

OZVOR · METHODOLOGY

The LLM Citation Tracker

A transparent methodology for measuring whether AI engines cite your brand.

2026 Edition

Contents

1. A Spreadsheet + Methodology for Monitoring When ChatGPT, Claude, Perplexity & Gemini Mention Your Business — in 10 Minutes a Week

2. PART A — The Methodology

3. PART B — The Spreadsheet Build Spec

4. Sources & references

A Spreadsheet + Methodology for Monitoring When ChatGPT, Claude, Perplexity & Gemini Mention Your Business — in 10 Minutes a Week

By Ozvor · 2026 · ozvor.com

The one-sentence version: AI answer engines now hand customers two or three business names instead of ten blue links — and if yours isn't one of them, you lost the customer before they knew you existed. This tracker is the cheapest, most honest way to find out whether you're on that shortlist, starting this week, with nothing but a spreadsheet and ten minutes.

PART A — The Methodology

Why this matters (read this part even if you skip the rest)

For about twenty years, getting found online meant one thing: rank on Google. You earned a spot on page one, the customer scrolled ten blue links, and the click was yours to win. That world is closing.

AI answer engines — ChatGPT, Gemini, Perplexity, Claude, and Google's own AI Overviews — now sit *on top of* search. They read the web and hand back a finished answer before a single link is shown. The scale is no longer niche: **OpenAI reported ChatGPT reached roughly 900 million weekly active users in February 2026** (OpenAI, via TechCrunch & Search Engine Land, Feb 2026), and **Google's AI Overviews reach more than 2 billion people a month across 200+ countries** (Alphabet/Pichai Q2 earnings, Jul 2025). The question your customer used to type into Google — "best dentist near me," "which bookkeeper for a small business" — they now ask an AI.

And the AI doesn't return ten options. It returns two or three names.

This is the part that's invisible to the people it hurts most. You can keep your Google rankings, keep doing everything your SEO agency told you to, and still watch your phone ring less — because the AI is now answering the question, and the searcher never scrolls down to your website. Across the open web, **68% of US Google searches ended without a click in early 2026, up from about 60% in 2024** (SparkToro / Rand Fishkin, Jun 2026). The clicks aren't going to page two. They're not going anywhere.

What you're losing isn't "traffic." It's the shortlist — and the shortlist is leads, and leads are money. When an AI names two businesses like yours and you're not one of them, you didn't lose a ranking. You lost the customer.

How bad is the gap for small businesses specifically? The hard numbers:

- **ChatGPT recommends only about 1.2% of local business locations** when asked for a business like yours (SOCi 2026 Local Visibility Index, ~350k locations, Jan 2026). It names *one*. Almost everyone else is invisible.
- In home services, **roughly 87% of independent HVAC and plumbing contractors have effectively zero AI citation share** while a handful of national franchises absorb the leads (5WPR HVAC & Plumbing AI Visibility Index, Q1 2026).
- The flip side — the upside — is just as real: **brands cited inside a Google AI Overview earn about 120% more organic clicks than uncited brands on the same query** (Seer Interactive, 5.47M queries, Apr 2026), and **Microsoft's own data across 1,277 websites found visitors arriving from AI sign up at up to 11x the rate of search visitors** (Microsoft Clarity, Nov 2025). (In fairness, conversion lift isn't universal — a separate 54-site study by Amsive found no statistically significant overall difference between AI and organic conversion, 4.87% vs 4.60% (p=0.794), though most sites and the higher-traffic ones did convert AI traffic better. The honest read: AI visitors *often* convert better, not always.)

Being named isn't a vanity metric. It's pre-qualified demand. The reason this tracker exists is that **most businesses have no idea whether they're being named or not** — only around 16% of brands systematically track AI search performance (HubSpot, citing Sep 2025 data). You're about to be in the minority that actually knows.

The honest caveat (we'd rather you trust us than overclaim)

There is no public API that tells you when an AI names your business. ChatGPT does not notify you. Perplexity sends no alerts. Claude and Gemini offer no citation-monitoring tools to the public. So the only way to know whether you're being cited is to **ask the engines yourself, manually, using the questions your customers would actually type**.

That has real limits, and we're going to name them so you read your own data correctly:

- **You only see the queries you think to run.** If you don't test a question, you have no idea how you appear for it.
- **Answers vary run to run.** Ask ChatGPT the same thing 100 times and it surfaces about 44 different brands, but only ~5 of them appear 80%+ of the time (Search Engine Land / Fishkin, Feb 2026). One sample is a snapshot, not a verdict. The *trend over weeks* is the signal.
- **Personalization and location skew results.** Log out, or use a fresh/incognito session, so you're seeing closer to what a stranger sees — not what the AI has learned about *you*.
- **"Mentioned" and "cited with a link" are different things.** An engine can name you without linking, or link a source without naming you in the prose. Track both.

This is not a perfect instrument. It is the best *free, honest, do-it-yourself* instrument that exists — and it gives you genuine, unmediated feedback no dashboard can fake. Run it with discipline and the pattern becomes undeniable within a month.

The four engines (and why you test all of them)

Each engine pulls from different places and updates on its own clock, so being named in one tells you little about the others.

Engine	What to know	Where it tends to pull from
ChatGPT	The giant — ~900M weekly users (OpenAI, Feb 2026). Runs a live web search on only ~34.5% of queries; the rest answer from training "memory" (per marketing research from Siana Marketing, Q1 2026).	Wikipedia, Reddit, homepages, news, plus live search results when triggered.
Perplexity	Built answer-first; almost always shows its sources, so it's the easiest place to <i>see which URL</i> it cited.	Reddit is its #1 cited domain; heavy on fresh web sources (Profound, 2026).
Claude	Conversational, increasingly used for research and B2B decisions.	Authoritative, well-structured sources; cites less visibly than Perplexity.
Gemini	Google's app — over 750M monthly users (Alphabet, Q4 2025). Grounded in Google's index and Maps for local.	Google Search results and Google Business Profile data (most accurate for local).
Google AI Overview	Not a chatbot but the AI answer box on top of normal Google results — 2B+ monthly users (Alphabet, Jul 2025). Worth tracking as a 5th "engine."	The live Google index; favors content that already ranks.

Why this is winnable for a small business: the discipline is real science, not snake oil. A peer-reviewed study from Princeton, Georgia Tech and the Allen Institute for AI (KDD 2024, arXiv:2311.09735) tested 9 tactics across 10,000 queries and showed you can lift a page's visibility in AI answers by up to 40% — driven by adding credible quotations (+41%), statistics (+33%), and citing authoritative sources (+28% on average, and up to +115% for an underdog page ranked 5th). Keyword stuffing was the one tactic that *backfired* (-8.7%). Translation: AI rewards being the clearest, best-sourced, most current answer — something a small, focused business can absolutely win.

How to build a query bank that actually reflects your customers

The single biggest mistake people make is testing the wrong queries — usually their own business name. **Customers who already know your name aren't the customers you're worried about.** You want to know if you show up for the people who *don't* know you yet. Those people ask questions.

The core rule: **write each query the way a customer would phrase it when they don't know you exist.** If your business name or tagline creeps into the query, rewrite it from the customer's side of the table.

Build your bank across three intent stages, because they behave differently:

1. **Discovery** — broad, "who even does this" questions. "*Best bookkeeper for a small e-commerce business.*" These are the hardest to win and the most valuable when you do.
2. **Consideration** — "how do I choose / what should I look for" questions. "*What should I ask a photographer before hiring them for brand work?*" These are GEO gold: if you've published a post that *is* the answer, you can get cited even with a tiny following.
3. **Decision** — specific, ready-to-act, local or product-specific. "*Same-day cosmetic dentist consultation in [your city].*" High commercial intent; commercial-intent phrasing drives meaningfully more brand mentions (BrightEdge, 2026).

A practical recipe for your bank (aim for 10-15 live queries you actually run each week, drawn from a larger reference list):

- 3-4 **discovery** queries that name your category + your city or niche.
- 3-4 **consideration** queries phrased as "how to choose / what to ask / what's included."
- 3-4 **decision** queries with your specific service + location or specialism.
- 1-2 **competitor-shaped** queries ("alternatives to [the big chain in my category]") to see who the AI defaults to.

The read-aloud test: say the query out loud. If it sounds like something a stranger would type to find a service like yours, it belongs in your bank. If it sounds like a press release about your company, cut it.

Part B of this document ships you a ready-made **Query Bank of 30+ customer-style queries across ten verticals** (dentist, accountant/bookkeeper, law firm, HVAC/plumber, restaurant, SaaS consultant, photographer, designer, fitness/PT, real estate) so you can start from proven phrasing and adapt, instead of staring at a blank cell.

How to read what you find: citations, position, sentiment

When you run a query, look for four things and record each one. (Part B turns every one of these into a spreadsheet column with dropdowns.)

1. Were you cited at all? (Y/N) This is the primary signal — the binary that drives everything on your dashboard. "Cited" means your business is named in the prose *or* your content is linked as a source. Be strict and honest: a generic mention of your *category* is not a citation of *you*.

2. What position? Being named first is worth far more than being buried fourth in a list. Use a simple scale: **1 / 2 / 3 / list / none**. "1" means you're the lead recommendation. "list" means you appear but in a long, unranked rundown. Position trend over weeks tells you whether you're climbing toward the default answer or slipping.

3. What's the sentiment — and is it accurate? Read *how* the engine describes you, using four buckets: **positive / neutral / negative / inaccurate**. The "inaccurate" flag is the most actionable one. If an engine says you're in the wrong city, offer a service you dropped, or get a specialty wrong, it's pulling from stale or thin content — and you can fix the record by publishing accurate, specific, current information. (Note: business-profile info was only ~68% accurate on ChatGPT and Perplexity in one large study, vs ~100% on Gemini, which is grounded in Google Maps — SOCi, 2026. Inaccuracy is common and correctable.)

4. Which source got cited? When an engine links a URL (Perplexity and ChatGPT Search do this most), record it. This is the highest-value data you'll collect: it tells you *exactly which piece of your content the AI is actually retrieving* — so you know what to publish more of. If a competitor's URL is cited instead, record that too. You're mapping the real battlefield.

The competitor column is not optional. Knowing you weren't cited is half the picture. Knowing *who was* — the same three franchises every week, or a rival who just published a great FAQ — tells you what "winning" looks like in your category and how far you have to climb.

The weekly cadence: 10 minutes, same day, every week

Weekly is the right rhythm. Retrieval indexes change constantly — Reddit's share of ChatGPT citations swung from ~60% of responses to ~10% inside six weeks in 2025 (Semrush, 230k+ prompts). Monthly testing is too coarse to act on; daily testing is noise. One disciplined weekly session beats both.

The routine:

1. **Open a fresh/incognito browser session** (so results aren't personalized to you). Log out of the AI accounts where you can.
2. **Run your live queries across the engines** that matter to your audience. You don't have to run all five engines for all queries every week — a sustainable default is your top ~8-10 queries across ChatGPT + Perplexity + Gemini, rotating Claude and Google AIO in.
3. **Log each result as one row** in the Weekly Log (Part B): query, engine, Cited Y/N, position, source URL, sentiment, competitor cited, a one-line note.
4. **Glance at the Dashboard.** Your citation rate, per-engine breakdown, share-of-voice vs competitors, and sentiment mix update automatically.
5. **Done.** Close the laptop. Ten minutes.

After **four weeks**, you have a trend you can trust. A realistic, healthy pattern for a business doing consistent GEO content looks like the example below — nothing for the first couple of weeks, then a first appearance once your content gets indexed and trusted:

Week	Date	Query	Engine	Cited?	Position	Source cited	Sentiment	Competitor cited	Notes
1	2026-05-20	Best Invisalign consultation in [city] centre	Perplexity	N	none	—	—	3 clinics named	None of my content surfaced
2	2026-05-27	Best Invisalign consultation in [city] centre	Perplexity	N	none	—	—	Same 3 clinics	Still no citation
3	2026-06-03	Best Invisalign consultation in [city] centre	Perplexity	Y	list	linkedin.com/in/[profile]/posts	positive	2 clinics	Linked my LinkedIn post from 2026-05-25 on the Invisalign assessment process
4	2026-06-10	Best Invisalign consultation in [city] centre	Perplexity	Y	2	mysite.com/invisalign	positive	1 clinic	Climbed from list to #2; now citing my own page

That climb — none → list → #2 — is exactly what working GEO looks like.

How to act on what you find

The tracker is only worth ten minutes a week if you *do something* with the pattern. Here's the decision tree:

- **Cited consistently, position 1-2, accurate, positive.** You own this query. Protect it: keep the cited content fresh (AI-cited URLs are ~26% fresher on average than top organic results — Ahrefs, ~17M citations, Dec 2025), and reuse the format that won on adjacent queries.
- **Cited sometimes / stuck in "list" position.** You're in the door but not the default. Strengthen the exact content being retrieved: add a credible quotation, a concrete statistic, an authoritative citation (the three tactics the Princeton paper proved move visibility most). Answer the question more directly, in a self-contained 50-150 word section.

- **Never cited; a competitor always is.** Study what *they* publish for that query and out-answer it — more specific, more current, better-sourced. Note where they're cited from (often Reddit, LinkedIn, a review profile, or an FAQ page) and show up credibly in those same places.
- **Cited but inaccurate.** Treat as urgent. Publish accurate, specific, dated content that states the correct facts plainly, and make sure your Google Business Profile (which feeds Gemini and AI Overviews) is correct. The engine is repeating what it found; give it something better to find.
- **Cited but negative.** Read the surrounding context, address the underlying issue, and build a counterweight of accurate, positive, well-sourced content. Don't argue with the model; change its inputs.

The throughline of all five: **AI cites the clearest, most credible, most current answer it can find.** Your job is to be that answer for the queries that turn into customers.

When to graduate to automated monitoring

This spreadsheet is genuinely useful, and for a single-location business testing a handful of queries it may be all you need for a while. But it has a ceiling, and you'll hit it:

- **Coverage.** Ten minutes covers ~10 queries. Real categories have hundreds of phrasings, and you can't manually test enough of them to be confident.
- **Memory and consistency.** One person, once a week, eyeballing answers, is subject to mood, fatigue, and "did I run that incognito?" Answers vary run-to-run, so a single weekly sample is statistically thin.
- **Competitors and benchmarking.** Tracking share-of-voice by hand across multiple rivals and engines gets unmanageable fast.
- **Scoring and content.** The spreadsheet tells you *that* you're invisible. It doesn't tell you your overall standing or exactly what to publish next.

That's the line where you graduate to **Ozvor** — the AI Search Trust Intelligence platform this tracker comes from. Ozvor audits how your brand appears across **ChatGPT, Claude, Perplexity, Gemini and Google AI Overview**, computes a 3-vector **TrustIndex Score** (Brand 30% / Performance 35% / AI 35%), benchmarks you against your competitors, and builds a **GEO content plan** so you're not guessing what to publish.

- **Free** — 1 brand, 3 competitors, 50 prompts, a monthly audit and your TrustIndex Score. No credit card. (Start here — it's the automated version of this spreadsheet's first week.)
- **Growth** — **\$99/mo** (or **\$831/yr** with founder pricing) — 1 brand, 10 competitors, 250 prompts, **weekly automated monitoring, citation tracking, and GEO content.** This is the upgrade for everyone who finds this spreadsheet working but wants their 10 minutes back and 25x the coverage.
- **Agency** — **\$149/mo** (or **\$1,251/yr** founder) — up to 25 brands, white-label reports, client workflow.
- **Founder discount:** 30% off, first 100 customers, annual only.

Done-for-you execution — the consistent, specific, well-sourced posts that earn citations — is available through **OrganicPosts by Ozvor**, our consultancy arm.

Why act now: AI engines learn to trust the sources they already cite. Today's mention becomes tomorrow's default, and the winnable "citation seats" per category are scarce — across repeated queries only ~5 brands per category get named consistently while ~72% sit in a long tail the AI rarely mentions (Search Engine Land, Feb 2026). This is the same setup SEO had around 2010: high upside, low competition, compounding trust. Plant your flag while it's cheap.

Start free at ozvor.com. Questions: hello@ozvor.com.

PART B — The Spreadsheet Build Spec

This section is a precise, unambiguous specification for building the companion .xlsx . It names exact sheets, columns (with column letters), data-validation dropdowns, conditional-formatting rules, and formulas referencing other sheets by name. A builder (human or script) should be able to produce the workbook directly from this spec with no further decisions. Default assumptions: 1-based rows, header in row 1, data starting row 2, formula fill-down range rows 2-1000 unless stated. Engine names and dropdown values must be used verbatim so the Dashboard COUNTIFS match.

Workbook overview

Four sheets, in this tab order:

1. **Read Me**
2. **Query Bank**
3. **Weekly Log**
4. **Dashboard**

Global styling: header row bold, white text on brand navy fill (#0B1F3A), frozen at row 1. Body font Arial/Calibri 11. Column widths set per sheet below. Freeze panes as noted.

Sheet 1 — "Read Me"

Single-column layout (merge A:F per content row, or just put text in column A and widen A to ~110). No formulas. Content blocks, in order:

Title block (A1): LLM Citation Tracker — by Ozvor (18pt bold). **A2:** ozvor.com · hello@ozvor.com (10pt grey).

Block — "What this is" (A4 heading, A5 body):

A simple, honest way to track whether ChatGPT, Claude, Perplexity, Gemini and Google AI Overview name your business when customers ask. ~10 minutes a week, no tools required.

Block — "The four engines + AIO" (A7 heading; A8:B12 small table): two columns, Engine | What to know, populated from the engine table in Part A (ChatGPT, Perplexity, Claude, Gemini, Google AIO).

Block — "Your weekly routine" (A14 heading; A15:A20 numbered list):

1. Open a fresh / incognito browser session and log out of the AI accounts (so results aren't personalized to you).
2. Run your queries from the **Query Bank** across the engines that matter to your audience.
3. Add one row per query-per-engine in **Weekly Log**.
4. Set **Cited?** to Y or N, then fill Position, Source URL, Sentiment, Competitor, Notes.
5. Check the **Dashboard** — it updates automatically.
6. Repeat same day each week. Review the trend after 4 weeks.

Block — "How to read it" (A22 heading; A23 body): short restatement of Cited / Position / Sentiment / Source meanings (one line each).

Block — "When to automate" (A28 heading; A29 body):

When 10 queries isn't enough coverage, start free at ozvor.com. Ozvor monitors all five engines automatically, scores your AI visibility, benchmarks competitors, and tells you what to publish next. Growth plan \$99/mo (or \$831/yr founder pricing).

Column A width: 110. Wrap text on. Row heights auto.

Sheet 2 — "Query Bank"

Purpose: ready-made, customer-style queries to adapt. Header row 1, freeze at row 1. Columns:

Col	Header	Type / validation
A	Vertical	free text (pre-filled)
B	Query	free text (pre-filled)
C	Intent	dropdown: discovery, consideration, decision
D	Notes	free text

Column C data validation: List source = `discovery,consideration,decision` . Show dropdown in cell. Apply A2:D200.

Column widths: A 22, B 70, C 16, D 45. Wrap on B and D.

Pre-fill these 32 rows (rows 2-33). Replace bracketed placeholders like `[your city]` when adapting. Format: `Vertical | Query | Intent | Notes` .

- Dentist | Best cosmetic dentist for Invisalign in `[your city]` | decision | Local + service-specific
- Dentist | What should I look for when choosing a family dentist? | consideration | "How to choose" content wins this
- Dentist | Emergency dentist open on weekends near `[neighborhood]` | decision | High urgency, high intent
- Accountant / Bookkeeper | Best bookkeeper for a small e-commerce business | discovery | Niche + category
- Accountant / Bookkeeper | Do I need an accountant or a bookkeeper for my startup? | consideration | Educational, very citable
- Accountant / Bookkeeper | Affordable small-business tax accountant in `[your city]` | decision | Local + price intent
- Law firm | Best small-business attorney for contract review in `[your city]` | decision | Local + specialism
- Law firm | What questions should I ask before hiring a business lawyer? | consideration | FAQ-style, high citation potential
- Law firm | How much does it cost to form an LLC with a lawyer? | discovery | Data-answerable question
- HVAC / Plumber | Best emergency plumber near `[neighborhood]` | decision | Home-services money query
- HVAC / Plumber | How do I choose a reliable HVAC contractor? | consideration | Trust-driven; checklist content wins
- HVAC / Plumber | Why is my AC blowing warm air and who can fix it in `[city]`? | decision | Problem + local
- Restaurant | Best `[cuisine]` restaurant in `[neighborhood]` | decision | Classic local AI query
- Restaurant | Where can I book a private dining room for a business dinner in `[city]`? | decision | Specific use-case
- Restaurant | Good restaurants in `[neighborhood]` for a group with dietary restrictions | consideration | Long-tail, specific
- SaaS consultant | Best HubSpot implementation consultant for SMBs | discovery | Software-specific niche
- SaaS consultant | How long does a typical CRM implementation take for a 20-person company? | consideration | Data-answerable, very citable
- SaaS consultant | What should I look for in a SaaS onboarding consultant? | consideration | "How to choose" content
- Photographer | How to find a commercial photographer for a product launch in `[city]` | discovery | Use-case + local
- Photographer | What should I ask a photographer before hiring them for brand work? | consideration | Strong GEO query — answer it in a post
- Photographer | Average cost of a brand photography session for a small business | discovery | Price/data question
- Designer | How to find a freelance designer for a small-business brand identity | discovery | Niche + use-case
- Designer | What's included in a brand identity package from a freelance designer? | consideration | Publish your package = get cited
- Designer | Freelance logo designer for a startup under `[budget]` in `[city]` | decision | Local + budget intent
- Fitness / PT | Best personal trainer for beginners in `[neighborhood]` | decision | Local + audience-specific
- Fitness / PT | How do I choose a personal trainer for weight loss? | consideration | Goal-driven, citable
- Fitness / PT | Personal trainer that offers online coaching for busy professionals | discovery | Niche differentiator
- Real estate | Best real-estate agent for first-time buyers in `[city]` | decision | Local + audience

9. Real estate | What questions should I ask a realtor before listing my home? | consideration | FAQ content wins this
0. Real estate | Real-estate agent who specializes in [neighborhood] condos | decision | Hyper-local specialism
1. (Cross-vertical) | Alternatives to [the big chain / competitor in my category] | consideration | Reveals the AI's default names
2. (Cross-vertical) | Who are the top-rated [your category] in [your city]? | discovery | The blunt discovery test

Sheet 3 — "Weekly Log"

Purpose: one row per query-per-engine-per-week. This is the data engine for the Dashboard. Header row 1, freeze at row 1. Columns:

Col	Header	Type / validation
A	Week	number (1, 2, 3 ...)
B	Date	date (format yyyy-mm-dd)
C	Query	free text (paste from Query Bank)
D	Engine	dropdown: ChatGPT, Claude, Perplexity, Gemini, Google AI0
E	Cited?	dropdown: Y, N
F	Position	dropdown: 1, 2, 3, list, none
G	Source URL cited	free text (URL)
H	Sentiment	dropdown: positive, neutral, negative, inaccurate
I	Competitor cited instead	free text
J	Notes	free text

Data validation (exact list sources — values must match verbatim for Dashboard formulas):

- D2:D1000 → ChatGPT, Claude, Perplexity, Gemini, Google AI0
- E2:E1000 → Y, N
- F2:F1000 → 1, 2, 3, list, none
- H2:H1000 → positive, neutral, negative, inaccurate

Conditional formatting (apply to A2:J1000):

1. **Cited = Y → green row.** Rule type: formula. Formula: `=E2="Y"` . Format: fill `#D5F5E3` , text `#1E6B3A` . (Applies whole row when the row's Cited? is Y.)
2. **Cited = N → soft red row.** Formula: `=E2="N"` . Format: fill `#FBE3E1` , text `#9B2C2C` . (Lower priority than rule 1.)
3. **Sentiment = inaccurate → bold orange cell** on column H only. Range H2:H1000. Formula: `=H2="inaccurate"` . Format: fill `#FFE8B3` , bold text `#B45309` .
4. **Position = 1 → bold cell** on column F only (highlight the wins). Range F2:F1000. Formula: `=F2="1"` . Format: bold, fill `#D5F5E3` .

Pre-fill rows 2-5 with the four example rows from the Part A trend table (the Invisalign example) so the user sees a worked example, then leave rows 6+ blank for their data. Mark example rows with (example – delete me) appended in Notes.

Column widths: A 7, B 13, C 55, D 14, E 9, F 11, G 40, H 14, I 28, J 50. Wrap C, G, I, J.

Sheet 4 — "Dashboard"

Purpose: formula-driven summary, zero manual entry. All formulas reference 'Weekly Log' . Use whole-column-ish ranges `=E2:E1000` etc. so new rows are captured automatically. Header row 1. Suggested layout by block (cell anchors are explicit so a builder can place them exactly).

Define a helper anywhere unobtrusive (e.g. L1): total logged rows =COUNTA('Weekly Log'!\$C\$2:\$C\$1000) → label L0: "Total queries logged".

Block 1 — Overall citation rate (A1:B6)

- **A1:** OVERALL (section header).
- **A2:** Total queries logged · **B2:** =COUNTA('Weekly Log'!\$C\$2:\$C\$1000)
- **A3:** Times cited (Y) · **B3:** =COUNTIF('Weekly Log'!\$E\$2:\$E\$1000,"Y")
- **A4:** Citation rate · **B4:** =IFERROR(B3/B2,0) — format as **percentage**.
- **A5:** Times not cited (N) · **B5:** =COUNTIF('Weekly Log'!\$E\$2:\$E\$1000,"N")
- **A6:** Position-1 wins · **B6:** =COUNTIFS('Weekly Log'!\$E\$2:\$E\$1000,"Y",'Weekly Log'!\$F\$2:\$F\$1000,"1")

Conditional format B4: ≥0.30 green, 0.10–0.30 amber, <0.10 red (3-color scale or three formula rules).

Block 2 — Citation rate per engine (A8:D14)

Header row A8: Engine | Logged | Cited | Citation rate . Then one row per engine (verbatim names), rows 9-13:

Cell	Formula
A9	ChatGPT
B9	=COUNTIF('Weekly Log'!\$D\$2:\$D\$1000,\$A9)
C9	=COUNTIFS('Weekly Log'!\$D\$2:\$D\$1000,\$A9,'Weekly Log'!\$E\$2:\$E\$1000,"Y")
D9	=IFERROR(C9/B9,0) (percentage)

Rows 10-13 identical with \$A10 = Claude , \$A11 = Perplexity , \$A12 = Gemini , \$A13 = Google AI0 . Fill B/C/D formulas down by changing only the row number. Add a data-bar conditional format on D9:D13 to visualize per-engine rate.

Block 3 — Trend by week (A16:C24)

Header A16: Week | Cited | Citation rate . For weeks 1-8 (rows 17-24), with the week number in column A:

Cell	Formula
A17	1 (then 2,3,... down)
B17	=COUNTIFS('Weekly Log'!\$A\$2:\$A\$1000,\$A17,'Weekly Log'!\$E\$2:\$E\$1000,"Y")
C17	=IFERROR(B17/COUNTIFS('Weekly Log'!\$A\$2:\$A\$1000,\$A17,'Weekly Log'!\$C\$2:\$C\$1000,">"),0) (percentage)

Insert a **line chart** on C17:C24 (citation rate by week) titled "Citation rate trend." This is the chart the user watches climb.

Block 4 — Share of voice vs competitors (A26:B33)

Header A26: Share of Voice .

- **A27:** My citations (Y) · **B27:** =COUNTIF('Weekly Log'!\$E\$2:\$E\$1000,"Y")
- **A28:** Queries where a competitor was cited · **B28:** =COUNTIF('Weekly Log'!\$I\$2:\$I\$1000,">") (non-blank Competitor cell).
- **A29:** My share of voice · **B29:** =IFERROR(B27/(B27+B28),0) (percentage).
- Optional named-competitor leaderboard (A31:B33): user types a competitor name in A31:A33 and B uses =COUNTIF('Weekly Log'!\$I\$2:\$I\$1000,""&\$A31&"") to count mentions (wildcards catch partial/multi-name cells). Sort/format as a small bar.

Insert a **bar/pie** comparing B27 vs B28 titled "Share of voice."

Block 5 — Sentiment mix (A35:C40)

Header A35: Sentiment | Count | % . Rows 36-39:

Cell	Formula
A36	positive · B36 =COUNTIF('Weekly Log'!\$H\$2:\$H\$1000,\$A36) · C36 =IFERROR(B36/SUM(\$B\$36:\$B\$39),0)
A37	neutral · B37 =COUNTIF('Weekly Log'!\$H\$2:\$H\$1000,\$A37) · C37 (same pattern)
A38	negative · B38 =COUNTIF('Weekly Log'!\$H\$2:\$H\$1000,\$A38) · C38
A39	inaccurate · B39 =COUNTIF('Weekly Log'!\$H\$2:\$H\$1000,\$A39) · C39

Conditional format A36:A39 by category color (green/grey/red/orange) matching the Weekly Log sentiment colors. Insert a **pie/doughnut** on B36:B39 titled "Sentiment mix." Add a callout cell A41: =IF(B39>0,"▲ "&B39&" inaccurate citation(s) – fix your content/profile","No inaccuracies flagged").

Footer (A43)

Outgrown the spreadsheet? Ozvor monitors all 5 engines automatically, scores your AI visibility, and benchmarks competitors. Start free at ozvor.com. (italic, brand navy.)

Dashboard formatting: section headers bold on #0B1F3A fill / white text; percentages to 0 decimal places; freeze nothing; column A width 34, B-D 12.

Build checklist (for the .xlsx generator)

- [] 4 sheets in order: Read Me, Query Bank, Weekly Log, Dashboard.
- [] Dropdowns verbatim: Engine (ChatGPT, Claude, Perplexity, Gemini, Google AI0), Cited? (Y, N), Position (1,2,3,list,none), Sentiment (positive,neutral,negative,inaccurate), Intent (discovery,consideration,decision).
- [] Weekly Log conditional formatting rules 1–4 applied with correct priority (green-Y beats red-N).
- [] Query Bank pre-filled with the 32 rows above.
- [] Weekly Log pre-filled with the 4 example rows (marked "example — delete me").
- [] Dashboard formulas reference 'Weekly Log'! ranges \$2:\$1000 ; all rates wrapped in IFERROR(...,0) and formatted as %.
- [] Three charts: weekly trend (line), share of voice (bar/pie), sentiment mix (pie/doughnut).
- [] Header row 1 frozen on Query Bank, Weekly Log; brand navy #0B1F3A header fill throughout.
- [] Footer CTA on Read Me and Dashboard pointing to ozvor.com.

*Ozvor — the AI Search Trust Intelligence platform for SMBs. Audit how you appear across ChatGPT, Claude, Perplexity, Gemini & Google AI Overview; get your TrustIndex Score; benchmark competitors; build a GEO content plan. Start free at **ozvor.com** · **hello@ozvor.com**. Done-for-you execution via OrganicPosts by Ozvor.*

Sources & references

The methodology in this tracker is grounded in named, dated, public research. Key sources below, with full URLs so every figure is verifiable.

Why measurement is needed (the scale of AI search)

- OpenAI / ChatGPT — 900M weekly active users (Feb 2026), via TechCrunch — <https://techcrunch.com/2026/02/27/chatgpt-reaches-900m-weekly-active-users/>
- Google AI Overviews — 2B monthly users, via TechCrunch — <https://techcrunch.com/2025/07/23/googles-ai-overviews-have-2b-monthly-users-ai-mode-100m-in-the-us-and-india/>
- Pew Research Center — AI summaries and link clicks (Jul 2025) — <https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>
- SparkToro / Rand Fishkin — 2026 zero-click study — <https://sparktoro.com/blog/in-2026-less-than-one-third-of-google-searches-still-send-a-click/>

What to measure, and why answers vary run-to-run

- Aggarwal et al., "*GEO: Generative Engine Optimization*", KDD 2024 — arXiv:2311.09735 · <https://arxiv.org/abs/2311.09735>
- Search Engine Land — repeated ChatGPT runs & brand visibility (~5 brands per category) — <https://searchengineland.com/repeated-chatgpt-runs-brand-visibility-468552>
- Cyrus Shepard / Zyppy — AI citation ranking factors — <https://signal.zyppy.com/p/ai-citation-ranking-factors>

Where citations come from (what to watch by source)

- Siana Marketing — where ChatGPT gets its information (2026) — <https://www.sianamarketing.com/resources/where-does-chatgpt-get-its-information-2026-report>
- Peec AI — most-cited domains, via Search Engine Land — <https://searchengineland.com/ai-search-engines-cite-reddit-youtube-and-linkedin-most-study-473138>
- Semrush — most-cited domains, 3-month tracking — <https://www.semrush.com/blog/most-cited-domains-ai/>
- Profound — Reddit & AI search — <https://www.tryprofound.com/blog/the-data-on-reddit-and-ai-search>