


How to Tell If AI Already Recommends Your Competitors

SEO | 6 October 2026 | 11 min read

Read online: <https://www.accordtechsolutions.com/blog/how-to-tell-if-ai-recommends-your-competitors>



By the Accord Tech Solutions team · Last reviewed October 2026

Which competitors keep appearing when buyers ask AI tools for recommendations, and where is your company missing? This guide helps you answer that specific question with a repeatable monthly check. Write 10 buyer questions, run each one in ChatGPT, Perplexity, Gemini and Google AI Overviews where available, and log the companies named, their order, the reasons given and any sources cited. Compare the same questions over time, then choose three important pages to review first. The first complete run may take longer than an hour, especially if you check the cited sources. Treat each answer as a snapshot, not as proof of how every buyer sees your market. For a broader first look at citations, accuracy and website readiness, use our 30-minute AI visibility audit; this guide focuses on competitor

recommendations over time.

What does it mean when AI recommends your competitors?

When a buyer asks an AI assistant for a shortlist, the answer usually names a few companies, often with a short reason attached to each one. A company that appears early and gets a clear description has a head start in that conversation. A company that is missing may still win work, but the buyer then has to find it through another route, such as a search, a referral from a peer or a directory the buyer already trusts. Your competitors gain from every answer where they are named and you are not.

AI answers are not a ranking in the traditional sense. They are generated text, built from the sources a tool can reach and from patterns it learned during training. The same question can return different companies on different days, and two tools can disagree on the same afternoon. That is why this check looks at patterns across several runs instead of one screenshot. The useful questions are simple: which competitors keep appearing, which sources the tools seem to rely on, and where your company is missing or described in vague terms.

It helps to separate two problems before you start. The first is absence, where a tool does not name your company at all. The second is weak positioning, where the tool names you but describes you in general words, or places you after competitors whose offers look less relevant to the question. Both problems show up in the log you will build in this guide, but the log alone does not establish their cause. For absence, inspect the pages and third-party sources relevant to the question. For weak positioning, compare the descriptions in the answer with what your own pages actually say, then check whether the offer and audience are clear.

How do you write 10 buyer questions to test?

Start with the questions a real buyer would type before they contact a vendor. Good questions name a problem, a category or a use case, and they often include a constraint such as team size, region, budget range or the tools the buyer already uses. Avoid brand-

only questions at first. A question such as "What is one company known for?" tells you what a tool knows about a single firm. A question such as "Which firms help a mid-sized company solve this specific problem?" tells you what a buyer will see while they are still deciding.



Write exactly 10 questions and keep them in a shared sheet. Use a mix of question types so the results are balanced. Some questions should ask for a category of service, some should describe a problem the buyer wants solved, some should ask for a comparison between two approaches, and some should ask for a shortlist for a specific situation. Ask each question in plain language, the way a buyer would speak to a search tool, and avoid industry jargon that only insiders use.

Keep the wording fixed for the whole test and for every month after. If you rewrite a question between runs, you change the test, and the results can no longer be compared. Save the exact text of each question, including punctuation, so that anyone on your team can run the same check later. A small change in wording can change the companies a tool names, so the sheet is the record you can rely on.

How do you run each question across the four tools?

Run every question in ChatGPT, Perplexity, Gemini and Google AI Overviews. Each tool behaves differently. Some show sources as links under the answer, some summarize without

visible citations, and Google AI Overviews appear inside regular search results only for some queries. The guidance on AI features in Google Search explains how these experiences are shown to people and what site owners can and cannot control. Read it before you interpret your results, so you know what an overview is and is not: [Google Search guidance on AI features](#).



Use a clean session for each run. Open a private window or a new chat, and avoid asking the question inside a conversation that already holds details about your company or a competitor. Do not add follow-up questions unless your method calls for them, because a follow-up changes the context. Ask each question once per tool, in the same order every month. Record the date and the time, because tools update their models and sources often, and an answer from one week may not match the answer from the next.

Note the settings you used. If a tool asks for a location, a language or an account, write it in your notes so later runs use the same options. For Google AI Overviews, run the question as a normal search and check whether an overview appears above the results. If none appears, log that result as no overview. A missing overview is still useful data. It tells you that the query did not trigger a summary in that period, and it stops you from assuming the tool showed an answer when it did not.

Imagine this fictional example: a small software firm asks Perplexity which tools suit a seven-person marketing team. The answer names three competitors, gives a short reason for each, and cites two review sites and one industry guide. The firm is not named at all. The team notes the sources, then checks whether its own product pages explain the same use case in plain words. This is only one run and only one snapshot, but it shows the team exactly where to look next.

What should you log from every answer?

The log is the core of the check. Use one row for each question in each tool, so a full monthly run with 10 questions across four tools produces 40 rows, including any no-overview entries. Record the date, the tool, the question, every company named in the order it appears, whether your company is named, and the sources cited. If Google AI Overviews does not appear for a query, record that as no overview rather than treating it as a company omission. In your notes, capture any reason the answer gives for naming a competitor and flag descriptions of your own company that are inaccurate or vague. Keep the notes factual and short.

The table below shows the layout with sample values. Every value in it is a sample for illustration only. None of it is a result from a real company.

Date	Tool	Question	Companies named	Your company named?	Sources cited
2026-10-01 (sample)	ChatGPT (sample)	Which tools help a small team plan content?	Competitor A, Competitor B, Competitor C (sample)	No (sample)	Review site (sample)
2026-10-01 (sample)	Perplexity (sample)	Which tools help a small team plan content?	Competitor B, Your company, Competitor D (sample)	Yes (sample)	Company blog (sample), industry guide (sample)

Date	Tool	Question	Companies named	Your company named?	Sources cited
2026-10-01 (sample)	Gemini (sample)	What should a mid-sized firm consider when choosing a partner?	Competitor A, Competitor D (sample)	No (sample)	Directory listing (sample)
2026-10-01 (sample)	Google AI Overviews (sample)	What should a mid-sized firm consider when choosing a partner?	No overview (sample)	Not applicable (sample)	Not applicable (sample)

Read the table row by row, then build a short summary in a separate tab. For each competitor, count how many eligible rows name it. For your own company, count how many eligible rows name you. Keep no-overview rows visible but separate from those counts. These numbers describe your fixed question sample for that month; they are not a market-share score or a benchmark. Compare them with earlier runs using the same questions, tools and settings.

The sources column can point you toward useful evidence. When the same review site, directory or guide is cited for a question across runs, open it and check what it says about the companies named. Repeated citations are a research clue; they do not prove that one source caused a particular recommendation. If a competitor is present there and you are not, note the gap and investigate whether that source is relevant to your buyer. For background on answer engines and why sources matter, read [what is AEO \(answer engine optimization\)](#).

Keep the notes neutral. Write what the tool said, not what you wish it had said. If a tool gives a wrong fact about your company, record it exactly, mark it as an accuracy issue, and move on to the next row. Accuracy problems and absence problems need different responses, so

separate them in the sheet.

How do you spot a competitor pattern in the log?

Start with one buyer question and compare its rows across the tools. Which companies are named more than once? Are they described for the same use case, and which sources are actually cited? Keep company mentions, cited pages and the tool's stated reasons as separate observations. Open a cited page before deciding whether it supports the answer; do not assume every mention came from a visible citation.

Next month, compare that exact question again under the same recorded settings. Write a short note on what repeated, what changed and whether your company was absent, inaccurately described or clearly positioned. Prioritize gaps tied to important buyer questions. This comparison guides what to investigate; it does not establish a universal AI ranking or prove what caused a tool's answer.

How often should you repeat the check?

Repeat the full check once a month, using the same 10 questions and the same four tools where available. Monthly repetition helps you separate a one-off answer from a pattern worth investigating. Tools change their models, sources and layouts, so one result can mislead you. A competitor who appears once is a weaker signal than one that appears for the same relevant questions across tools or several months. Record when an overview does not appear and keep those rows separate from answers that name no companies.



Keep every earlier log. A short comparison from one month to the next is more useful than a long report on a single run. Add a note when you change your website, publish a major article, earn a new listing or start a new partnership, so you can connect changes to results later. Do not treat a change in AI answers as proof that one action caused it. Several things can shift at once, and the monthly comparison only shows direction over time.

Choose a fixed day each month for the run, and give it a named owner. A check that depends on memory tends to slip, and a gap in the record makes it hard to see trends. Reserve enough time to run the questions, inspect relevant citations and finish the log. The first run may take longer while you build the sheet; later runs should follow the same method and settings.

Which three pages should you fix first?

After the log is complete, look for the high-value questions where competitors appear and your company does not, or where your company is described inaccurately or vaguely. Group those questions by the page a buyer would most likely land on, such as a service page, a comparison page, a guide or a case study page. Then choose the three pages that serve the highest-value questions for your business. Review those three first and note what you change. The next monthly run can show whether the observed answers moved, but it cannot

prove that a page edit alone caused the change. Keeping the first work list short makes it easier to review and document.



For each page, check a short list of points. Does the opening paragraph answer the buyer question directly, within the first few sentences? Does the page name the category, the problem and the audience in plain words? Can every claim on the page be supported by a source a reader can verify? Does the page link to the supporting material that explains the offer in more depth? Does the page avoid repeating generic claims that every competitor also makes? Each yes or no tells you where the page needs work. For a shorter look at what to check on your own site, read the [30-minute AI visibility audit](#).

Write the changes in ordinary language. Clear sentences help human readers, and they also help tools that read and summarize pages. Avoid stuffing pages with keywords, because repeated phrases make a page harder to read and do not explain the offer. Add a clear section that answers the buyer question in a few short paragraphs, and keep the rest of the page focused on the same problem. Pages that already rank in regular search have a useful head start, so check those first.

Common mistakes

Most failed checks fail for the same few reasons. Watch for these:

- Running the test once and treating the result as a final verdict.
- Changing the wording of the questions between months.
- Running questions inside a chat that already holds company details.
- Logging only the companies you like and skipping the ones that surprise you.
- Ignoring the sources column, which shows where the tools look for answers.
- Fixing ten pages at once, so the team cannot tell which change helped.
- Writing new pages that repeat the same generic claims your competitors make.

Frequently asked questions

Can I run this check without a paid tool?

Yes. The method uses the standard access that each tool offers and a spreadsheet for the log. Paid plans can show different answers or features, so write the plan you used in your notes and keep it the same each month.

Why does the same question return different companies on different days?

AI tools generate answers from models and sources that change over time, and many answers vary slightly on each run. Treat a single difference as normal variation. Look for patterns over several months before you change your strategy.

Does a "yes" in the log mean my company is winning?

No. A yes means your company was named in one answer, for one question, in one tool. The wording, the order and the description matter as well. Being named early with an accurate, specific description is more useful than a bare mention at the end of a long list.

What should I do if none of the four tools names my company?

Treat that as a useful finding, not a verdict. Start with the three-page fix list, check whether those pages answer the buyer questions in plain words, and review the sources cited for those questions. Then repeat the check next month to see whether anything changes.

If your first log shows gaps and you would like help reading it, book a 30-minute working session with the Accord team: [book a 30-minute working session](#). Bring your 10 questions and your first monthly log, and use the time to decide which pages to look at first.

Accord Tech Solutions

AI-powered B2B growth agency. We turn SEO, email and lead generation into booked sales calls for US B2B companies.

<https://www.accordtechsolutions.com/contact-us> · contact@accordtechbd.com · +1 307 370 1787