

AI VISIBILITY · FUNDAMENTALS

What is AI Visibility?

AI visibility measures how often and accurately your brand appears in AI-generated answers from platforms like ChatGPT, Perplexity, and Gemini. It differs from SEO because AI engines synthesise recommendations rather than list links, making third-party brand consensus and structured content the primary citation signals.

AIVIARA RESEARCH · MAY 2026

For most of the past decade, "can buyers find us?" had a reasonably clear answer. You checked rankings. You looked at organic traffic. If you ranked on page one for the right terms, you were visible.

That answer is no longer complete. According to G2's April 2026 survey of 1,076 B2B software decision-makers, 51% now begin their research with an AI chatbot more often than with a search engine, up from 29% eleven months prior. One in three of those buyers purchased from a vendor they had no prior awareness of, after an AI system recommended them. The shortlist is increasingly assembled before any vendor website gets visited.

51%

of B2B software decision-makers now begin vendor research with an AI chatbot more often than a search engine — up from 29% eleven months prior.

G2 · APRIL 2026 · N=1,076

What determines whether your brand is in that conversation is AI visibility. It's the measure of how frequently your brand appears in responses generated by AI platforms (ChatGPT, Perplexity, Google AI Overviews, Gemini, Claude) when users query those platforms for information relevant to your category.

The distinction between AI visibility and SEO is structural, not just semantic. Search engine optimisation determines where your page ranks in a list of links. AI visibility determines whether your brand appears in a synthesised answer the user may act on without clicking anything at all. An AI system doesn't hand the user a list and let them choose; it processes sources, synthesises a response, and presents something that already resembles a

recommendation.

One more distinction that most definitions skip: visibility and accuracy are separate problems. A brand can appear frequently in AI responses with wrong pricing, outdated features, or attributes belonging to a competitor. Appearing inaccurately may be worse than not appearing at all. Both require attention; neither automatically fixes the other.

Your Google rank is not a reliable predictor of your AI citation rate

SEO work carries over, partially. AI engines index the web, so a brand that has built strong SEO foundations is better positioned than one starting from nothing. But the correlation between organic rank and AI citation is much weaker than most teams expect.

Chatoptic ran a controlled study in September 2025 examining 15 brands across five categories, running 1,000 queries. Guided by Omer Ben-Porat, a researcher at the Technion, the study found that brands on Google's first page appeared in ChatGPT responses only 62% of the time. The Spearman correlation between Google rank and ChatGPT position was 0.034 with web browsing enabled, and 0.022 with browsing off. Enabling ChatGPT's web browsing feature improved Google-ChatGPT alignment by exactly one percentage point.

Brands on Google's first page appeared in ChatGPT responses only 62% of the time. The Spearman correlation between Google rank and ChatGPT position was 0.034.

CHATOPTIC / TECHNION — 15 BRANDS, 5 CATEGORIES, 1,000 QUERIES

Two systems, two different objectives. Google ranks pages for click-through, asking which page the user is most likely to want to visit. The correlational evidence from multiple studies suggests AI models cite sources with structured, answer-first content more reliably. A page that builds carefully toward its conclusion across twelve paragraphs may rank well in Google and get rarely cited by an AI model. A page that states its answer in the first sentence, with specific verifiable claims, appears more likely to be cited regardless of where it ranks organically.

The gap across platforms is just as pronounced. The Digital Bloom's December 2025 synthesis of 680 million citation records reported that only 11% of sites get cited by both ChatGPT and Perplexity. The two platforms draw from substantially different source pools. A brand tracking only one AI platform is likely to have a systematically distorted picture of its actual presence.

11%

of sites are cited by both ChatGPT and Perplexity. The two platforms draw from substantially different source pools.

THE DIGITAL BLOOM · DECEMBER 2025 · 680 MILLION CITATIONS

Brand consensus, not backlinks, is the strongest AI citation signal

No AI platform has published a specification of its citation selection algorithm. What's available is correlational research from multiple independent studies, and the directional findings are consistent enough to be actionable.

The strongest signal the available research surfaces is brand consensus across independent sources. Entities.org, synthesising an Ahrefs analysis of 75,000 brands and a Yext analysis of 17.2 million citations, found that brand mentions correlate with AI visibility at $r=0.66$. Backlinks correlate at only $r=0.10$.

A brand cited consistently across review sites, comparison platforms, and independent publications is significantly more likely to appear in AI answers than one with strong backlinks but thin third-party presence. The research describes a corroboration threshold: AI systems appear to require two or three independent authoritative sources citing a brand before recommending it with any consistency. Below that threshold, responses get hedged or the brand gets omitted.

Content structure is the second major lever. Per The Digital Bloom's synthesis of the Princeton GEO research framework, adding statistics to content is associated with a 22% improvement in AI visibility, and adding direct quotations with a 37% improvement.

Adding statistics to content is associated with a 22% improvement in AI visibility. Adding direct quotations from named sources: 37%.

THE DIGITAL BLOOM SYNTHESIS OF PRINCETON GEO RESEARCH FRAMEWORK

To be precise: these specific percentages come from The Digital Bloom's analysis and attribution rather than being directly confirmed figures from the Princeton paper itself. The broader finding from the GEO research is that citation-optimised content structure can improve AI visibility substantially overall. The directional signal is clear even where the per-strategy numbers should be treated as approximate.

Cross-platform presence compounds. The Digital Bloom found that sites appearing on four or more platforms are 2.8 times more likely to appear in ChatGPT responses. Once a brand is consistently referenced across authoritative platforms, AI systems encounter it more frequently in both training and retrieval contexts, which reinforces further citation. This is a self-reinforcing dynamic that benefits brands already in the conversation and disadvantages those trying to enter it.

Take two B2B project management platforms of comparable size. One has a well-optimised blog and 80 organic backlinks. The other has been listed in G2's top 10, quoted in three TechCrunch articles, and mentioned regularly in Reddit discussions about project management tools. The second will likely appear in AI responses for project management queries even if its organic search footprint is smaller. The mechanism isn't that AI systems ignore SEO signals. It's that brand consensus across independent sources is weighted more heavily than backlink density when an AI model is deciding which sources to draw from.

The scale of the problem most B2B teams underestimate

ChatGPT reached 900 million weekly active users as of February 2026, per OpenAI's announcement. Gemini reached 750 million monthly active users by Q4 2025, according to Sundar Pichai's prepared remarks at Alphabet's Q4 2025 earnings call. Perplexity's CEO Aravind Srinivas reported 780 million queries processed in May 2025, stated publicly at the Bloomberg Tech Summit.

Those numbers are measuring different things. ChatGPT's 900 million is weekly active users. Gemini's 750 million is monthly active users. Perplexity's 780 million is monthly queries, not monthly users. That figure says something about query volume but cannot be compared directly to user counts. All three platforms are large enough to matter for B2B purchasing decisions. They cannot be stacked as equivalent metrics.

G2's survey found that 69% of B2B software buyers selected a different vendor than originally planned based on AI chatbot guidance. One in three bought from a vendor they had no prior awareness of before the AI introduced them. The AI isn't just answering questions buyers already had. It's shaping the consideration set before buyers form preferences.

69%

of B2B software buyers selected a different vendor than originally planned based on AI chatbot guidance.

G2 · APRIL 2026 · N=1,076

DerivateX ran 1,400 buyer-intent prompts across 50 B2B SaaS companies on ChatGPT, Perplexity, Claude, and Gemini. 44% of those companies scored below 50 out of 100 on composite AI presence. That figure comes from a study with a specific sample and methodology; it is not a universal distribution across all B2B categories. But Loganix's April 2026 synthesis of six independent studies, which collectively covered hundreds of millions of data points, reached a consistent directional conclusion: 73% of B2B buyers now incorporate AI tools into their research process. The underlying study methodologies vary, and the 73% is a synthesis headline rather than a single controlled survey result. G2's primary survey of 1,076 buyers is more methodologically precise, and its figures point the same direction.

Brands invisible to AI aren't losing one channel among many. They're absent from the moment in the buying process where the shortlist gets formed.

Measuring AI visibility requires tracking across platforms and query types

The standard approach is prompt-based scanning. You build a library of queries representing how target buyers actually phrase questions (category queries, comparison queries, problem-focused queries, purchase-intent queries) and run those prompts systematically across AI platforms. You measure how often your brand appears, in what role within the response, and with what sentiment. You track changes against a baseline over time.

The core metrics: mention rate (brand mentions per prompts tested), share of voice within your category (your mentions as a share of all tracked brand mentions), citation rate (prompts where your URL is explicitly linked), and position within the response. Being the primary recommendation in a comparative query carries different weight than a brief mention in a long list.

Coverage needs to span platforms. Because only 11% of sites appear in both ChatGPT and Perplexity results, tracking a single platform gives a systematically incomplete picture. Query type also affects citation patterns significantly. The brands that appear in informational queries are not always the same brands that appear in purchase-intent queries

for the same category. Both stages matter, and they often require different content investments.

Aiviara is building infrastructure for monitoring AI brand citations and factual accuracy across LLM platforms. Early access information is available at aiviara.com.