





5 AI engines. Different sources. Same brands.



BrightEdge AI Catalyst analyzed citations and brand mentions across 5 AI search engines. What they cite is wildly different. What they recommend is largely the same.

BRIGHTEDGE




Same Questions. Completely Different Sources.



.2%


Share of Gemini's citations that come from UGC sources. It barely touches forums, video, or social content..



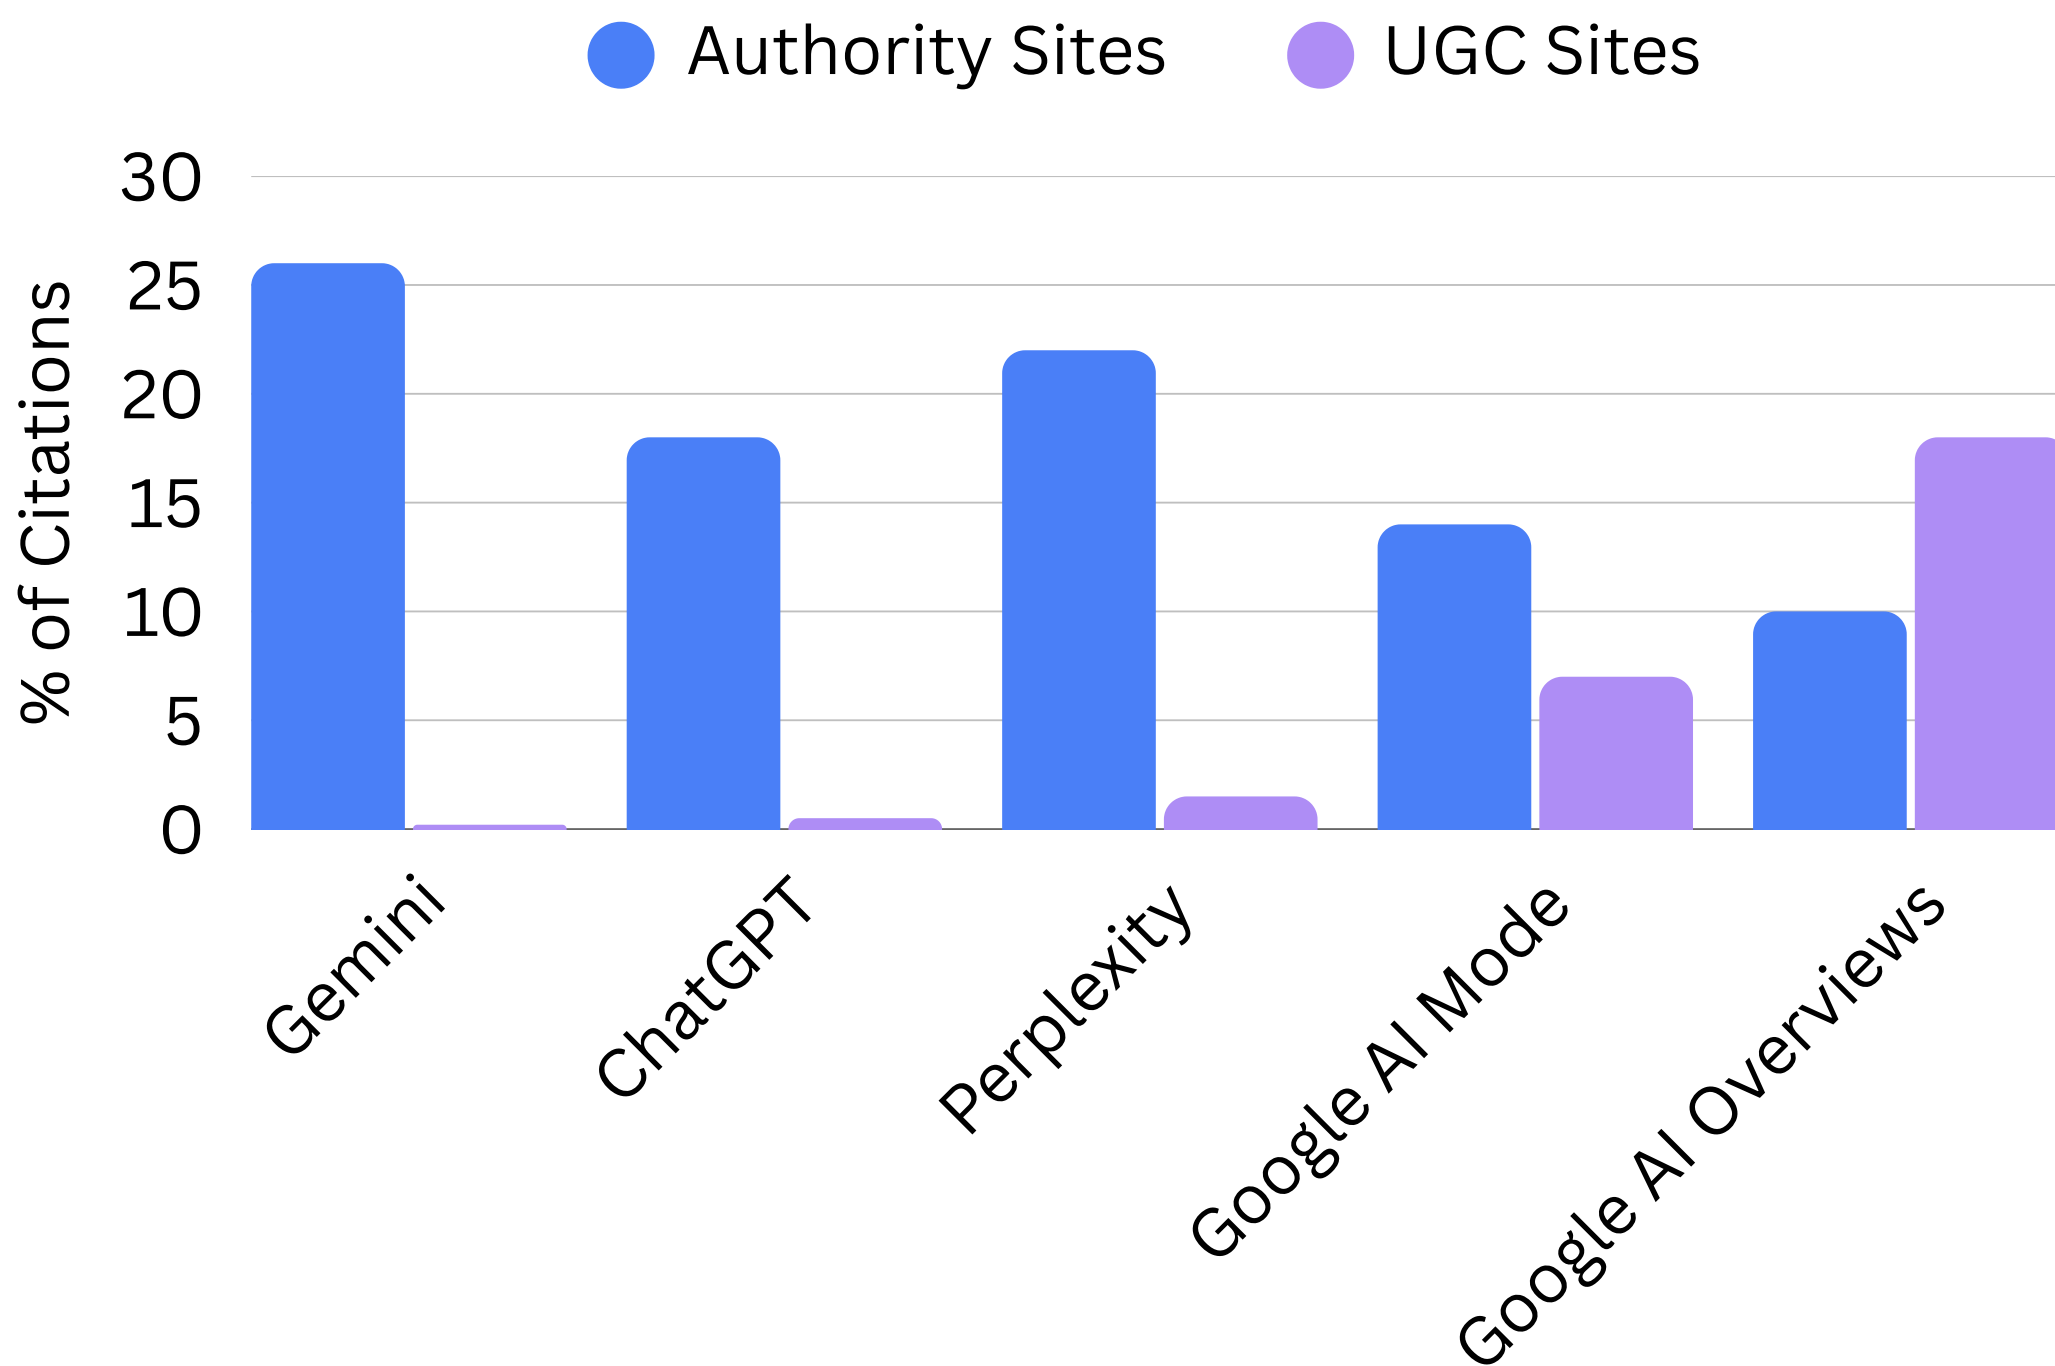
18%

Share of Google AI Overviews' citations that come from UGC. On AIO, UGC actually outweighs authoritative sources.


Both are Google AI products with a significant difference in user-generated content reliance. Each AI engine, even from the same company, utilizes a distinct source mix. Do you know which layer AI accesses in your category?



Engines Sit at a Different Points on the Spectrum



Every engine functions uniquely, yielding different outputs from identical prompts. This prompts the question: why do they consistently recommend the same brands despite their diverse source mixes?



On sources, engines wander. On brands, mentioned, they hold.


16-59%

Source overlap between engines Caption: A 43-point spread. Some engine pairs barely cite the same things. Others overlap heavily. There's no consistent pattern.

35-55%

Brand overlap between engines Caption: A 19-point spread. Across every pair of engines, the brands they name land in a tight, predictable band.

The engines disagree about sources in wildly different ways. But the brands they recommend cluster inside a narrow range, no matter which two engines you compare. That consistency is what makes a single strategy viable across all 5.



What Marketers need to know



The engines disagree on sources but agree on brands.

Citation overlap across engines runs 16% to 59%. Brand mention overlap runs 36% to 55%, consistently higher. The same brand can win across all 5 engines at once, even though each pulls from a completely different slice of the web to get there.



Build for three source layers, not five engine playbooks.

Every engine draws from authoritative sources, commercial and editorial content, and UGC. They just weight the mix differently. Covering all three matters more than building a bespoke strategy for each engine.



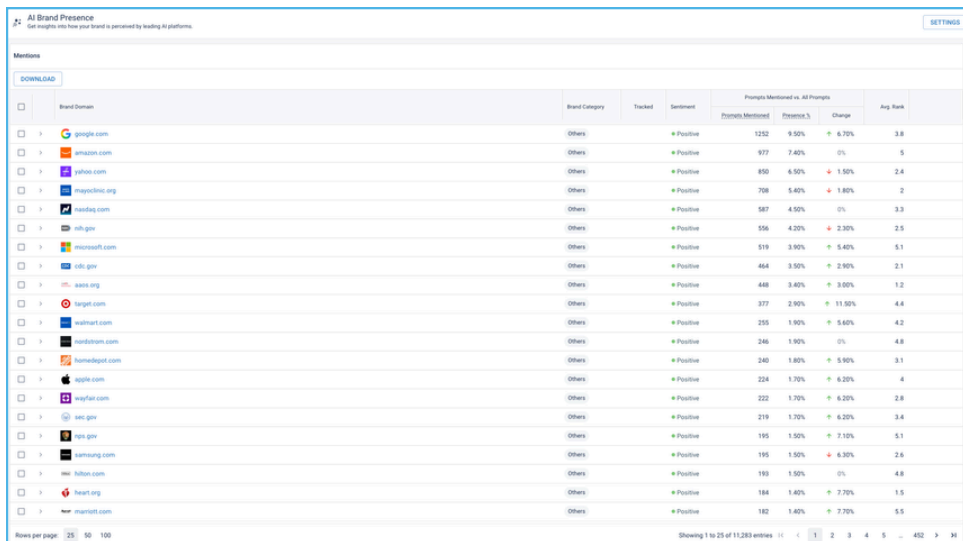
Authority is category-relative, not category-universal. Not every brand belongs on .gov or .edu, and you shouldn't force it. Every category has its own authoritative layer: trade associations, analyst firms, expert publishers, standards bodies. Be covered by the ones that are trusted in your vertical.



Weight investment toward the engines that matter most to your buyers. Perplexity and Gemini lean authority. ChatGPT and AI Mode lean commercial. AI Overviews leans UGC. Every layer is still a factor in every engine, but the weighting tells you where to put your chips first.

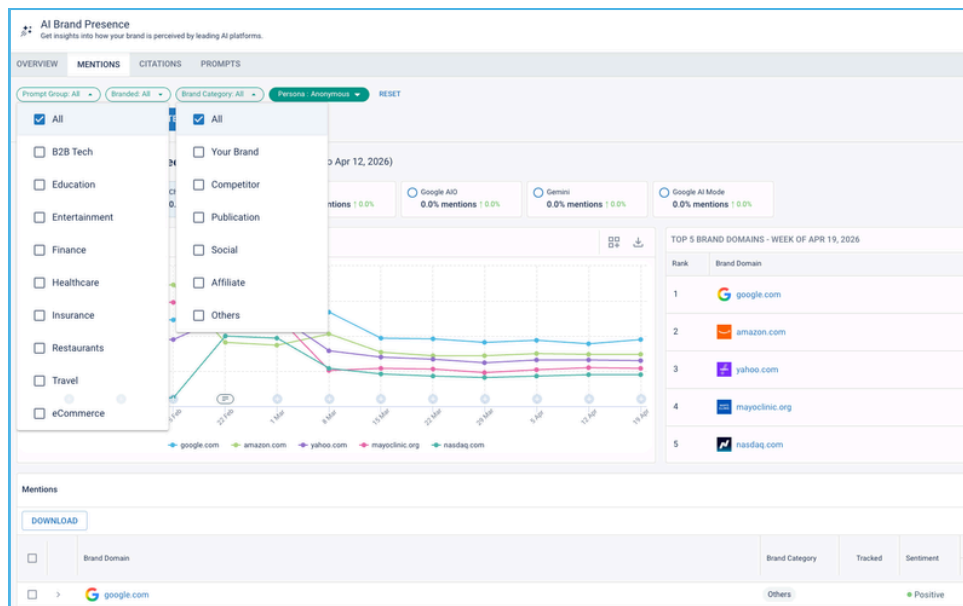
AI Catalyst defines your strategy

See all the top cited sources for the prompts that matter



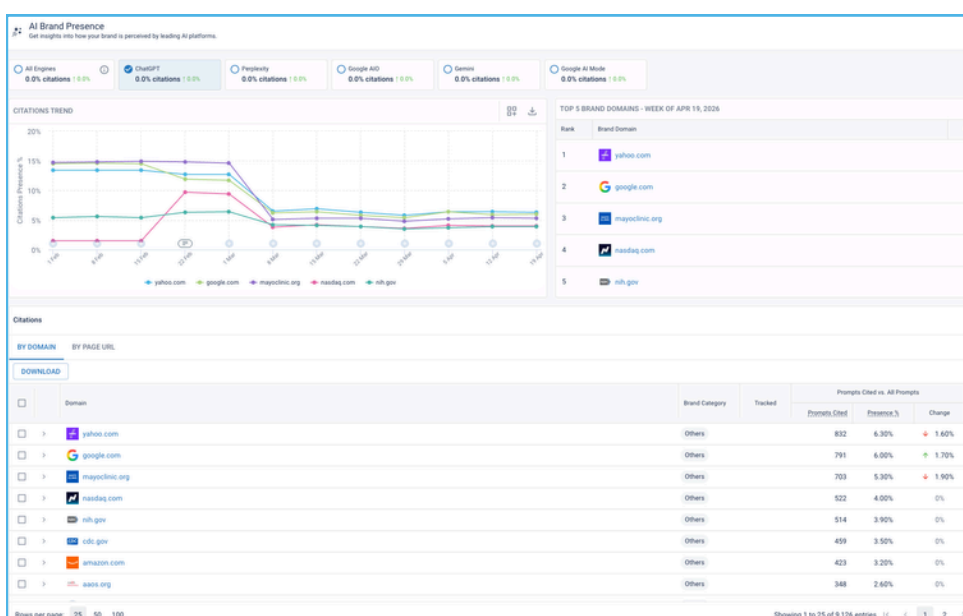
AI Catalyst examines and gathers the most cited sources from the AI engines you monitor, allowing you to quickly assess how they value different sources. This insight enables you to tailor your strategy effectively.

Prompt Groups and Site Filters make granular analysis simple



Organize Prompt Groups by Business Unit using AI Catalyst to automatically categorize site types. This allows for quick analysis of how publications, competitors, and social media influence AI responses to potential customers.

Track market activity across AI search engines



AI Catalyst streamlines the exploration of engines and collective examination of citations and mentions, revealing how nuances vary across different AI engines and emphasizing what is significant to them.