


37%5 1%29%99%38%7 6%54 250M
105%4616 35%1 4013 2402 1254
3924 5502 3961 3236 5047 3863
9840 2864 5010 7562 4680 9713
M90% 1000 0026 K24K 14K3 5K7K
21K4 K46K 451K 12M4 3%24%14%
45%5%296 123%05%1 21%2 86%3
87%3 5%42%285 7%38 67%7 33%1
041% 2456 0303 3791 03%3 160%
4371 1%194 4%42 49%3 254%3779
%320 1%24 13%4 6190 5082 8570
5280 7665 0284 5448 4938 5604
1663 7943 9203 9139 4414 5717
9018 3737 8503 3489 5286 7635
72%6 45%6 8%73%21% 23%3%976
8587 5651 6458 26%1 2%31%8%2
%540%510%20%31%18%41%29%
3x30 -40%37%5 1%29%99%38%7
6%54 250M 105%4616 35%1 4013
2402 1254 3924 5502 3961 3236
5047 3863 9840 2864 5010 7562
4680 9713 M90% 1000 0026 K24K
14K3 5K7K 21K4 K46K 451K 12M4
3%24%14%45%5%296 123%05%1
21%2 86%3 87%3 5%42%285 7%38
67%7 33%1 041% 2456 0303 3791
03%3 160% 4371 1%194 4%42 49%3
254%3779%320 1%24 13%4 6190
5082 8570 5280 7665 0



brainlabs

A NEW PATH TO AI DISCOVERY.



37%5 1%29 %99% 38%7 6%54 250M
105% 4616 35%1 4013 2402 1254
3924 5502 3961 3236 5047 3863
9840 2864 5010 7562 4680 9713
M90% 1000 0026 K24K 14K3 5K7K
21K4 K46K 451K 12M4 3%24 %14%
45%5 %296 123% 05%1 21%2 86%3
87%3 5%42 %285 7%38 67%7 33%1
041% 2456 0303 3791 03%3 160%
4371 %194 4%42 49%3 2

SETTING THE STAGE

01	The Current State of AI Visibility	03
02	Measuring AI Visibility	05
03	Understanding the AI Visibility Pyramid: What Actually Creates AI Visibility	10
04	Turning Measurement Into Strategy: Where to Invest and What to Do	12

Insights provided by:



Aziza
Associate Director, SEO



Adam Edwards
Chief Product Officer



Vic Spero
Sr. Manager, SEO



Elliot Glass
Sr. Director, AI Visibility





01

THE CURRENT STATE OF AI VISIBILITY

Search is becoming a much bigger idea than Google. Today, 37% of consumers start their search with an AI tool instead of Google. The shift is even more pronounced in B2B, where 51% of software buyers now turn to an AI chatbot first, up from 29% a year ago.

That matters because an AI answer plays a very different role from a traditional search result. Instead of presenting ten links and asking someone to decide where to go next, an AI tool can compare the options, summarize what other people think and recommend what it believes best fits the question.



In other words, discovery is increasingly happening inside the answer.

Because an AI answer reads like a personal recommendation, buyers treat it like one. That creates a crucial distinction between being mentioned and being recommended. Mentions show that an AI system knows a brand exists. Recommendations tell us whether that brand is actually being put forward for consideration.

At the same time, the relationship between traditional rankings and AI visibility is weakening. In 2024, there was a 99% overlap between AI Overview citations and the top 10 Google results. By 2026, that overlap had fallen to 38%, meaning traditional rankings alone can no longer tell us whether a brand will surface in AI.

76% → 38%

AI Overview overlap with Google's top 10 results in a single year. The assumption that ranking well on Google translates to AI visibility is eroding fast.



Organic discovery is splitting

We can see the consequences in traffic. Across 54 client accounts representing over 250 million sessions, traditional organic traffic fell 10.5% in the eight months after AI Overviews expanded across US desktop searches. That decline affected 46 of the 54 clients analyzed.

Over the same period, AI referral sessions grew 163.5%.

Those two trends belong together. People still search, research and evaluate products. Increasingly, however, some of that activity is happening before a traditional organic click ever takes place.

The result is a discovery journey that is harder to understand through website traffic alone. Someone may encounter your brand in an AI Overview, ask ChatGPT to compare you with a competitor, check Reddit for validation and only then arrive on your website. By the time that visit happens, much of the work of shaping their decision has already taken place elsewhere.



ORGANIC SEARCH				AI REFERRAL TRAFFIC		
METRIC	PRE-AIO	POST-AIO	Δ	PRE-AIO	POST-AIO	Δ
Sessions	140,132,402	125,439,245	▼ 10.5%	502,396	1,323,650	▲ 163.5%
Key Events	47,386,398	40,286,450	▼ 15%	107,562	468,097	▲ 335%
Key Events Rate	10%	10.4%	▲ 0.3pp	8.2%	13.4%	▲ 5.2pp
New Users	48,274,863	42,205,020	▼ 12.6%	161,553	430,575	▲ 167%
Returning Users	30,792,655	28,988,278	▼ 6%	79,693	222,997	▲ 180%


Research done by Samantha Paez



Different AI surfaces are doing different jobs

Our analysis of more than 1.3 million AI referral sessions makes that clear.

ChatGPT accounted for more than 90% of AI referral sessions in our client sample. Perplexity, Claude and Gemini combined contributed fewer than 100,000. But traffic volume alone misses an important part of the story.



Perplexity traffic barely grew during the period we analyzed, yet its conversion rate nearly doubled. Claude showed almost the opposite pattern: rapid traffic growth, but conversion remained below organic. ChatGPT combined scale with stronger commercial performance.

Traffic volume alone, then, doesn't tell us how valuable an AI surface is. One platform may be introducing people to your brand. Another may be helping them research. Another may be showing up much closer to a decision. Treating all of those interactions as one AI channel hides the behavior marketers actually need to understand.



This fragmentation extends to the information shaping those answers. AI systems can draw on brand websites, publishers, Reddit threads, YouTube videos, review platforms, retail listings and other sources across the open web. The result is a discovery environment in which a brand's visibility is increasingly shaped by an ecosystem of information rather than a single owned property.

02 MEASURING AI VISIBILITY

You can measure AI visibility today, but you have to look at the right signals. The problem is that most brands are not looking at all.

While 43% of marketers name AI optimization a core strategic priority, only 24% track AI visibility generally, and just 14% track citations specifically. Part of that gap comes from the measurement itself. AI visibility can be measured, but it cannot be measured like traditional search.

Traditional SEO gave marketers relatively deterministic signals. A keyword had a position. A page received an impression. A user clicked. Google Search Console recorded the event.

AI systems don't offer the same certainty. There is currently no first-party, query-level dataset showing every prompt people use, every answer they receive and every brand or source included. Instead, AI visibility tools approximate that experience by identifying likely prompts, running them through AI platforms and recording what comes back.


That makes the data **probabilistic rather than deterministic**. Different tools use different prompt sets and methodologies, while the answers themselves can change between runs. Citations can disappear within weeks, major platforms agree on which brand to recommend first only around 45% of the time, and fewer than 5% of query sets show full consensus across major AI engines.

Every measurement platform runs prompts and records mentions. Where they differ is in how they estimate prompt volume, and there are broadly four approaches in market.

METHODOLOGY	HOW IT WORKS	STRENGTH	WEAKNESS
Panel & survey estimation	Derives prompt volume from consumer panels or survey data	Attempts to reflect real human behaviour	Meaningful margin of error, especially in niche or B2B verticals
Clickstream & traffic inference	Uses anonymized browsing behaviour to infer query activity across AI platforms	Useful for platform-level comparisons (e.g. ChatGPT vs Gemini growth)	Less reliable at individual prompt or topic level
Keyword-to-prompt modeling	Maps existing keyword search volumes to estimated AI prompt frequency	Widely available; leverages familiar SEO data	Conversion factor largely assumed; doesn't account for how differently people search on LLMs
Direct API sampling	Runs a fixed prompt set on a scheduled cadence and records outputs	Fully transparent, you know exactly what was asked	Makes no claim about real-world prompt volume



Ask an AI tool the same question 100 times. The chances of any two responses returning the same brand list? Less than 1%.

**2,961**

Prompt runs across ChatGPT, Claude, and Google AI in the SparkToro study

<1 in 100

Chance of seeing the same brand list in any two runs of the same prompt

~1 in 1,000

Chance of seeing the same brand list in the same order across two runs

23%

Citations still active after 14 days, across models

45%

Agreement between platforms on top recommended brand

<5%

Query sets where all platforms cite the same brands

The answer isn't to abandon measurement. It's to understand what the data can reliably tell us, and build a measurement framework around it.



Start by separating AI from organic

The first change is simple: AI referral traffic deserves its own reporting line.

Last year, we believed it could reasonably sit within organic reporting because analytics platforms were already categorizing much of it that way. A year of watching the traffic behave has changed our view.

AI referral traffic often carries stronger intent. Someone searching a broad category term on Google may still be at the beginning of their research. Someone asking an AI assistant to recommend the best platform for a specific use case has already articulated the problem, added criteria and asked for a recommendation.

By the time that person clicks through, part of the consideration journey has already happened.

That helps explain why relatively small amounts of AI traffic can have an outsized commercial impact. Ahrefs, for example, reported that AI search visitors represented just 0.5% of its traffic but drove 12.1% of signups.



Blending those visitors into organic traffic makes both channels harder to understand. At minimum, marketers should be able to see AI referral sessions, engagement and conversion separately, and then break those results down again by platform.

But referral traffic only captures the people who click. Much of AI discovery doesn't. Understanding that part of the journey requires looking inside the answers themselves.



Measure the quality of visibility, not just presence

Mentions and citations are useful starting points, but neither tells us on its own what role a brand is playing in an AI answer.

We think a combination of measures gives marketers a clearer picture.

Share of Mentions tells you how often your brand appears across the prompts that matter to your category. The important distinction is whether those prompts are branded or non-branded. If someone asks specifically about your company, appearing is expected. If they ask for the best solution to a problem and the model introduces your brand without being prompted, that is a much stronger discovery signal.

Within those mentions, **citations** tell you how often the model also points people toward your content. One useful way to track that relationship is **Citation-to-Mention Ratio**: the percentage of brand mentions accompanied by a citation. A brand may be frequently mentioned but rarely cited, suggesting that awareness inside the model isn't translating into an opportunity to drive traffic.

Share of Recommendations raises the bar again. It measures how often the model actively recommends your brand rather than simply including it among several options. If mentions tell us whether a brand is in the conversation, recommendations tell us how often it is being chosen from that conversation.

Finally, **Share of Narrative** looks at what the model actually says. Which use cases does it associate with you? Which strengths or weaknesses appear repeatedly? How are you positioned against competitors? Is the information positive, negative, outdated or simply inaccurate?

That context matters because a strong visibility score can still hide a brand problem. You might be mentioned frequently while AI consistently associates you with the wrong use case, outdated information or a negative perception.

Together, these measures move AI visibility away from a binary question of *Did we show up?* and toward a more useful one: **What role is our brand playing in the answer?**



There is no single source of truth

Referral data can show which AI platforms are sending people to your site and what those visitors do next. Search reporting can capture some AI-driven impressions and clicks. Prompt tracking can show when your brand is mentioned, cited or recommended inside answers that never result in a visit.

Each captures a different part of the journey, so measurement has to combine them. The same principle applies across AI platforms: combining them into one visibility score can hide important differences.



We learned this through our own AI visibility program. Brainlabs increased Share of Voice from 28.6% to 38.7%, a 35% relative lift, alongside a 42% increase in Mention Rate. On the surface, that looks like one clear upward trend.

METRIC	PRIOR 3 MO (DEC 1 – FEB 28)	LAST 3 MO (MAR 1 – MAY 31)	CHANGE
Share of Voice	28.57%	38.67%	+10.10 pts (+35.4%)
Mention Rate	7.33%	10.41%	+3.08 pts (+42.0%)
Answers analysed	24,560	30,337	+23.5%

Underneath it, the platforms told different stories. Google AI Mode gained 12.1 points. ChatGPT gained 5.3. Perplexity fell 7.9.

PROVIDER	PRIOR SOV	RECENT SOV	Δ
Google AI Overview	—	91.03%	New
Google AI Mode	31.60%	43.71%	+12.11 pts
Gemini	19.44%	42.49%	+23.05 pts
ChatGPT	32.54%	37.79%	+5.25 pts
Perplexity	32.01%	24.13%	-7.88 pts

An aggregate score would have hidden both where the growth was coming from and where performance was moving in the opposite direction.

That becomes even more important when platforms have different levels of usage and commercial value. A five-point improvement on a major discovery surface may matter considerably more to the business than the same movement somewhere your audience rarely searches.

So measurement needs to answer two questions at once: **How visible are we overall, and where specifically is that visibility being won or lost?**

The second question is what makes the first actionable.



Measure trends, then connect them to the business

Because AI responses are volatile, individual movements should be treated carefully. A small week-to-week change may reflect model variation rather than a meaningful shift in brand performance.

The trend matters more.

For most teams, tracking the core metrics weekly and reporting them monthly creates enough frequency to see meaningful movement without encouraging people to react to every fluctuation.



WHAT TO TRACK	TRACK	REPORT
AI referral sessions + engagement	Weekly	Monthly
Conversion rate by AI source	Weekly	Monthly
Mentions + citations	Weekly	Monthly
Citation-to-Mention Ratio	Weekly	Monthly
Share of Mentions, branded vs. non-branded	Weekly	Monthly
Share of Recommendations	Weekly	Monthly
Share of Narrative + sentiment	Weekly	Monthly + alert on material changes

Narrative requires a slightly different approach. A fluctuation in Share of Mentions can wait for the trend to become clear. A material change in how AI describes your brand, particularly inaccurate or damaging information, may require action immediately.

From there, the harder job is connecting visibility to something the wider business already values.

Direct attribution is still difficult. AI can influence a decision without producing a referral click, and the eventual conversion may happen days or weeks after the original interaction. Trying to force perfect attribution onto an imperfectly observable journey creates more confidence than the data deserves.

A better starting point is correlation. If Share of Recommendations improves consistently, look at what happens next. Does branded search demand rise two or three weeks later? Do direct visits increase? Do assisted conversions or sales move with it? More importantly, does that relationship persist?



That is where AI visibility begins to graduate from an SEO metric into a business signal.

But measurement creates another question. If one brand performs well in ChatGPT and poorly in AI Mode, or appears frequently but is rarely recommended, **what actually determines those outcomes?**

To answer that, we need to understand what AI visibility is built on.



UNDERSTANDING THE AI VISIBILITY PYRAMID: WHAT ACTUALLY CREATES AI VISIBILITY

AI vs Traditional SEO

Pyramid Structure:

- Top Layer (AI Search):** Grok, Claude, Gemini, Copilot, Perplexity
- Middle Layer (AI SEO):** YouTube, App Store, Reddit, LinkedIn, News Sites, Review Sites, Retail Sites
- Bottom Layer (Traditional SEO):** Google, Bing, Blog, Backlinks, Website, Technical

Comparison Matrix:

	AI SEARCH ERA	TRADITIONAL SEO
IMPROVES	- In Depth Content - EEAT - Complete Product Info - Positive Sentiment	- No Human Involvement - PBN Backlinks - Poor Schema
DETRACTS	- Internal Linking - Backlinks - Optimized Page Titles & Descriptions	- Disorganized Site Map - Spam Content - Missing Anchor Text - Meta Keywords



Layer 1: Build the foundation

At the base of the pyramid are the SEO fundamentals that make a brand accessible and understandable in the first place. **AI visibility builds on this foundation rather than replacing it.**

Google's AI features still draw from the standard web index and existing quality signals such as E-E-A-T, which means the same fundamentals that support traditional search visibility continue to matter. Clear site architecture, internal linking, technically sound pages and useful, authoritative content all increase the likelihood that information can be found and retrieved.

But strong traditional rankings don't guarantee AI visibility. While many AI Overview citations come from pages that already rank well, cited sources can also appear outside the top 50. AI systems have another layer of selection once content has been retrieved, creating opportunities for genuinely useful, authoritative content beyond the highest-ranking pages.

That's why the fundamentals remain the base of the pyramid, rather than the whole strategy. **They make AI visibility possible; the layers above determine how effectively that foundation translates into being cited, represented and ultimately recommended.**

We have seen that relationship play out in our own testing. Working with a leading North American nutritional beverage brand, we tested a relatively simple SEO intervention: updating website metadata.

METRIC	BEFORE (APRIL 2026)	AFTER (JUNE 2026)	% IMPROVEMENT
Clicks	46,190	50,828	+10%
Impressions	5,705,280	7,665,028	+34%
New Users	45,448	49,385	+9%
Sessions	60,416	63,794	+6%
AIO Mentions	3,920	3,913	-.17%
AIO Citations	944	1,457	+54%
LLM Mentions	1,790	1,837	+3%
LLM Citations	378	503	+33%
LLM Traffic	348	952	+174%
LLM New Users	86	763	+787%

The change was followed by improvements across both traditional and AI search metrics. Mention Share increased from 57.2% to 64.5%, while Citation Share rose from 6.8% to 7.3%. Those gains also occurred outside the brand's historically high-traffic periods, making seasonality a less likely explanation for the uplift.

We've seen the foundation influence not only whether a brand appears, but **what AI says about it**. In another test, we identified inaccurate claims appearing in AI responses and corrected the underlying information on the brand's website. Those corrections were picked up quickly, alongside an increase in citation visibility.



Together, the tests show why the foundation matters. The information a brand publishes can influence whether AI systems find it, cite it and represent it accurately. **But brands aren't the only ones contributing information to that picture.**

Layer 2: Build external authority

Above the foundation sits the broader information environment around the brand: publishers, YouTube, Reddit, LinkedIn, review platforms, retail sites and other third-party sources. **AI systems use these surfaces to build context around what a brand says about itself, adding external evidence of its authority, reputation and relevance.**

This makes off-site visibility a much bigger part of the equation. Community discussions, reviews, creator content and editorial coverage can all influence how a brand is understood and represented in AI answers. Strong topical authority, consistent product information and positive sentiment across these surfaces reinforce the foundation below.

Layer 3: Earn AI visibility

At the top are the AI experiences that synthesize those signals into answers: platforms such as Gemini, Claude, Copilot and Perplexity. **This is where the work happening across the layers below becomes visible to the user, shaping which brands are surfaced, how they are described and which are ultimately recommended.**

The important point is that these layers are connected. You cannot optimize the top of the pyramid in isolation.

AI visibility is the outcome of what a brand publishes, how the wider web understands and validates it, and how AI systems ultimately interpret those signals.

If an AI system cannot reliably access or understand your product information, there is less for it to work with. But owned content is only part of the picture. When third-party sources tell a different story, or there is little external evidence about your brand at all, another optimized landing page can only solve so much.

That is why the question "**How do we optimize for AI?**" can be misleading. There isn't one AI optimization lever. There is an information ecosystem that needs to be understood and influenced.

The challenge is deciding which parts of that ecosystem are actually worth influencing for your brand. That's where measurement turns into strategy.



TURNING MEASUREMENT INTO STRATEGY: WHERE TO INVEST AND WHAT TO DO

The temptation with any new channel is to look for a universal playbook. Publish more Reddit content. Invest in digital PR. Add FAQs. Refresh every article. Produce more video. Fix your schema. Some of those tactics may work but none of them tells you whether they are the right place for *your* next dollar or hour.

Our client data shows why generic prescriptions fall apart quickly. In one category, Reddit didn't appear among



the top 100 cited sources. In another, it represented 21% of all citations, three times the next-largest domain. Even for the same brand, its importance changed by AI surface: Reddit represented 23% of ChatGPT citations but only 3% of AI Overview citations.



So there is no universal AI visibility mix. There is your mix.



The Organic Media Mix

Enter the Organic Media Mix: a framework for allocating resources across organic channels based on where AI systems are actually citing your category and where you can realistically influence those citations.

The output is a one-page strategic view of which organic channels you're prioritizing, why, and what you expect that investment to deliver. It makes the trade-offs explicit and gives teams a shared plan to work from.

The principle comes from media planning. A paid media team wouldn't divide budget equally between every available channel simply because those channels exist. It considers the audience, performance, cost and role of each, then makes deliberate trade-offs.

AI visibility needs the same discipline.



Step 1: Pull your citations report

The starting point is your citation data. Look at the prompts that matter to your business and identify which sources AI systems repeatedly use to answer them. Then group those sources into channels: your owned site, publishers, Reddit and forums, review platforms, YouTube, retail listings, Wikipedia and whatever else genuinely appears in your category.

Top 10 Cited Domains

Domains cited most often in answers

Domain	% of Total Citations	Citations
1.  Reddit	2.3%	976
2.  Youtube	2.0%	858
3.  Clutch	1.8%	756
4.  LinkedIn Jobs	1.2%	516
5.  Directive Consulting	1.1%	458

Illustrative example of a citation report



Step 2: Score each channel across four dimensions

Raw citation volume tells you what's happening. It doesn't tell you where to invest. For that, you weight each channel across four dimensions.

1.

Influenceability: How much ability do you realistically have to change what appears there?

2.

Difficulty: What would it cost in time, resources, relationships or development work to make that change?

3.

Commercial proximity: Are these sources influencing broad informational questions or prompts much closer to a purchase decision?

4.

Sentiment: Is the channel helping the brand, hurting it or simply mentioning it?

Citation share alone can send a team in the wrong direction. A source might appear constantly but be extremely difficult to influence. Another might account for fewer citations but sit much closer to high-intent prompts and offer a more realistic opportunity to change what customers see.



The OMM forces that trade-off into the open.



Step 3: Build the mix

With those scores in place, you can build the mix. Here's a hypothetical for a CPG brand. Before running the citations analysis, the team assumed owned content was doing most of the work and had allocated accordingly. The data told a different story

ORGANIC MEDIA MIX Channel Scoring Matrix						
CHANNEL	CITATION SHARE	INFLUENCE	DIFFICULTY	PROXIMITY	SENTIMENT	PRIORITY
Owned Site	26%	●●●●	●○○○	●●○○	Positive	1
YouTube	12%	●●●○	●●○○	●●○○	Positive	2
3P / Editorial	31%	●●○○	●●○○	●●○○	Positive	3
Reddit / Forums	21%	●●○○	●●○○	●●○○	Mixed	4
Review Platforms	8%	●●○○	●○○○	●●●●	Mixed	5
Wikipedia	2%	●○○○	●●●●	●○○○	Neutral	6
●●●● = Very High ●●●○ = High ●●○○ = Medium ●○○○ = Low						
For illustration only						

Anonymized example: Consumer health brand (illustrative)



The finding that shifted the strategy was third-party editorial. It was the largest citation driver, had high commercial proximity and consistently positive sentiment, yet the brand had almost no structured investment behind it. Digital PR had been treated primarily as a brand awareness play. The OMM showed it was also influencing AI visibility, creating a case to reallocate budget.

Prioritizing Prompts and Applying Proven Tactics

Once you know where to invest through your Organic Media Mix, the next step is deciding **which prompts are actually worth influencing**.

Not all prompts carry the same value. And unlike traditional search volume, prompt volume is directional at best, based on assumptions that vary by platform. That makes volume alone a weak way to decide what matters.

Instead, we use a weighted scoring framework that combines **business value, current visibility and how the brand is being represented**. This allows hundreds of tracked prompts to be filtered down to the opportunities that matter most to the business.

SCORE	WHAT IT MEASURES	WHERE IT COMES FROM	SCALE
Value	How much this prompt is worth to the business if you win it	Scored against a fixed rubric	1-5, rescaled to 1-100
Mention %	How often the brand is actually named in the engine's answer	Platform data (ArcAI, Peec.ai , Profound , or equivalent)	0-100
Citation %	How often the brand is cited as a source, named or not	Platform data (ArcAI, Peec.ai , Profound , or equivalent)	0-100
Priority	A ranking that combines Value with the unclaimed part of Mention % and Citation %	Calculated from Value, Mention %, and Citation %	Roughly 1-100
Sentiment	Whether the brand appears well and accurately, not just often	Platform sentiment data where available, otherwise scored while reading the response	Positive / Neutral / Negative or inaccurate

Apply tactics against the gaps

Once you've identified the prompts that matter most, the next question is **what needs to change to improve performance against them**.

Our testing so far shows that relatively targeted changes can influence both how often a brand appears and, just as importantly, what AI systems say when it does.

Fixing technical SEO can improve accessibility

In our testing, we found that hidden JavaScript rendering dependencies were making critical content harder for traditional search engines and AI systems to access. Resolving those dependencies improved content accessibility and LLM performance.



MENTION SHARE		CITATION SHARE	
Pre-Experiment	57.2%	Pre-Experiment	6.8%
Post-Experiment	64.5%	Post-Experiment	7.3%

Correcting misinformation can improve answer accuracy

Visibility is only valuable if the information being surfaced is accurate. In another test, we identified incorrect information appearing in AI responses and updated the underlying website content.

The correction was picked up quickly across multiple AI platforms. While overall mentions decreased during the same period, the remaining mentions were more accurate, reinforcing an important distinction: **the goal isn't simply to maximize mentions. It's to make sure the mentions you earn are useful and accurate.**

MENTION SHARE		CITATION SHARE	
Pre-Experiment	54.0%	Pre-Experiment	2.0%
Post-Experiment	51.0%	Post-Experiment	3.1%

Strengthen the content AI has to work with

Sometimes the opportunity is in the content itself. Our testing and wider research show that how information is structured, supported and kept up to date can influence whether AI systems use it.

The goal isn't to produce more content for the sake of it. It's to strengthen the information available around the priority prompts you've identified.

OPTIMIZATION	OUTCOME	SOURCE
Adding brand USP models to key pages	LLM traffic increased 18%	Omio
Applying direct-answer headings ("How to fix X" vs. "About X")	41% average citation rate for direct-answer headings vs. 29% for vague ones	Brainlabs
Refreshing content every three months	Refreshed content is 3x more likely to be cited by LLMs	AirOps
Adding statistics and quotations to content	One of the features that lifted AI visibility roughly 30-40%	Princeton GEO Study

Together, these tests point to a broader principle: there is no single tactic for improving AI visibility. The right intervention depends on the gap measurement reveals. **Make information accessible when AI can't find it, correct it when AI gets it wrong, and strengthen it when the content isn't doing enough.**