SEO.STREAM · OPERATOR SOP #000

The AI Search Baseline

A 30 minute audit that tells you exactly where you, or your client, actually stand inside ChatGPT, Google's AI answers and Perplexity. No tools required. Repeatable every month.

By Charles Floate. Eighteen years of ranking sites, and this is the first exercise I run on anything now, before I touch a single page. You cannot fix what you have not measured, and almost nobody in this industry has measured this yet. That is the whole opportunity.

---

0 The 60 second version

You are going to ask the three big AI engines the questions your customers actually ask before they buy, score how visible you are in each answer from 0 to 3, and turn that into one number out of 100. That number is your **AI Visibility Score**. Then you do it again next month and watch it move.

What you need:

- About 30 minutes and the tracking sheet (linked above, or build your own from the spec in Step 6)
- ChatGPT with search on, Google (US view), and Perplexity logged out
- Optional but powerful: Google Search Console access, and any AI visibility tool you already pay for (Ahrefs Brand Radar and similar) to sanity check what you find by hand

> Why by hand, when tools exist? Because the tools sample. Doing 15 of your real money queries with your own eyes, once, teaches you more about how these engines treat your brand than a month of dashboard averages. After that, let the tools scale it. This is how I still start every audit.

1 Build your 15 queries

This is the step everyone rushes and it is the one that decides whether the whole audit is worth anything. You are not listing keywords you want to rank for. You are listing **the exact questions a buyer types or speaks right before money changes hands.** Get this wrong and you will measure vanity. Get it right and the score becomes a map of where your revenue is leaking.

10 money queries

The commercial ones. Mix these shapes so you are not just testing one type of intent:

- 2 to 3 category head terms: "best [category]", "[category] for [audience]"
- 2 to 3 comparison and decision queries: "[you] vs [competitor]", "best [thing] for [use case]"
- 2 to 3 problem queries your money pages exist to answer
- 1 to 2 local or vertical specific queries, if that is your game

5 brand queries

These test whether the machines trust you specifically: your brand name, "[brand] reviews", "is [brand] legit", "[brand] vs [main competitor]", and one category term ("best [category]") where you only watch for whether **you** get named at all.

Stuck for the money ten? Have an AI draft them, then cut ruthlessly. You want revenue queries, not traffic queries:

```
Act as an SEO strategist. My business: [what you do, who for, where].
My money pages: [list 3 to 5 URLs or short descriptions].
Give me 10 commercial-intent queries my ideal customer asks an AI
assistant in the hours before they buy. Mix category, comparison and
problem queries. Rank them by likely revenue impact, highest first.
```

> **Charles's move:** doing client work? Run this per client and hand them the finished score before you have pitched a single deliverable. A one page "here is exactly where you are invisible in ChatGPT, and here is who is eating your lunch" is the strongest cold audit in the game right now. The Baseline is not just measurement. It is the best sales asset you own.

2 Run the sweep

Now you ask all 15 queries in all 3 engines. That is 45 checks. Budget about 20 minutes and keep it clean, because a dirty run gives you a flattering number that helps nobody.

ENGINE	HOW TO RUN IT CLEAN
ChatGPT	Fresh chat for every query, search turned ON. Ask it plainly, do not lead it toward yourself. Note what it says AND what it links in the citations.
Google AI Overviews / AI Mode	Incognito window. Force your market with a country code (add `&gl=us` to the results URL for the US, or your own). Screenshot anything that surprises you.
Perplexity	Logged out. Watch the sources row as closely as the answer text. Perplexity shows its working, so it is the fastest read on who the engines trust.

For every one of the 45 checks, you are looking for four things and writing down three of them:

1. **Are you named?** Does your brand appear in the answer text at all.
2. **Are you cited?** Is your site actually linked as a source.
3. **Are you recommended?** Are you the pick, the "best", the first name out.
4. **Who owns it?** Log the 2 to 3 brands or domains that ARE named and cited. This column is gold. It is your real competitor set in AI search, and it is very often not who you think it is.

> **The one rule that matters:** score the answer a stranger gets. Do not argue with the model, do not re-roll until it flatters you, do not run it signed into an account that already knows you and your site. Incognito, logged out, first answer. If the answer is unstable and changes when you run it twice, score the worse one. You want the truth, not a trophy.

3 Score every check, 0 to 3

As you sweep, drop a number into each of the 45 cells. The scale is deliberately simple so you are consistent across all 45:

0 · Absent

You are nowhere in the answer.

1 · Named

Mentioned in the text, but no link.

2 · Cited

Your site is linked or cited as a source.

3 · Recommended

You are the pick: first, "best", the clear recommendation.

The tracking sheet does the adding for you. If you are doing it by hand, the maths is in the next step and it is genuinely simple.

4 Turn it into your score

Here is the whole calculation, in plain numbers. There is nothing to be scared of.

How many points are on the table:

Every check can score up to 3. You have 15 queries across 3 engines, so:

`15 queries × 3 engines × 3 points =` **`135 points maximum`**

Those split into two groups you will want to read separately:

- **Money score:** 10 money queries × 3 engines × 3 = `90 points`
- **Brand score:** 5 brand queries × 3 engines × 3 = `45 points`

90 + 45 = 135. That is why the tracker shows three totals: money, brand, and the two added together.

Your AI Visibility Score:

(your total points ÷ 135) × 100

Add up every number you wrote, divide by 135, multiply by 100. That is your headline percentage. Do the same with your money points over 90, and your brand points over 45, to see which side is dragging.

A worked example so it is concrete. Say you scored 27 points across your money queries and 9 across your brand queries:

- Money: `27 ÷ 90 = 0.30 = 30%`
- Brand: `9 ÷ 45 = 0.20 = 20%`
- Overall: `(27 + 9) ÷ 135 = 36 ÷ 135 = 0.267 = 27%`

That is a "Present" score with a weak brand signal, which tells you something specific: the engines see your content but do not trust your brand enough to hand you the click.

Different fix from a low money score. This is why you keep the two apart.

5 Read your band, then the scissors

Drop your overall percentage into one of these four bands. Be honest about which one you are in.

OVERALL %	BAND	WHAT IT ACTUALLY MEANS
0 to 14%	Invisible	AI search does not know you exist. You are losing by default. The good news: this is the fastest band to climb out of, because there is nowhere to go but up.
15 to 34%	Present	You get named, rarely cited. The machines see you but do not trust you enough to send anyone your way yet.
35 to 54%	Contender	Real citations on real money queries. Now it is a share of voice fight with the 2 to 3 names you logged in the "who owns it" column.
55%+	Dominant	You are the answer. Your job flips to defending it and expanding it, while the models keep changing under you every month.

The scissors check (3 minutes, Search Console)

One more number, and it is the one that makes decision makers move. Open Search Console, go to Performance, set the range to the last 16 months, and look at impressions against clicks. If impressions are climbing while clicks flatten or fall, you are watching the scissors open: **Google is increasingly answering with your content without sending you the visit.** Screenshot it and date it. For client work, that single chart is the most persuasive slide you will ever put in a pitch.

6 What moves the number

Your score is not luck. It is set by three levers, and each one is a full playbook. This is the map, not the territory. The territory is what we run inside SEO.Stream every week.

1. Consensus alignment. AI answers synthesise what the corpus already agrees on. If your pages fight the consensus grammar of your topic, you do not get pulled into the answer. Align to the consensus first, then differentiate. Most sites lose here without knowing it.

2. Entity depth. The models cite brands they can resolve as real: consistent naming, a genuine footprint, third party corroboration. A thin entity does not get recommended no matter how good the page is. This is why your brand score can lag your money score.

3. Citable structure. Answers lift chunks, not whole pages: tight definitional blocks, clean claims, extractable formatting. Most sites do not have a single paragraph an engine can safely quote. Fixing that is often the fastest points on the board.

Re-run the whole Baseline every month. Same 15 queries, same clean method. Log it in the tracker's Monthly tab and watch the movement. The score climbing is your proof the work is working, and it is the receipt you show a client or a boss.

7 Your tracking sheet

The download at the top gives you the whole thing built and self calculating: paste your queries, drop in the 0 to 3 scores, and it totals your money, brand and overall percentages and names your band automatically. It has three tabs:

- **Baseline:** the 45 cell grid with an example row, live totals, and your band.
- **Monthly Tracker:** log each re-run and it shows the month on month movement.
- **Scoring Guide:** the 0 to 3 scale, the bands, and the clean run checklist, always on hand.

Prefer to build your own? Here is the column spec. One row per query, 15 rows.

COLUMNS	WHAT GOES IN THEM
A, B	Query · Type (Money or Brand)
C, D, E	ChatGPT: score 0 to 3 · cited URL · who owns it
F, G, H	Google AI: score · cited URL · who owns it
I, J, K	Perplexity: score · cited URL · who owns it
L	Row points: `= C + F + I`
M	Notes: phrasing that changed the answer, screenshots

COLUMNS	WHAT GOES IN THEM
Totals	Money points (sum of the 10 money rows' L) ÷ 90 · Brand (5 rows) ÷ 45 · Overall (all 15) ÷ 135

Ran it? Then you are already ahead of about 95% of the industry, because they are still arguing about whether AI search matters instead of measuring it. Inside SEO.Stream, this Baseline is week one homework: members post their score, and the weekly live tests and the downloadable machines exist to move it. 100 founding seats at $179/mo, locked for life. Doors open August 17 at **seo.stream**.

Charles Floate · SEO.Stream · Found this useful? Send a friend to seo.stream for their own copy, do not forward the file. That is how the room stays worth being in.