

Turnover.

Rankings rotate, sources disagree, and citations often skip your name.

Here's what the data actually looks like.

52%

of ChatGPT's #1 rankings
rotate on rerun

3.8%

of sources every engine
agrees on

40%

of citations never name the
brand

Every metric in AI search moves. We measured six of them at once.

At Writesonic, we track 2 billion+ AI conversations across 10+ platforms and 50+ markets. That scale let us test a few questions most of the industry has only guessed at. So we set out to answer them ourselves.

The six studies ahead cover what moves, what doesn't, and what to do about it: which rankings hold, which sources engines trust, who gets displaced when a brand breaks in, and why a citation still isn't a mention.

IN THIS REPORT

-
- | | | |
|----|---|-------|
| 01 | Ranking stability. The brand at #1 changes 52% of the time. | p. 03 |
|----|---|-------|
-
- | | | |
|----|--|-------|
| 02 | Brand displacement. 90% of entries happen alongside an exit. | p. 04 |
|----|--|-------|
-
- | | | |
|----|--|-------|
| 03 | Citation overlap. No two engines agree on more than 24%. | p. 05 |
|----|--|-------|
-
- | | | |
|----|---|-------|
| 04 | Source influence. Reddit and YouTube outlift your own site 2 to 3x. | p. 06 |
|----|---|-------|
-
- | | | |
|----|---|-------|
| 05 | Growth drivers. Citation share predicts growth, sentiment does not. | p. 07 |
|----|---|-------|
-
- | | | |
|----|--|-------|
| 06 | Ghost citations. 40% of citations never say your name. | p. 08 |
|----|--|-------|
-

The brand at #1 changes 52% of the time.

52%

of ChatGPT's #1 rankings rotate to a different brand across reruns. Only 8% are locked.

We ran 631,999 prompt-model pairs across seven AI platforms, each prompt repeated at least ten times. The brand holding #1 today is 1.08 to 1.18x more likely to hold it tomorrow than random chance. That is the entire incumbency advantage AI search currently offers.

THE MOST VOLATILE PLATFORM VS. THE STICKIEST, SIDE BY SIDE

 ChatGPT	MOST VOLATILE	 Grok	STICKIEST
Distinct #1 brands / query	9.2	Distinct #1 brands / query	5.6
Mean persistence	0.519	Mean persistence	0.655
Raw stay-rate	45.2%	Raw stay-rate	57%
Incumbency lift	1.18x	Incumbency lift	1.08x



Gemini, Perplexity, Copilot, AI Overviews, and AI Mode all land inside this 1.08–1.18x range. No platform clears meaningful lock-in. Grok's higher raw stay-rate comes from running fewer competing brands per query, not from a stronger moat.

The top answer keeps rotating. Every prompt still produces a #1 answer, it just rotates among 7 to 9 brands. By slot 5, that rotation is up to 4.8x more volatile than at slot 1.

Getting mentioned matters more than getting linked. Across all seven platforms, how often a brand gets mentioned online correlates with AI ranking 3x stronger than how many sites link to it.

TAKEAWAY

Run your key prompts ten times, not once. The distribution of outcomes is your real visibility, a single result tells you almost nothing on its own.


When a new brand enters an AI answer, another one drops out **90%** of the time.

90%

of brand entries into an AI answer happened alongside at least one brand exit.

Across 10.7 million tracked changes between repeated AI answers, we found that AI answers behave like a fixed-size shelf. The answer rarely expands to fit a new name. It rotates instead, at an average of about 3.4 brands entering and 3.4 leaving per change.

AVERAGE BRANDS LEAVING AN ANSWER



Brands are removed **2.6x more often** when a new brand shows up. Entry and exit move together far more than chance would predict.

At least 26% of swaps are internal. One in four add/drop events is a brand trading places with its own alias, product, or parent company, not a competitive loss at all.

Brands swap back and forth. The strongest pairs replace each other repeatedly across reruns. OpenAI and Gemini traded places 4,184 times in one 30-day window.

TAKEAWAY

Build your competitor set from observed swaps, not assumptions. Then check whether your own products or aliases are the ones actually replacing you.

No two AI engines agree on more than 24% of their sources.

3.8%

is the share of sources ChatGPT, Gemini, Perplexity, and AI Overviews all cite in common for the same prompt.

Across 161,286 prompts and four platforms, we measured citation overlap using Jaccard similarity. Of the 70,879 prompts where all four engines returned citations, 72–73% of every cited domain appeared on exactly one engine and nowhere else.

JACCARD OVERLAP BY ENGINE PAIR, SHARE OF CITED SOURCES BOTH AGREE ON



No pair clears 24%. Perplexity, AI Overviews, and Gemini cluster higher since all three are search-grounded. ChatGPT is the outlier, averaging ~13% overlap against the other three, the lowest in the study.

Two citation ecosystems. AI Overviews, Perplexity, and Gemini form one related, search-grounded ecosystem. ChatGPT draws from a structurally different pool of sources.

The consensus layer is thin. Chasing the sources all engines agree on covers just 3.8% of the picture. The remaining 72–73% of citation decisions happen in engine-exclusive territory.


TAKEAWAY

Audit citations per engine, not in aggregate. A strategy built for Perplexity and AI Overviews will not transfer to ChatGPT.

Reddit and YouTube move brands 3x more than your own website.

We measured Influence Lift, how much more likely a brand is to appear in an AI answer when a given source mentions it, across 2,000+ brands and 14 sources. A mention on your own site raises your odds by 1.30x. On Reddit or YouTube, by more than 3x.

INFLUENCE LIFT ON AI ANSWER VISIBILITY, BY SOURCE



Own website (outlined) is the weakest of all 14 sources measured.

WHERE MOST BUDGET GOES VS. WHERE LIFT ACTUALLY IS

Universal lifters	3.18x
B2B review sites	2.40x
Marketplaces	1.46x
Your own website	1.30x

Frequency and influence are different questions.

Wikipedia is cited constantly but lifts visibility only 1.42x. Reddit is cited far less often, yet lifts 3.17x when it does.

Reddit isn't a niche play. It covers 1,984 of the 2,007 brands studied, at a 3.17x lift. That combination of reach and lift is rare.

A big lift means little if barely anyone sees it. A few niche sources post lifts of 4 to 5x but cover only a few dozen brands. YouTube and Reddit lift 3x and cover nearly all 2,000 brands in the study. That combination is what matters.

TAKEAWAY

Rebalance budget away from owned content. A GEO plan that starts and ends with your own site is investing in the 1.3x channel while the 3x channels sit untouched.

Citation share predicts growth. Sentiment doesn't.


We tracked 4,670 brands from March through June 2026 to find which signals moved with AI visibility growth. The brands gaining ground were not described more warmly. They were cited more often.

CORRELATION WITH VISIBILITY GROWTH

Citation share change $r = 0.41$



Sentiment change $r = 0.04$



Citation share's relationship with growth was roughly ten times stronger than sentiment's.

SHARE THAT INCREASED CITATION SHARE

60%



Top
visibility
growers

26%



Bottom
visibility
decliners

Growers were 2.3x more likely to land in the top quartile for visibility growth.

A CLEARER MEASUREMENT HIERARCHY

PRIMARY KPI

Citation share

Track monthly, and against competitors.

OUTCOME METRIC

AI visibility

Shows whether the brand is gaining ground.

DIAGNOSTIC

Sentiment

Watch for reputation risk, not growth.

Sentiment barely separated the groups. Top growers averaged a sentiment score of 70.2; bottom decliners, 70.1. A difference of one-tenth of a point.

The movements matched. Top growers gained 1.55 points of citation share on average; decliners lost 1.66, nearly a mirror image.

TAKEAWAY

Track citation share as a primary monthly KPI. Keep sentiment as a reputation diagnostic, not the metric you optimize for growth.


A citation is not the same as being named.

40%

of citations, across roughly 16 million brand appearances, never named the brand in the answer.

A citation confirms an AI engine used your content. It does not confirm the reader saw your name. Writesonic and Search Engine Land call the gap a "ghost citation": cited as a source, but absent from the answer text.

GHOST CITATION RATE, BY ENGINE



Engines split into namers and citers. Gemini and Copilot name brands more often but link to pages less. Perplexity and Google's engines link often but name brands less.

Proximity may matter. A finding buried three paragraphs from the brand name that produced it is easier for a model to anonymize than one built into the same sentence.

TAKEAWAY

Measure citations and mentions separately. Pair your brand name directly with the claim it produced, on your own pages and in the coverage you can influence.

A quarterly check misses all of this. AI search moves faster than that.

The brands staying ahead measure per engine, not in aggregate, and track citations and name recognition as two separate numbers, not one blended score.

01 Measure every engine separately.

Citation pools overlap by as little as 12% between engine pairs. A blended score hides which platform actually has a problem.

02 Track repeated runs.

Rankings rotate on rerun, and brand displacement is routine. Consistency across repeated runs is the signal that matters, not any one screenshot.

03 Separate citations from mentions.

A link is not the same as recognition. Nearly 4 in 10 citations never say the brand's name. Track the two as separate numbers.

04 Prioritize citation share and third-party mentions.

Citation share correlates with growth 10x stronger than sentiment. Community, video, and review sources outlift owned content 2 to 3x.

This report is Writesonic, applied to itself.

Writesonic is the AI visibility platform behind this research. It tracks brand mentions, citations, and rankings across ChatGPT, Gemini, Perplexity, and ten other AI platforms, then prioritizes the fixes most likely to move the needle. Every study in this report ran on that same infrastructure.

13,000+

marketing teams and agencies

4.8 / 5

average rating on G2 and Capterra


2B+

real AI conversations tracked

BRANDS AND AGENCIES ON WRITESONIC

BRANDS			
			
			
			

AGENCIES			
			
			
			



HOW THE PLATFORM WORKS

Track where a brand shows up across every AI engine, prioritize the gaps by impact in the Action Center, act with AI agents and the Article Writer, then verify and iterate as the answers change again. It's a loop, not a one-time report, because AI search never holds still.

SEO tools ignore AI. AI tools ignore SEO. **Writesonic** does both.

Explore Writesonic

See where you actually stand.

Writesonic tracks AI visibility, citation share, and brand mentions across ChatGPT, Gemini, Perplexity, and seven other platforms. Run distributions, not single snapshots, so you see the pattern before a competitor does.

[Start free](#)[Book a demo](#)

AI Visibility Tracker

Share of voice across 10 platforms

Citation & Mention Split

See ghost citations by engine

Action Center

Prioritized gaps, ranked by impact

Read the full studies.

The Brand at #1 in AI Search Changes 52% of the Time

631,999 prompt-model pairs across 7 platforms

[Read the study →](#)

Who Loses When You Win in AI Search?

10.7 million AI answer changes analyzed

[Read the study →](#)

Do AI Engines Cite the Same Sources?

161,286 prompts across 4 platforms

[Read the study →](#)

Which Sources Influence AI Answers?

2,000+ brand study across 14 sources

[Read the study →](#)

What Separates Brands Gaining AI Visibility From Those Losing It?

4,670-brand study, March–June 2026

[Read the study →](#)

Ghost citations: Why AI search cites your content, not your brand

~16M brand appearances · Search Engine Land

[Read the study →](#)
