective digital marketing strategies. It allows businesses to reach their customers directly and build strong relationships."

Good:

"We A/B tested 47 subject lines last quarter. The ones that started with a question outperformed statements by 34%. Here's the exact framework we use now, why most of the 'industry best practices' are wrong, and the three subject line patterns we won't touch again."

Move 6 — Answer the implied question completely

A page titled "Best CRM for solopreneurs" should answer that question AND the 2-3 obvious follow-ups: pricing tiers, integrations, when to switch, what breaks at scale. If the reader has to search again, AI cites the page that doesn't make them search again.

Bad: A 400-word page that says "the best CRM depends on your needs" and links to a comparison tool.

Good: A 2,000-word page with named recommendations for 5 user types, real pricing, real integration limits, real downsides, and a decision matrix at the bottom.

Move 7 — Specific examples, not generic claims

"Increased conversion" is invisible. "Took our signup rate from 3.1% to 5.8% by removing the second password field" is citable. AI engines are quote machines — they preferentially cite content with quotable specifics because those specifics make the AI's summary more credible.

Bad: "We helped them grow their revenue significantly."

Good: "We grew their MRR from \$14k to \$87k in 9 months by switching from a free trial to a \$1 onboarding offer. Trial conversion went from 3.2% to 18%."

Category 3 — Content Organization & Extractability (20 pts)

Move 8 — One H1 that matches the implied query

Not "Welcome" or "Our Services." If the page is about "the 5 mistakes founders make on landing pages," your H1 says that. Exactly that. Multiple H1s confuse the parser; generic H1s waste the strongest signal you have.

Bad: `<h1>Welcome to our blog</h1>`

Good: `<h1>The 5 landing page mistakes that cost founders their seed round</h1>`

Move 9 — Headers as questions, not labels

Instead of `<h2>Pricing</h2>` write `<h2>How much does Linear cost?</h2>`. AI looks for question-answer pairs. Give it exactly that, and your H2 plus the paragraph below it becomes the answer block AI quotes. This single change can double your citation rate.

Bad:

- `<h2>Pricing</h2>`
- `<h2>Features</h2>`
- `<h2>Support</h2>`

Good:

- `<h2>How much does Linear cost?</h2>`
- `<h2>What features ship in the Free plan?</h2>`
- `<h2>How do I contact Linear support?</h2>`

Move 10 — Direct answer in the first paragraph, under 80 words

Treat your opening paragraph like a featured snippet. Define the thing, name the answer, then expand for the rest of the page. AI engines lift this paragraph as their summary — make sure it's a summary worth lifting.

Bad:

"In today's fast-paced digital landscape, founders face unprecedented challenges when it comes to landing page optimization. With attention spans shrinking and competition increasing..."

Good:

"The 5 landing page mistakes that kill founder conversion are: long forms, vague headlines, no social proof above the fold, multiple CTAs, and a hero image that doesn't show the product. Fix these in order — each one moves your conversion rate 5-15%."

Move 11 — Real FAQ section with 3+ Q&A pairs

Use `<details>` accordion elements or a `<dl>` Q&A list. Each question is a separate query that AI can match. A page with 5 FAQ items is effectively 6 pages of citation surface. Most sites skip this because they think FAQs are tacky — AI engines love them.

Bad: A "Questions?" link to your contact page.

Good:

```
<details>
  <summary>How long does setup take?</summary>
  <p>Most teams are live in under 20 minutes. The longest step is
    importing existing tickets from Jira/Linear/Asana — that takes
    5-15 min depending on volume.</p>
</details>
<details>
  <summary>Can I switch plans mid-month?</summary>
  <p>Yes. Upgrades are prorated to the day. Downgrades take effect at
    the next billing cycle.</p>
</details>
```

Move 12 — Internal links ≥ 3 to related content

Pages with zero internal links read as orphans. AI engines weight "is this woven into a real site?" as a quality signal. Three contextual links — not a sidebar of recent posts — are the minimum.

Bad: A blog post that ends with no links anywhere.

Good: A blog post that mentions "we wrote about this in [our pricing teardown](#)" and "this builds on [the SPCL framework](#)" in the body, with 1-2 more contextual links sprinkled through.

Category 4 — Technical Structure (15 pts)

Move 13 — Use `<main>` and `<article>` semantic HTML

Not `<div class="main">`. Browser-based AI agents (Google AI Mode with browsing, ChatGPT Atlas, Claude with browsing) look at the accessibility tree first. `<main>` tells them "the page content starts here." Without it, they're guessing — and guessing often returns wrong content.

Bad: `<div id="content"><div class="article">...</div></div>`

Good: `<main><article>...</article></main>`

Move 14 — Title tag 50-60 chars + meta description

Both are still pulled into AI summaries as fallbacks when the page itself doesn't have a clean answer paragraph. Title that names the thing, meta description that previews the answer.

Bad: `<title>Home</title>` (8 chars, useless)

Good:

```
<title>The GEO Playbook — 26 fixes to get cited by AI (2026)</title>
<meta name="description" content="The exact website fixes that get AI
engines (ChatGPT, Perplexity, Gemini) to cite your site. 26
specific changes, ranked by impact.">
```

Move 15 — Alt text on every image

Multimodal AI is reading your images now. Decorative images get `alt=""`. Content images get a real description. Charts, screenshots, and product photos especially — those are the ones AI engines describe back to users.

Bad: ``

Good: ``

Move 16 — Open Graph tags

`og:title`, `og:description`, `og:image` — used when AI engines render link cards in chat interfaces. ChatGPT's link previews use OG. Perplexity uses OG. Without these, your shares look bare and your AI-rendered link cards have no image.

Good:

```
<meta property="og:title" content="The GEO Playbook by Artem Novitckii">
<meta property="og:description" content="26 fixes to get cited by AI
engines.">
<meta property="og:image" content="https://novitckii.com/og-geo-
playbook.png">
<meta property="og:type" content="article">
```

Category 5 — Entity & Enhancement Signals (15 pts)

Move 17 — JSON-LD schema that matches what's on the page

Article schema on articles, Product schema on products, Course schema on course pages, FAQPage schema on FAQ pages. Don't fake it (Google detects this) — but don't skip it either. Schema is how AI engines confirm "this page is actually about what it claims to be about."

Good for an article:

```
{
  "@context": "https://schema.org",
  "@type": "Article",
  "headline": "The GEO Playbook",
  "author": { "@type": "Person", "name": "Artem Novitckii" },
  "datePublished": "2026-05-31",
  "dateModified": "2026-05-31"
}
```

Move 18 — Named author + visible publication/updated date

Undated content loses to dated content in AI citation, period. "Updated: May 2026" beats no date every time. Named authors beat "Admin" or "Staff Writer" by a wide margin. AI engines are trained to prefer recent, attributed sources.

Bad: A blog post with no author and no date.

Good: "By Artem Novitckii · Published April 8, 2026 · Updated May 31, 2026"

Move 19 — Statistics with source attribution

Per the Princeton GEO research (KDD 2024), adding cited statistics increases AI citation visibility by 37%. Adding sources increases it by 40%. Real numbers from named sources are AI engine catnip.

Bad: "Most founders fail."

Good: "Of the 2,847 startups that raised pre-seed in 2024, 67% shut down before raising a Series A (CB Insights, 2025 Year in Review)."

Move 20 — Don't block cross-platform AI bots

GPTBot, PerplexityBot, ClaudeBot, Bingbot. These are how ChatGPT/Perplexity/Claude/Copilot crawl. Block them in `robots.txt` and you're invisible across half the AI internet. Many sites block these "to protect IP" — that protection costs you every citation those engines would've sent.

Bad:

```
User-agent: GPTBot
Disallow: /
User-agent: PerplexityBot
Disallow: /
```

Good:

```
User-agent: *
Allow: /
```

(Default-allow is the right call unless you have a very specific legal reason to block.)

Part 2 — Entity & Offer Extractability (15 pts, separate score)

The first 20 moves get the AI engines TO your site. These last 6 make sure they can identify **who or what** you are when they get there — and surface your offer when someone asks for it. This is where most personal-brand and SaaS sites lose the citations they should be winning.

Category 6 — Entity & Offer Extractability

Move E1 — Add entity schema (Person for creators, Organization for companies)

Most sites have nothing here. Add a `Person` JSON-LD block for personal brands or an `Organization` block for companies. This is THE single signal that tells AI engines "this site represents [identified entity]." Without it, you're an unnamed page even if you have great content.

Personal brand example:

```
{
  "@context": "https://schema.org",
  "@type": "Person",
  "name": "Artem Novitckii",
  "url": "https://novitckii.com",
  "image": "https://novitckii.com/headshot.jpg",
  "jobTitle": "AI Content Systems for creators & founders",
  "sameAs": [
    "https://instagram.com/artem.novitckii",
    "https://youtube.com/@artem.novitckii",
    "https://www.linkedin.com/in/artem-novitckii/"
  ]
}
```

Company example:

```
{
  "@context": "https://schema.org",
  "@type": "Organization",
  "name": "Linear",
  "url": "https://linear.app",
  "logo": "https://linear.app/logo.png",
  "description": "The issue tracker for fast-moving software teams.",
  "sameAs": [
    "https://linkedin.com/company/linear",
    "https://x.com/linear"
  ]
}
```

Move E2 — Add `sameAs` links to all your platforms

LinkedIn, X, IG, YouTube, GitHub, Crunchbase, Wikipedia if you have one. This is the identity bridge — it tells AI engines "the @username on X, the LinkedIn page, the GitHub org, and the author of this site are the same entity." Without `sameAs`, your identity is fragmented across the internet and AI can't stitch it together. Three or more external URLs is the threshold.

Bad: Person schema with just `name` and `url`.

Good: Person schema with `sameAs: ["https://instagram.com/...", "https://youtube.com/@...", "https://x.com/...", "https://www.linkedin.com/in/..."]`

Move E3 — Name your primary offer in your hero section

Don't write "Book a demo" or "Subscribe." Write "Book a demo of *Linear* — the issue tracker used by Vercel, Cash App, and Ramp" or "Join *The Weekly Drop* — a newsletter for AI builders, 47k+ subscribers." When ChatGPT is asked "how do I get [brand]'s [thing]?", it can only answer if the offer is named and extractable in your visible text. Generic CTAs are invisible to AI summaries.

Bad:

[Subscribe] [Get Started] [Learn More]

Good:

*Join The Weekly Drop — 47k builders get one tactical AI workflow every Friday.
[Subscribe free]*

Move E4 — Put a "Who is [name] / What is [brand]" block on every important page

Not just on [/about](#). AI engines crawling a single URL won't follow internal links to figure out who you are. They need it inline on every page that should be citable for entity queries. One sentence plus one specific proof point on every important page is enough.

Bad: A homepage that says "Welcome — see what we do" with no identifying content.

Good: A homepage hero that says: "Hi, I'm Artem. I help creators, founders, and operators build AI content systems. 10K+ creators have used my templates. 100+ systems shipped. 3 years building in public."

Move E5 — Use named voice with a visible byline

First-person "I" for personal brands or first-person "we" for companies — both beat passive third-person institutional voice. "I tested 12 CRMs" beats "We tested 12 CRMs" beats "12 CRMs were tested." Pair the named voice with a visible author or attributable team so AI knows whose opinion this is.

Bad: "12 CRMs were evaluated using a standardized methodology developed by our research team."

Good: "I (Artem) tested 12 CRMs over 6 weeks with my agency clients. Here's what actually worked." — with a visible author byline at the top.

Move E6 — Put at least one original artifact on every page that matters

A customer case study with named results. A named methodology. A transcript. Show notes. Original data. These are what AI engines lift as direct quotes. A SaaS landing page about "automated billing" that doesn't include a single customer story is generic. The same page WITH "Ramp scaled from \$1M to \$100M ARR on our billing infrastructure" becomes citation surface for every related query.

Bad: A landing page that lists 8 generic benefits with no proof.

Good: A landing page with one named customer case study, one named framework ("The 3-Step Billing Migration"), and one piece of original data ("We process \$2.4B/year across 1,200 customers").

What does NOT work (Google says so directly)

These show up in every GEO blog post but Google has explicitly debunked them in their official AI optimization guide. Don't waste time on:

- **llms.txt files.** Google states they have no effect on AI search. Don't bother.
 - **Special "AI-targeted" schema.** Standard schema.org types are what matter. There's no `AIOverviewOptimized` schema. The schema you already know is the schema you need.
 - **Chunking your content into bite-sized pieces for AI.** Write for humans. AI ingests what humans read.
 - **Rewriting everything specifically for AI.** Google's exact words: *"Success often requires no overt SEO at all."*
 - **Treating `Google-Extended` as an AI Overview block.** `Google-Extended` is a model-training directive, not an AI Overviews gate. AI Overviews are governed by `Googlebot + noindex / nosnippet`. Blocking `Google-Extended` does nothing to your AI search visibility.
-

How to use this playbook

Don't try to do all 26 moves at once. The fastest path:

1. **This week** — fix the eligibility blockers (Moves 1-4). Most sites pass these by default, but the 5% that don't are losing 100% of their AI traffic. 30 minutes.
2. **Next week** — add entity schema (Moves E1, E2). This is the single highest-leverage fix for any site with a brand or person behind it. 1-2 hours.

- 3. Week 3** — restructure your top 3 pages: H1 + first-paragraph answer + FAQ + named offer (Moves 8, 10, 11, E3). 3-4 hours.
- 4. Week 4** — fill in technical gaps: Open Graph, alt text, dates, author bylines (Moves 14-18). 2-3 hours.
- 5. Re-audit and prioritize what's left.**

Re-run the audit 4-6 weeks after shipping fixes. That's roughly how long AI engines take to re-crawl, re-index, and re-evaluate.

Track citations, not just rankings. Search Console shows AI Overview traffic under the standard Web search type — there's no separate AI report. Look at impressions + clicks for your target queries over 90-day windows.

Get the audit skill

Reading the playbook is one thing. Running it on your actual site and getting a scored report with the exact fixes ranked by impact is another. That's what the `geo-audit` Claude skill does.

It opens any URL in a real browser, checks all 26 signals from this playbook, and hands back a 100-point base score plus a 15-point entity score with:

- The top 3 highest-impact fixes (do these first)
- Every failing check with the exact recommended fix
- An evidence bundle so you can verify nothing was hallucinated
- A 4-week action plan tiered by effort

Comment SEARCH on [@artem.novitckii](#) on Instagram and I'll DM you the skill + this playbook as a PDF.

If you've already got the skill, run it on your homepage and one important inner page. DM me your scores — I'm tracking which categories most sites fail.

Sources

[Google's AI optimization guide](#) · [Google's AI features docs](#) · [Princeton GEO research \(KDD 2024\)](#).

— Artem Novitckii, May 2026