

How to Track Brand Mentions in Perplexity AI (10 Best Tools for 2026)

Metehan Yesilyurt

AI Search & Information Retrieval Researcher · metehan.ai

Published: September 06, 2026 · Canonical: <https://metehan.ai/articles/track-brand-mentions-perplexity-ai/>

Perplexity often answers with just one or two sources. That makes it brutal and simple at the same time. You are either the cited source or your competitor is, with very little middle ground. So tracking your brand in Perplexity is not a vanity metric, it is the difference between being the answer and being invisible in a channel that now handles hundreds of millions of queries a month.

There is one trap most guides miss. In Perplexity, a mention and a citation are not the same thing. Perplexity will often cite your page as a source while describing the content without ever writing your brand name in the answer. So you need to track mention rate and citation rate separately, or you will misread your own visibility. Below I cover both, compare the 10 best tools to track Perplexity brand mentions, and back the strategy with peer-reviewed research instead of vendor claims.

TL;DR

- Best overall tool to track Perplexity brand mentions: Peec AI. Daily Perplexity tracking, mentions and citations side by side, sentiment, and unlimited seats.
- Best for pure mention-share: Gumshoe AI. Best for agencies and multi-country: SE Visible.
- Track mentions and citations as two different metrics, because Perplexity frequently cites a page without naming the brand.
- Track all the Sonar models, since the same prompt can surface different brands across Sonar, Sonar Pro, and the Reasoning variants.
- Perplexity is non-deterministic, so measure on a schedule and read trends over weeks, not a single run.

Perplexity by the numbers

- Perplexity now serves more than 400 million queries a month^[1], with reported monthly visits around 159 million and roughly 15 million monthly active users.
- It cites the most sources per prompt of the major engines, yet each source carries less weight than a ChatGPT citation, per the citation selection and absorption study^[2].
- Its sources barely overlap with Google's normal results, because generative engines pull from a different source ecosystem^[3] and favor fresher, more niche pages.
- The same prompt returns different sources across runs^[4], which is why a one-time check is not tracking, it is guessing.

Mentions vs citations in Perplexity, and why it matters

A citation means Perplexity pulled a URL as a source. This is the distinction that separates a real Perplexity tracker from a dashboard. A mention means your brand name actually appears in the written answer. Perplexity regularly does one without the other. It can cite your pricing page while describing the product in generic terms, so you earned the traffic signal but not the brand recall. Or it can name your brand from its model memory while citing a third-party review instead of your site, so you have the mindshare but none of the click.

Track both. Citation rate tells you which of your pages Perplexity trusts as evidence. Mention rate tells you whether the answer actually says your name. Sentiment tells you whether that mention helps or hurts. A tool that collapses these into one "visibility score" hides the exact problem you need to fix.

How to track your brand mentions in Perplexity

You cannot see Perplexity in Google Search Console, so you need a dedicated AI visibility tracker that scrapes the real Perplexity answer. Three things make the difference between accurate tracking and noise. First, track every Sonar model, because Sonar, Sonar Pro, and the Reasoning variants can surface different brands for the same prompt. Second, build your prompt set from real demand instead of guessing, the way I describe in [stop guessing which prompts to track](#)^[5]. Third, measure on a schedule, because Perplexity is non-deterministic and the research is clear that you should [not rely on a single measurement](#)^[6] and should treat AI visibility as a number with [real uncertainty](#)^[7]. If you want to understand why a given page gets picked, I reverse-engineered [Perplexity's ranking patterns](#)^[8] and its [likely algorithm weights](#)^[9].

Perplexity brand mention trackers compared


Every tool here tracks Perplexity. The columns that matter are whether it separates mentions from citations, whether it reads sentiment, and who it fits best.

Tool	What it tracks in Perplexity	Sentiment	Best for	Starts at
Peec AI	Mentions, citations, share of voice	Yes	All-round teams and agencies	€89/mo
Scrunch AI	Mentions, citations, personas	Yes	Enterprise brand safety	\$250/mo
SE Visible	Mentions, citations, cached answers	Yes	Agencies, multi-country	\$99/mo
Gumshoe AI	Share of LLM, mention share	Yes	Pure mention-share, pay as you go	Free, then usage
AthenaHQ	Mentions, citations, content gaps	Yes	Tracking plus pitches	\$295/mo
Hall AI	Mention and citation alerts	Yes	Slack-first alerting	Free tier, then paid
Ahrefs Brand Radar	Share of voice, citations	Limited	Existing Ahrefs users	\$199/mo
Semrush AI Visibility Toolkit	Mentions, citations, sentiment	Yes	Existing Semrush users	\$99/mo
Rankscale	Page-level citations, position	Yes	Budget, hands-on SEOs	€20/mo
Omnia	Citations, answer share, briefs	Yes	Action-first scaleups	€79/mo

The 10 best tools to track brand mentions in Perplexity AI

To determine the most reliable software for monitoring Perplexity citations, I evaluated 10 leading platforms against live Perplexity Pro searches. We compared prompt execution frequency, citation chip extraction accuracy, source URL mapping, and sentiment scoring capabilities. Here is the hands-on ranking:


Peec AI



Screenshot placeholder: Peec AI Perplexity dashboard

Peec AI^[10] is the tracker I reach for first, because it shows your Perplexity mentions and citations side by side, adds sentiment and competitor share of voice, and tracks ChatGPT, Gemini, Google AI Mode, Claude, and Copilot in the same view. Setup takes minutes, the dashboard reads cleanly for non-technical stakeholders, and every plan ships with unlimited seats so the whole team and your clients fit. Honest limit: it is a tracking and insight layer, so it surfaces the gaps but you still do the fixing. Starts at €89 per month.


Scrunch AI



Screenshot placeholder: Scrunch AI dashboard

Scrunch AI^[11] is the enterprise pick, tracking Perplexity mentions and citations by persona, funnel stage, and geography, with sentiment and a real-time bot feed that shows when crawlers hit your pages. Its Agent Experience Platform can serve an AI-friendly version of your site to lift citations. Honest limit: it leans toward monitoring over action, and the depth comes with enterprise pricing. Starts at \$250 per month.


SE Visible



Screenshot placeholder: SE Visible dashboard

[SE Visible](#)^[12] tracks Perplexity alongside ChatGPT, Gemini, AI Overviews, and AI Mode, and stores cached snapshots so you can open the actual answer and verify every mention. All plans include unlimited seats, which makes it economical for agencies running many client brands across countries. Honest limit: it is monitoring only and exports are CSV only. Starts at \$99 per month.


Gumshoe AI



Screenshot placeholder: Gumshoe AI dashboard

[Gumshoe AI](#)^[13] is the most focused on the mention question, built around a Share of LLM metric that measures how often Perplexity and other models name your brand against competitors. It simulates real buyer personas to predict what they ask, then points to the actions most likely to lift your share. Honest limit: the pay-as-you-go model is great for testing but can scale unpredictably at high volume. The first report runs are free.


AthenaHQ



Screenshot placeholder: AthenaHQ dashboard

[AthenaHQ](#)^[14] tracks Perplexity mentions and citations alongside the other major engines and pairs it with content gap suggestions and pitch workspaces, so it leans more toward action than most monitors. Honest limit: some of its content recommendations still feel early, so review them before you publish. Starts at \$295 per month.


Hall AI



Screenshot placeholder: Hall AI dashboard

[Hall AI](#)^[15] is the alerting layer, with a Slack-first dashboard and citation heatmaps that catch a Perplexity mention or a lost citation the day it changes. Its agent analytics show how AI crawlers move through your site. Honest limit: data depth can be inconsistent for long-term trend lines, so pair it with a deeper tracker. It has a free tier.


Ahrefs Brand Radar



Screenshot placeholder: Ahrefs Brand Radar dashboard

[Ahrefs Brand Radar](#)^[16] tracks your Perplexity share of voice and citations and ties them back to your backlink and ranking data, which is the real reason to use it if you already live in Ahrefs. Honest limit: it is lighter on sentiment and step-by-step fixes, and it needs an Ahrefs subscription plus a per-engine add-on. Starts at \$199 per month per engine.


Semrush AI Visibility Toolkit



Screenshot placeholder: Semrush AI Visibility Toolkit dashboard

[Semrush](#)^[17] brings Perplexity mention, citation, and sentiment tracking into its AI Visibility Toolkit, so existing users get share of voice and the exact pages cited without leaving the platform. Honest limit: the entry toolkit covers only 25 prompts and one domain, so multi-brand teams scale cost quickly. Starts at \$99 per month.


Rankscale



Screenshot placeholder: Rankscale dashboard

[Rankscale](#)^[18] is the budget, AI-native option, with page-level Perplexity citation tracking, a visibility score, competitor benchmarking, and schema audits for hands-on SEOs. Honest limit: prompt input is mostly manual and the interface is technical. Starts at €20 per month.

Omnia



Screenshot placeholder: Omnia dashboard

[Omnia](#)^[19] tracks Perplexity citations and answer share, then turns the gaps into a brief that says what to build and which sources to pitch, which is the point for a lean team that needs to ship fixes. Honest limit: it is built for scaleups, so very large enterprises with deep governance needs may want something heavier. Starts at €79 per month.

What the research says about getting cited in Perplexity

Perplexity runs a live search for almost every query, so fresh, well-structured, evidence-rich pages win. The science backs this. Perplexity cites broadly, more sources per answer than other engines, but the pages it actually leans on are longer, better structured, and rich in extractable evidence^[2] like definitions, numbers, and comparisons. Structural cleanup alone lifted citation rates by 17.3% across six engines^[20], and a study that

included Perplexity found that [metadata and freshness, semantic HTML, and structured data](#)^[21] had the strongest link to being cited. In head-to-head tests, [topical relevance and position drive the first citation](#)^[22], while formatting-only edits barely move it.

Two cautions. Generative search engines have a long history of citing pages that do not fully support their statements, as the foundational [verifiability study](#)^[23] showed, so verify that the answer represents you correctly, not just that you were cited. And because answers are non-deterministic, read trends over weeks. I like to confirm findings in the code and in the lab, which is the whole point of [finding it in the code and watching science prove it](#)^[24].

How to get mentioned more in Perplexity

Start by auditing the prompts where Perplexity answers your audience but never names you. Those are your topical gaps. Then write content that answers the question directly near the top, with question-style headings, clear definitions, comparisons, and schema, since that is the structure Perplexity lifts from. Keep pages fresh, because Perplexity prefers recently updated sources, and earn placements on the third-party sites it already trusts, because earned media wins more AI citations than your own pages. Build depth across a topic rather than one post, so the model treats you as authoritative. For the full toolkit, see my [AI visibility tools guide](#)^[25] and [AEO and GEO tools guide](#)^[26], and for the citation mechanics, [the RRF method for getting cited](#)^[27].

Frequently asked questions

Here are direct, authoritative answers to frequently asked questions about tracking brand citations, analyzing source domains, and monitoring brand sentiment on Perplexity AI:

Which tool is best for tracking brand mentions in Perplexity?

Pee AI for most teams, because it shows mentions and citations side by side, adds sentiment and competitor share, and includes unlimited seats. For mention-share specifically, Gumshoe AI; for agencies and multi-country tracking, SE Visible.

What is the difference between a mention and a citation in Perplexity?

A citation means Perplexity pulled your URL as a source. A mention means your brand name appears in the written answer. Perplexity often does one without the other, so track both. Citation rate shows which pages it trusts as evidence, and mention rate shows whether the answer actually says your name.

Can I track Perplexity brand mentions in Google Search Console?

No. Search Console only reports Google's own results, so it cannot see Perplexity. You need a dedicated AI visibility tracker that scrapes the real Perplexity answer. You can sanity-check the basics first by [seeing whether the models even recognize your brand](#)^[28].

Do I need to track all the Perplexity Sonar models?

Yes, if you want an accurate picture. Sonar, Sonar Pro, and the Reasoning variants can surface different brands for the same prompt, so tracking only one model can hide a gap or a win.

How often should I track Perplexity?

Daily if you publish often, weekly at the minimum. Perplexity runs live searches and is non-deterministic, so a single check is unreliable. Track trends over weeks, not one run.

How do I get Perplexity to mention my brand?

Answer the question directly and early, use question-style headings and schema, keep pages fresh, build topical depth, and earn placements on the third-party sources Perplexity already cites. Then track whether your mention rate and citation rate move, and double down on what works.

Sources and further reading

- Zhang, He, Yao, [From Citation Selection to Citation Absorption: A Measurement Framework for GEO Across AI Search Platforms](#)^[2], arXiv:2604.25707.
- Chen, Wang, Chen, Koudas, [Navigating the Shift: A Comparative Analysis of Web Search and Generative AI Response Generation](#)^[3], arXiv:2601.16858 (EDBT/ICDT 2026).
- Grossman et al., [How Generative AI Disrupts Search: Google Search, Gemini, and AI Overviews](#)^[4], arXiv:2604.27790 (SIGIR 2026).
- Vishwakarma, Kumar, Jamidar, [What Gets Cited: Competitive GEO in AI Answer Engines](#)^[22], arXiv:2605.25517.
- Yu, Yang, Ding, Sato, [Structural Feature Engineering for Generative Engine Optimization](#)^[20], arXiv:2603.29979.
- Kumar, Palkhouski, [AI Answer Engine Citation Behavior: An Empirical Analysis of the GEO-16 Framework](#)^[21], arXiv:2509.10762.
- Liu, Zhang, Liang, [Evaluating Verifiability in Generative Search Engines](#)^[23], arXiv:2304.09848.
- [Quantifying Uncertainty in AI Visibility](#)^[7], arXiv:2603.08924, and [Don't Measure Once](#)^[6], arXiv:2604.07585.

Perplexity brand mention tracking is one input into AEO. Pair it with ChatGPT citation data and Gemini citation data and you have the minimum viable answer engine optimization picture. Tracking only Perplexity gives you a single AI engine in a multi-engine world.

REFERENCES & SOURCES

- [1] **more than 400 million queries a month** — <https://developer.nvidia.com/blog/spotlight-perplexity-ai-serves-400-million-search-queries-a-month-using-nvidia-inference-stack/>
- [2] **citation selection and absorption study** — <https://arxiv.org/abs/2604.25707>
- [3] **pull from a different source ecosystem** — <https://arxiv.org/abs/2601.16858>
- [4] **different sources across runs** — <https://arxiv.org/abs/2604.27790>
- [5] **stop guessing which prompts to track** — <https://metehan.ai/blog/stop-guessing-which-prompts-to-track/>
- [6] **not rely on a single measurement** — <https://arxiv.org/abs/2604.07585>
- [7] **real uncertainty** — <https://arxiv.org/abs/2603.08924>
- [8] **Perplexity's ranking patterns** — <https://metehan.ai/blog/perplexity-ai-seo-59-ranking-patterns/>
- [9] **likely algorithm weights** — <https://metehan.ai/blog/advanced-technical-insights-into-perplexity-ais-potential-algorithm-weights-a-case-study/>
- [10] **Peec AI** — <https://peec.ai/>
- [11] **Scrunch AI** — <https://scrunch.com/>
- [12] **SE Visible** — <https://visible.seranking.com/>
- [13] **Gumshoe AI** — <https://gumshoe.ai/>
- [14] **AthenaHQ** — <https://athenahq.ai/>
- [15] **Hall AI** — <https://usehall.com/>
- [16] **Ahrefs Brand Radar** — <https://ahrefs.com/brand-radar>
- [17] **Semrush** — <https://www.semrush.com/ai-seo/overview/>
- [18] **Rankscale** — <https://rankscale.ai/>
- [19] **Omnia** — <https://www.useomnia.com/>
- [20] **17.3% across six engines** — <https://arxiv.org/abs/2603.29979>
- [21] **metadata and freshness, semantic HTML, and structured data** — <https://arxiv.org/abs/2509.10762>
- [22] **topical relevance and position drive the first citation** — <https://arxiv.org/abs/2605.25517>
- [23] **verifiability study** — <https://arxiv.org/abs/2304.09848>
- [24] **finding it in the code and watching science prove it** — <https://metehan.ai/blog/i-found-it-in-the-code-science-proved-it-in-the-lab-the-recency-bias-thats-reshaping-ai-search/>
- [25] **AI visibility tools guide** — <https://metehan.ai/articles/best-ai-visibility-tools>
- [26] **AEO and GEO tools guide** — <https://metehan.ai/articles/best-aeo-geo-tools>
- [27] **the RRF method for getting cited** — <https://metehan.ai/blog/the-rrf-top-n-playbook-how-to-get-cited-by-chatgpts-web-mode/>
- [28] **seeing whether the models even recognize your brand** — <https://metehan.ai/blog/does-chatgpt-recognize-your-brand/>