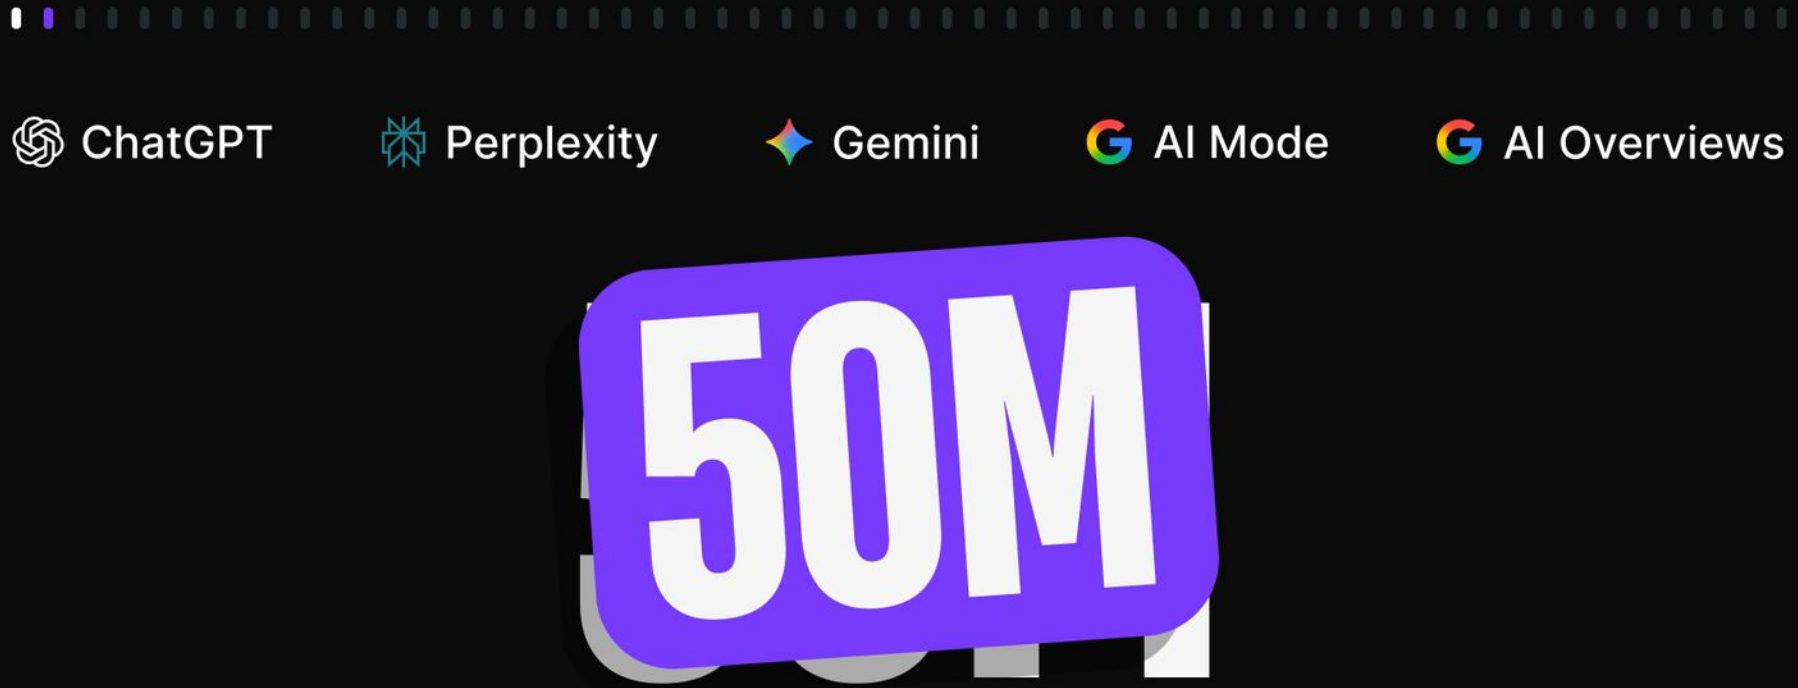
THE STORY SCRAPED 50M LLM ANSWERS

Story by Michał Suski

Surfer


ChatGPT
Perplexity
Gemini
AI Mode
AI Overviews
50M
SCRAPED AI ANSWERS.
5 engines. Every day.

2025

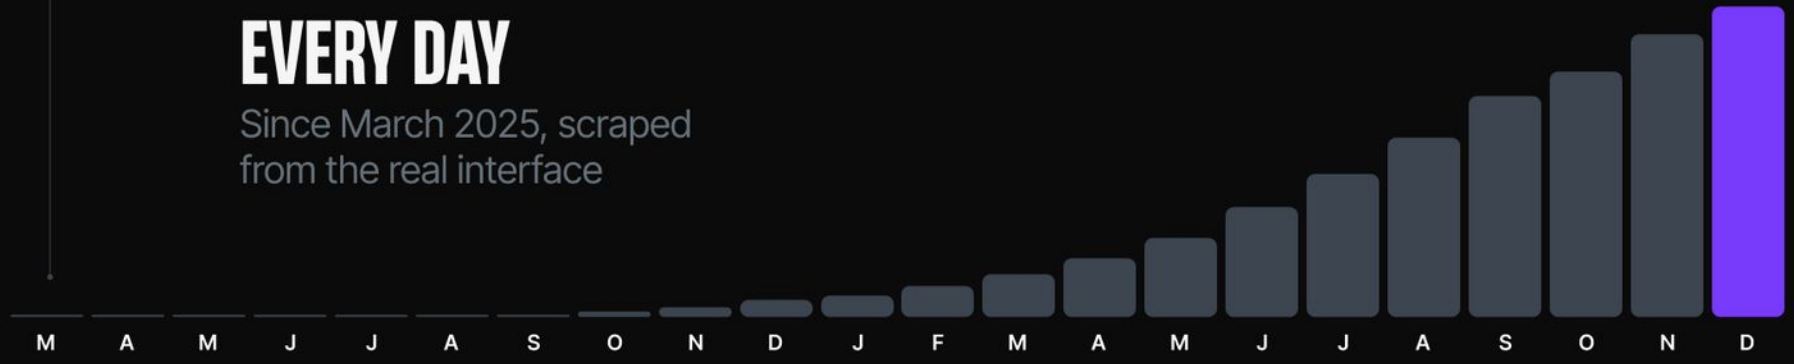
100K PROMPTS A DAY

The questions brands
actually track, across niches

EVERY DAY

Since March 2025, scraped
from the real interface

2026



Cumulative AI Tracker answers by month, Mar '25 - Dec '26.

Data From Surfer AI Tracker

HOW MUCH OF WHAT YOU HEAR ABOUT AI SEO SURVIVES CONTACT WITH DATA?

THREE ACTS

I WHAT MOVES AN AI ANSWER

How to measure without fooling yourself.

II FALSE SIGNPOSTS

The proxies and checklists everyone uses.

III WHAT GETS YOU CITED

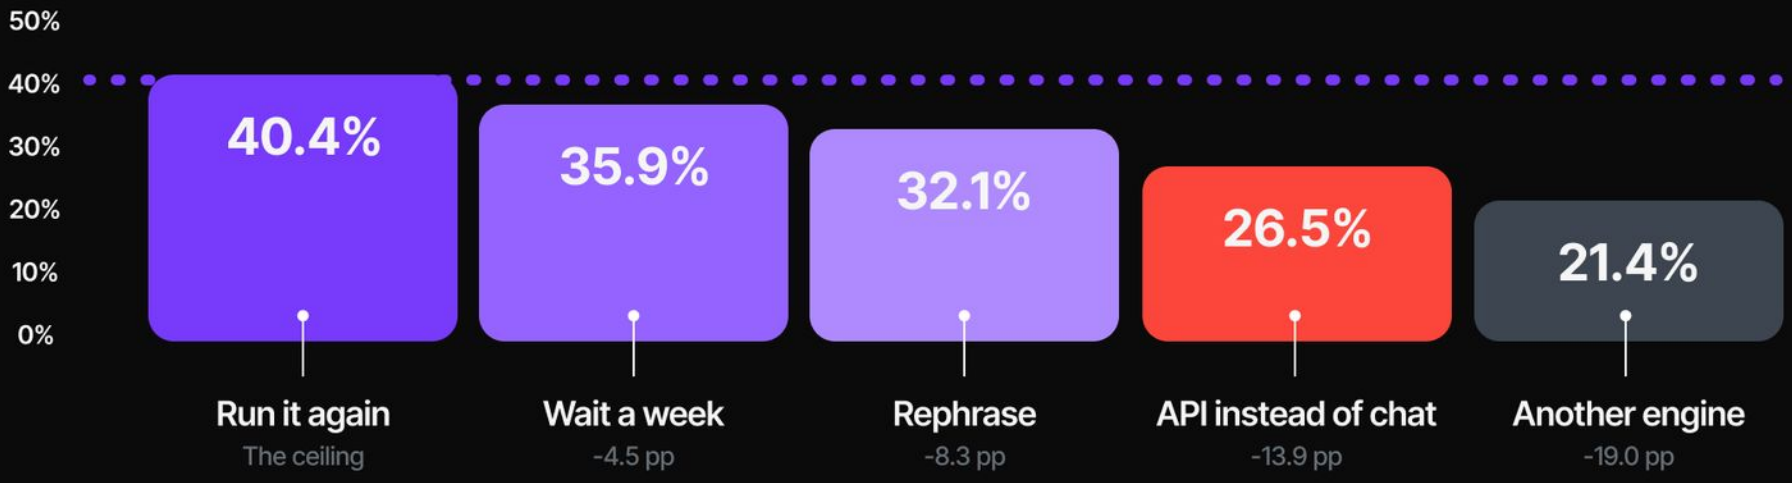
The levers that hold up, incl. the full listicle story.

ACT I

WHAT MOVES AN AI ANSWER?

Engine, channel, calendar, wording or pure chance.

CHANGE ONE THING. WATCH THE ANSWER CHANGE.



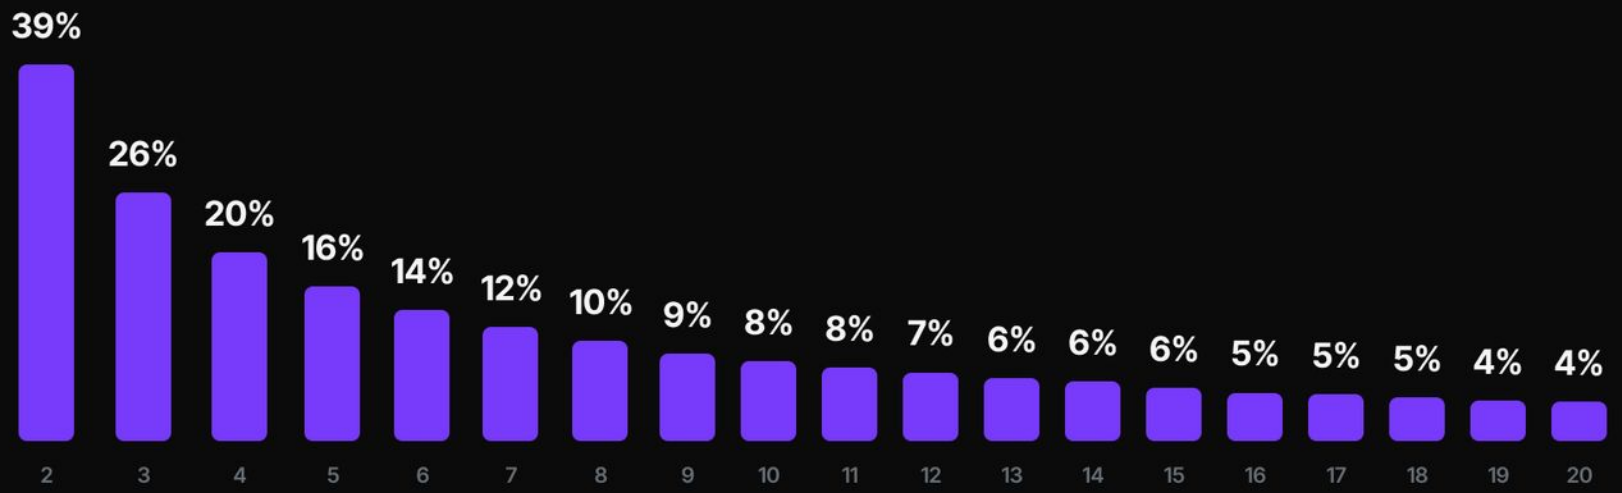
Jaccard of canonical brand sets, mean of engines. Jul-Sep 2026. Rephrase and API from separate studies

Data From Surfer AI Tracker

WHAT IF YOU JUST ASK AGAIN?

NOISE

Find everyday winners.



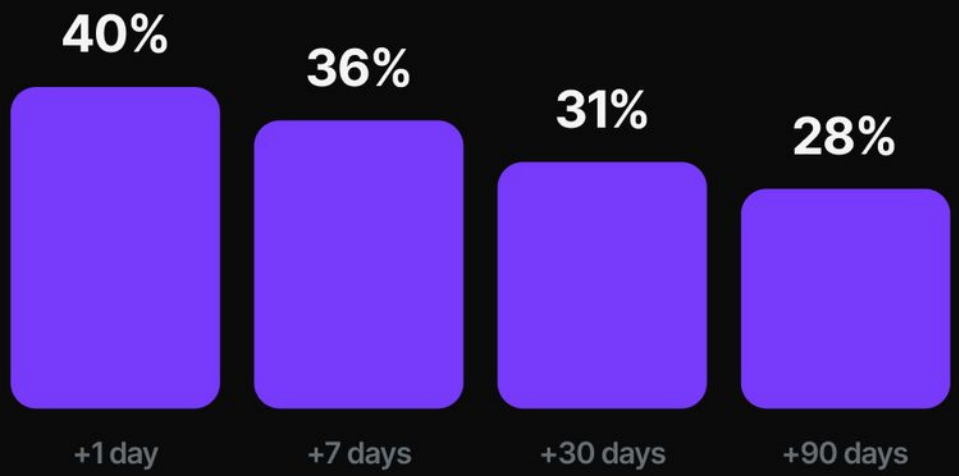
Brands present in all n daily runs, %. Mean of 5 engines. Jul-Sep 2026

Data From Surfer AI Tracker

NOISE

WHAT IF YOU WAIT A WEEK?

Brands stay, sources drift.



Overlap (Jaccard) of answers d days apart. Mean of 5 engines. Jul-Sep 2026

Data From Surfer AI Tracker

REWORD IT ALL YOU WANT

NOISE



The words barely matter. The intent does.

Brand overlap as a share of run-again overlap. Average of ChatGPT, AI Mode, AIO and Perplexity. June 2026

NOISE

REWORD IT ALL YOU WANT

TRACK INTENTS, NOT REWORDINGS.

Related intent
A neighbouring question

Different intent
Another question

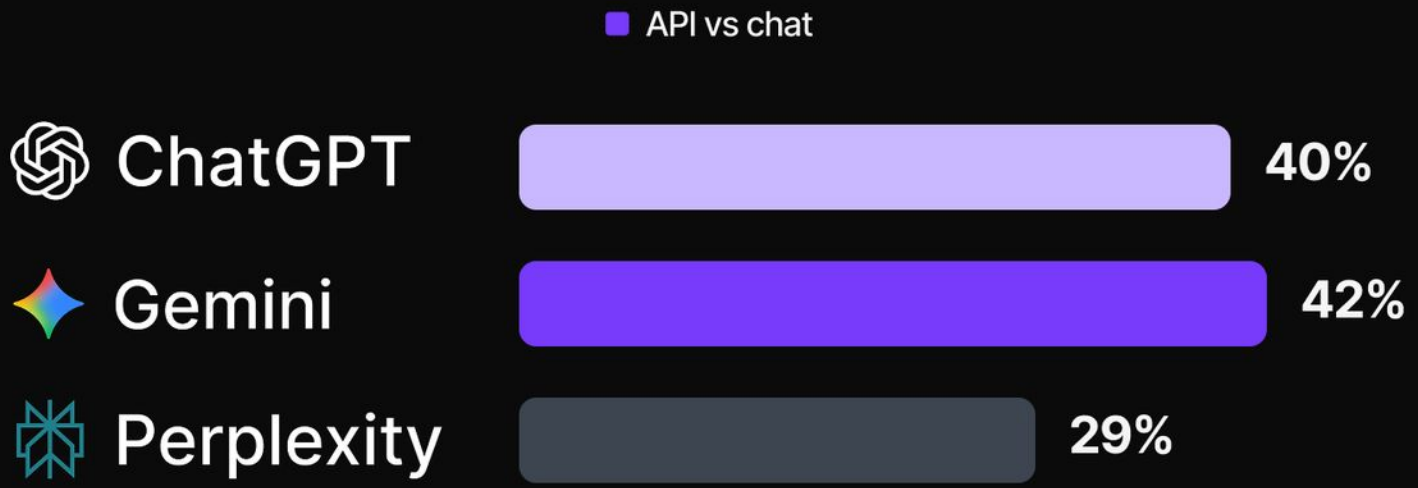
36%

→ The words barely matter. The intent does.

Brand overlap as a share of run-again overlap. Average of ChatGPT, AI Mode, AIO and Perplexity. June 2026

WHAT IF YOU ASK THE API?

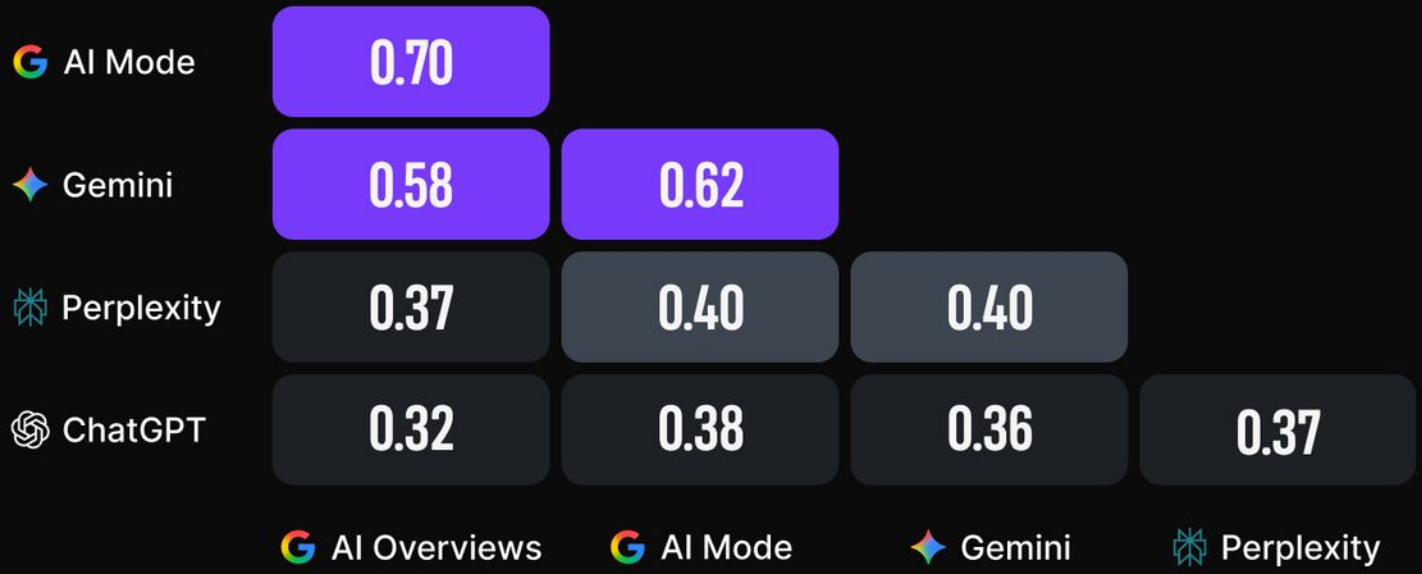
NOISE



Top-3 brands kept, %. API-vs-scraped study + AI Tracker floor

Data From Surfer AI Tracker

WHAT IF YOU ASK NOISE ANOTHER ENGINE?



Spearman, within-report ranks, scale 0-1. Jul-Sep 2026

Data From Surfer AI Tracker

NOISE YOU CAN'T REMOVE.

ERROR YOU CAN STOP ADDING.

HOW TO MEASURE HONESTLY

1. **Many runs**
   A probability, not a position
2. **Every day**
   Drift is real and silent
3. **Intents, not strings**
   Paraphrases collapse into one answer
4. **The real interface**
   The API is a different product
5. **Each engine separately**
   No single AI visibility number

ACT II

FALSE SIGNPOSTS

Measuring is hard, so everyone grabs a proxy.

MYTH

GLOBAL TOP DOMAINS AREN'T YOUR TOP DOMAINS

Only 8.9% of a niche's core sources come from the global top 20.

3,720 reports. Jul-Sep 2026

Data From Surfer AI Tracker

MYTH

GLOBAL TOP DOMAINS AREN'T YOUR TOP DOMAINS

Only 8.9% of a niche's core sources come from the global top 20.

Find your niche's core sources before you chase anyone else's.

3,720 reports. Jul-Sep 2026

Data From Surfer AI Tracker

MYTH

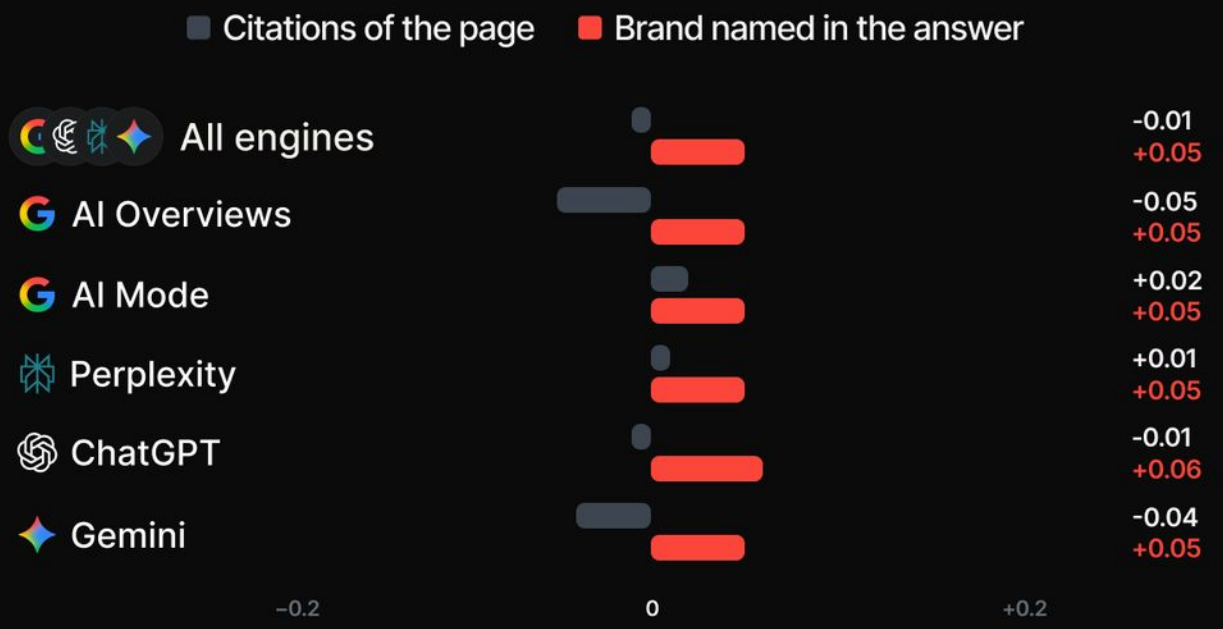
AUTHORITY IS A PROXY, NOT A FACTOR

AUTHORITY

AGE

PAGERANK DOESN'T GET YOU CITED

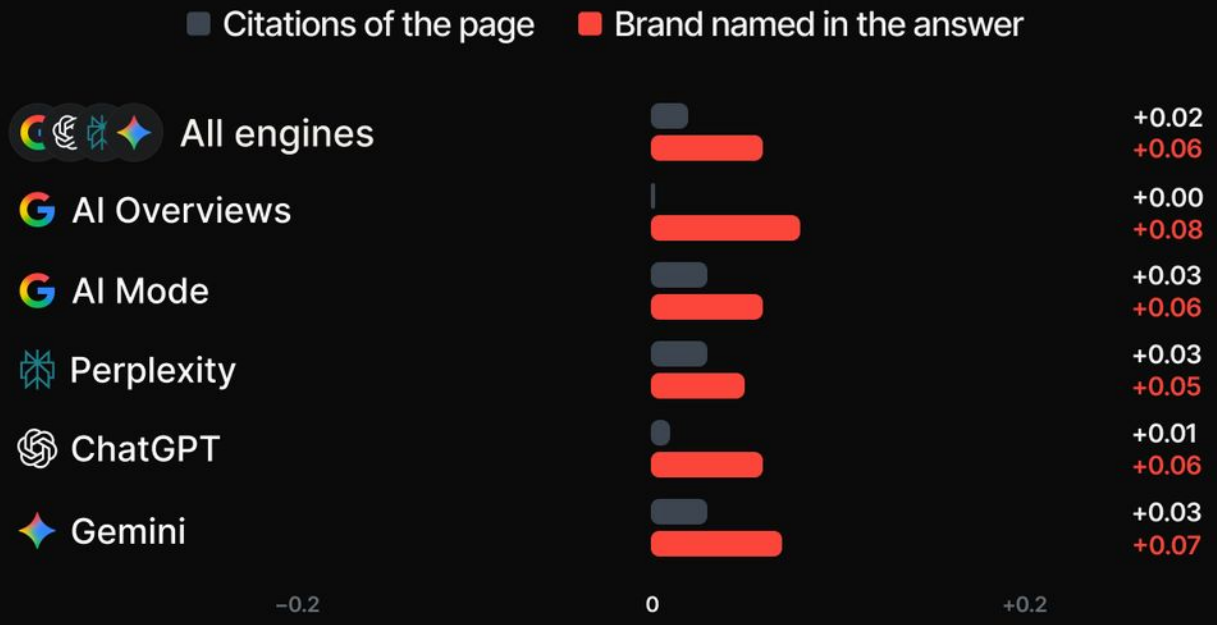
MYTH



Jul-Sep 2026. Scale -0.2 to +0.2

Data From Surfer AI Tracker

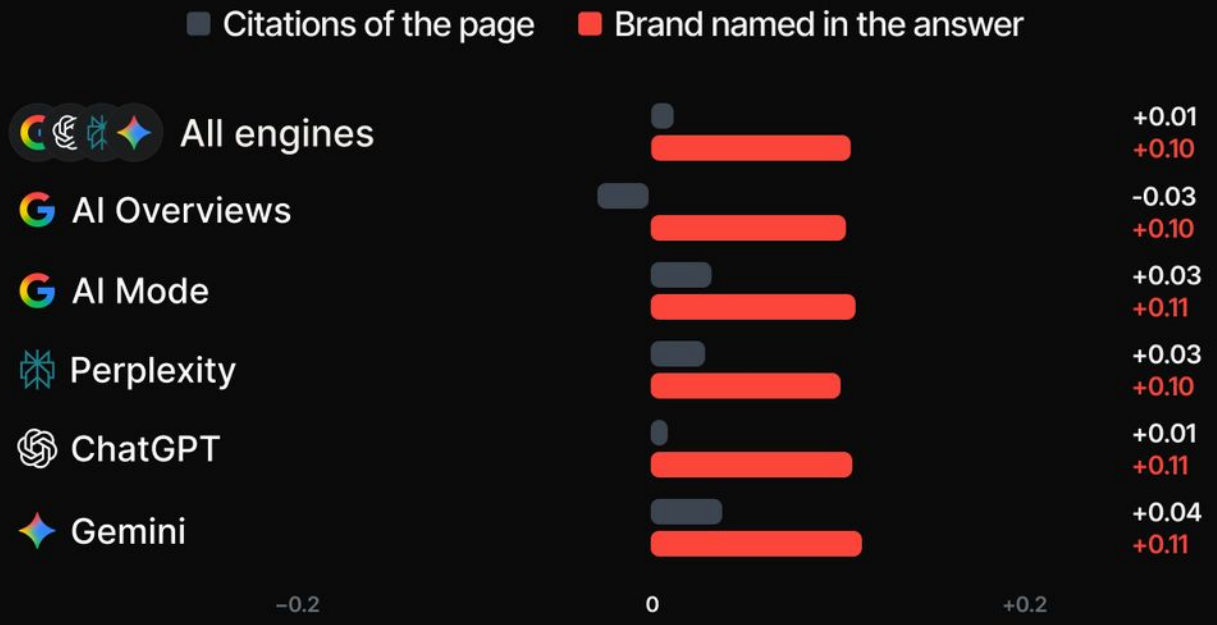
MYTH
CENTRALITY HELPS, BARELY



Jul-Sep 2026. Scale -0.2 to +0.2

Data From Surfer AI Tracker

MYTH
TRAFFIC GETS YOU NAMED



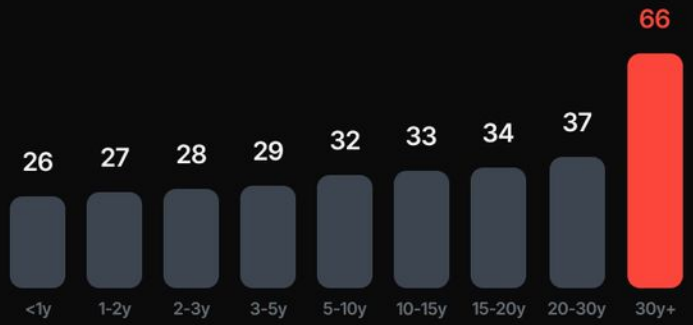
Jul-Sep 2026. Scale -0.2 to +0.2

Data From Surfer AI Tracker

MYTH

DOMAIN AGE

Old sites win on volume, not on merit.





2.5X more citations for 30y+ domains than for domains under 1 year (66 vs 26)

0.06 at most, within one prompt (ρ). Negative on ChatGPT and Perplexity.

Median over 3 months

Data From Surfer AI Tracker

MYTH

THE GEO CHECKLIST IS DECORATION

What checklists say

Add stats, quotes, an FAQ and a TOC - and AI will cite you.

What the data says

≈4–5% each: the bottom 5 of 20 page features.

4 AI models × 1,500 prompts. ~90 runs each. ~500k AI answers

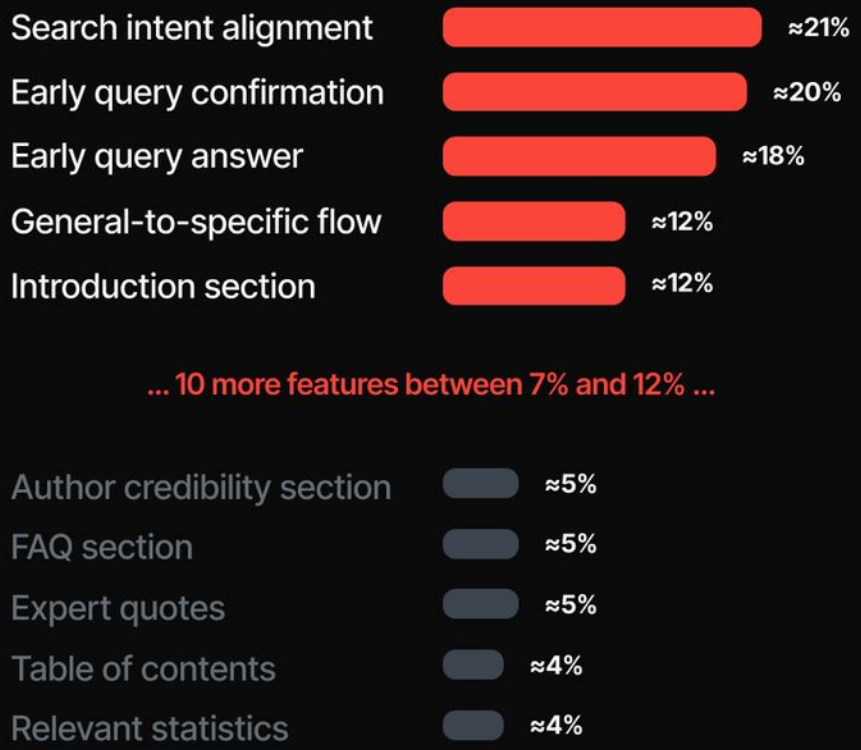
Data From Surfer AI Tracker

MYTH

THE CHECKLIST SITS AT THE BOTTOM OF 20

Writing the answer Checklist items

Fix intent and the first answer before decorating.



How much more often a page with the feature is cited. Scale 0-25%. Values read from chart, approximate

Data From Surfer AI Tracker

MYTH

LINKING TO WIKIPEDIA WON'T SAVE YOU

The Advice

Link out to trusted sites to look credible.

The Data

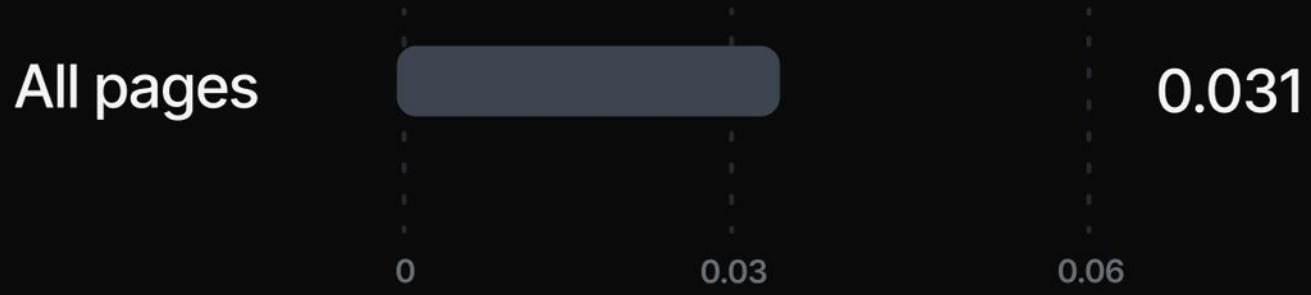
External links in blog posts vs being cited: ≈ 0.007.

4 AI models × 1,500 prompts. ~90 runs each. ~500k AI answers

Data From Surfer AI Tracker

MYTH

LINKING TO WIKIPEDIA DOES NOTHING

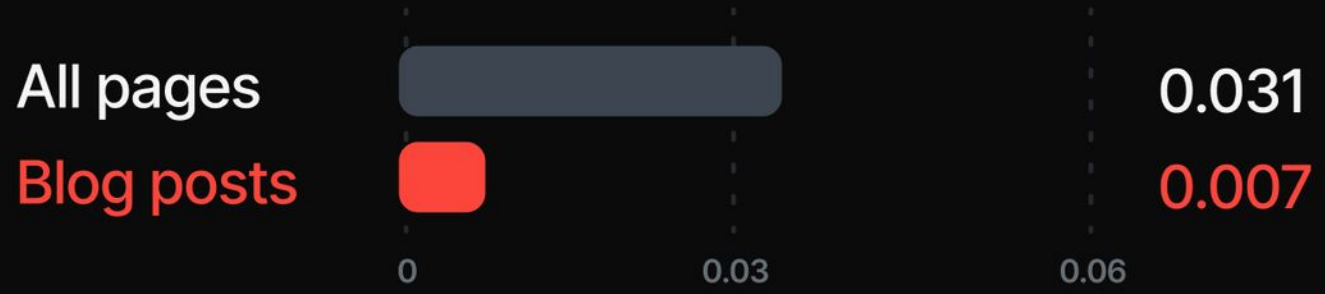


Correlation, scale 0-0.06. Values read from chart, approximate

Data From Surfer AI Tracker

MYTH

LINKING TO WIKIPEDIA DOES NOTHING



Correlation, scale 0-0.06. Values read from chart, approximate

Data From Surfer AI Tracker

MYTH

LLMS.TXT IS THE NEW META KEYWORDS. SCHEMA ISN'T.



LLMS.txt



Schema (JSON-LD)

Jul–Sep 2026. Association, not effect

Data From Surfer AI Tracker

MYTH

LLMS.TXT FLIPS BY ENGINE



Sites with llms.txt minus without, pp. 4,366 prompt × engine contexts. Jul–Sep 2026

Data From Surfer AI Tracker

MYTH

SCHEMA IS UP ON EVERY ENGINE



Skip llms.txt. Keep structured data on pages you actually maintain.

Pages with JSON-LD minus without, pp. 40,712 crawled pages. Jul–Sep 2026. Association, not effect


BIG ISN'T
THE SAME
AS CITED

ACT III

WHAT GETS YOU CITED

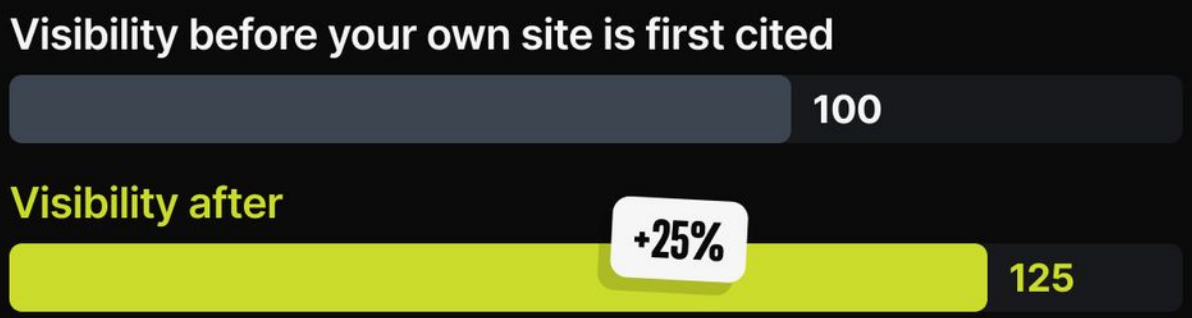
LEVER

GOOGLE AI IS STILL GOOGLE

95% of cited domains rank in Google's Top 10 for at least one keyword.

95%

THE FIRST CITATION LEVER IS THE BIG ONE



First wave only: later citations add little on top

Get your own site cited once - that first citation moves the needle most.

Visibility index, before = 100. 2,685 brands, 17,299 newly cited own pages. Jun 2025 – Sep 2026

Data From Surfer AI Tracker

LEVER

HOW TO WRITE FOR AI

Format beats gimmicks

1. Answer in the first 50 words
2. Match the intent
3. Write chunks AI can lift
4. Cover the facts
5. Be specific, not original

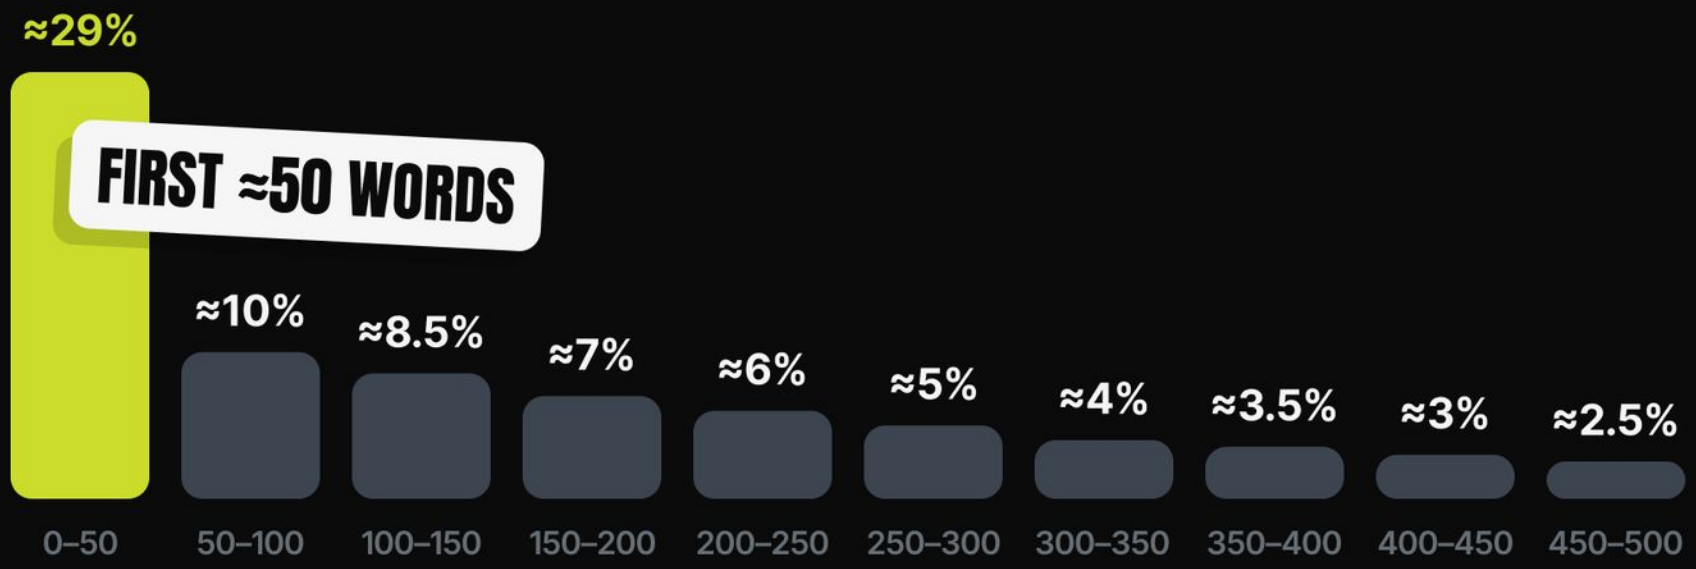
~21% More citations when the page matches the intent - the strongest of 20 page features

20 page features scored by LLM. Values read from chart, approximate

Data From Surfer AI Tracker

LEVER

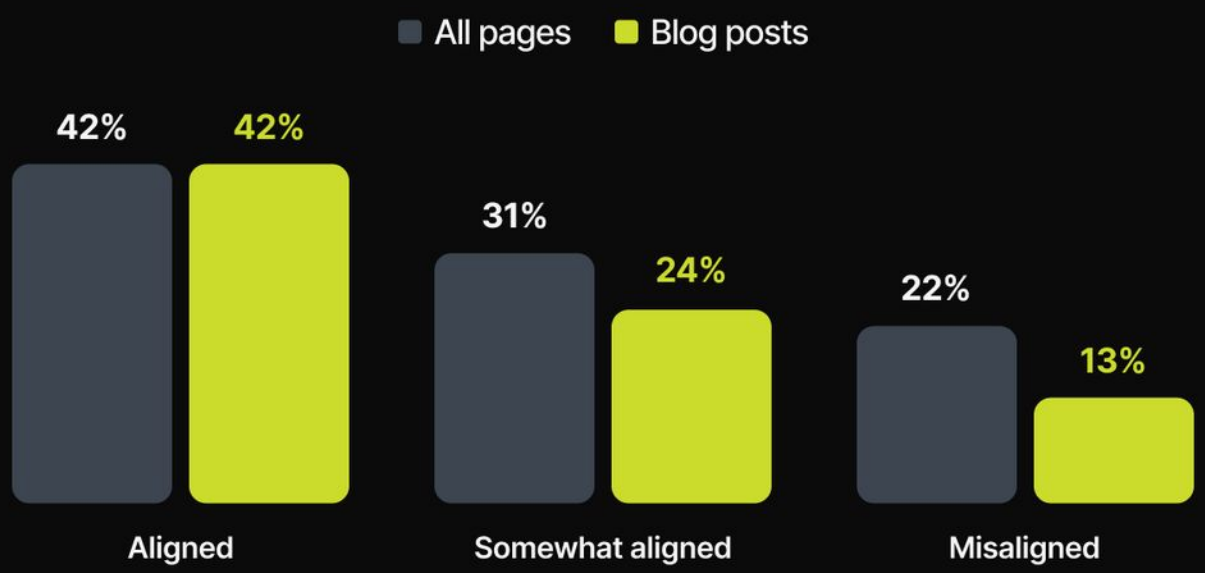
ANSWER IN THE FIRST 50 WORDS



Share of citations by position of the cited passage in the page, words. Values read from chart, approximate

Data From Surfer AI Tracker

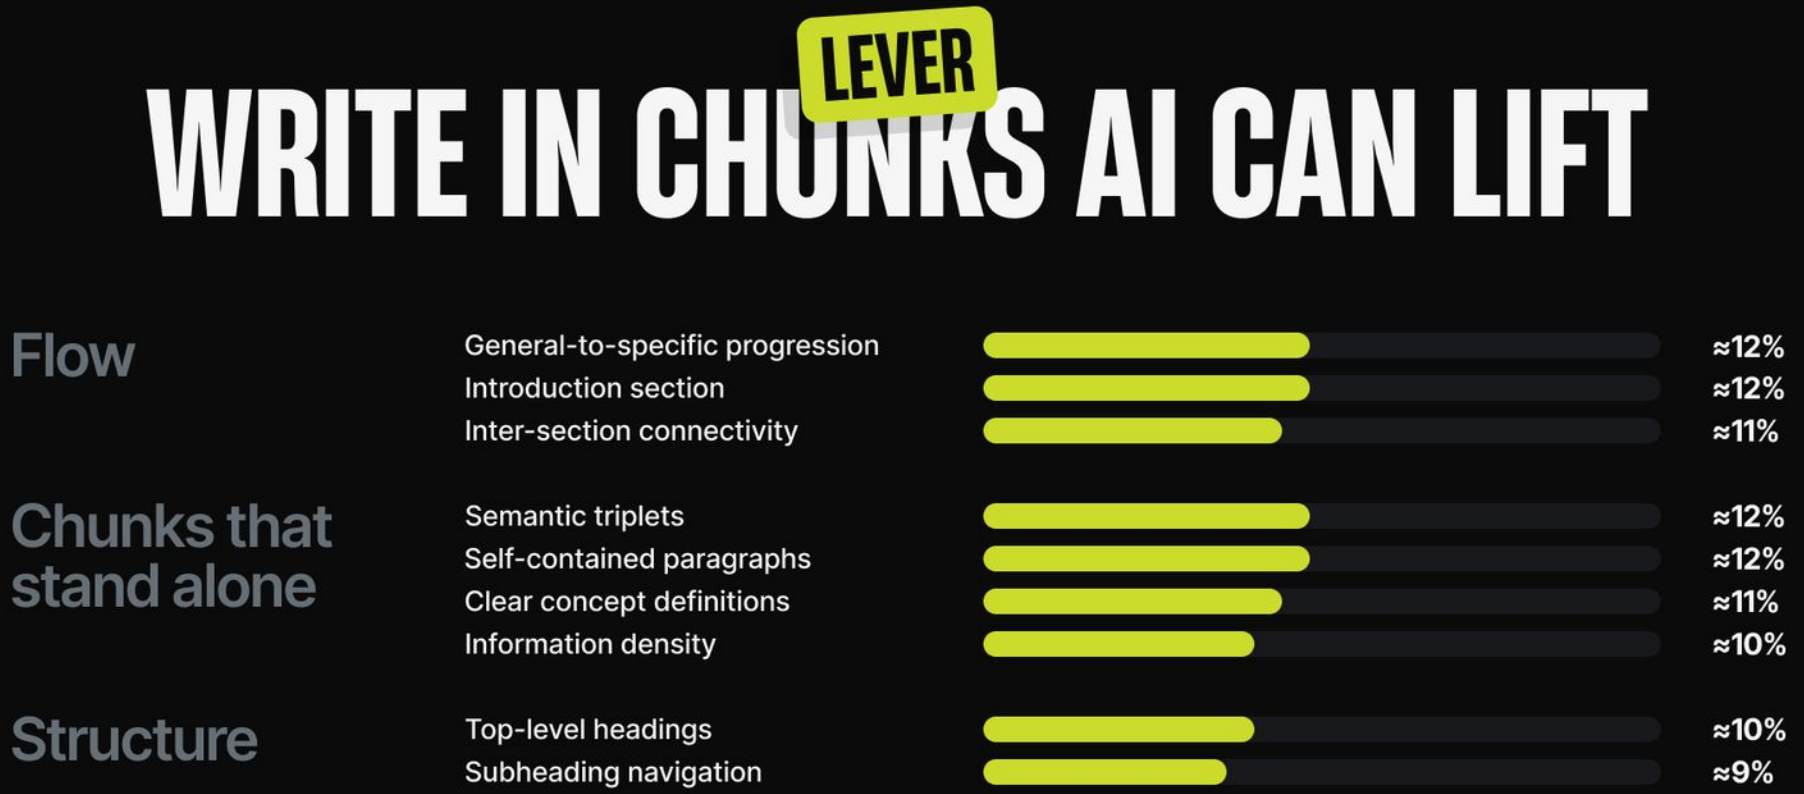
MATCH INTENT OR LEVER DON'T BOTHER



How often a page was a source, by intent alignment scored by LLM. All pages vs blog posts

Data From Surfer AI Tracker

LEVER
WRITE IN CHUNKS AI CAN LIFT

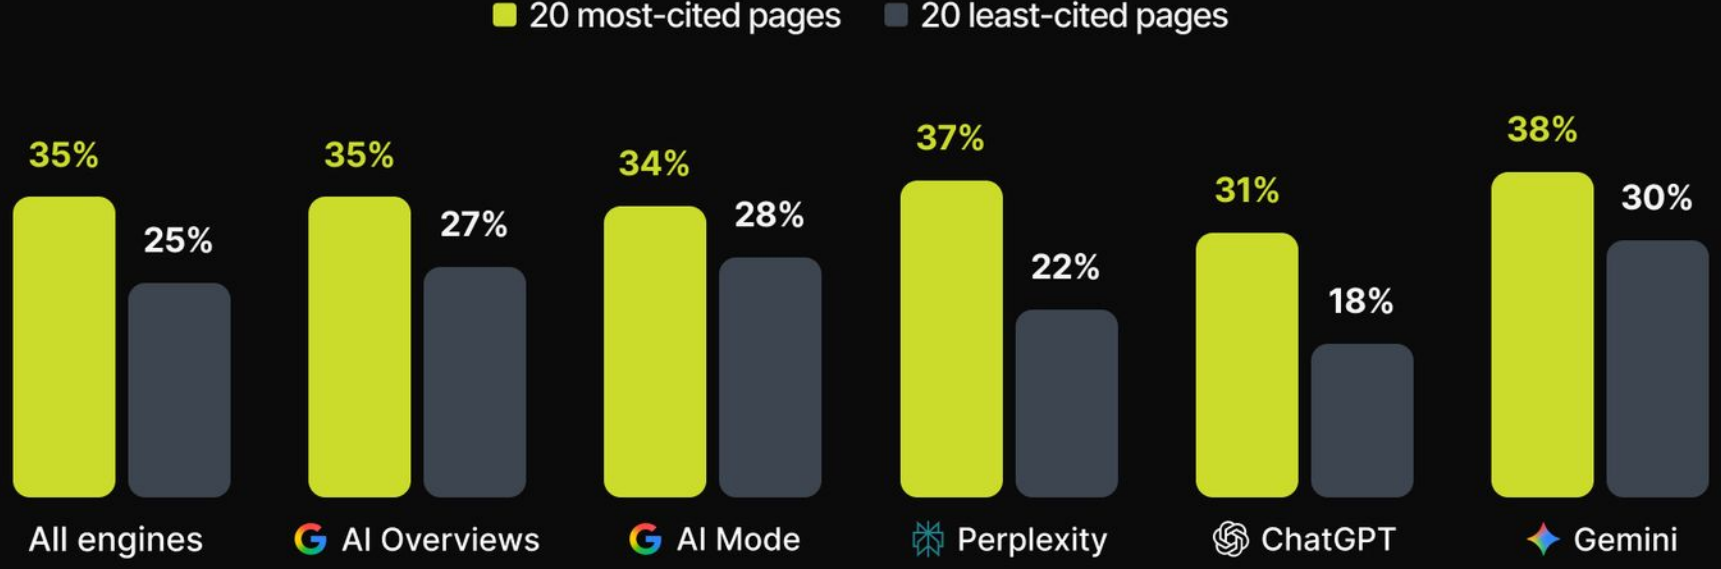


How much more often a page with the feature is cited. Scale 0–25% (same as #8). Values read from chart, approximate

Data From Surfer AI Tracker

FACTS WIN THE LEVER SOURCE GAME

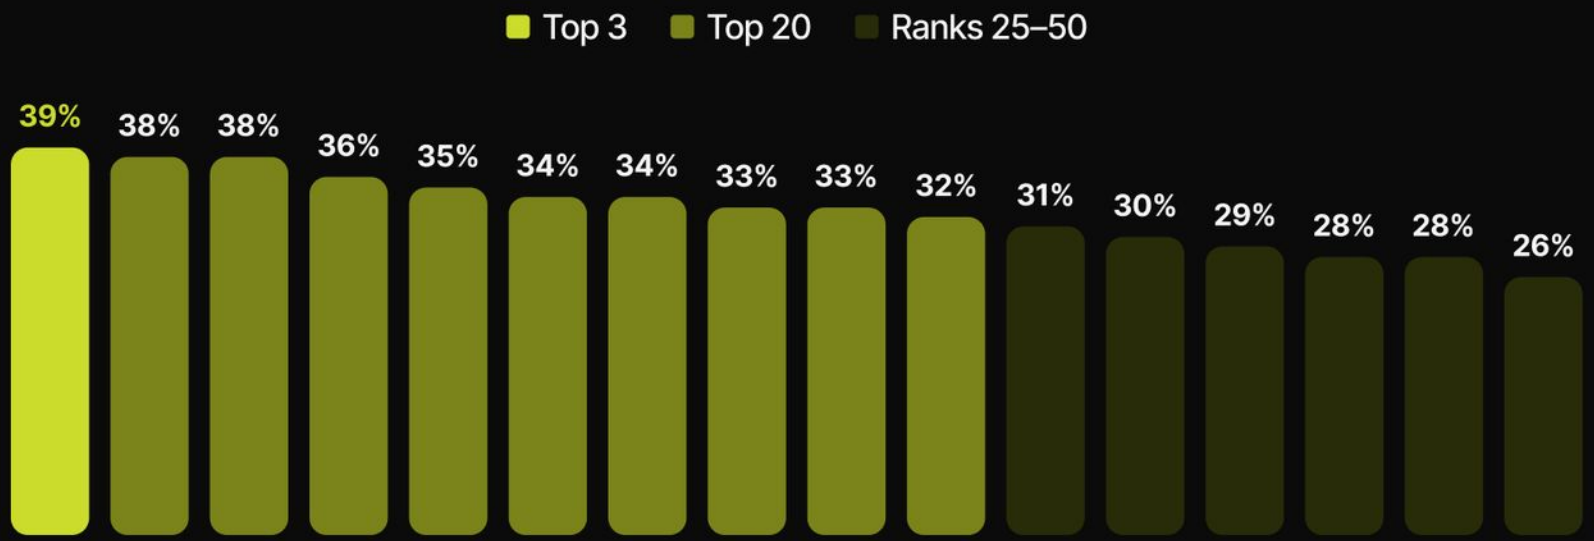
NEW STUDY



Facts coverage (Content Editor metric), %. 100 prompts per engine. Jul-Sep 2026. Page length doesn't explain it

Data From Surfer AI Tracker

EVERY RANK DOWN, LEVER FEWER FACTS



Mean facts coverage by citation-frequency rank within prompt × engine, %. 5-rank moving average. Jul–Sep 2026

Data From Surfer AI Tracker

LEVER

UNIQUE ANGLE

0.12 Correlation between a unique angle and being cited - small, but positive on every engine



Share of high-uniqueness pages showing each trait (multi-label), scale 0–50%

Data From Surfer AI Tracker

TWO KINDS OF LISTS

THEIR LIST

Paid or earned

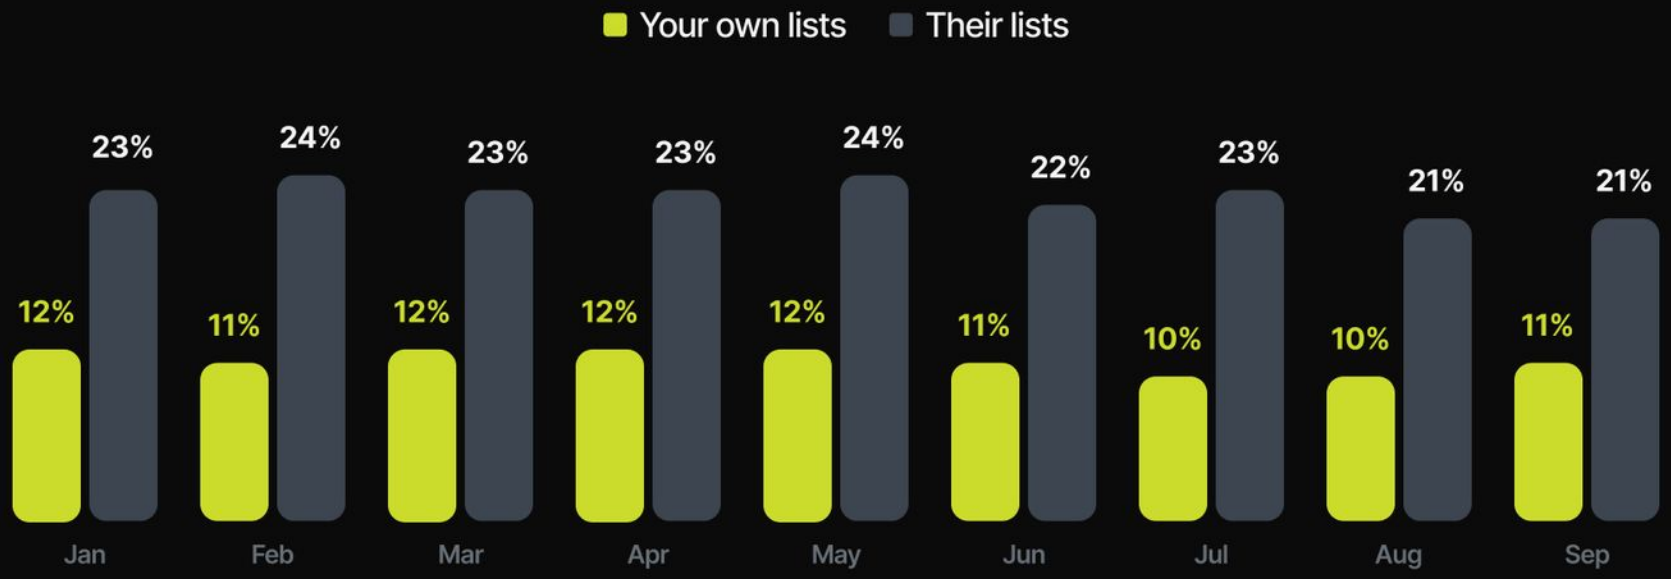
YOUR OWN LIST

Under control and free

Listicle study. 8,677 reports, 238.6M citations. Jan–Sep 2026

Data From Surfer AI Tracker

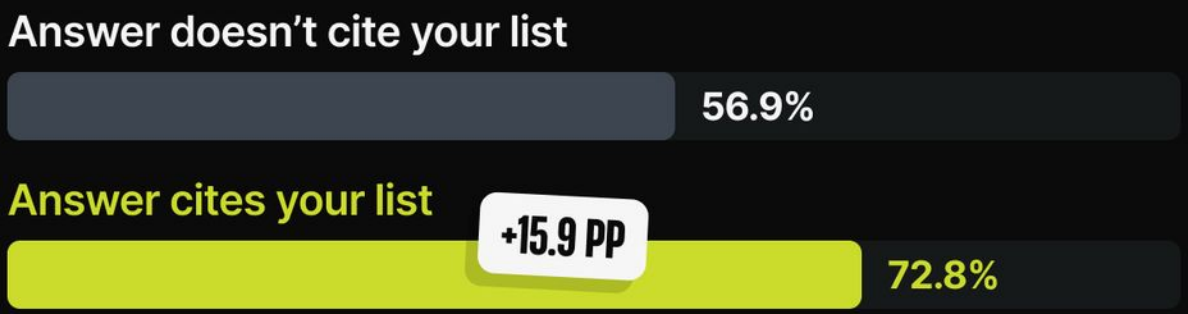
LISTS AREN'T LOSING GROUND



Share of citations. Same 546 reports, Jan–Sep 2026

Data From Surfer AI Tracker

YOUR LIST GETS LEVER YOU NAMED



And you're named higher: average position 3.08 → 2.77

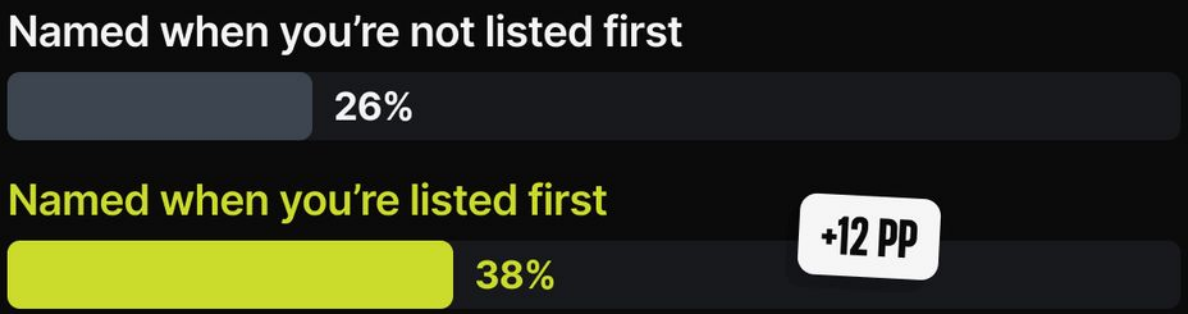
Every competitor on your list also gets +3.7 pp.

Listicle study. 8,677 reports, 238.6M citations. Jan–Sep 2026

Data From Surfer AI Tracker

LEVER

PUT YOURSELF FIRST



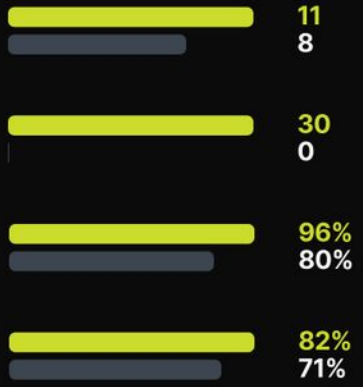
Listicle study. 8,677 reports, 238.6M citations. Jan–Sep 2026

Data From Surfer AI Tracker

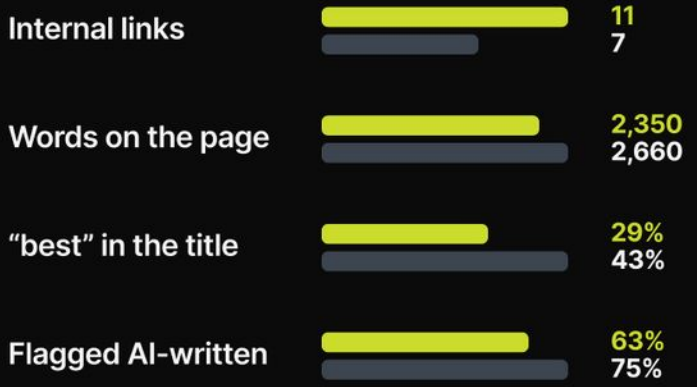
CITED VS. NAMED: TWO GAMES

LEVER

Gets the page cited



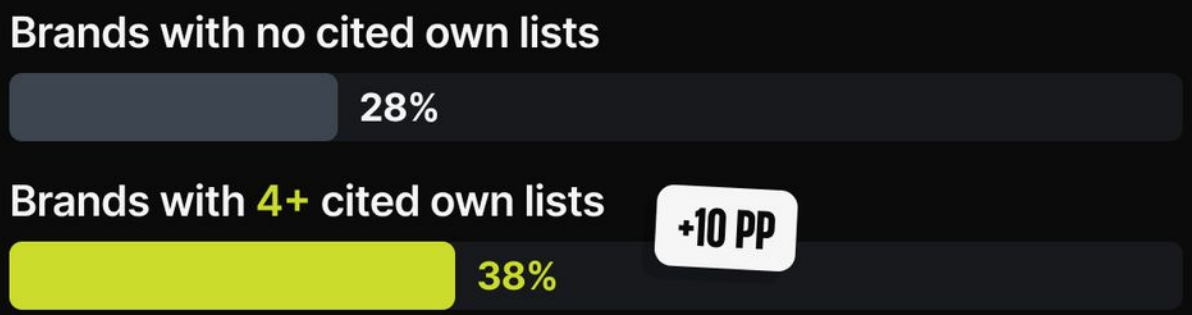
Gets your brand named



Writing-factors study. 201,534 self-promo pages, 5,162 reports. Jul–Sep 2026

Data From Surfer AI Tracker

MORE LISTS, LEVER MORE PRESENCE



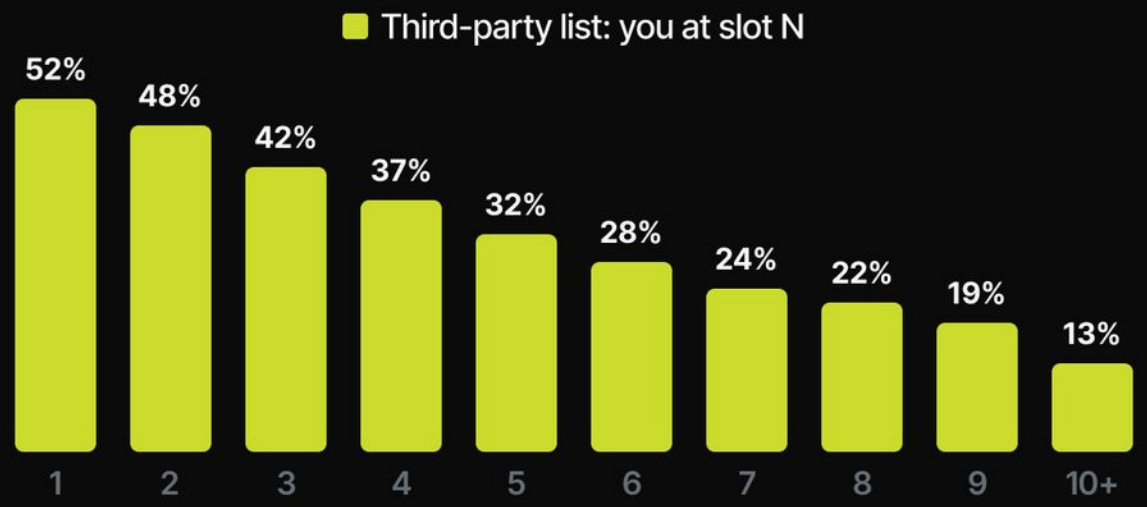
Build a set of comparison pages, not one.

Listicle study. 8,677 reports, 238.6M citations. Jan–Sep 2026

Data From Surfer AI Tracker

THE SLOT IS LEVER EVERYTHING

How often you're named in the answer that cites the list, by your slot on it.



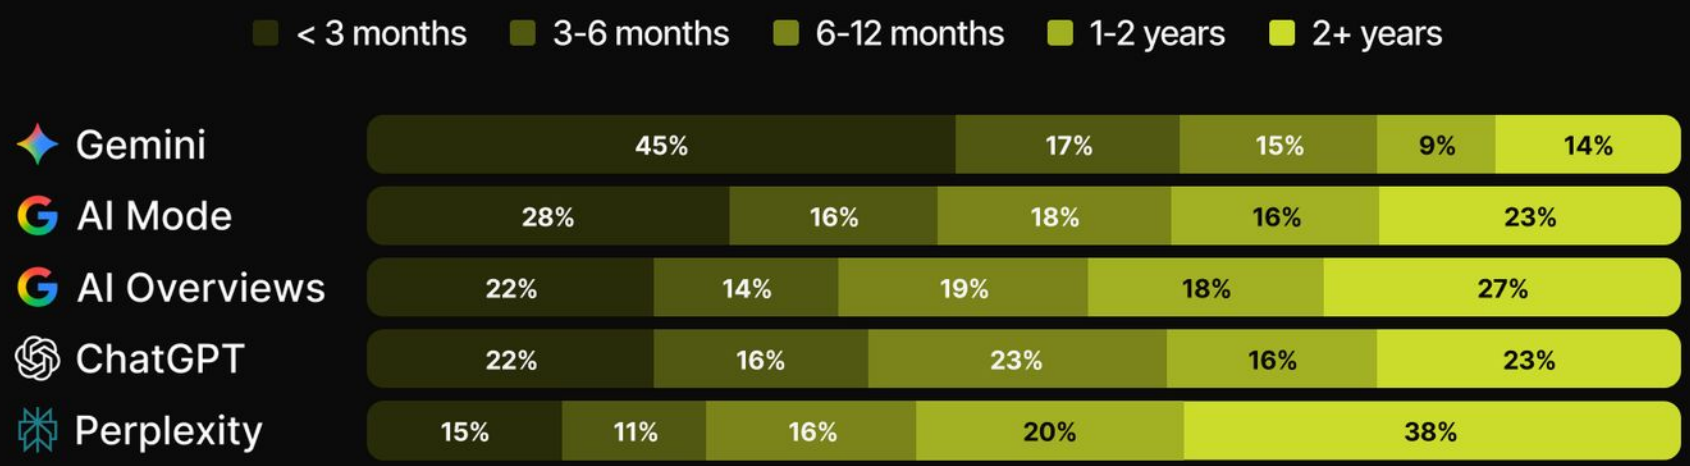
5% random sample of listicle citations. 8,677 reports, 238.6M citations. Jan–Sep 2026

Data From Surfer AI Tracker

YOUR LIST MAKES YOU IN CHARGE OF THE MESSAGE.

Their list makes the category visible.

OLDER CONTENT LEVER STILL GETS CITED



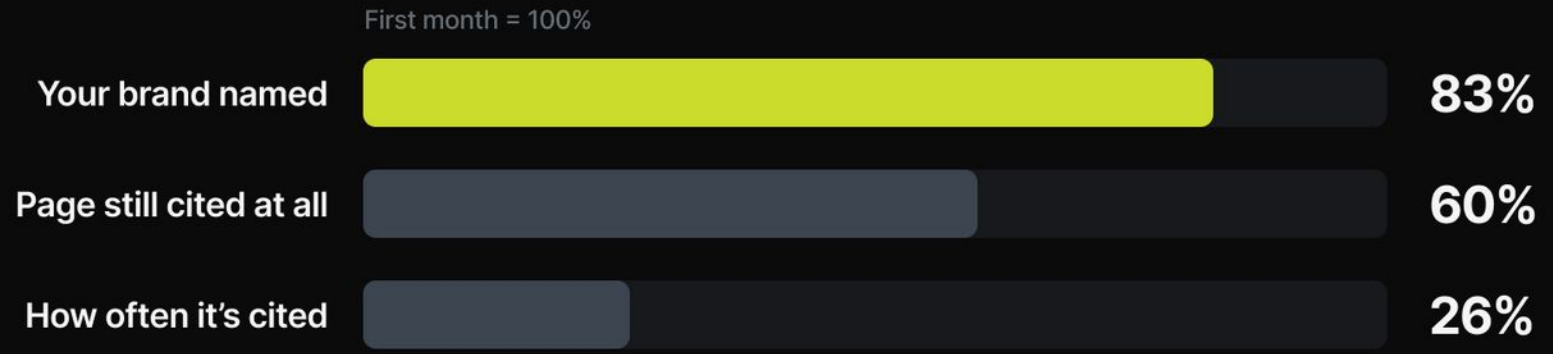
57–74% of citations in AI Mode, AI Overviews, ChatGPT and Perplexity go to pages older than six months.

Share of AI citations by page age (since publication date), Mar–May 2026. Pages with an explicit date only

Data From Surfer AI Tracker

THE PAGE FADES. THE NAME STAYS.

11 weeks after AI first cites your own list, as a share of the first month



Use the list to earn the mention; refresh it to keep the citations.

Writing-factors study. 74,718 + 146,151 page × report pairs. Jan–Jul 2026

Data From Surfer AI Tracker

WHAT 50M ANSWERS TOLD US

MEASURE HONESTLY

Ask again, it changes

Scrape the chat, not the API

Track intents, not rewordings

No single “AI visibility” number

IGNORE THE PROXIES

Global top domains aren’t yours

Authority gets you named, not cited

The GEO checklist is decoration

llms.txt does nothing

PULL THE LEVERS

Rank in Google

Answer first, cover the facts

Get on lists, win the top spot

The first citation is the big one


DATA
SHINY OBJECT

DON'T TRUST ME. CHECK YOUR BRAND.



THANK YOU!



Michał Suski

msu.ski

surfer.ai

Surfer