



STATE OF
AI SEARCH
2025

Foreword from the CEO

When I joined Google's Search team, I thought I was working on the future of information discovery. What I discovered instead was that **I was witnessing the end of an era.**

During my time on Google's information acquisition team, the team responsible for gathering and processing the data that powers Search, I had a front-row seat to one of the most significant shifts in digital history, akin to the shift from web to mobile. I watched as traditional search behavior began evolving into something entirely new. Users were no longer content with lists of links; they wanted direct, intelligent answers.

The "10 blue links" paradigm that had defined search for decades was giving way to AI-generated responses capable of understanding context, synthesizing vast amounts of information, and delivering tailored guidance to each user.

We carefully chose the term "generative engines" rather than "answer engines" to emphasize the dynamic and adaptive nature of these systems, as each response is uniquely shaped by the question, context, and data at hand.

As we saw the rise of ChatGPT and the subsequent Search Generative Experience to AI Overviews to now AI Mode on Google, **the question wasn't whether or not this shift would come, but how quickly it will,** and which companies would seize the opportunity. That realization led to a stark choice: **remain on the sidelines of this revolution, or step into the arena and help shape it.** I chose to act.

At AthenaHQ, we're not merely building another analytics platform. We're pioneering an entirely new discipline called **Generative Engine Optimization, because we understand that the rules of digital visibility have fundamentally changed.** The billions of dollars invested in traditional SEO strategies over the past two decades represent knowledge that must now be reimaged for an AI-first world.

Rather than relentlessly pursuing clicks and traffic, imagine LLMs acting as skilled salespeople who engage with users, understanding their needs, and guiding them to relevant solutions. Overall, this results in less traffic, but each one will be far more qualified and valuable to your business.

This report reflects months of research across millions of AI responses. More importantly, it represents our mission to ensure that businesses don't just survive this transition, they thrive in it. **We've seen companies go from zero AI search presence to capturing 25% share of voice in their category within months.** We've watched smaller companies outmaneuver competitors twenty times their size by mastering these new dynamics first, and have seen large companies protect their lead on AI Search from smaller, more nimbler competitors.

The data you'll find in these pages isn't just academic research. It's actionable intelligence born from our unique position at the intersection of traditional search expertise and cutting-edge AI technology. **Every insight has been tested with our 200+ customers** who trust us to navigate this new landscape. **The future belongs to organizations that become part of the answer, not just part of the results.**

This report is your roadmap to that future.

October 8th, 2025




Andrew Yan

Co-Founder & CEO, AthenaHQ

Former Google Search Product Manager

We Are AthenaHQ

The leading platform for Generative Engine Optimization (GEO).
Helping brands build authority, optimize content for AI search,
and gain visibility in the era of generative answers.

 Y Combinator

Built With Experience From



Andrew Yan CEO



Alan Yao CTO

servicenow

 DeepMind

Google

Trusted By

200+ LEADING COMPANIES

Backed By

MILLIONS OF AI RESPONSES

Featured In



Table of Contents

Executive Summary

Key Learnings & Trends	05
Performance Benchmarks	06
Strategic Recommendations	10

AI Search Landscape

The End of an Era	11
The Paradigm Shift	11
Why This Matters Now	12

Athena Playbook

Why Traditional SEO Falls Short	13
The Strategy	14
Bringing It All Together	15
How to Read & Methodology	15
About AI sources monitoring	17
General Data	18

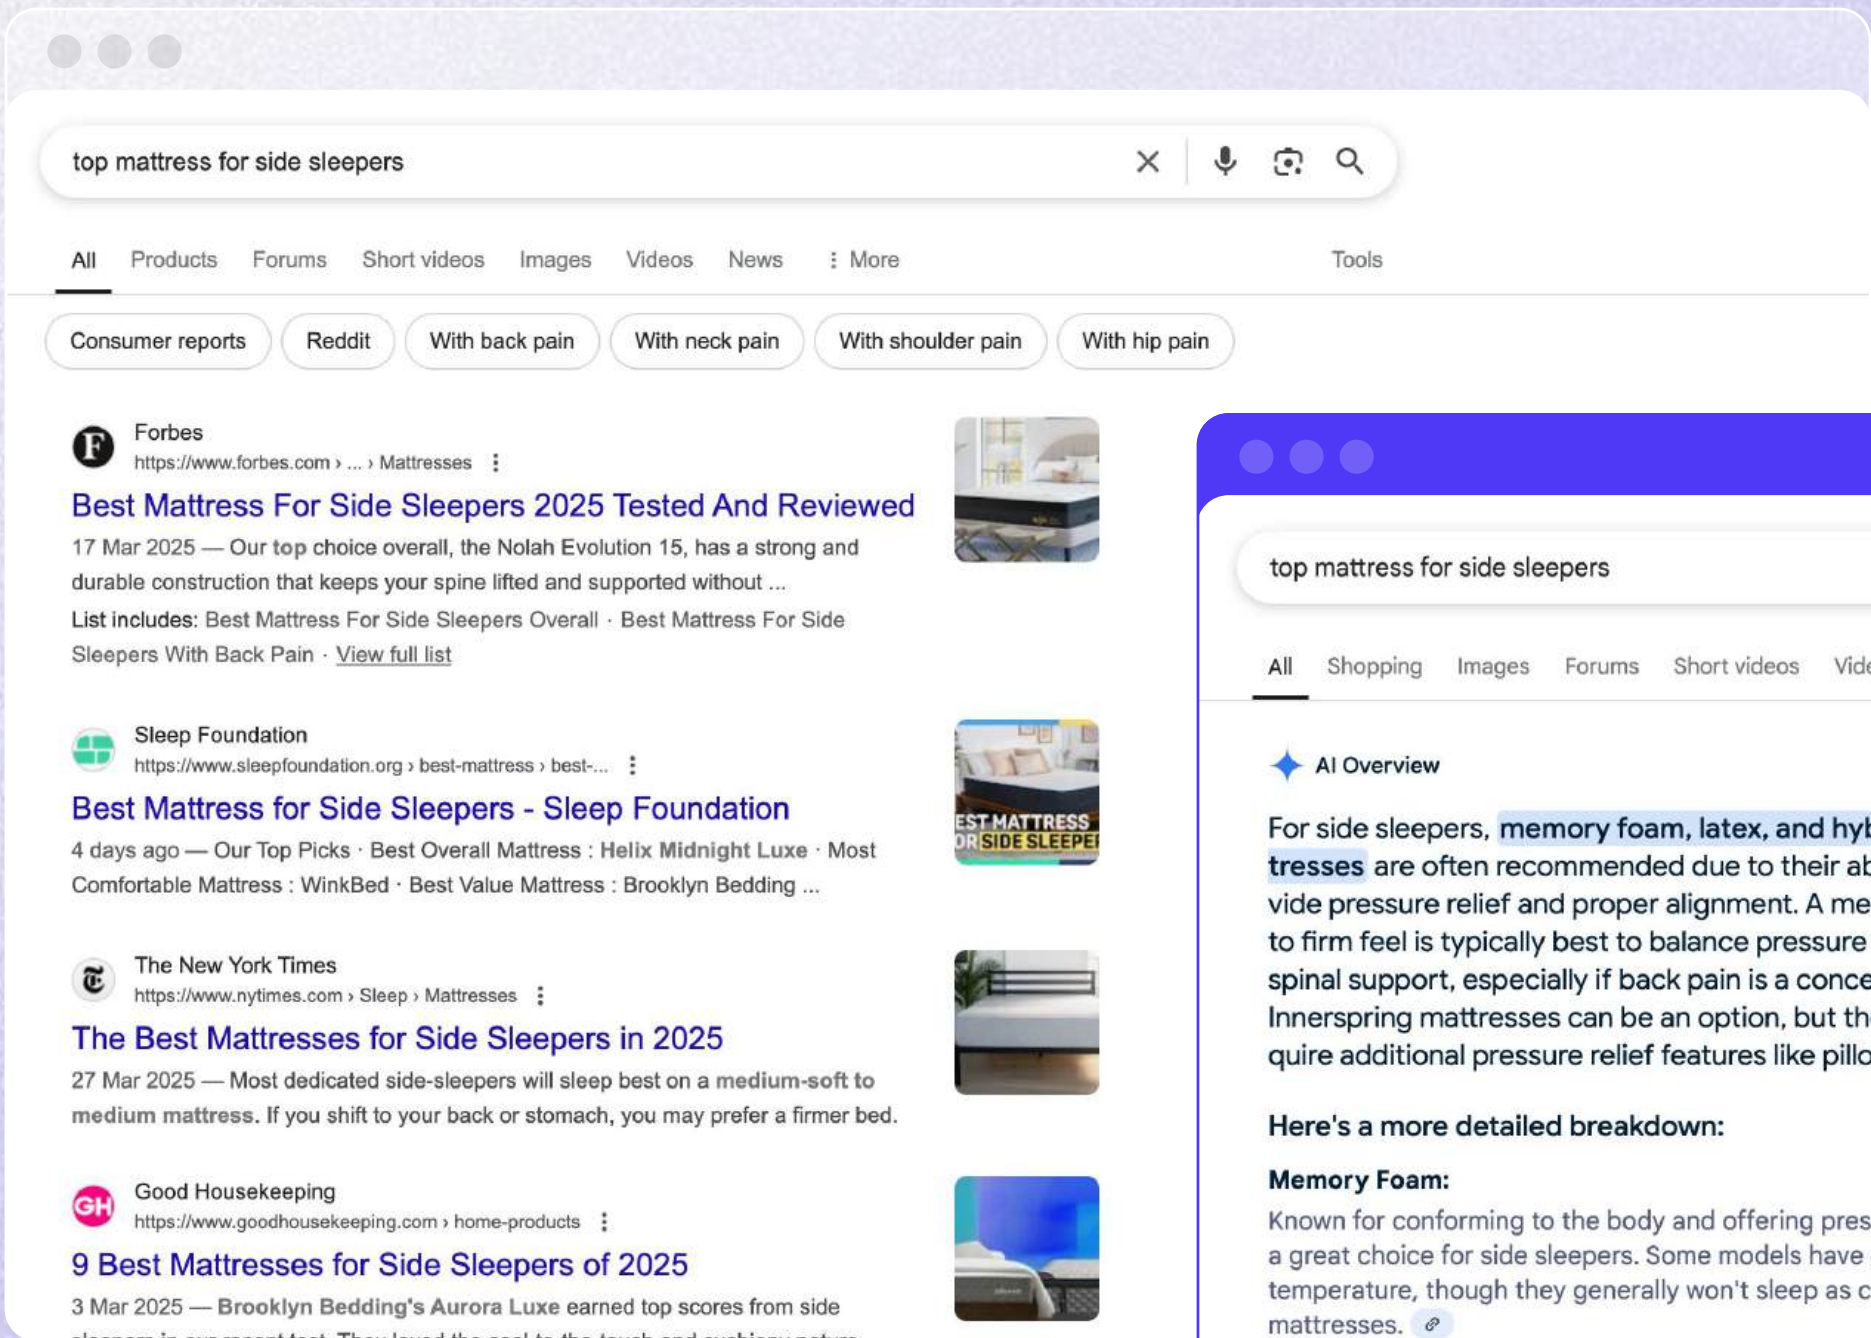


Executive Summary

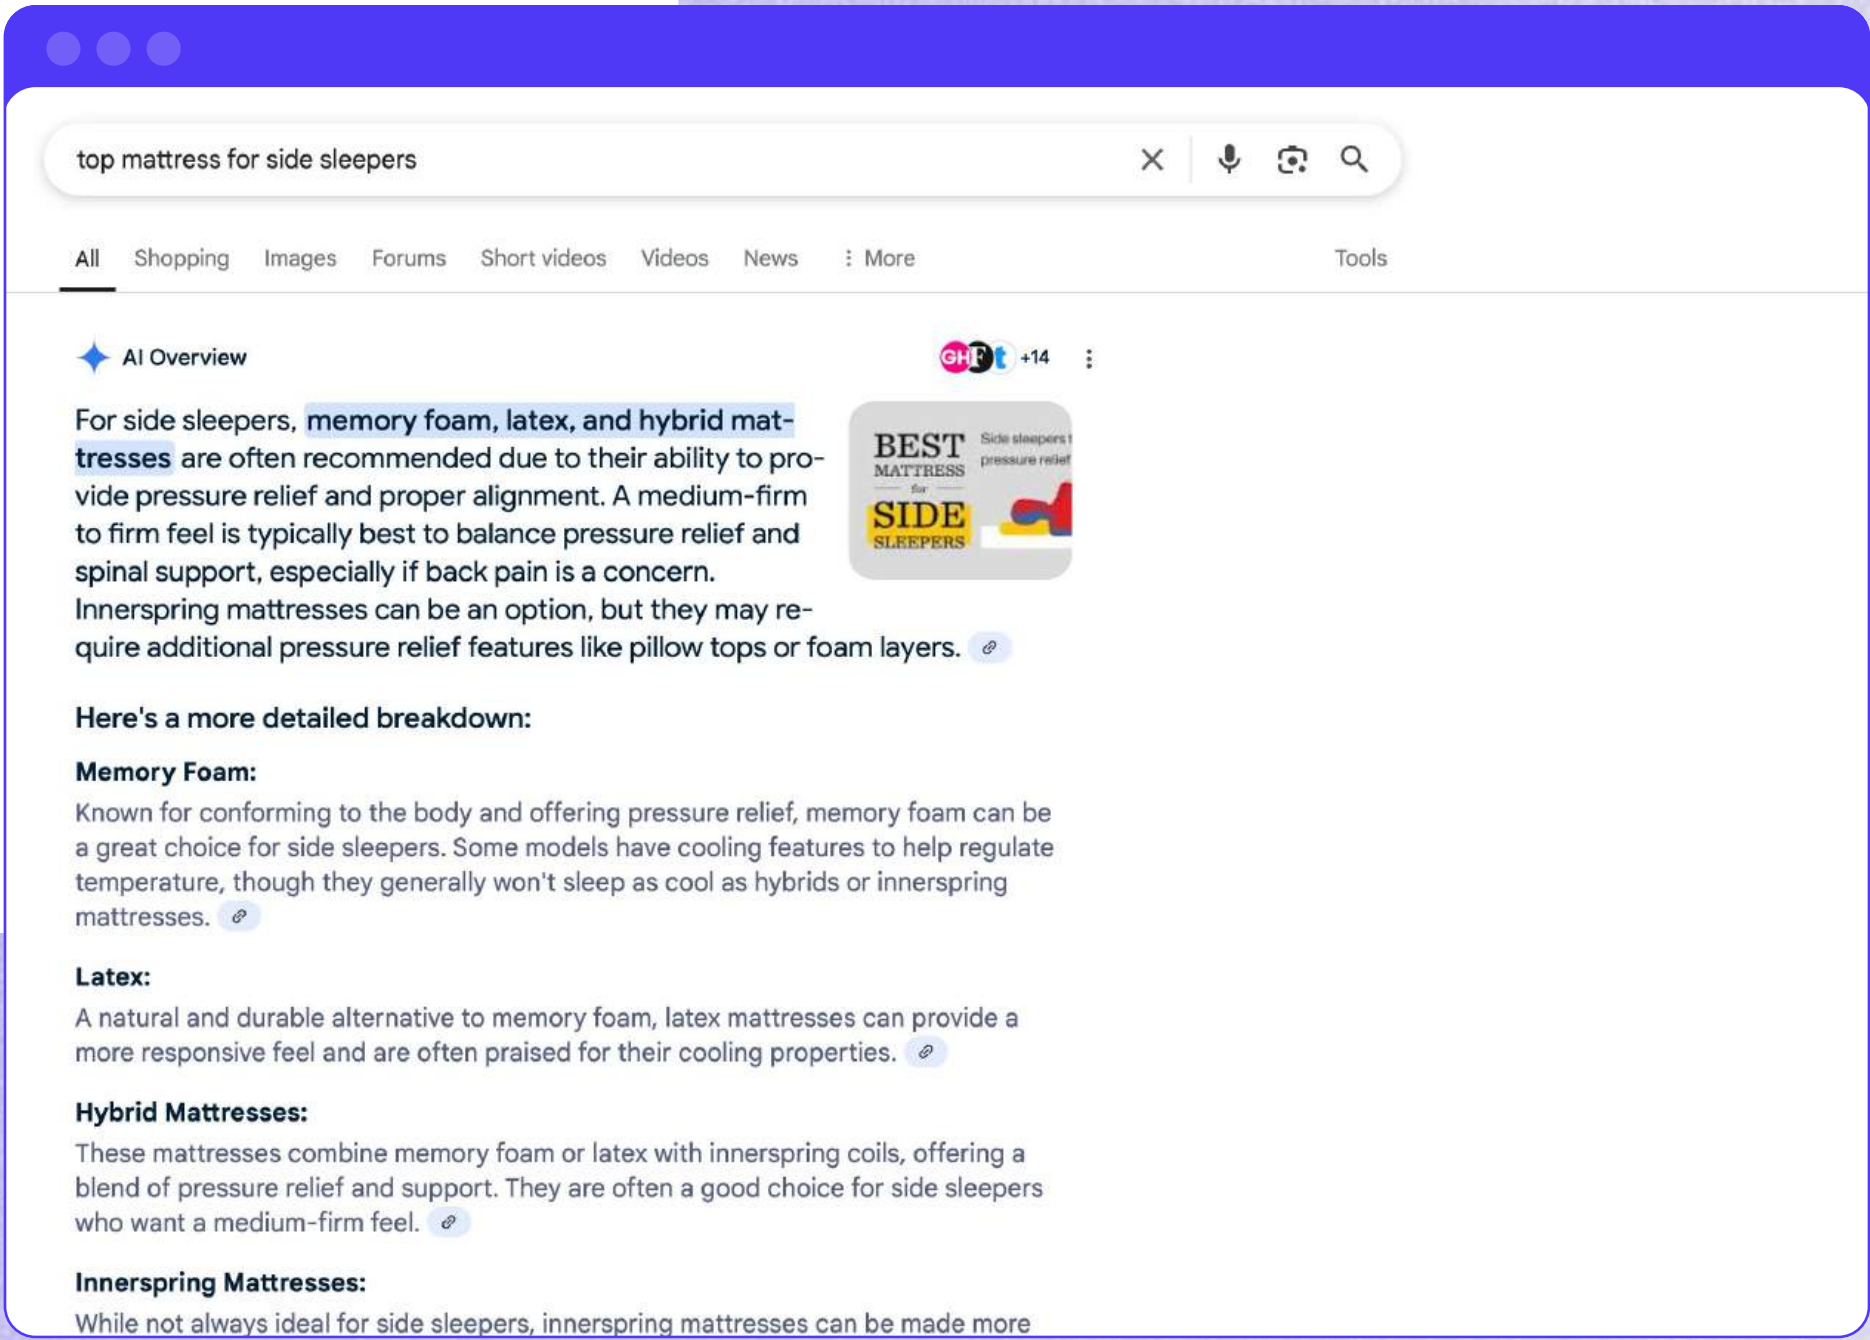
The arrival of generative models has transformed online search. Search engines like **Google have evolved beyond the "10 blue links,"** now delivering answers directly on the results page through AI Overviews and AI modes. **This has drastically reduced the need for clicks and slashed publisher traffic:** several leading sites report drops of more than 50% after these changes. Although the quantity of clicks is dramatically declining, the people that do click are better educated and much more likely to buy.

This is why it's so important to have a strong presence on LLMs like AI Overview and ChatGPT. In this new landscape, **traditional SEO is no longer enough.** Brands must rethink digital visibility strategies to ensure they become part of the answer, not just the results.

Traditional Search



AI Search



	Before	Today
SEO v GEO	Users click links and find answer on publisher sites	AI delivers answers directly on the result page
Traffic Impact	Sites received steady inbound traffic from search.	Publisher traffic has dropped 50% after AI rollouts.
User quality	Visitors arrive early in their journey, less informed, lower intent .	Fewer visits, but users arrive highly educated and closer to purchase.

Key Learnings & Trends

1. Search as an Answer, Not a Click: Users now get information directly in search, undermining impressions and clicks for external sites. The link-based web gives way to an **AI-driven direct answer model**, redefining how content is discovered and consumed.

2. Generative Engine Optimization (GEO) Changes Everything: Everyone knows they need to perform well on LLM searches, but many don't know how to do it. Everything changes with LLM search. There are **new metrics to track**, **new strategies** to implement and **new content to create**.

3. Understanding What Your Customers Are Asking: Customers are asking LLMs **much more sophisticated questions** than they ever asked Google Search. They are getting their basic questions answered quickly and then they are asking deeper questions. It's important that you **understand both what the initial and deeper questions** customers are asking.

4. New Content Strategies: Once you understand the new questions your customers are asking, your content strategy has to change. **It's no longer sufficient to write a blog saying your product is a good solution** for a particular use case. Now, for example, you need to explain **why your solution is a better solution** to that use case than the other alternatives the customer is considering.

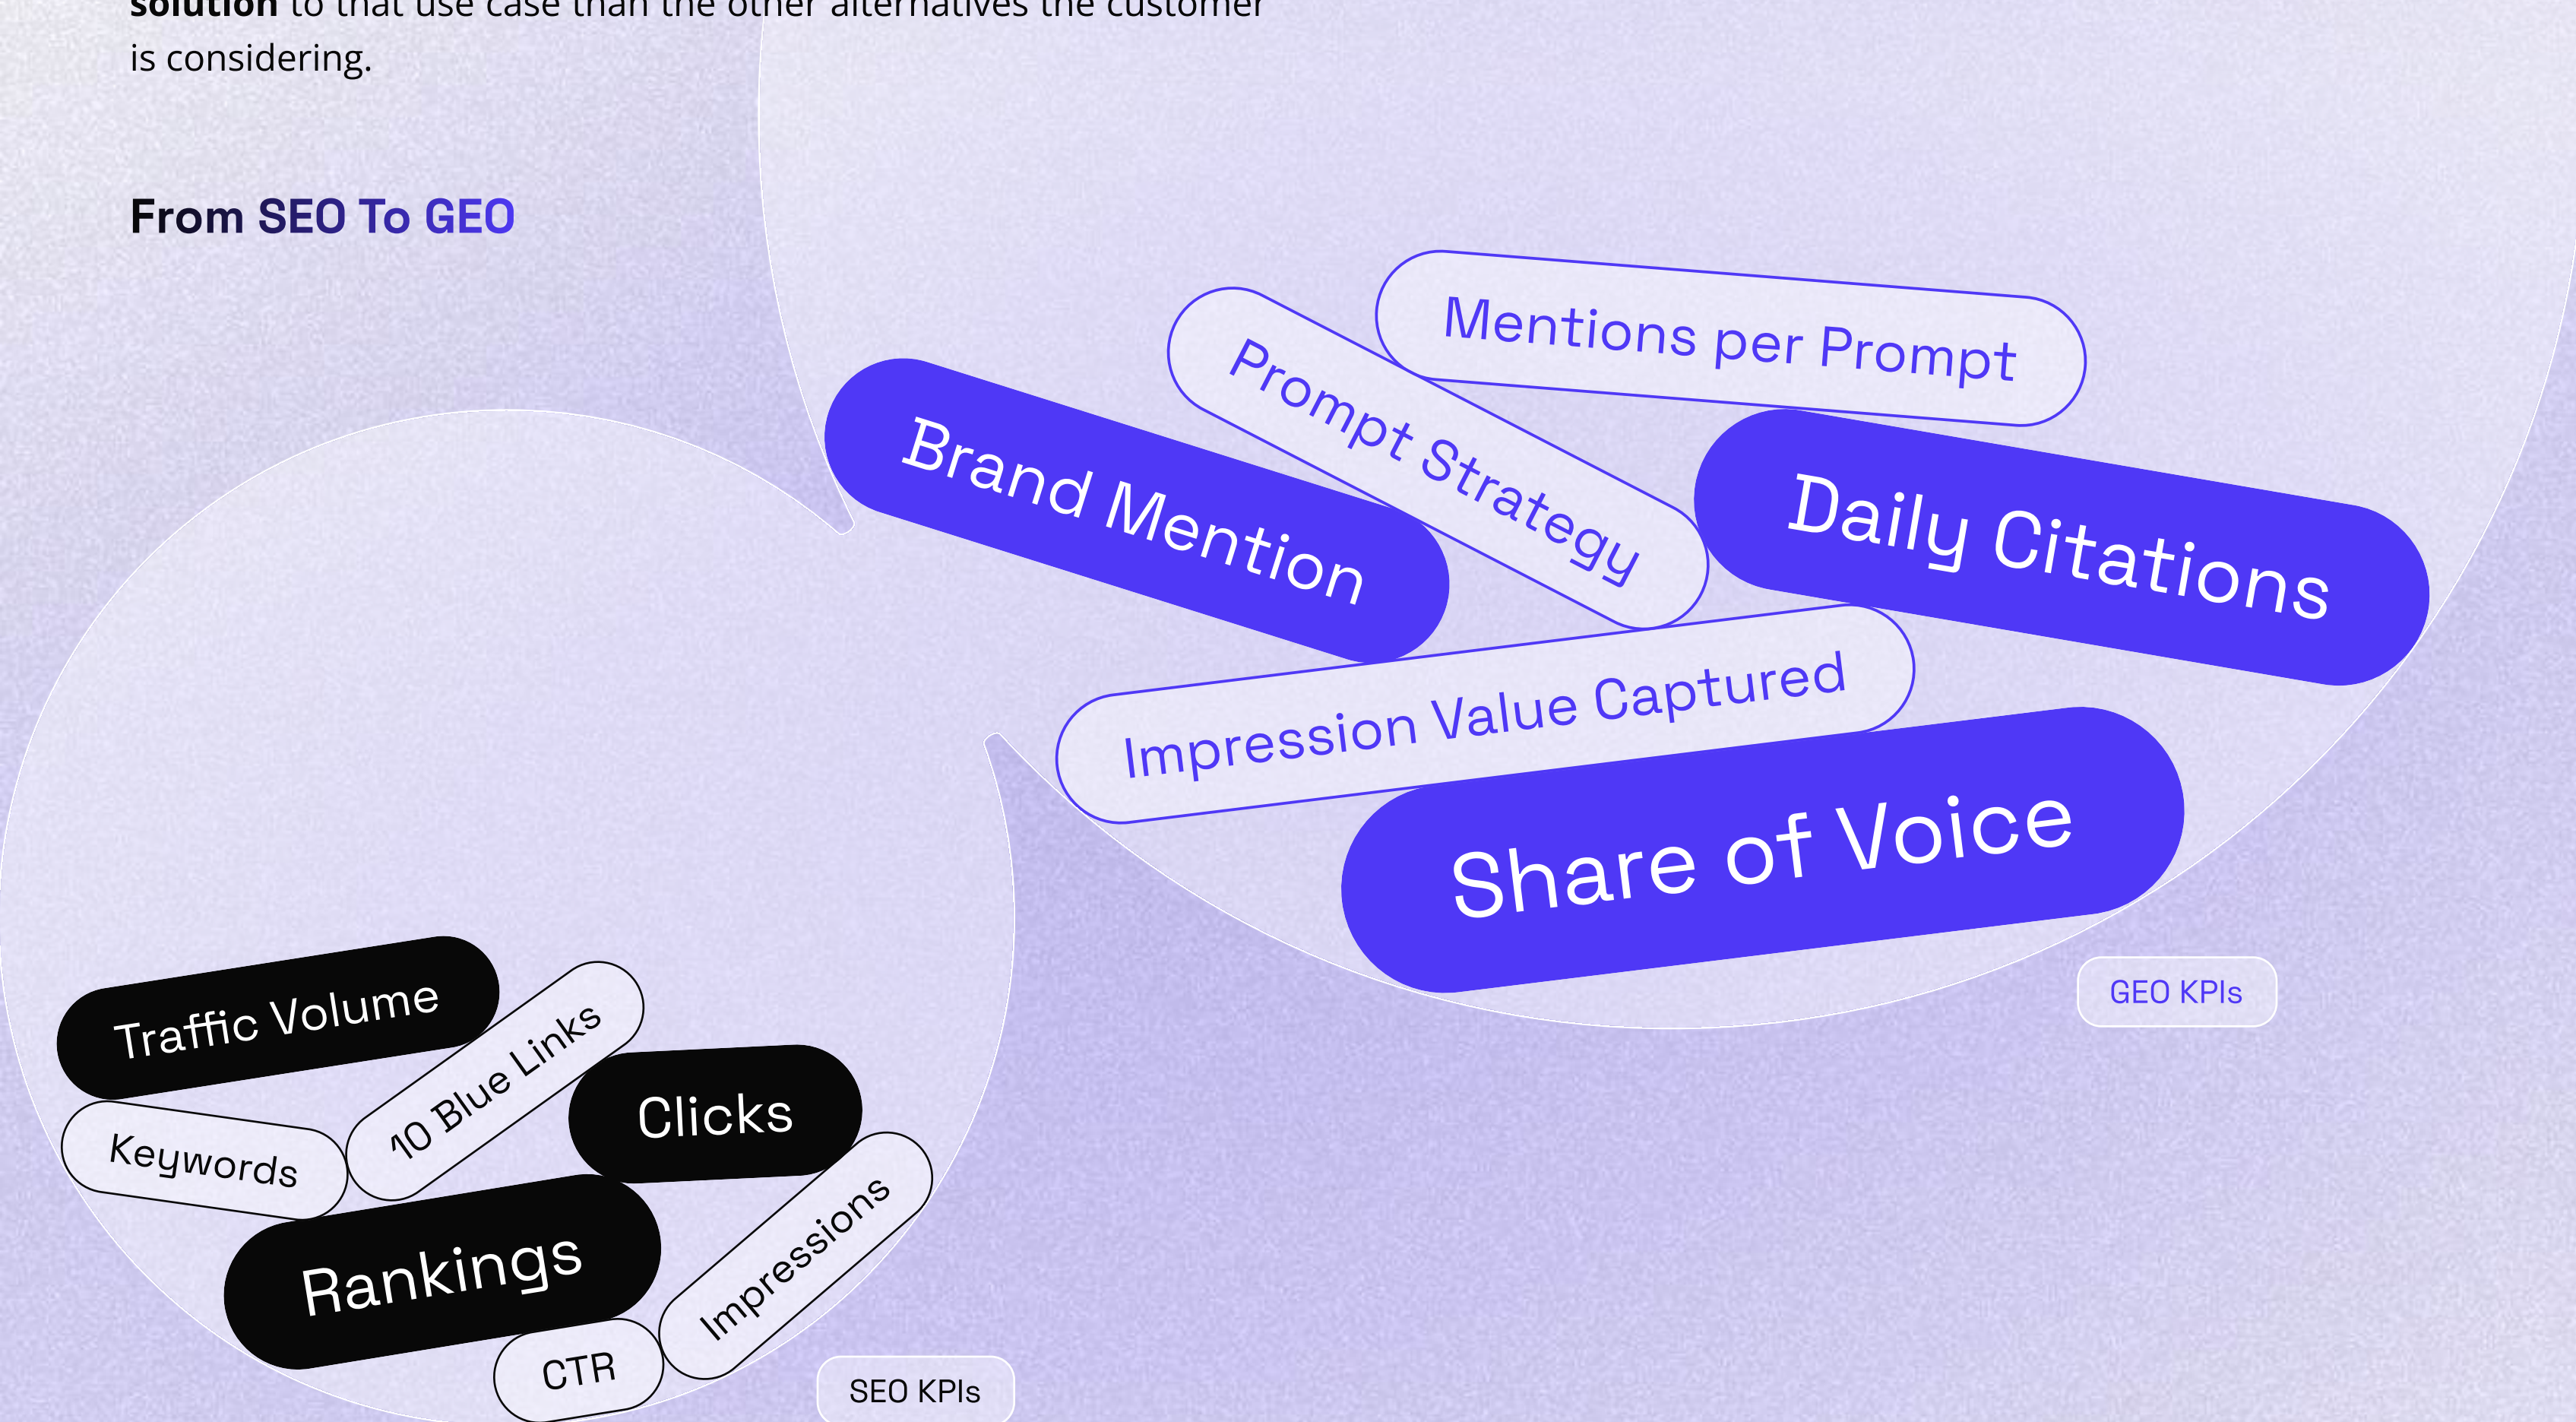
5. Writing for LLMs is Different from Writing for People: While LLMs will read content that you write for people, if you **write with LLMs in mind**, you're much more likely to rank. Having the right partner to help you write well for LLMs can have a huge impact.

6. Traffic Decline Requires Action: With sharp drops in organic traffic, marketing leaders must prepare for a future with fewer visits from search engines. **Early adopters** of generative AI **already secure a significant share of voice** within months, gaining a competitive edge.

7. From SEO to GEO: Enter Generative Engine Optimization (GEO). Unlike SEO, which focuses on page rankings, GEO ensures that brand data and expertise are integrated into LLM-generated answers. **Being discovered no longer depends on clicks, the goal to become a trusted source** cited by AI systems.

8. New Visibility Metrics: Traditional KPIs (rankings, clicks) lose relevance. New measures emerge, such as Share of Voice (SoV) - the percentage of brand mentions in answers versus competitors plus Mentions per Prompt, Daily Citations (how often AI references your content), and ROI per Prompt. **These metrics reveal who dominates the context of answers, not just clicks.**

From SEO To GEO



Performance Benchmarks (All Segments)

Before you start on your GEO journey, it's valuable to have **some context on how AI search works**, where it gets data from and **what "good" looks like**. Below are a few statistics to help with that.

Brand Presence in Answers:

On average, 16.9% of AI responses mention a given brand in discovery prompts. However, **leading companies reach up to 59.4%**; a huge gap in brand visibility. This shows that some players manage to insert themselves in responses **almost 3x more than the market average**. A similar pattern is seen in Share of Voice, where leading brands average 36%, while the 2nd and 3rd places reach 19.7% and 13.2%, respectively.

Average Brand Mentions

Average Brand Mentions Per Company

Top Brand
59.42%

Average
16.86%

TOP BRAND HAS
ON AVERAGE 3X
SHARE OF VOICE

Average Top Share Of Voice

Average SoV of Leading Companies

#1 Rank
36.55%

#2 Rank
19.70%

#3 Rank
13.21%

Domain Citations:

Als directly cite a brand's domain in only ~9.2% of responses (overall average). This means **most responses do not include the brand's own domain**, highlighting an opportunity to build greater authority. Additionally, each AI model references between 2.5 and 13.7 distinct domains per response-some engines rely on fewer sources, while others aggregate content from multiple sites, varying the diversity of references.

Average Citations

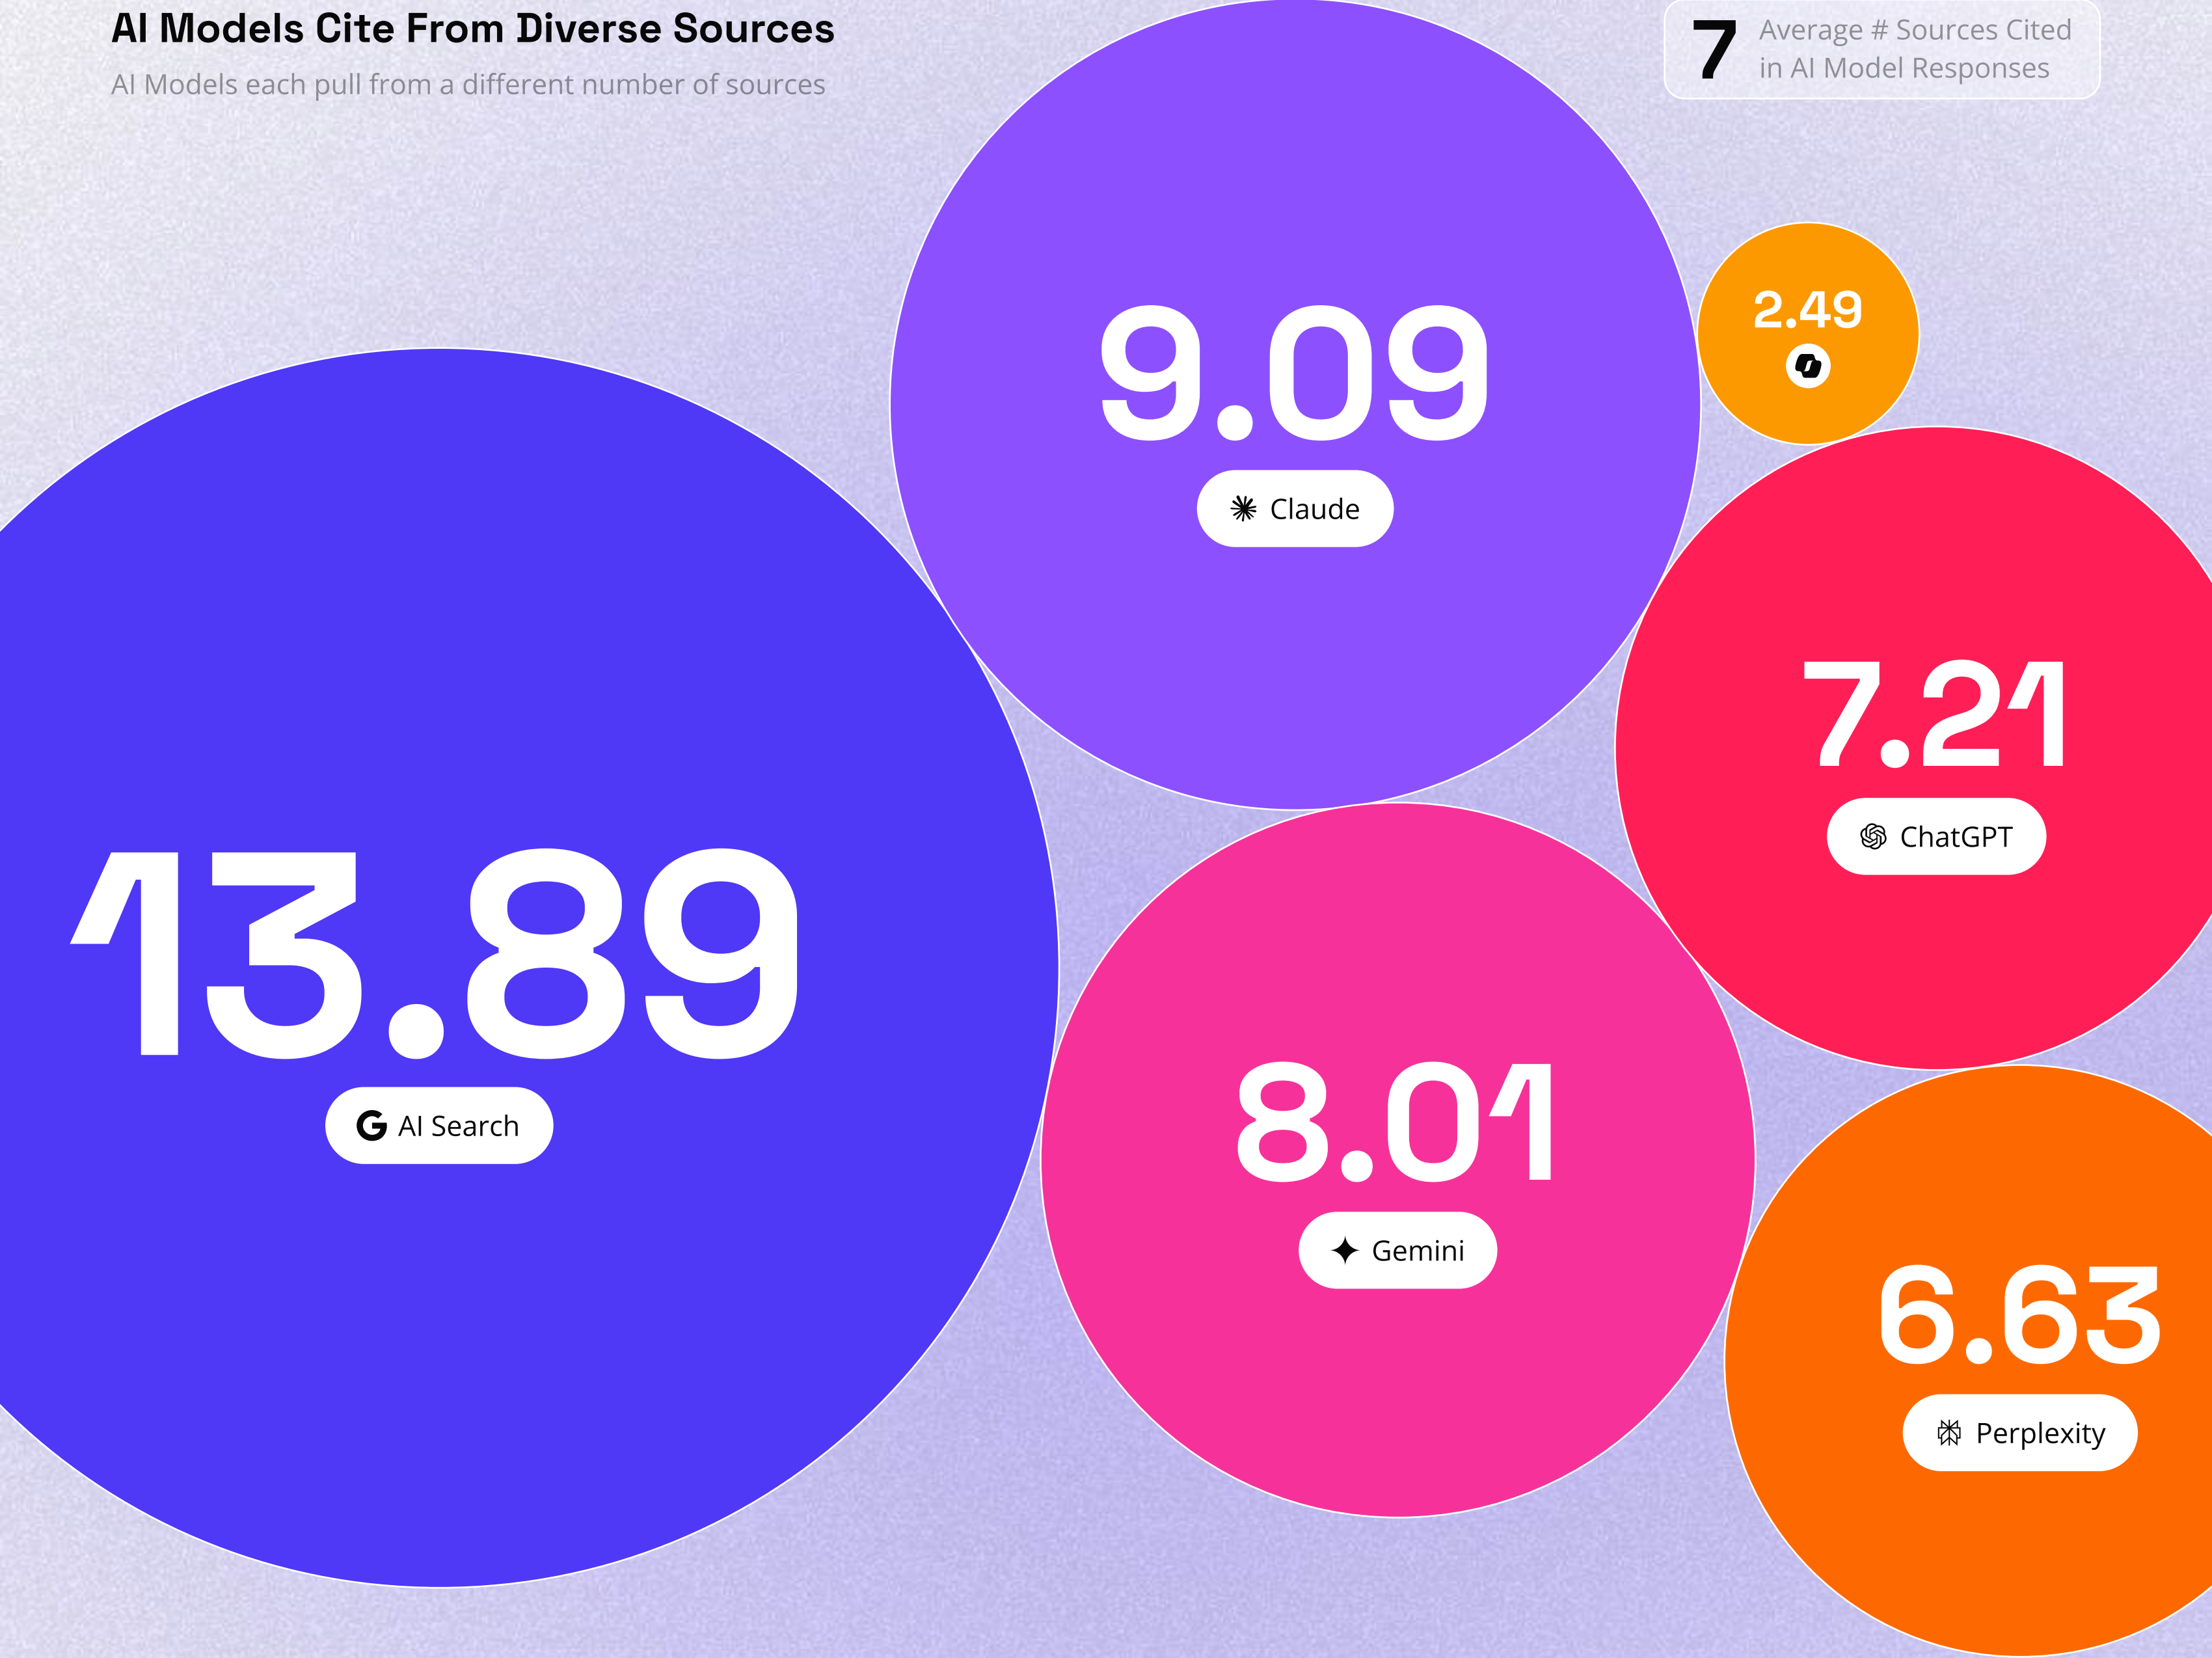
Percentage of Brands Domains Cited by Als as Sources

9.21%

AI Models Cite From Diverse Sources

AI Models each pull from a different number of sources

7 Average # Sources Cited in AI Model Responses



BRANDS MUST OPTIMIZE
ACROSS DIVERSE SOURCES
NOT JUST A
SINGLE CHANNEL

This diversity impacts brand visibility:
more sources = higher chance to be cited

LLM Favorite Sources:

LLMs, when responding to prompts, **do not rely on a fixed dataset but actively reference a mix of sources**. For a single prompt repeated daily over six months, the models may draw from hundreds to more than a thousand unique sources.

Some domains, such as **Reddit, appear disproportionately often, representing 10 to 40 percent of responses** and covering more than a hundred unique Reddit URLs. This shows that LLMs do not depend on a static reference set but continually expand and refresh their source pool. **In certain verticals, niche domains dominate**. For example, in finance, Investopedia and NerdWallet surface far more frequently than general sources. Smaller sites that appear in only one percent of responses can still play an important role.

Viewed collectively, these consistently surfaced sites may account for as much as 30 percent of all sources referenced over time. This makes them **valuable as strategic partners**.

When aggregating results across many prompts and models, it becomes clear which sources sustain visibility over longer periods. Increasing share of voice is not about maintaining a Wikipedia page or filling Reddit with mentions, since reliance on those sources is continuously decreasing. **What makes a difference is creating material that aligns semantically with the queries models receive**. Even when negative comments appear on Reddit, publishing a fresh article on your site or ensuring coverage on a smaller third-party domain that highlights improvements can still influence how LLMs shape their answers.

Top 10 Most Cited Sources (% Of Total)



DOMINATE
THE SOURCES,
OWN THE
ANSWERS



GET VERTICALS
ON SPECIFIC
MARKETS HERE!

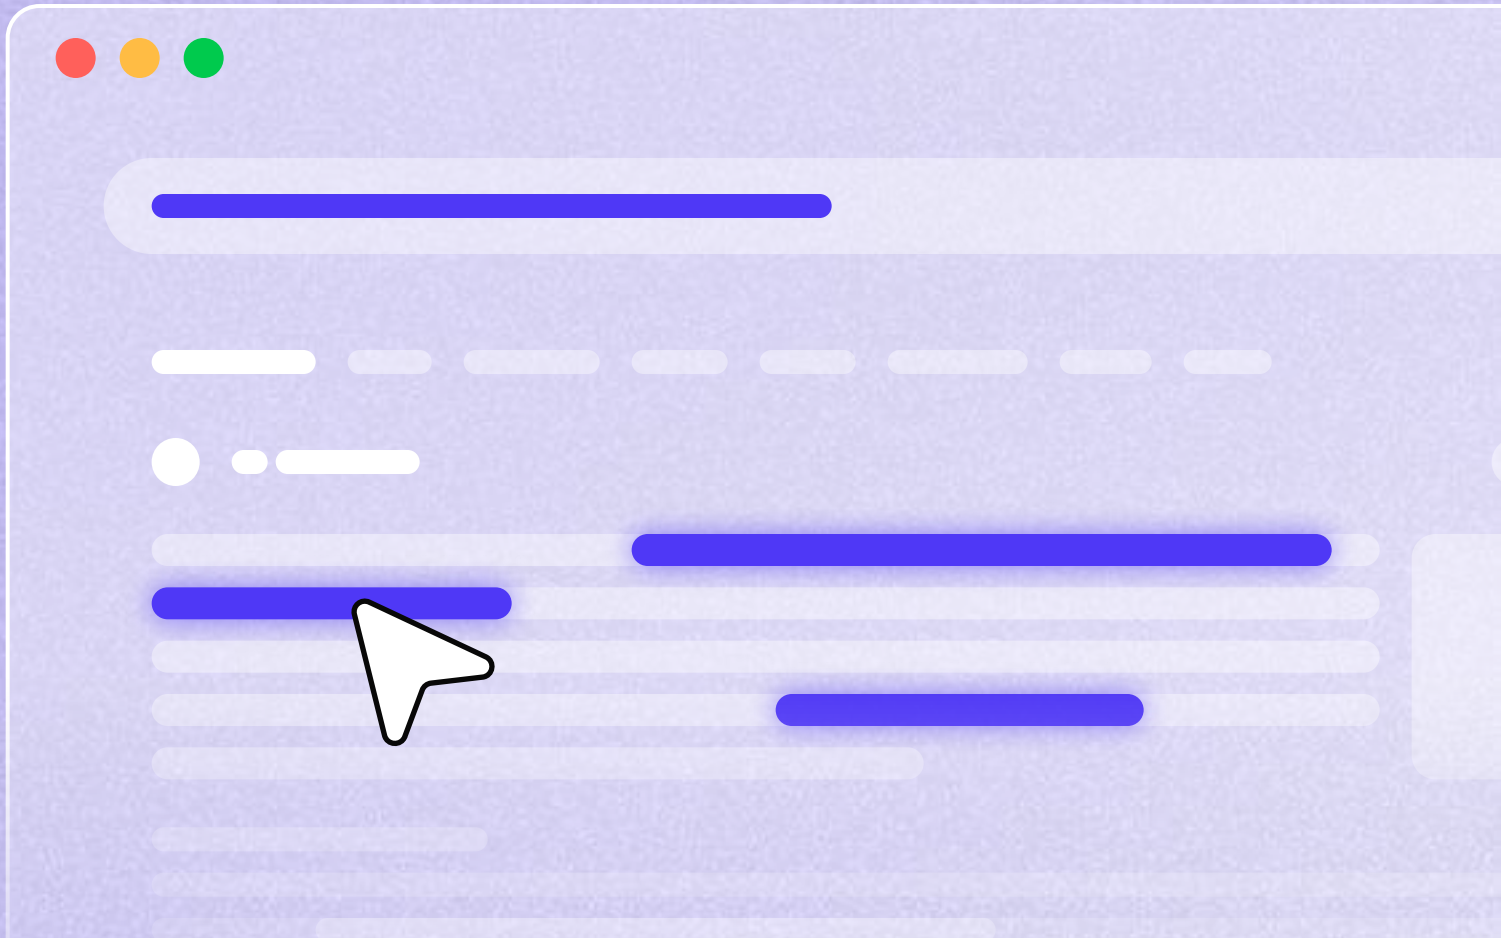
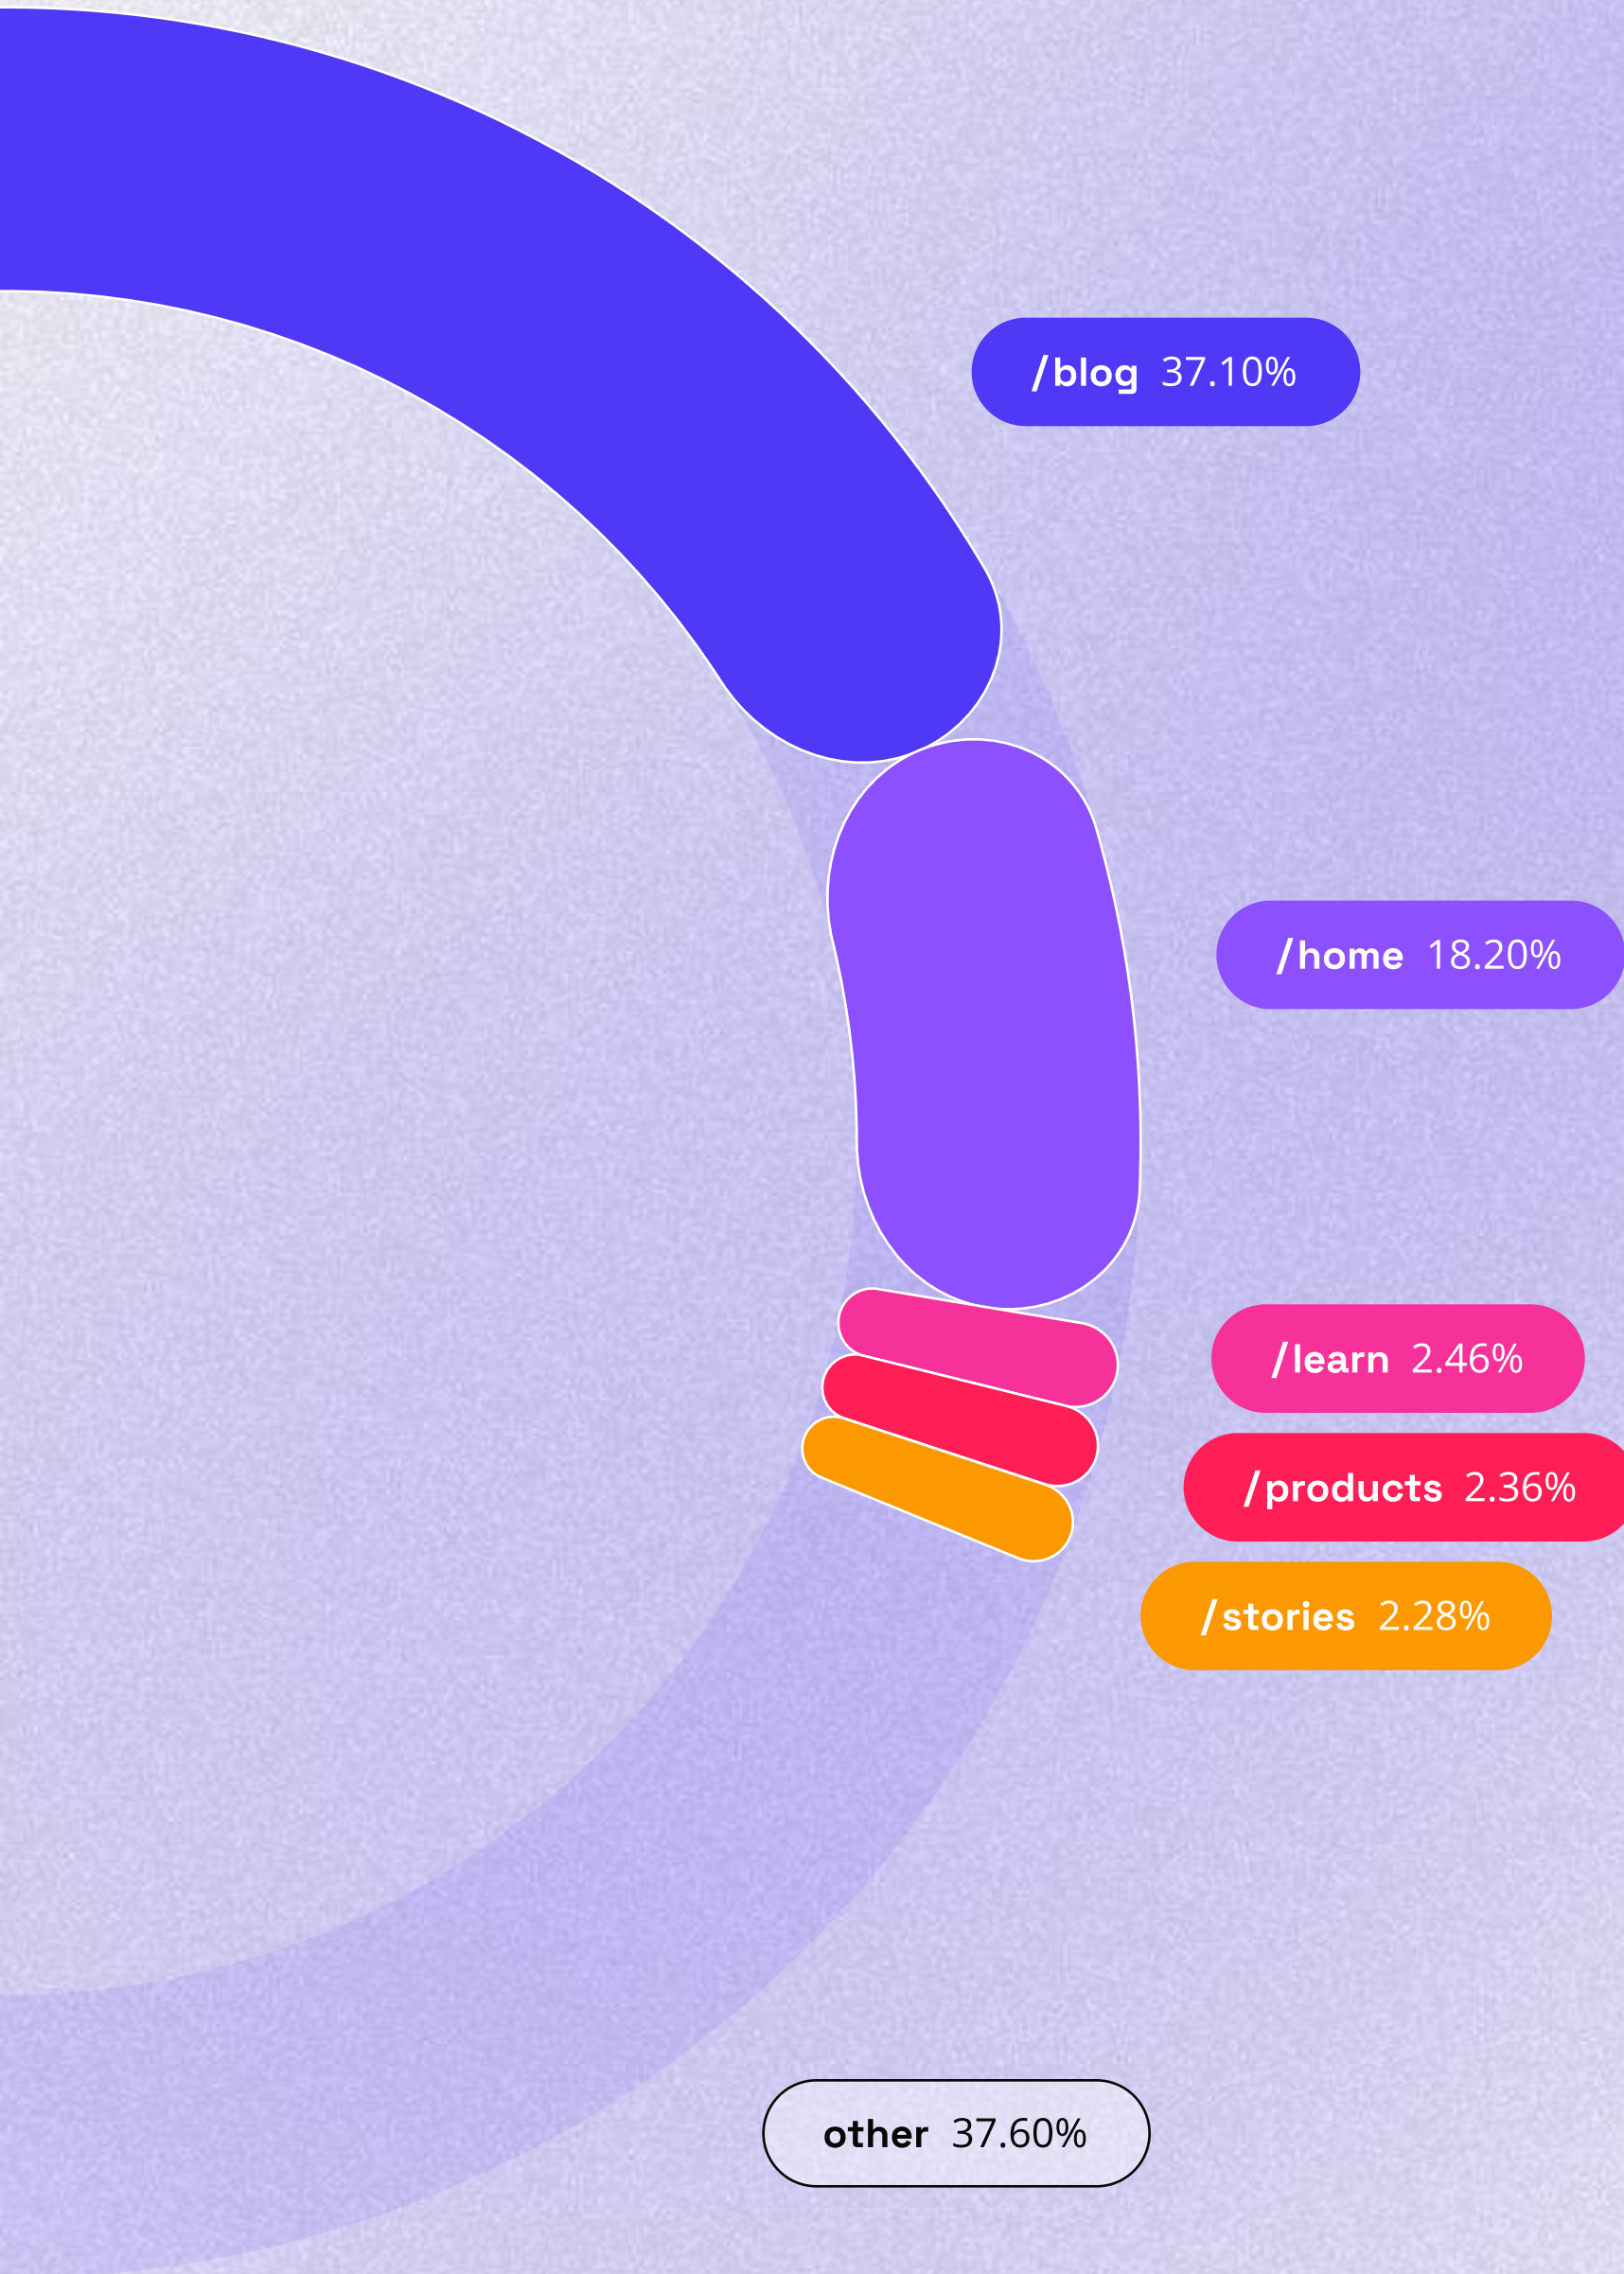
athenahq.ai/reports/state-of-ai-report-2025

Top Entry Points (Sites)

Within brand websites, **Als most often access open content sections first** such as /blog (37.1%) and /home (18.2%) dominate as the most common entry points for AI exploration, reinforcing how content hubs and homepage structures often shape the first layer of visibility. How AI models weight answers varies by segment, but a good rule of thumb is: 50% on-page factors (matching titles, slugs, and timely updates) and 50% sustained off-page reinforcement. On-page elements, especially those that directly match user prompts, invite initial visibility, but **long-term search presence depends equally on authority signals**: third-party PR, listicles, reviews, and active social proof. AI search optimization is not a siloed SEO task; it's a cross-functional 'team sport' where PR, content, product, and social efforts must align. **Those who break down silos and coordinate messaging consistently earn lasting visibility, while isolated teams lose traction regardless of budget**

Most Commonly Cited Paths

The 5 most common ways AI searches for information on your website



EVERY
CLICK
STARTS
WITH
CONTENT

Strategic Recommendations (GEO - Generative Engine Optimization)

1. Open, Structured, and Trustworthy Content: Publish information in an accessible, organized way for AIs. Centralize knowledge in open sections (e.g., /blog and /home pages, without content hidden behind forms, paywall or JavaScript-heavy) and provide **clear context so models can interpret and trust the data**. Reliable, up-to-date content supported by credible sources increases the likelihood of being cited by AIs. Avoid shiny objects like creating separate versions of your page for AI and for humans. These are short-term, low-impact efforts. **A good website for humans is what AI is evaluating.**

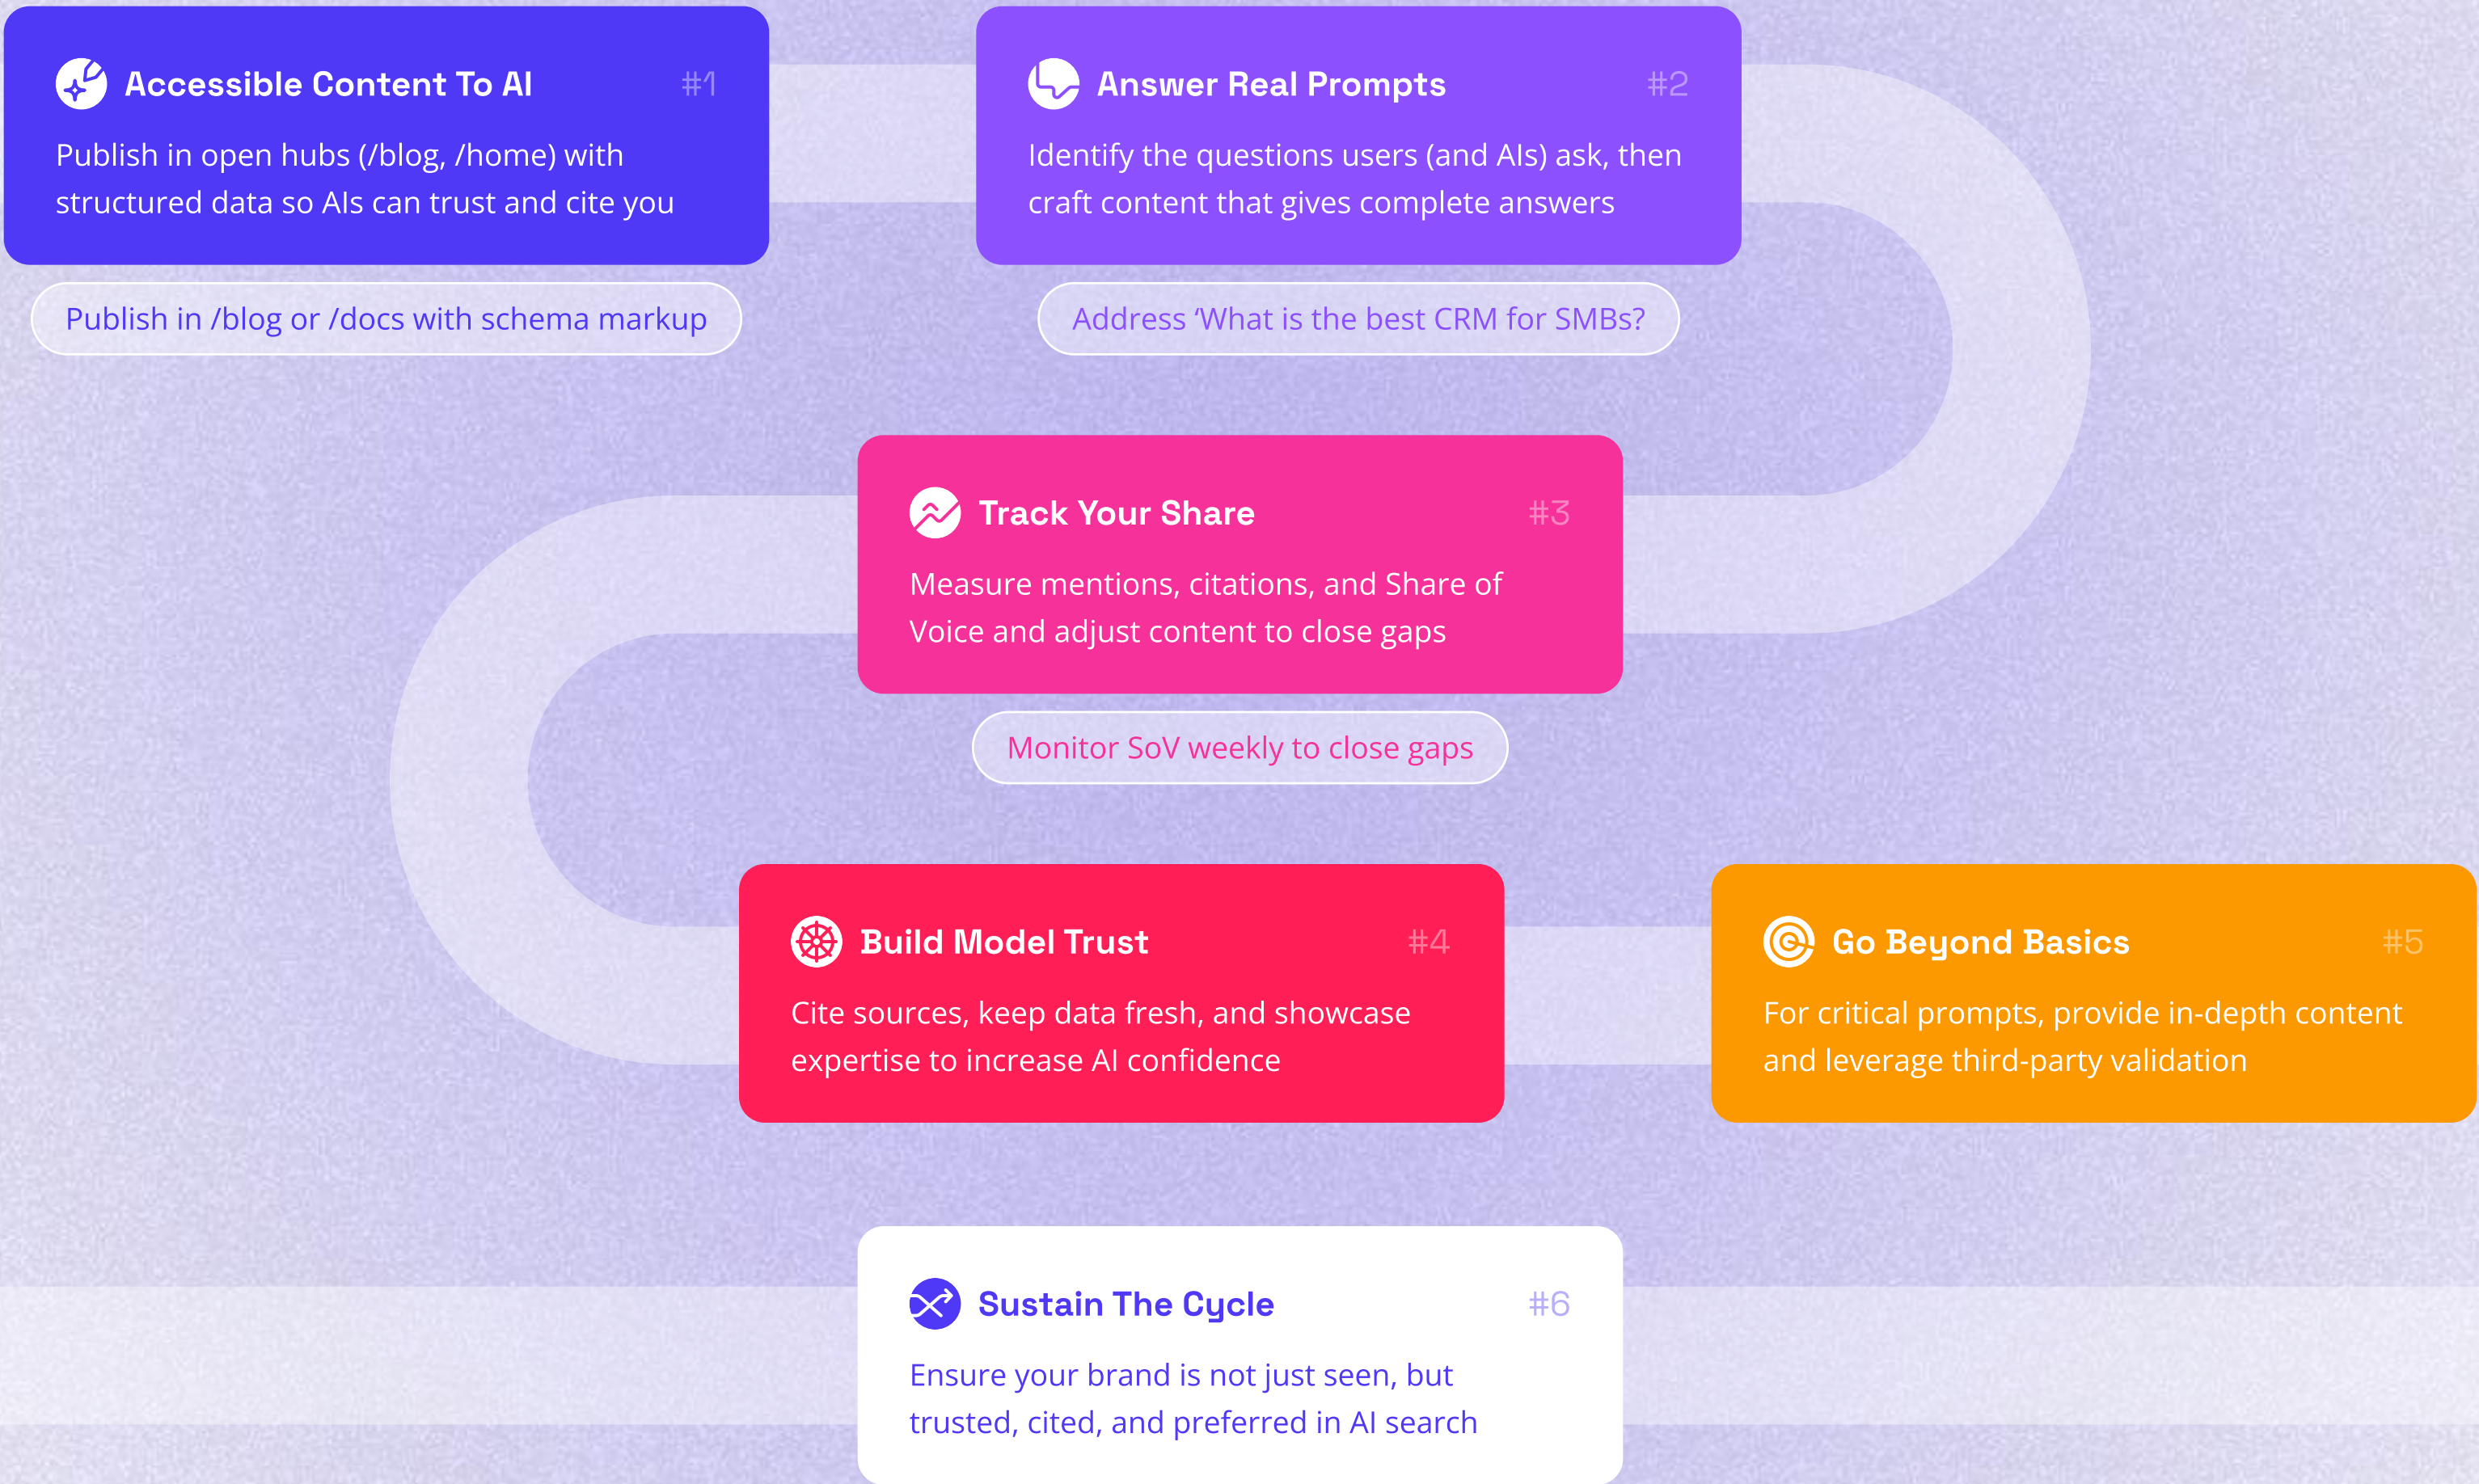
2. GEO and Prompt Strategy: Identify the key questions (prompts) your audience or industry asks chatbots and generative searches. Then, **create in-depth content that directly answers those questions**. Think beyond keywords: address specific topics and doubts clearly and comprehensively, ensuring the AI has **high-quality material from your brand** to include in answers.

3. Continuous Monitoring and Optimization: Actively track **when and how your brand appears in AI answers**. A change in brand mentions or share of voice over a week does not provide the full picture. **Models often highlight new content from competitors**, but this advantage will fade if they cannot back their claims with **coverage from highly reputable third-party sources** like Forbes.

If customer reviews fail to support the competitor’s positioning over time, that competitor will lose influence and credibility. **Measure your Share of Voice** in key prompts and **monitor daily citations** of your content. Use these new metrics to refine your strategy: for example, if a topic doesn’t mention your brand, create content to fill that gap. **Regularly adjust both content and format to strengthen your brand’s presence** in generated responses.

4. Authority Signals: Reinforce elements that build model confidence in your content. Cite respected sources, include updated data and references, and highlight your company's expertise. Building authority is part of GEO: **the more your site is perceived as a reliable source** (through a history of useful, well-referenced content), **the higher the likelihood of being cited in AI responses**. In short, become part of the knowledge base that generative engines rely on, positioning your brand as a preferred source in your category.

5. Going Deep: For your **most valuable prompts, you'll want to go deep**, both on your content and with third parties. On the most important topics, you'll want as many third parties as possible reinforcing your messaging and differentiation.



AI Search Landscape

The End Of An Era

50%+ Traffic Declines Signal a New Digital Landscape. The transformation of traditional search engines into generative answer engines marks a **fundamental shift in how information is accessed online**. Platforms like Google are evolving beyond the familiar list of blue links, using features such as AI Overviews and AI Mode to deliver synthesized answers directly to users.

The effects are already measurable:



This shift challenges the long-standing dynamics of the link-based web ecosystem. **As users increasingly receive complete answers** on the search results page, **the need to click through to external websites diminishes**, directly impacting impressions and referral traffic. Traditional SEO frameworks are no longer sufficient. The way content is evaluated, surfaced, and consumed is being redefined by generative AI models.

The Paradigm Shift

In response to this evolving environment, organizations must transition from Search Engine Optimization (SEO) to a new strategic approach: Generative Engine Optimization (GEO) - also known as Answer Engine Optimization. While SEO is designed to improve the ranking of individual URLs, **GEO focuses on ensuring that a brand's data, expertise, and insights become part of the foundational knowledge that generative engines rely on to construct their answers**.

The goal is to **become a trusted source of truth** - not just indexed, but actively cited and integrated into the responses delivered by AI systems.

Why This Matters Now: The Opportunity Ahead

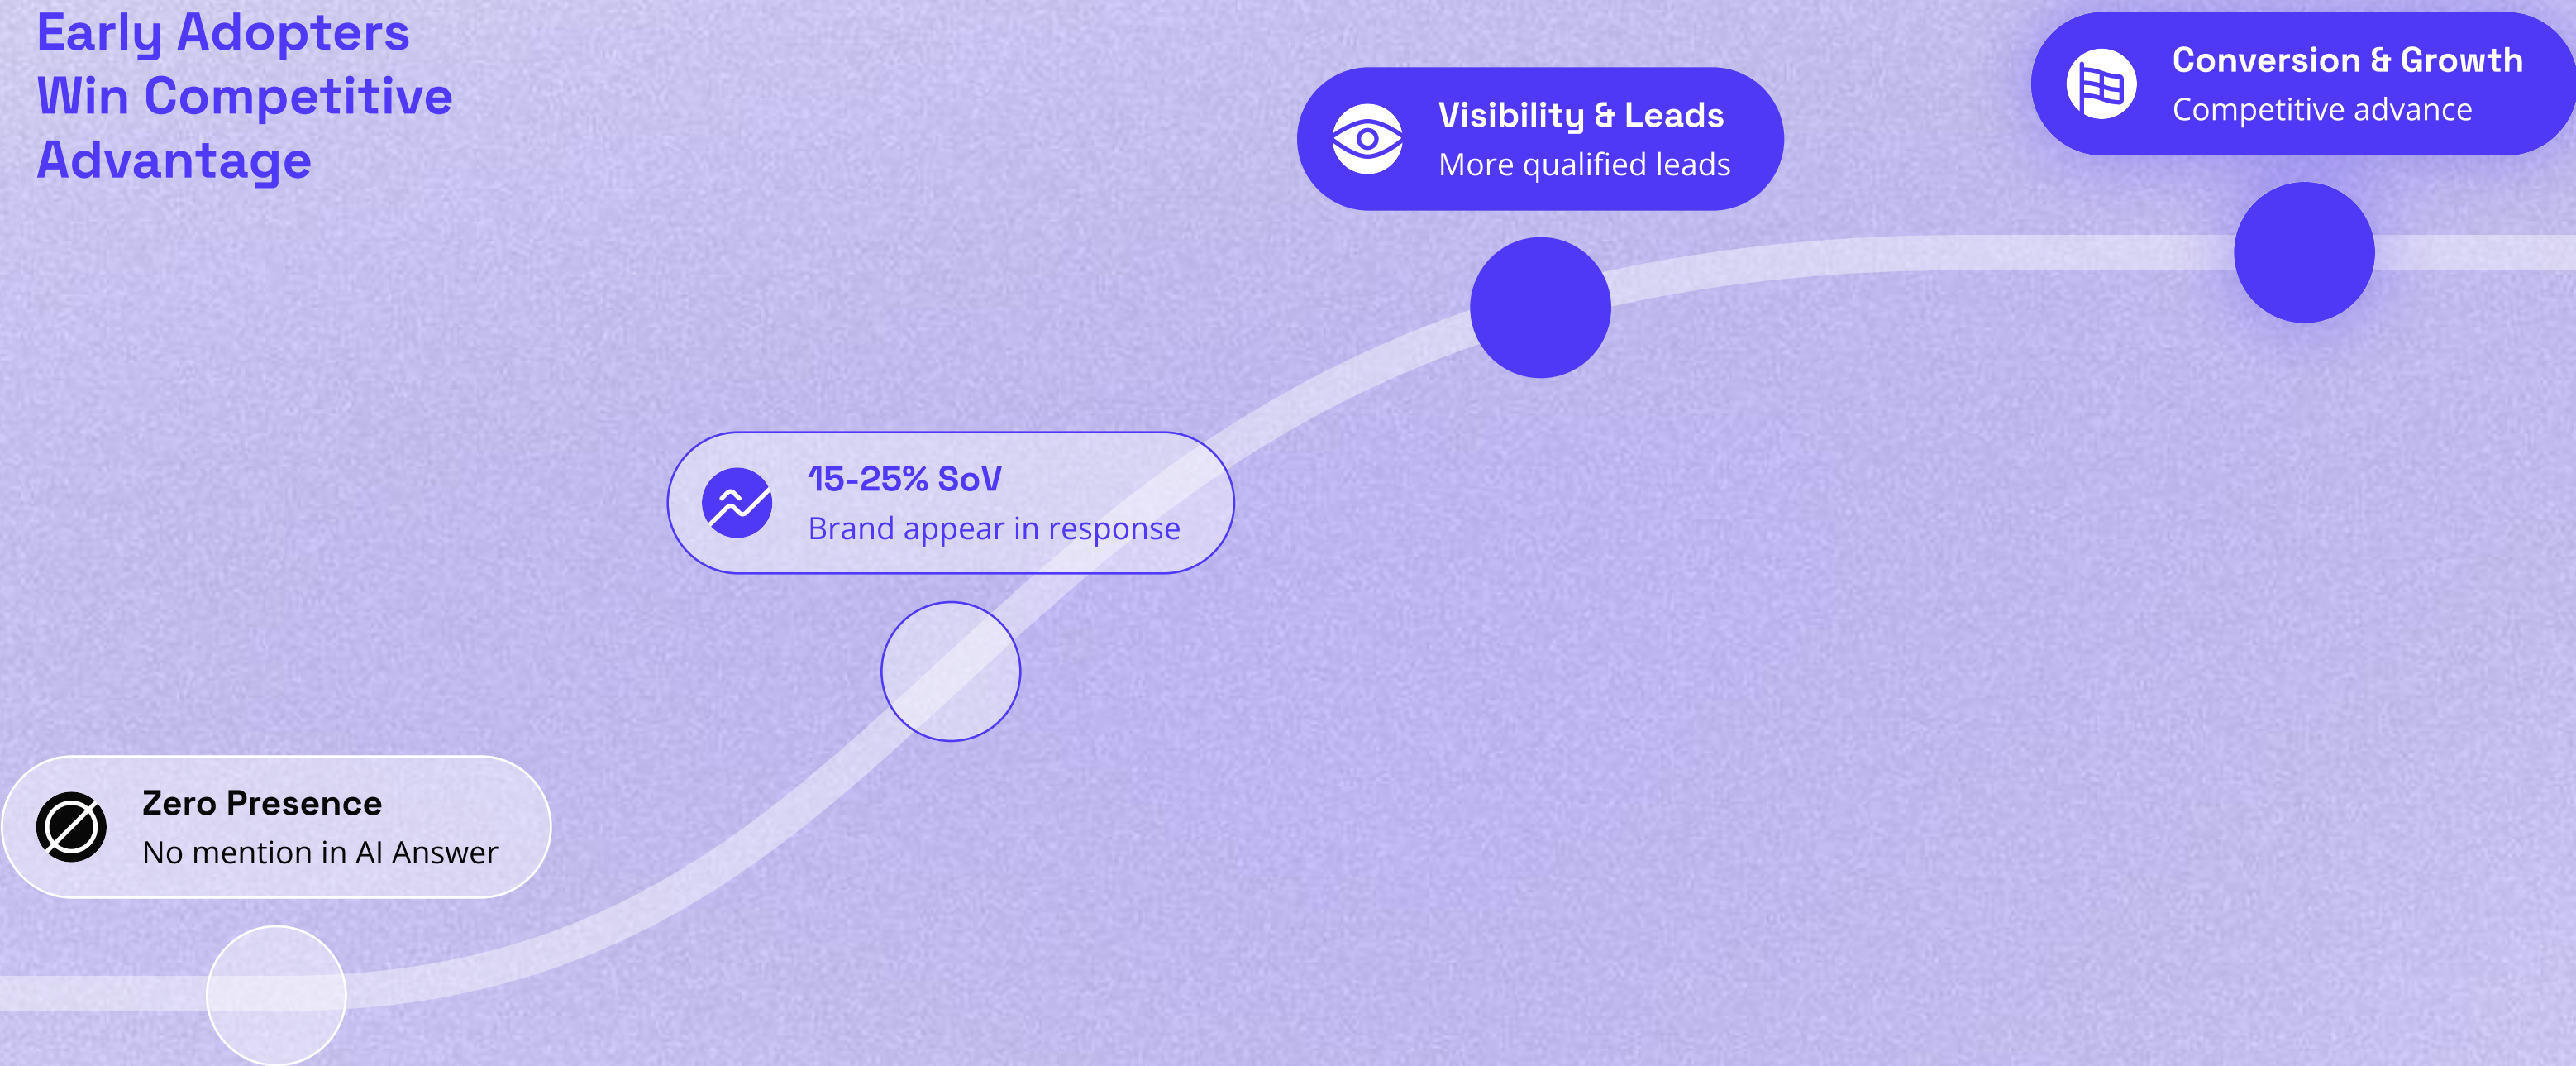
We are at the beginning of a new era. As search behavior shifts toward direct answers instead of link exploration, digital visibility is being redefined. In this new landscape, being discovered doesn't depend on being clicked - it depends on being part of the answer.

Brands that are proactively positioning themselves within generative ecosystems are already seeing exceptional results. In just a few months, it's common to see companies **go from zero presence** in AI-generated responses **to capturing 15-20-25% share of voice** in their category. This shift translates directly into greater visibility, more qualified leads, increased conversions, and ultimately - growth. The message is clear: Those who act early are not only adapting, they are gaining competitive advantage.

State of AI2025 serves as a clear signal: **organizations must reconsider how they structure, publish, and position their knowledge** to ensure relevance in a search experience driven by generative AI.

Being part of the result is no longer enough - you must become part of the response.

Early Adopters
Win Competitive
Advantage



BEING PART OF
THE RESULT IS
NO LONGER ENOUGH

YOU MUST
BECOME
PART OF
THE ANSWER

Athena Playbook: The GEO Operating Model

Why Traditional SEO Falls Short, And What To Measure Instead

As AI transforms search engines into answer engines, **traditional SEO KPIs lose their relevance**. Rankings, impressions, and click-through rates matter less in a world where users never see a link: they see a response. In this new environment, influence and authority within **AI-generated answers become the real drivers** of visibility and growth.



metric #1

SHARE OF VOICE

Percentage of brand mentions compared to competitors in response to key prompts.



metric #2

MENTIONS PER PROMPT

How consistently your brand appears across AI-generated queries.



metric #3

DAILY CITATIONS

The number of times AI engines cite or reference your content daily.



metric #4

AI- SOURCED LEADS

Leads or conversions triggered by discovery in AI answers.



metric #5

ROI PER PROMPT

The business impact generated from appearing in specific queries.



metric #6

AVERAGE BRAND MENTIONS

Your brand's presence in answers that include at least one competitor.

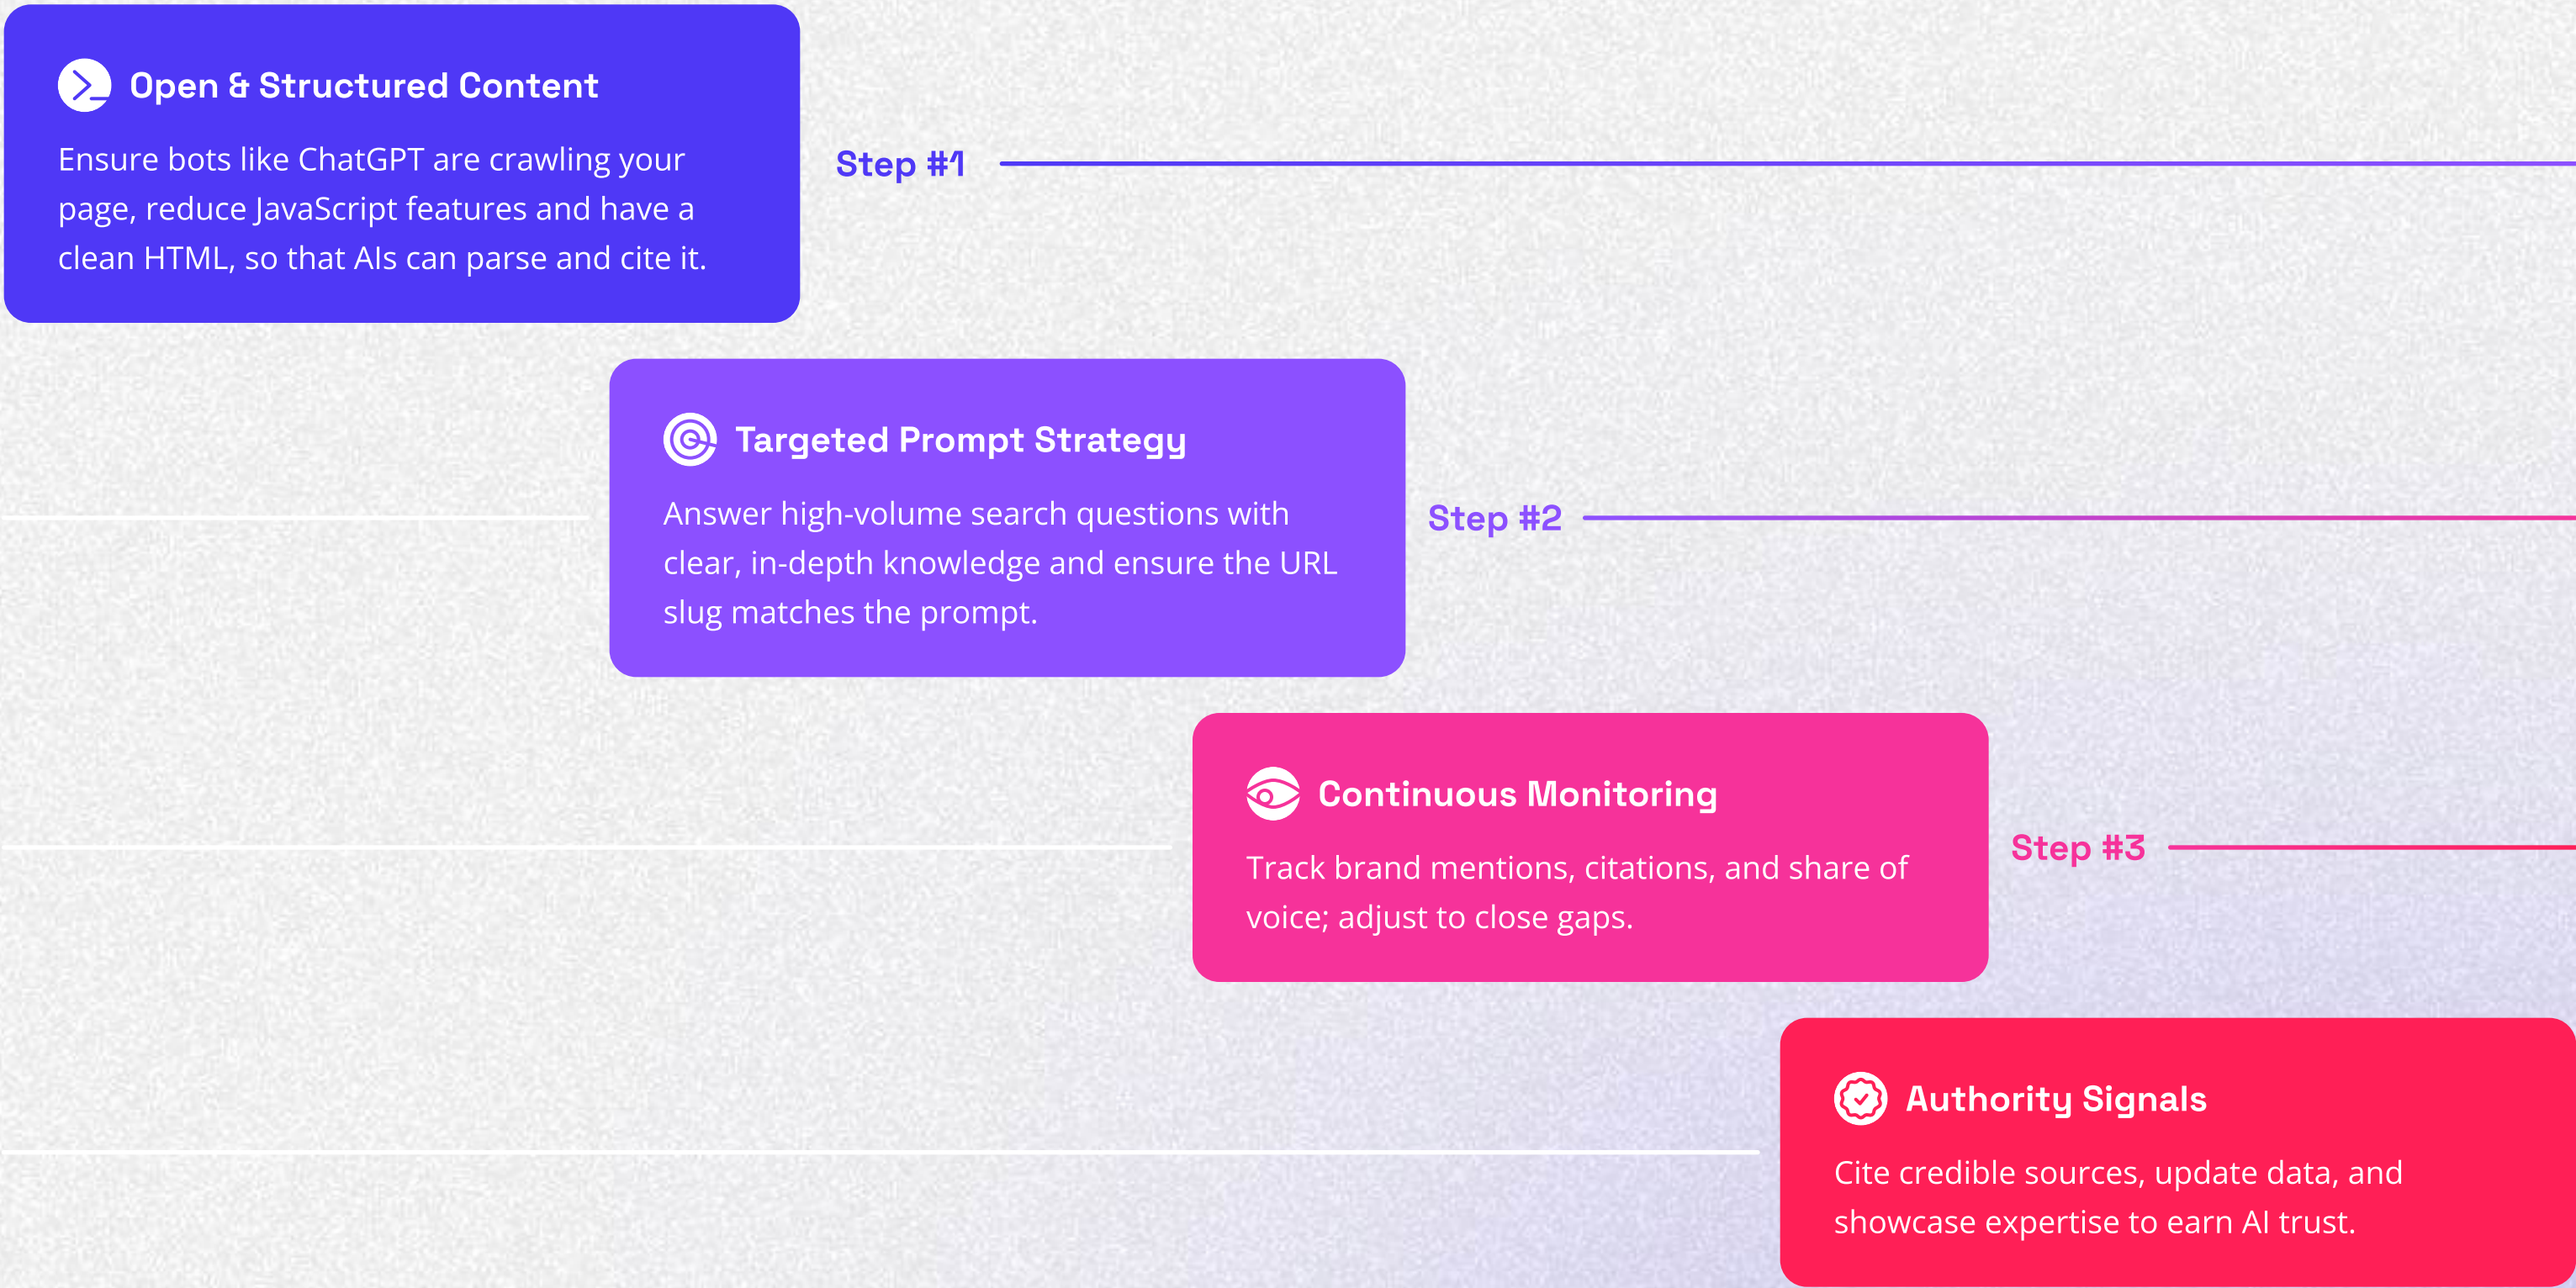
New metrics for AI-Native Discovery

These metrics provide a **true picture of visibility, authority, and conversion potential** in the generative search ecosystem: where the competition isn't for clicks, but for context.

The Strategy: Open Knowledge, Structured Content, And Targeted Action

At the core of a successful GEO strategy is a fundamental mindset shift: **treat generative engines as a primary audience, not just users**. That means making technical documentation, product details, and thought leadership publicly accessible and structured - rather than hidden behind forms or fragmented across platforms.

The 4 Steps to GEO Success



The goal is to become part of the knowledge base that AI engines draw from, ensuring that when they answer, your perspective and expertise are represented. When someone asks an AI: "What's the best solution for [specific problem]?" They are not looking for a browsing experience - they want a direct, trusted answer.

By openly sharing authoritative, structured, and context-rich content, you enable AI systems to:

- 1. Compare Alternatives**
Show customers how your offer stacks up against competitors so AI models can position you clearly.
- 2. Be Included In Answers**
Ensure your perspective is cited in AI-generated responses to key prompts.
- 3. Educate Before Clicks**
Reach potential customers with the information they need before they even visit your site.

This shortens decision-making and raises intent.
Leading to more qualified contacts, proposals, and conversions.

Bringing It All Together

Generative search is redefining how visibility, authority, and trust are earned online. Understanding the shift from SEO to GEO, the new success metrics, and the strategic steps to position content effectively is no longer optional, it is the foundation for digital relevance in the AI era.

And it is in this context that State of AI 2025 emerges. Drawing on insights from over 6 million AI-generated responses, **this report provides a clear, data-backed view** of how generative engines **shape brand visibility**, influence **customer decisions**, and create **new opportunities for growth**.

The goal is simple: To give you a practical, evidence-based resource that helps you navigate - and thrive in - this new search landscape.

How To Read & Methodology

Methodology

The insights in this report are based on data collected and analyzed by AthenaHQ, encompassing **6 million responses across six AI models over six months**.

After thorough evaluation, we realized that **the distribution follows a power law, and increasing the number of responses does not influence the end result**, so we set a cap of 1 million responses per model. However, we do see an influence over time. For example, some sources like Reddit are losing their citation relevance for certain prompts and topics because AI models are improving their search algorithms.

Benchmark Metrics

Serve as reference. They **represent the average behavior of each metric**, allowing you to assess **how your brand compares against the broader market**.

Example: if your brand mentions are above the benchmark, it indicates stronger visibility than the market average.

Actionable Metrics

Highlight specific situations and opportunities that can be turned into concrete actions.

They are designed to inspire strategies that increase your brand's presence within AI responses.

6 MILLION AI RESPONSES

Collected and analyzed **across 6 leading AI models**

#1 Data Collection

#2 Processing & Modeling

#3 Refinement

#4 Insight

How to Read this Report

This report was structured to give you maximum flexibility in navigation:

1. Explore Verticals

Each industry vertical is analyzed on its own dedicated page for clarity.

2. Choose Your Path

You can start directly with your sector of interest or review all sectors for a comparative view.

3. Use Key Metrics

Every section highlights core metrics (brand mentions, SoV, and citations) to guide interpretation.

1. Brand Mentions

The **percentage of times your brand is mentioned** in responses where at least one brand (yours or a competitor's) is referenced.

2. Share of Voice

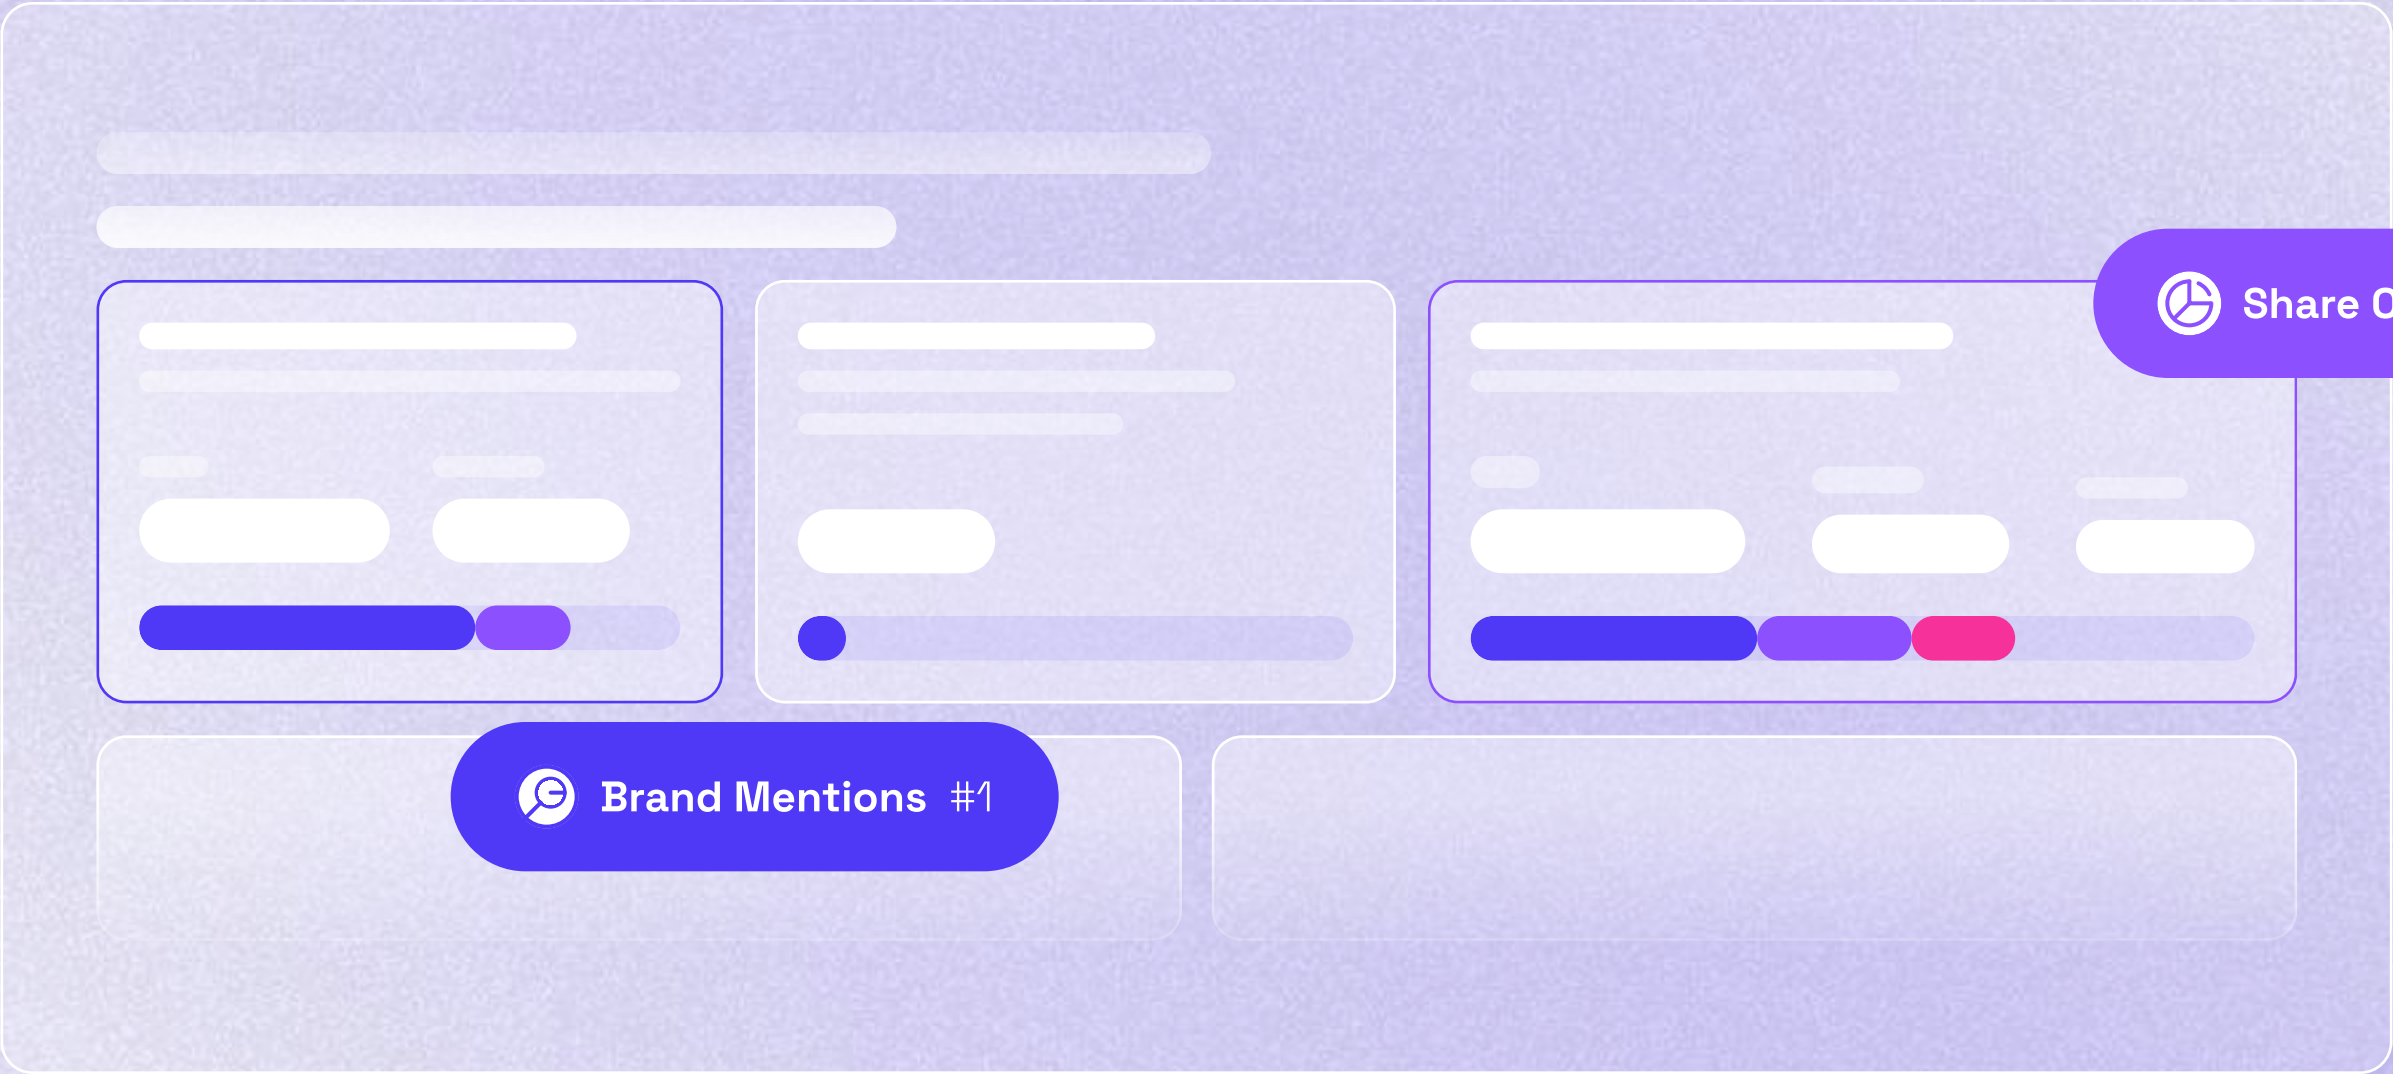
The **percentage of mentions your brand receives compared to** the total number of **all brand mentions** across tracked prompts. This metric helps you understand your relative visibility versus competitors in AI responses.

3. Citations

The **percentage of responses that cite your domain** as a source.

4. Sources

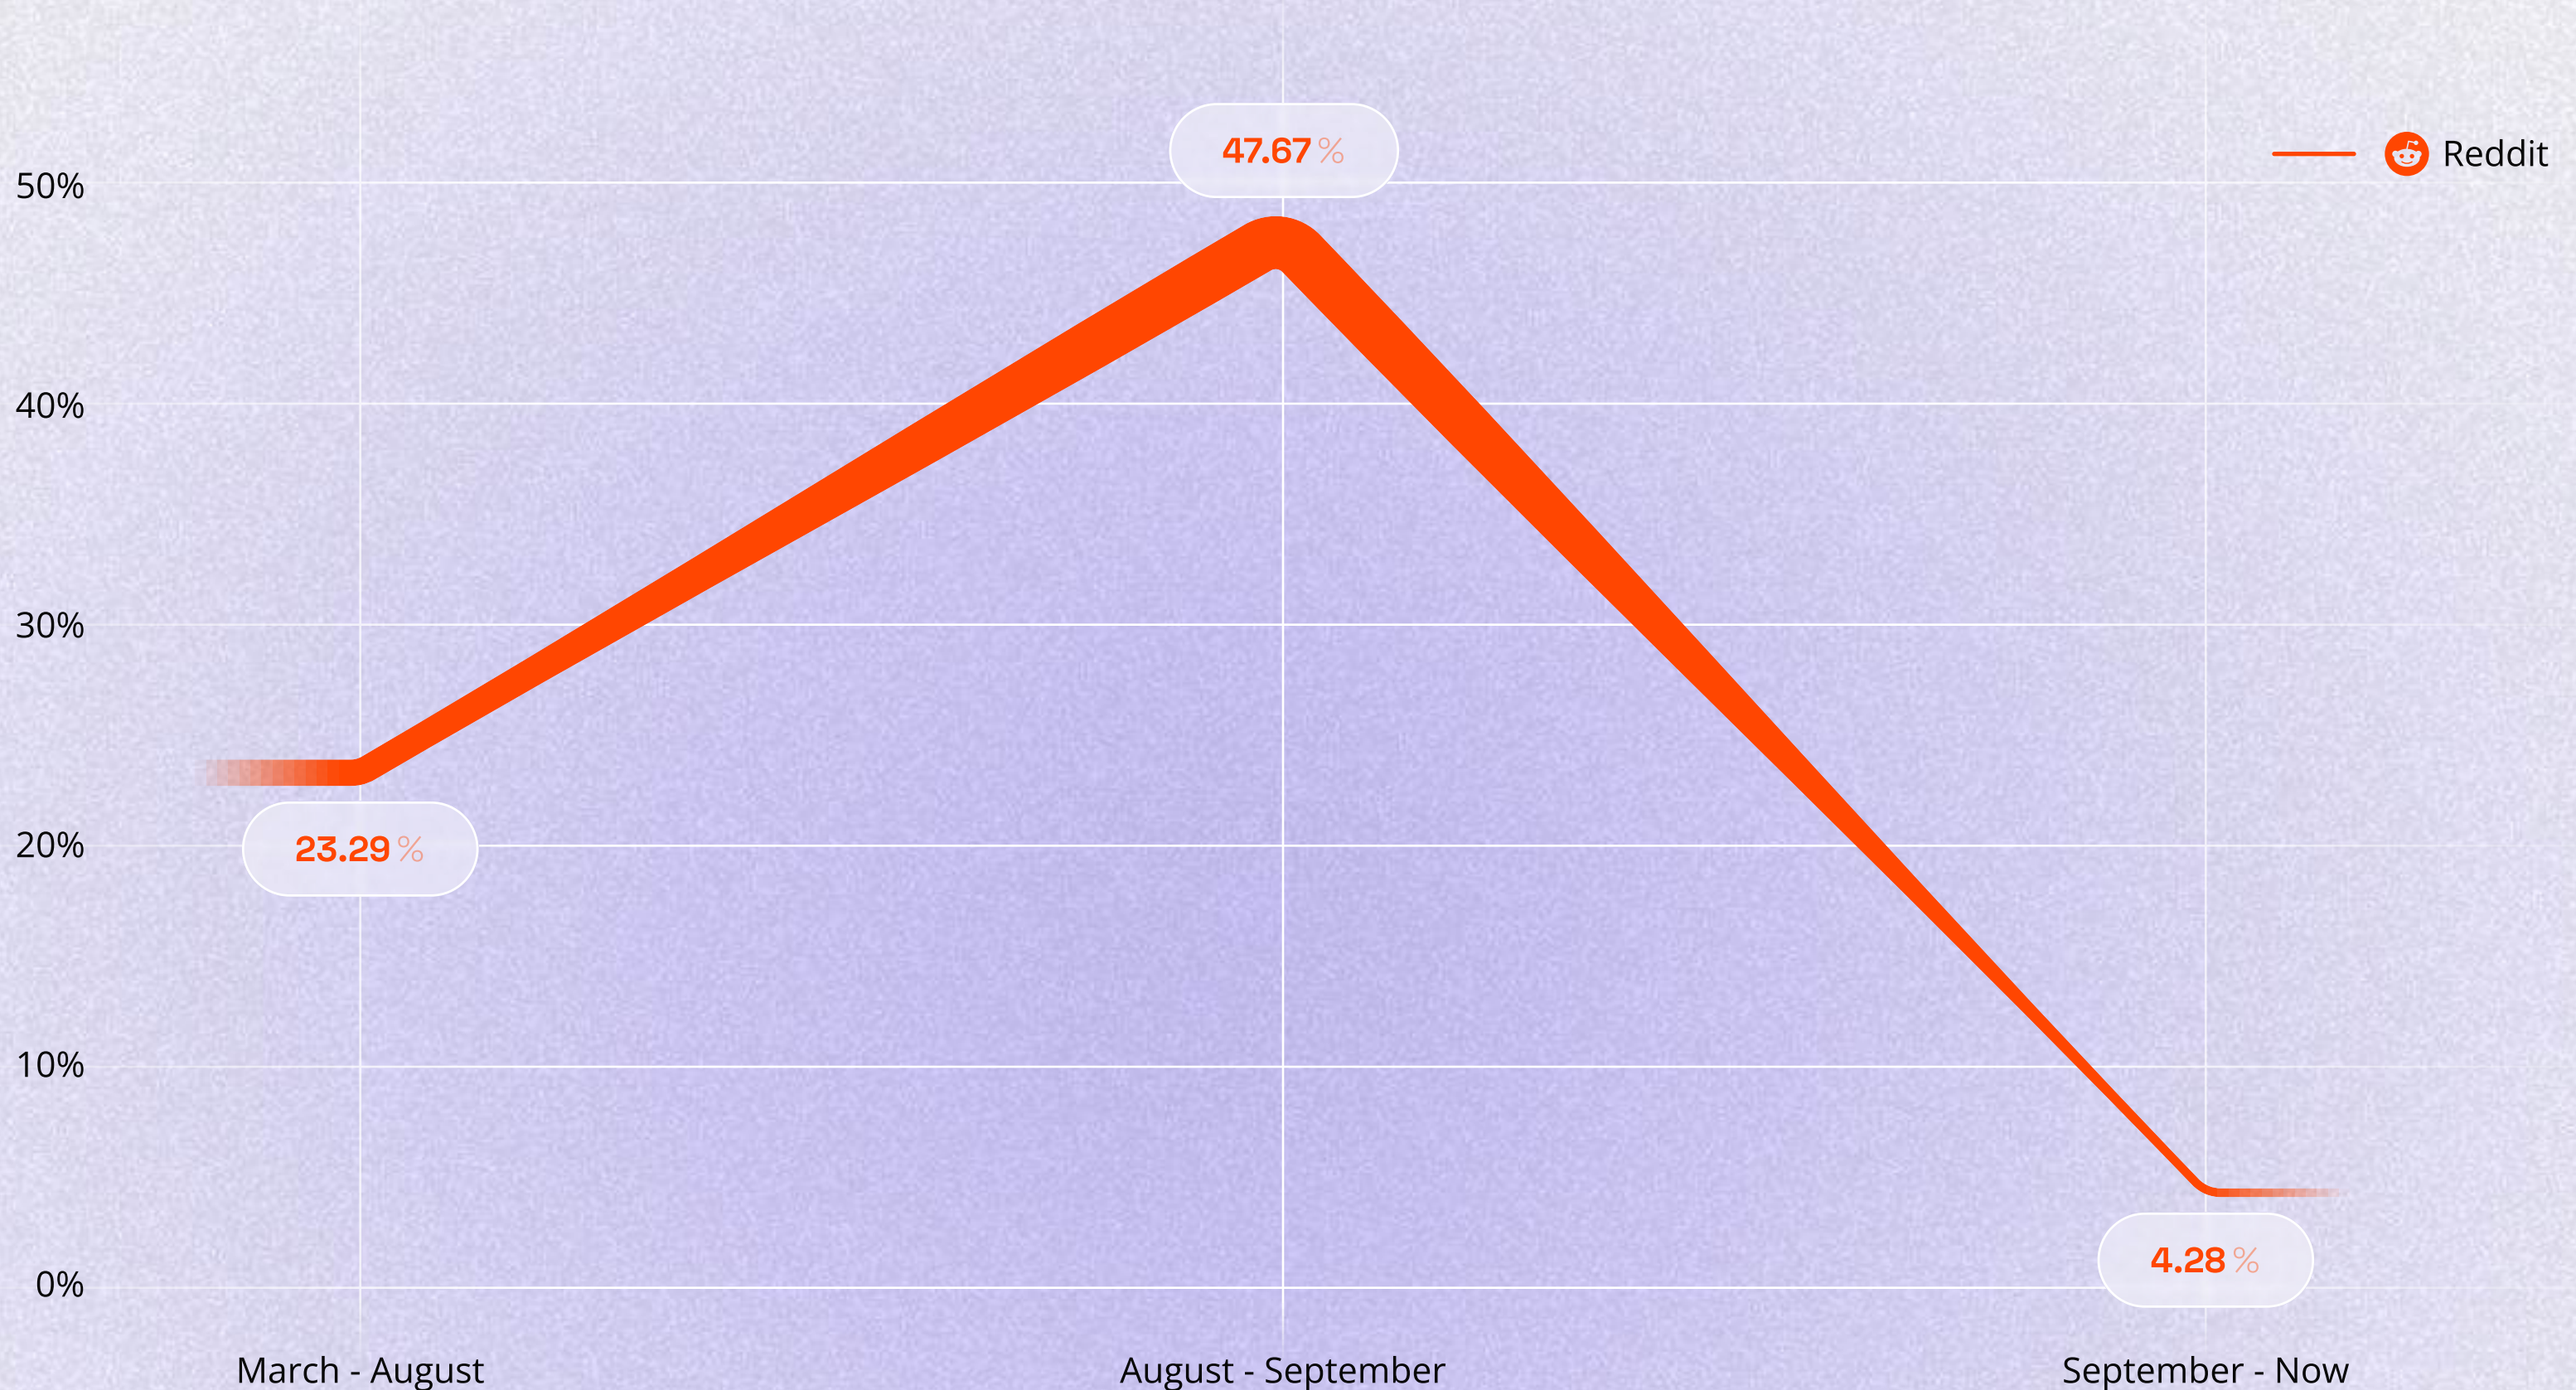
Analysis of the **origins used by AI models**, including diversity, concentration, top sources, and common paths.



About AI Sources Monitoring

It is important to highlight that the **ranking of sources is inherently volatile**. For this study, we consolidated data from March to August 2025 in order to smooth out short-term peaks and drops that result from shifts in the behavior of different AI models.

A clear example is **Reddit within ChatGPT**. While the overall average share during the period was 23.29%, we observed significant fluctuations: at certain points, **its share spiked to 70%**, while **in others it dropped to as low as 10-11%**. These swings are clearly visible in the charts below:



These variations do not invalidate the overall ranking but instead underscore the need for **continuous monitoring through generative engine optimization tools**. Such monitoring enables the detection of micro-trends and provides visibility into how the momentary volatility of sources can directly influence brand presence in AI-generated answers.

This also reinforces the importance of not chasing short-lived trends, but instead focusing on building a solid, consistent, and long-term content strategy. It is also important to have a lean team that is ready to adapt to algorithm changes.

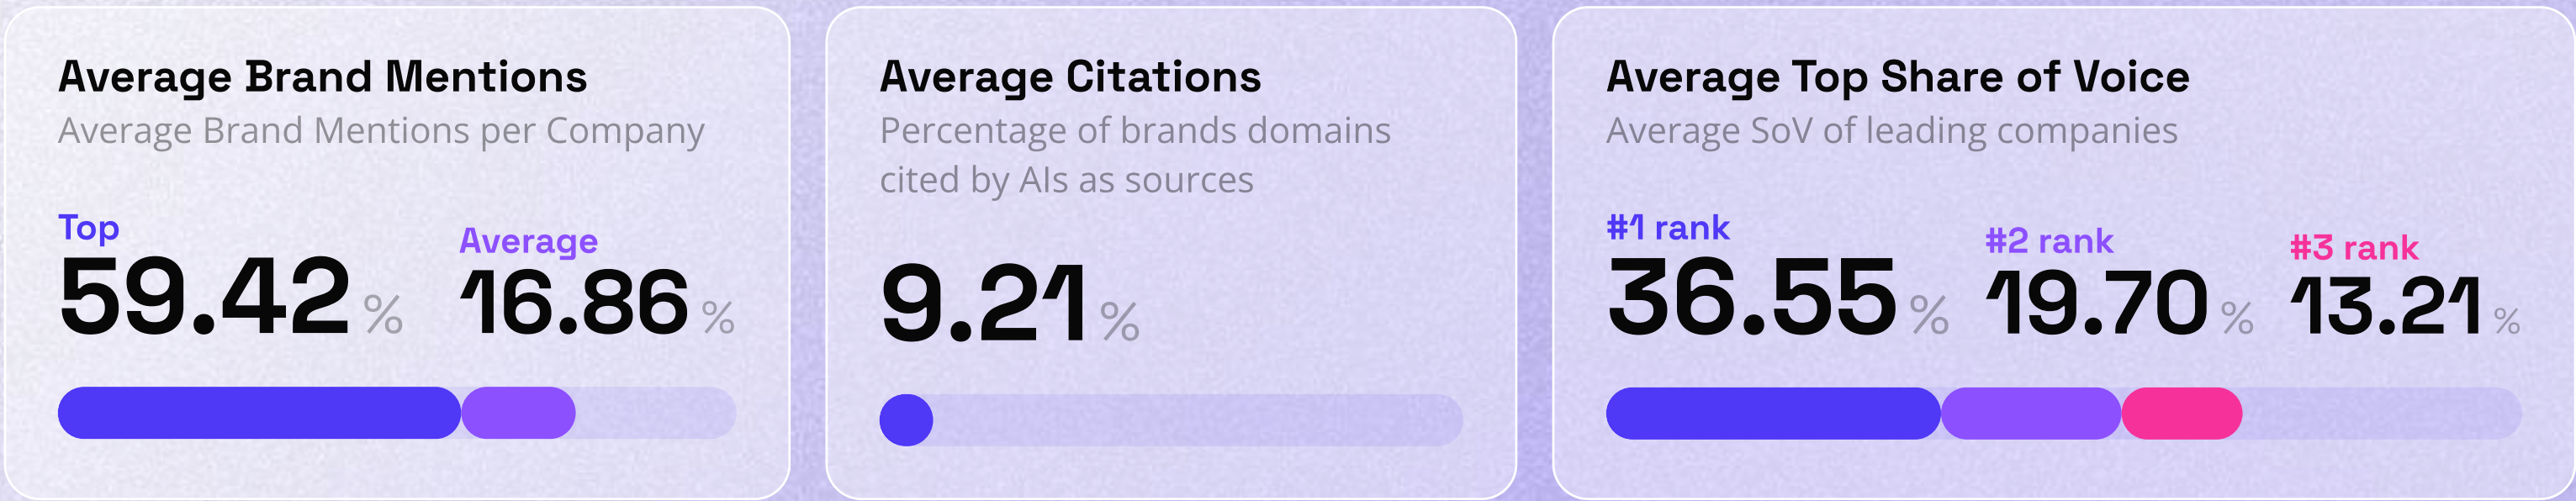
DON'T CHASE SPIKES
BUILD STABILITY

General (All Segments) AI Search Benchmarks

The overall view of AI search highlights both the scale of brand visibility and the role of authoritative sources. **On average, brands register 16.9% in mentions**, but **leaders achieve nearly 59.4%**, showing that presence can vary dramatically between companies. Citations: when a brand's domain is directly referenced as a source remain lower, with an average of 9.2%, while **each AI model cites between 2.5 and 13.7 distinct domains per response**, indicating differences in how models diversify their references.

Looking at where AIs draw information, a handful of domains consistently stand out: **Wikipedia, Reddit, YouTube, and Forbes rank among the most-cited across models**. In addition, site sections such as **/blog (37.1%) and /home (18.2%) dominate as the most common entry points for AI exploration**, reinforcing how content hubs and homepage structures often shape the first layer of visibility. Together, these findings suggest that brand exposure in generative engines is tied not only to overall mention volume but also to how well domains and content paths align with the sources AIs already prioritize.

Key AI Search KPIs: Brand Mentions, Share Of Voice, Citations



Average Top Share of Voice

Average SoV of leading companies

#1 rank

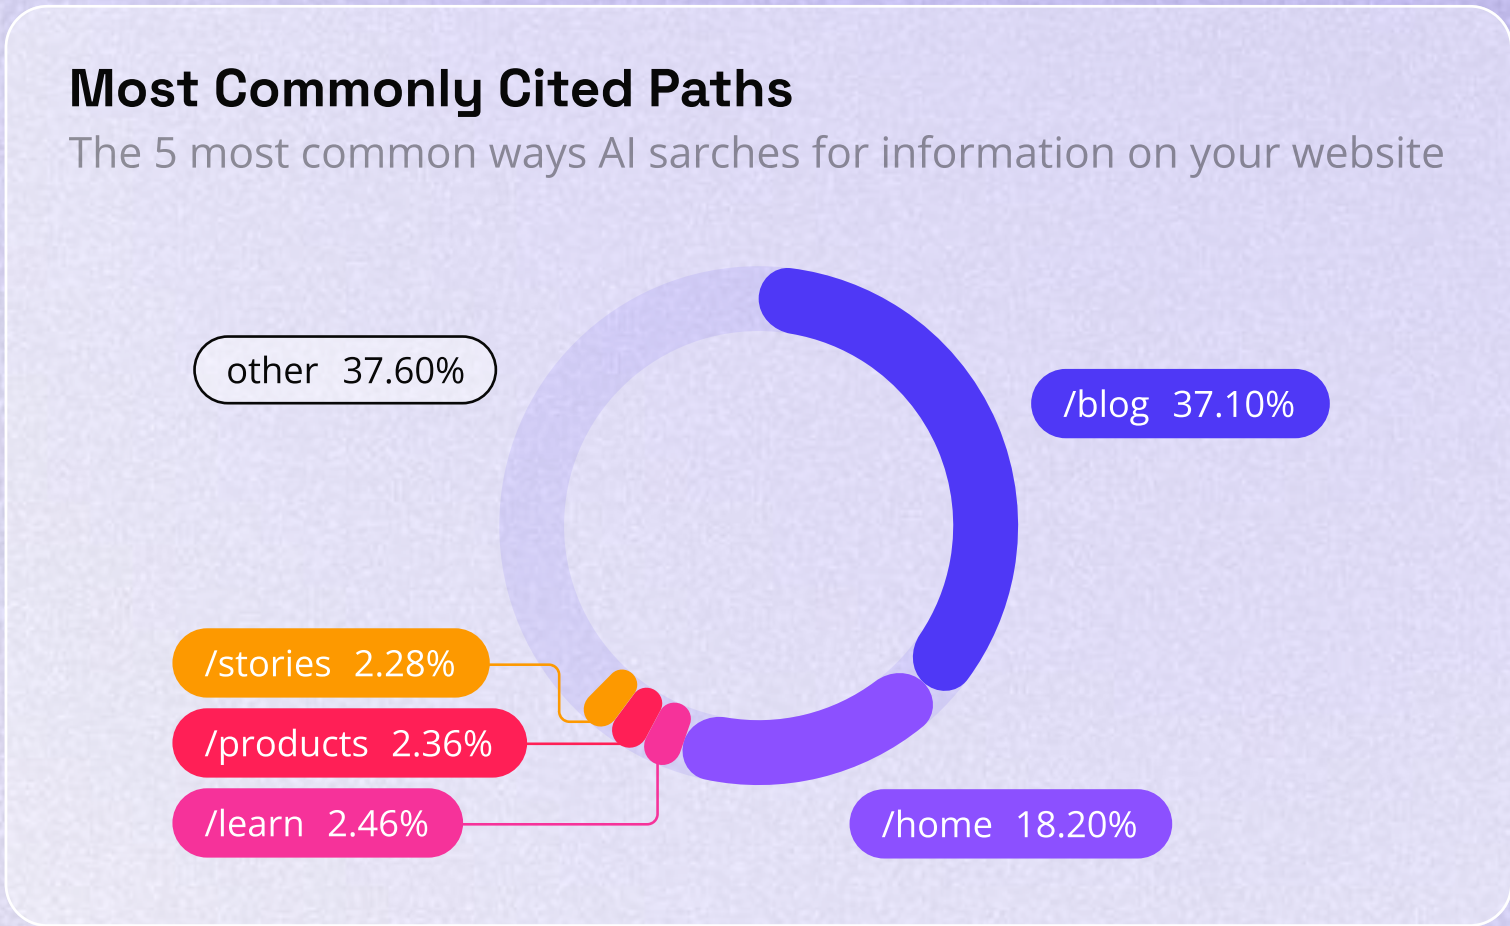
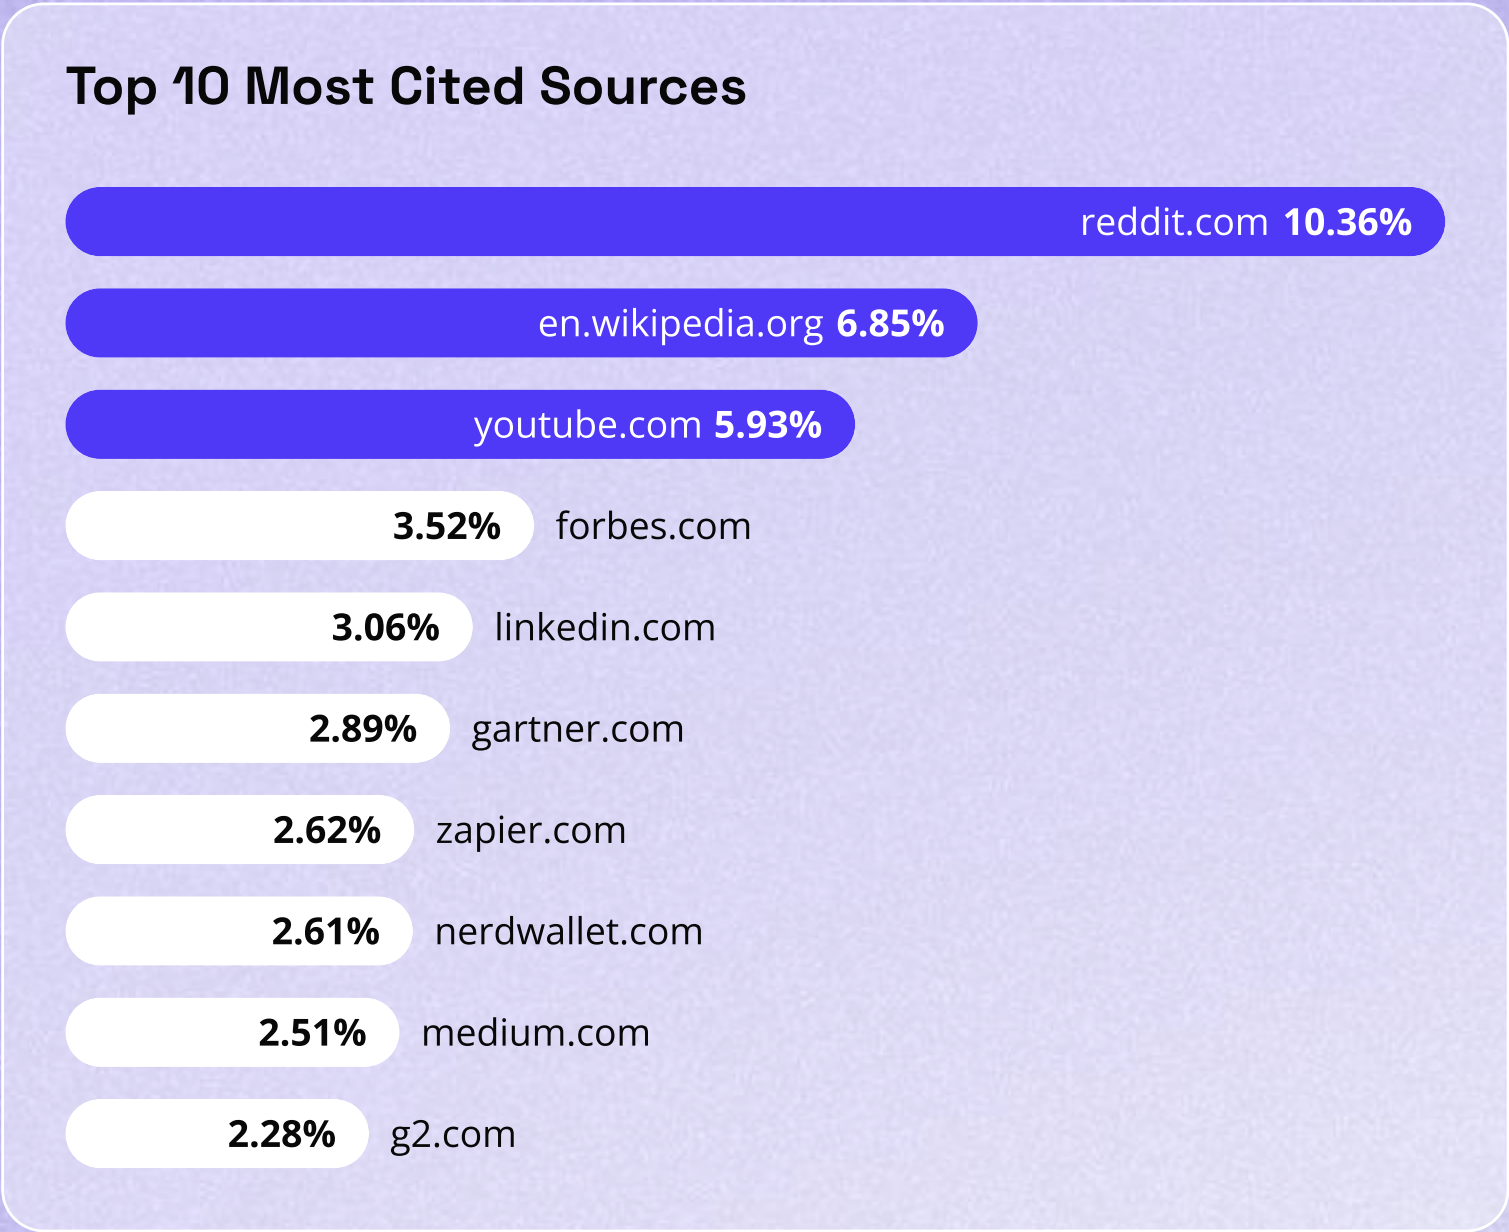
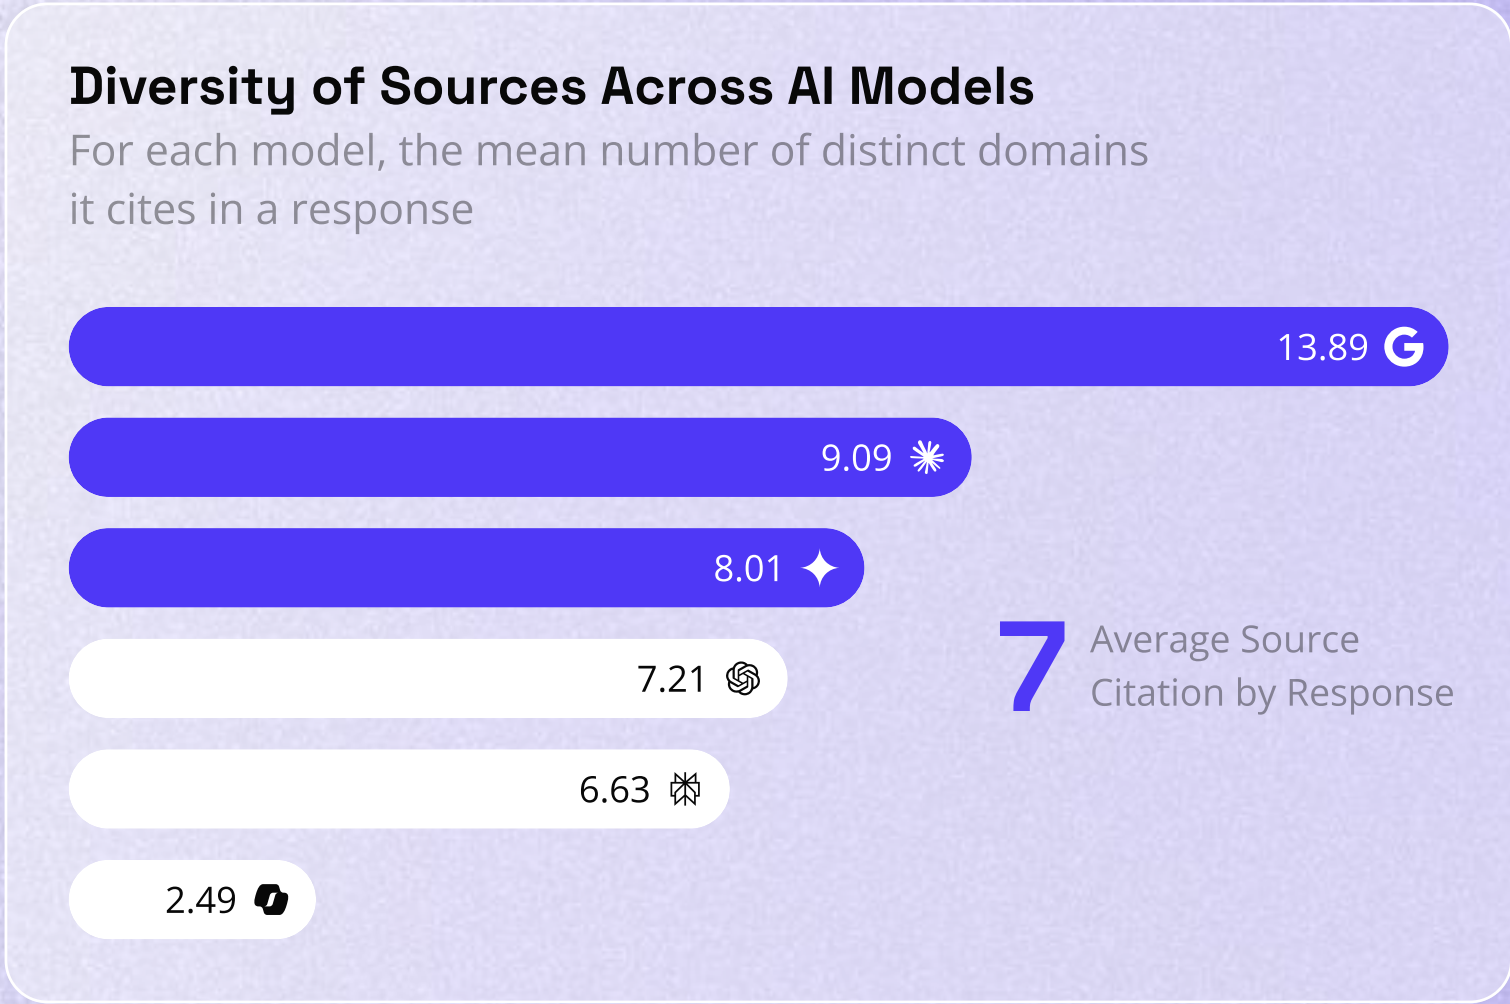
36.55 %

#2 rank

19.70 %

#3 rank

13.21 %



KEY TAKEAWAY

A handful of sources dominate AI answers. **Aligning content with these trusted domains**, plus **optimizing entry hubs** (/blog, /home), dramatically increases visibility.

AI Doesn't Search Everywhere: It Keeps Going Back To The Same Trusted Sources.

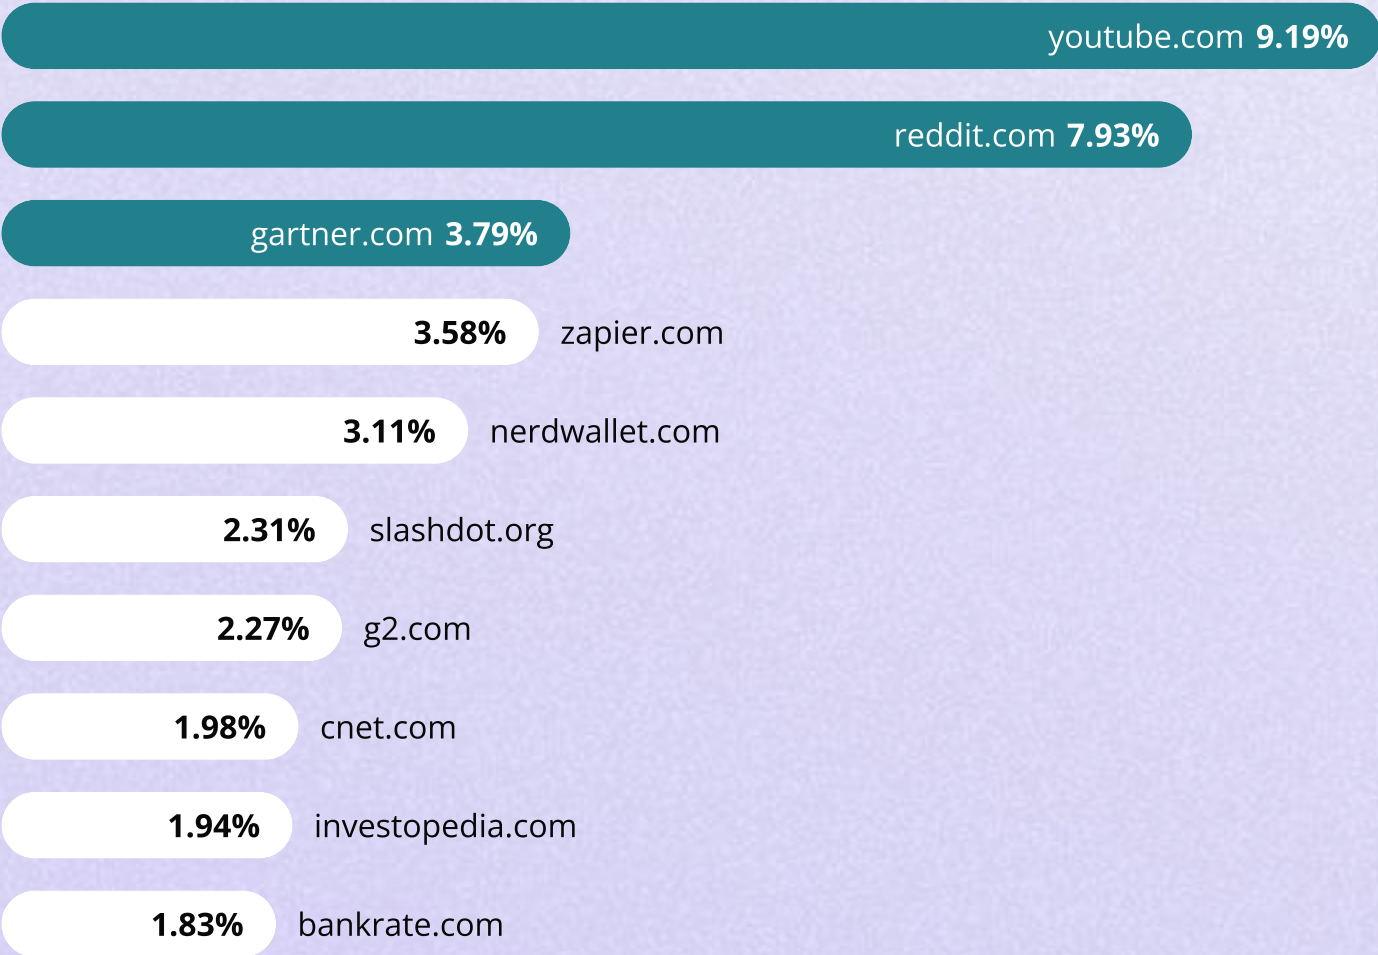
TOP AI SOURCES:
GENERAL

The top sources most cited by leading AI models, showing **which platforms dominate** AI-driven discovery

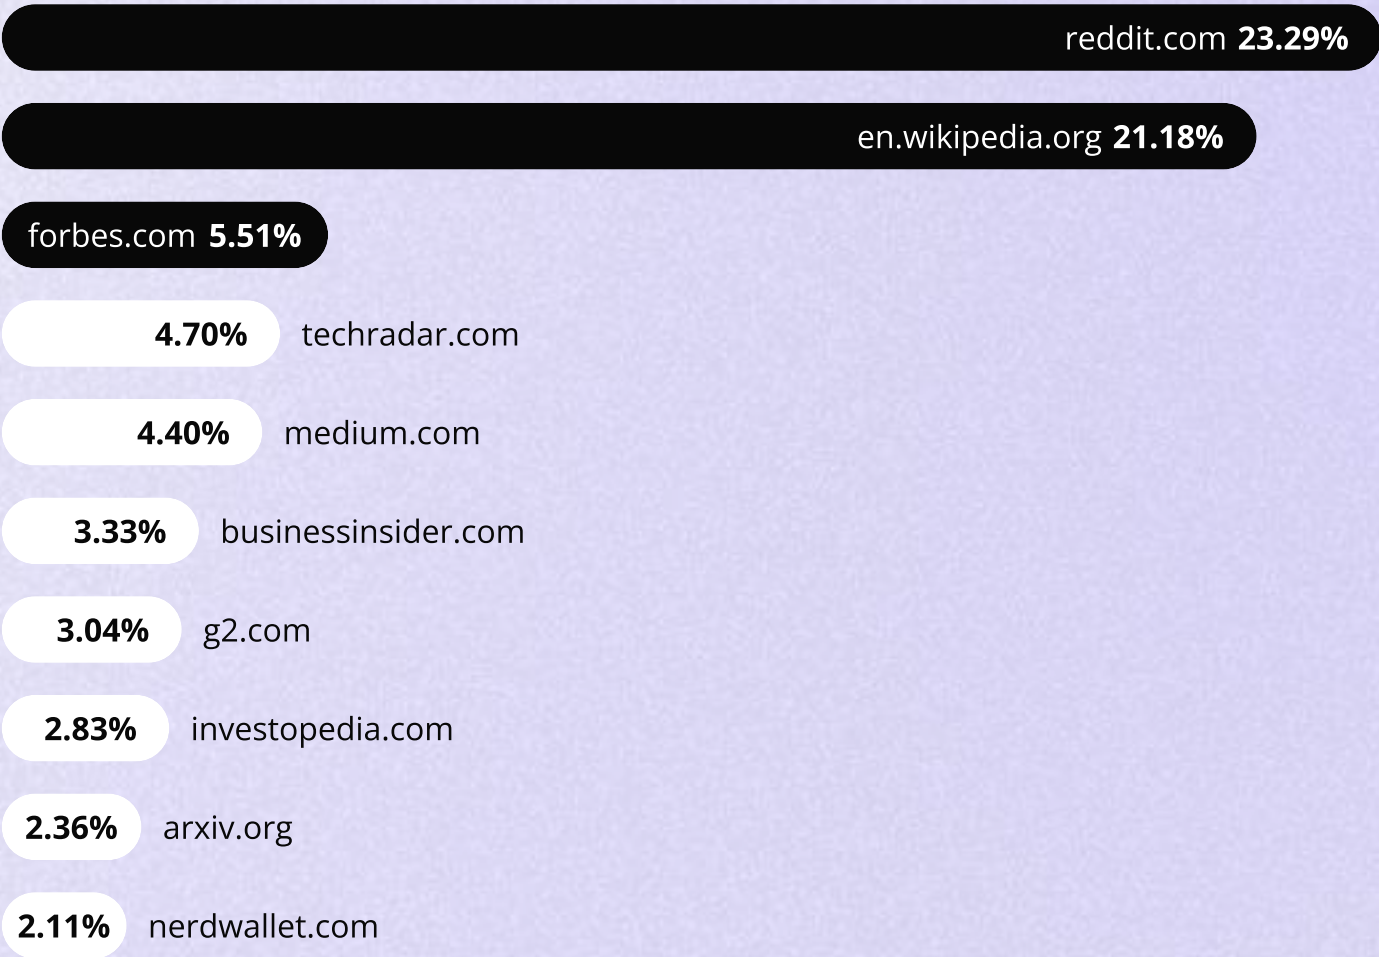
Community-Driven Sites Dominate AI Training Sources

- #1 Reddit
- #2 YouTube
- #3 Wikipedia

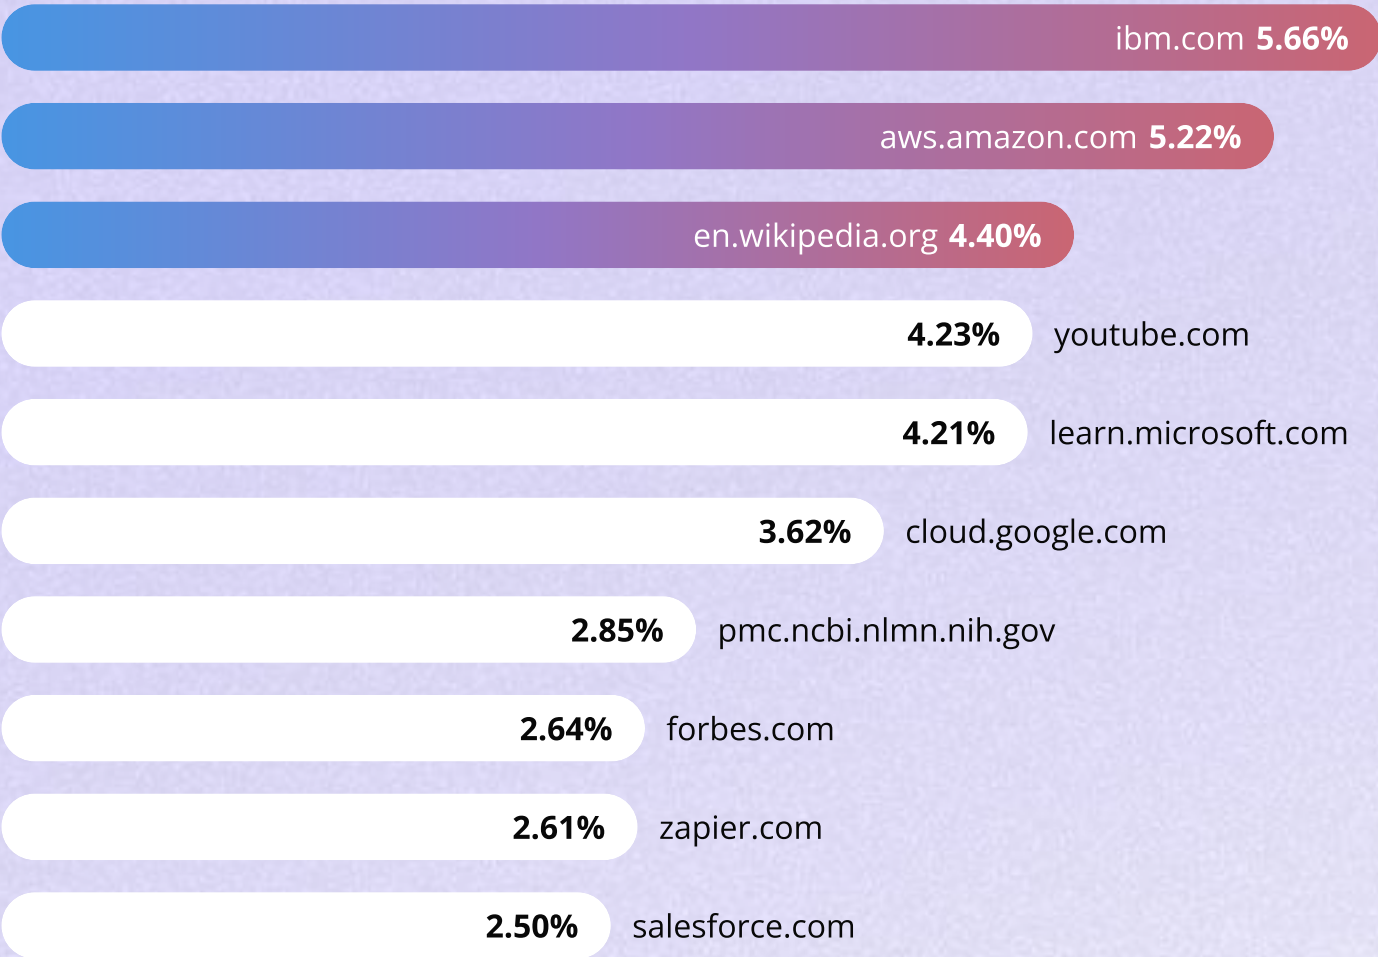
Perplexity



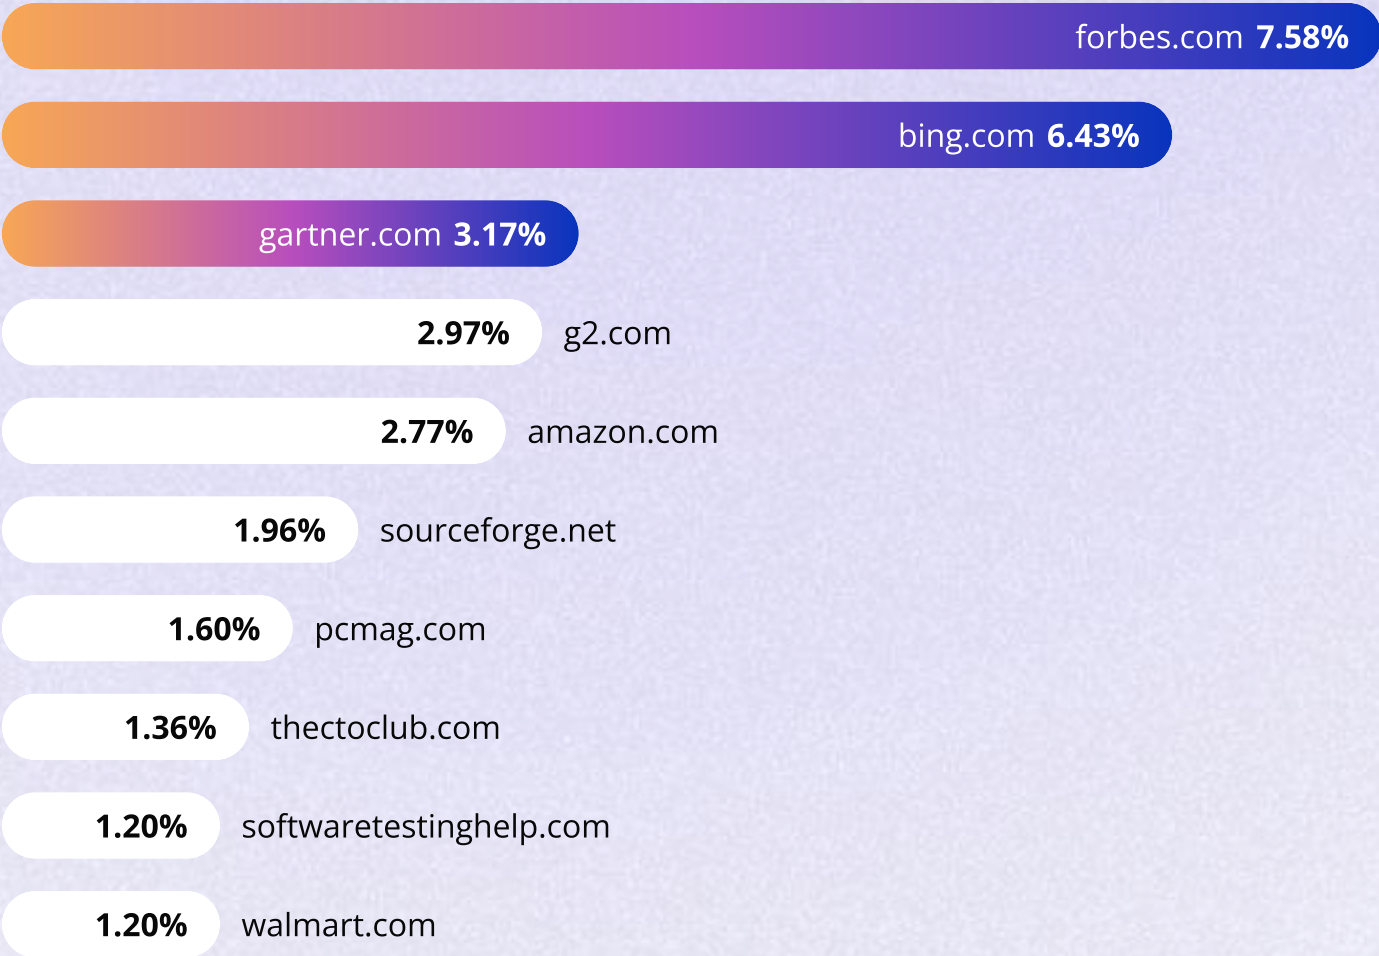
ChatGPT



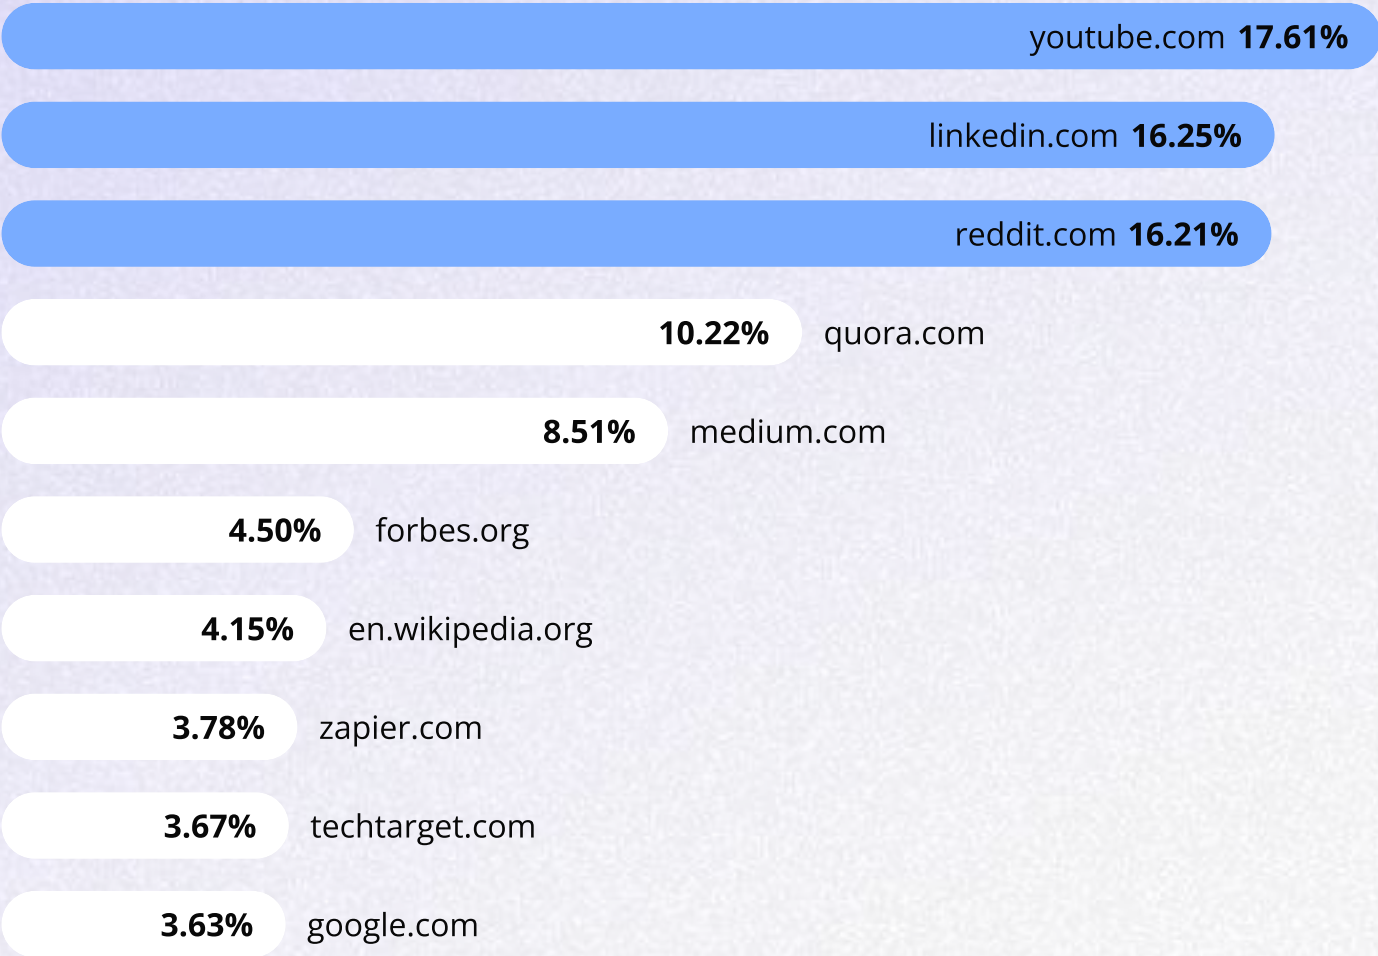
Gemini



Copilot



AI Overview



About The Authors



Renan Serrano
AthenaHQ Founding Team, Technical Program Manager, Strategic Research
[LinkedIn](#) [Email](#)



Natã Blum
AthenaHQ Data Specialist, Business Intelligence & Data Analysis
[LinkedIn](#) [Email](#)

About The Contributors



Ashley Stirrup
B2B Software Leader, former CMO at Algolia,
[LinkedIn](#)



Ethelbert Williams
Head of Retail Media Search & DTC Marketing, PepsiCo
[LinkedIn](#)

Research Team

This report was produced with the support of AthenaHQ’s research and strategy team, who contributed to data analysis, visualization, and editorial work.

Acknowledgments

We would like to thank our advisors, early customers, and industry partners who provided feedback and insights during the preparation of this study. Their perspectives helped ensure accuracy and relevance across industries.

Copyright

AthenaHQ is the leading platform for Generative Engine Optimization (GEO), empowering brands to understand and enhance their presence within AI-generated answers. Our advanced technology enables organizations to benchmark their digital visibility, gain actionable insights, and optimize content strategies for the new era of generative search.

Trusted by more than 200 companies and multinational organizations, AthenaHQ is redefining what it means to achieve digital visibility in an AI-driven landscape.

www.athenahq.ai

This publication is provided for general informational purposes only and is supplied "as is." It is not intended to serve as, nor should it be relied upon as, a substitute for professional advice.

Copyright © 2025 AthenaHQ. All rights reserved.

