

A PRACTICAL GUIDE FROM UPSKAILL

The AI Visibility Check

A fifteen-minute method for finding out whether AI engines mention your brand, and what to do about it if they do not.



15 MINUTES · NO TOOLS REQUIRED

Why this matters now

A growing share of buying research now ends inside an answer rather than on a results page. The buyer asks a question, reads what comes back, and shortlists from the handful of brands named in it. Every other brand in the category is absent at the exact moment consideration happens.

The difficult part is that this can happen while your reporting still looks healthy. Rankings hold. Traffic softens slowly enough to be explained away. Meanwhile the brands being named in those answers are compounding an advantage that no rank tracker will show you.

This guide gives you a baseline. Not a score to celebrate or panic over, but a measurement you can repeat in thirty days to see whether anything has moved. Every credible strategy in this area starts with one, and most brands have never taken it.

It takes about fifteen minutes and works for any brand in any category, whether you are checking your own or a client's.

WHAT YOU WILL NEED

- Fifteen minutes and a spreadsheet, or a printed copy of this guide
- Free accounts on ChatGPT, Perplexity, and Gemini
- A brand to check, which can be your own or a client's

No subscription, no software, and no technical work is required to complete the check.

1 Write your five buying questions

These are the questions a buyer asks immediately before choosing a product like yours. They do not contain your brand name, because nobody types your name at this stage. They are the questions someone asks when they know they have a problem and are deciding who solves it.

EXAMPLES FROM THREE CATEGORIES

- What is the best project management tool for a small agency
- How do I reduce customer onboarding time
- Which CRM works best for field sales teams

Write your own five below. Be specific about the situation your buyer is in, because vague questions produce vague answers and a useless baseline.

- 1 _____
- 2 _____
- 3 _____
- 4 _____
- 5 _____

2 Run them across three engines

Ask all five questions in ChatGPT, Perplexity, and Gemini. Record each engine separately, because they draw on different sources and do not behave the same way. That is exactly why a single-engine check tells you very little.

For each answer, note every brand named and every source cited. Results vary by phrasing, location, and the day you run them, so treat this as a directional baseline rather than a precise measurement. That is a limitation of the discipline, not of the method.

On Google. Gemini and Google's AI Overviews draw on the same underlying signals, so treat them as one column. If you prefer, run your questions in Google and record the AI Overview instead of using Gemini directly.

On other assistants. The same method works for any assistant your buyers use, including Claude and Microsoft Copilot. Add a column for each one you add, and remember that your maximum score rises by fifteen points per engine.

YOUR QUESTION	CHATGPT	PERPLEXITY	GEMINI
Question 1			
Question 2			
Question 3			
Question 4			
Question 5			

In each cell, write the brands that were named and the sources that were cited, and mark whether your own brand appeared.

STEP THREE

3 Score where you stand

Go back through what you recorded and score each question, on each engine, using the three outcomes below. Score each engine separately rather than adding them together, because the whole point is to see where you are strong and where you are absent.

- ☒ **Named · 3 points**
Your brand appears in the answer text
- ☒ **Cited · 1 point**
Listed only as a source or link
- ☐ **Absent · 0 points**
No mention of your brand at all

If your brand is both named and cited in the same answer, score it as named. Write the points in each cell below.

YOUR QUESTION	CHATGPT	PERPLEXITY	GEMINI
Question 1			
Question 2			
Question 3			
Question 4			
Question 5			
Total for this engine	out of 15	out of 15	out of 15

Now read each engine's total against the bands below. A brand can sit in a different band on each engine, and that difference is itself a finding worth acting on.

0–3 Effectively invisible

This engine does not associate your brand with these questions. Start with entity clarity and third-party presence before touching content structure.

4–8 A foothold


You appear, but inconsistently. The priority is content structure and making your description consistent everywhere this engine looks.

9–15 Real presence

You are part of the answer in your category here. The priority shifts to defending that position and watching what competitors do next.

4 Follow the citations

This is the step people skip, and it is the most useful one. The brands that appeared in those answers were pulled from somewhere. Follow each citation back to its source and write down where it came from.



What you end up with is a list of the publications, directories, review platforms, and communities where your category's answers are being decided. That list is your first roadmap, because it shows you where the engines are already looking rather than where you assumed they were.

NOTE THE SOURCES YOU FOUND

- Which publications appeared more than once
- Which listings or directories the engines trusted
- Which communities or forums were quoted
- Which competitor pages were cited directly

Three things worth doing this week

You do not need help or a budget for any of these. They are the highest-value moves available to almost every brand that has just run this check for the first time.

01 Move the answer to the top of one page

Take the single page that answers your highest-intent buying question and rewrite the opening so the answer appears in the first two sentences. Answers buried under several paragraphs of context do not get extracted, however good the rest of the page may be.

02 Make your description consistent everywhere

Check that what your website, your LinkedIn page, and the main third-party listings in your category say about you agrees. Engines cross-reference these sources, and disagreement between them is one of the most common reasons a brand fails to be trusted enough to cite.

03 Repeat this check in thirty days

Run the same five questions across the same three engines and compare each engine's total. A single measurement is a snapshot. Two measurements are the beginning of a monitoring habit, which is what turns this from an exercise into a discipline.

This is where our course begins

The check you have just run is the opening exercise of our AI Search Strategy cohort. The course takes it a great deal further across six live sessions, moving from a baseline to entity authority, content architecture, monitoring, and a prioritized strategy for your own brand.

Delegates work on their own brand throughout, and finish by presenting a strategy that the trainer and their peers pressure-test in the final session.

COURSE	AI Search Strategy (AEO and GEO)
FORMAT	Six live sessions, two hours each, twelve hours in total
DATES	11 to 27 August, Tuesdays and Thursdays, 11 AM ET
ROOM	Capped at fifteen delegates
TRAINER	Irma Karsten, twenty-five years in search
YOU KEEP	Playbook, monitoring tracker, prompt set, recordings, and a certificate anyone can verify



See the full syllabus

upskaii.com/course/ai-search-strategy-answer-generative-engine-optimization-aeogeo

Scan the code or type the address above.