 RENOWN RESEARCH

AI Visibility Report

AI Coding Tools

May 2026

THE THESIS

Cursor, Claude Code, and Codex are the bleeding edge developers actually use. AI search is recommending the leaderboard of 2023.

A data-driven analysis of how AI models perceive, rank, and recommend brands in the AI Coding Tools space.

Talk to engineers who ship code every day in 2026 and the bleeding edge is clear. Cursor is the IDE most of them have switched to. Claude Code is the agent they reach for when the work is non-trivial. Codex is the OpenAI-side answer to the same demand. None of this is contested in conversations, in developer surveys, or in usage data from the tools themselves.

The AI visibility picture of this same category looks nothing like that. We asked ten AI search engines a hundred questions about AI coding tools. They returned 4,209 brand mentions across thirty products. GitHub Copilot leads. Cursor is closing the gap. Tabnine, a tool whose practical relevance peaked years ago, sits comfortably at #4 propped up by training-data echo from a thousand "best of 2022" comparison articles. Claude Code (released February 2025) and OpenAI's Codex agent (released April 2025) are nearly invisible to half the AI surfaces your buyers consult.

This gap matters because AI-mediated discovery is now a real distribution channel, and the picture it shows isn't limited to people inside the AI news cycle. Developers who ask ChatGPT or Google AI Overview about "the best coding assistants" get back a 12-month-old map. Engineering managers vetting tools through Perplexity or Claude get answers shaped by training data, not what their team is actually evaluating. The disconnect shows up anecdotally too. Less-online developers, polled in surveys or in casual conversation, name the tools AI search shows them, not the ones the bleeding-edge crowd is using.

Three findings in this report should change how you think about AI visibility.

One. Model providers are **not** aggressively self-promoting. ChatGPT mentions Claude Code more often than Claude itself does. Gemini under-promotes Gemini Code Assist. The corporate self-interest theory of model bias does not hold up in this data.

Two. A single small vendor, Augment Code, has captured the citation game so thoroughly that **augmentcode.com** is the most-cited domain by four out of five AI search engines we studied. They are a tier-2 player by mention count, but a tier-1 force in what AI search knows about the category.

Three. Grok-4.3 mentions Cursor 81 times across 100 questions. It mentions OpenAI's Codex zero times. xAI has been openly attempting to acquire Cursor, and Elon Musk's conflict with Sam Altman is well-documented. That corporate feud is showing up in the recommendations your customers see.

These are not anomalies. They are the structural shape of AI-mediated buying. The brands winning this layer are doing it on purpose.

SCOPE

What we measured

30 brands and 10 AI models evaluated across 100 prompts, generating 1000 AI responses.

BRANDS EVALUATED

GitHub Copilot	Cursor	Claude Code	Windsurf
Cline	Replit	Codex	Amazon Q Developer
Gemini Code Assist	Continue	Aider	Tabnine
Sourcegraph Cody	JetBrains AI Assistant	Augment Code	Bolt.new
v0	Lovable	Devin	Zed
Roo Code	Plandex	Goose	OpenHands
Trae	Phind	Pieces	Qodo
	Kilo Code	Sweep	

AI MODELS EVALUATED

ChatGPT	Claude	Gemini	Perplexity
Google AI Overview	Google AI Mode	Grok	DeepSeek
	Mistral	Qwen	

55% #1 GitHub Copilot	7% Median Visibility	30 Brands Tracked	10 AI Models Tested
---------------------------------	--------------------------------	-----------------------------	-------------------------------


READING THIS REPORT

How we measure *AI Visibility*



Not all AI platforms are equal

Visibility scores are weighted by real-world reach. A recommendation from ChatGPT or Google AI Overview carries more weight than one from a niche model, because more people see it.



% AI VISIBILITY

Weighted mention rate across all platforms. Higher means more real users encounter this brand through AI. A brand at 78% appears in most responses on high-traffic platforms.

Range in this report: 0.5% to 78%

AVG POSITION

Where in the response a brand typically appears. Position #1 means mentioned first. Lower is better. Users tend to remember and act on whatever AI says first.

#1 = strongest top-of-mind signal










•• MODEL BREADTH

How many of the 10 AI platforms mention a brand at all. 10/10 = universal consensus. 2-3 = niche or emerging player in AI's view.

Top 7 brands: 10/10 models


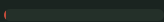


LEADERBOARD

AI Visibility Rankings

#	BRAND	AI VISIBILITY	AVG POSITION	MODELS
1	 GitHub Copilot	<div style="width: 55.3%;"><div style="background-color: #0070C0; height: 10px;"></div></div> 55.3	#2-3	10 / 10
2	 Cursor	<div style="width: 55.0%;"><div style="background-color: #0070C0; height: 10px;"></div></div> 55.0	#2-3	10 / 10
3	 Claude Code	<div style="width: 33.7%;"><div style="background-color: #FFC000; height: 10px;"></div></div> 33.7	#2-3	8 / 10
4	 Tabnine	<div style="width: 29.7%;"><div style="background-color: #FFC000; height: 10px;"></div></div> 29.7	#3-4	10 / 10
5	 Windsurf	<div style="width: 28.8%;"><div style="background-color: #FFC000; height: 10px;"></div></div> 28.8	#3-4	10 / 10
6	 Amazon Q Developer	<div style="width: 19.5%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 19.5	#4-5	10 / 10
7	 Sourcegraph Cody	<div style="width: 17.9%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 17.9	#4-5	10 / 10
8	 Codex	<div style="width: 16.4%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 16.4	#3-4	9 / 10
9	 Replit	<div style="width: 16.3%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 16.3	#4-5	10 / 10
10	 Augment Code	<div style="width: 14.4%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 14.4	#2-3	10 / 10
11	 Aider	<div style="width: 13.5%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 13.5	#4-5	9 / 10
12	 Continue	<div style="width: 13.0%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 13.0	#3-4	10 / 10
13	 Qodo	<div style="width: 11.9%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 11.9	#4-5	9 / 10
14	 JetBrains AI Assistant	<div style="width: 8.7%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 8.7	#4-5	10 / 10
15	 Bolt.new	<div style="width: 7.9%;"><div style="background-color: #FF4500; height: 10px;"></div></div> 7.9	#4-5	8 / 10

LEADERBOARD

AI Visibility Rankings (cont.)

#	BRAND	AI VISIBILITY	AVG POSITION	MODELS
16	 Gemini Code Assist	 7.2	#3-4	9 / 10
17	 Devin	 7.1	#3-4	10 / 10
18	 Cline	 6.7	#4-5	7 / 10
19	 Lovable	 5.4	#3-4	10 / 10
20	 v0	 4.2	#4-5	9 / 10
21	 Kilo Code	 2.8	#4-5	5 / 10
22	 Zed	 2.7	#5-6	8 / 10
23	 OpenHands	 2.0	#4-5	8 / 10
24	 Pieces	 2.0	#5-6	5 / 10
25	 Roo Code	 1.3	#4-5	7 / 10
26	 Goose	 1.1	#4-5	3 / 10
27	 Sweep	 0.3	#5-6	4 / 10
28	 Phind	 0.2	#3-4	3 / 10
29	 Plandex	 0.1	#9	1 / 10
30	 Trae	 0.1	#4	1 / 10

KEY OBSERVATION

GitHub Copilot leads with **55.3%** weighted AI visibility, recognized across **10** of 10 AI platforms. The top 5 brands capture the majority of AI mindshare, with a steep drop-off after Windsurf at **28.8%**.

MODEL INTELLIGENCE

What the models reveal *about themselves*

Each AI model has a distinct recommendation personality, shaped by its training data, retrieval architecture, and recency bias. The *Distinctive Pick* is the brand each model favors most relative to the overall average, revealing where its unique bias lies.

ChatGPT

The Default Channel

The highest-traffic AI surface and the closest to a balanced view of the AI coding market. Mentions Claude Code 53 times, more than Claude itself does, and only mildly over-promotes OpenAI's Codex (2.3 times the cross-model average). The category's neutral arbiter, with a slight Reddit-shaped tilt.

Cursor, Claude Code, GitHub Copilot
Top 3 Brands

Claude Code

Distinctive Pick

3.3

Avg Brands/Response

Claude

The Developer-Native

Cites GitHub, dev.to, and medium.com more than any other model. Treats developer-platform content as canonical. Surprisingly modest about Claude Code (42 mentions vs. ChatGPT's 53). The best surface for reaching technical buyers in their native dev-content habitat.

GitHub Copilot, Cursor, Claude Code
Top 3 Brands

Claude Code

Distinctive Pick

4.2

Avg Brands/Response

Gemini

The Stuck Encyclopedia

Training-only knowledge with no Codex, zero Claude Code, and zero Aider. Heavy on legacy completers: Tabnine 62, Amazon Q 47. Under-promotes Gemini Code Assist by 67% relative to peers. A reference desk for tools that existed in 2024.

GitHub Copilot, Tabnine, Amazon Q Developer
Top 3 Brands

Tabnine

Distinctive Pick

3.0

Avg Brands/Response

Perplexity

The Video-Reader

Cites youtube.com 70 times for AI coding queries. Trusts community forums and developer videos. The most likely surface to lift emerging tools (Augment, Qodo, Cline) above their mention-share, because content from the past six months gets through.

GitHub Copilot, Cursor, Tabnine
Top 3 Brands

**Augment
Code**

Distinctive Pick

4.3

Avg Brands/Response

MODEL INTELLIGENCE (CONTINUED)

What the models reveal *about themselves*

Google AI Overview

The Mass-Market Filter

The highest user-reach surface in the panel. Cites augmentcode.com 90 times, where a single vendor blog dominates Google's AI Overview citations for this category. Surfaces appear directly inside Google search results, making this the single biggest leverage point for any brand visible in 2026.

Cursor, GitHub Copilot, Windsurf
Top 3 Brands

Cursor

Distinctive Pick

3.9

Avg Brands/Response

Google AI Mode

The YouTube Aggregator

Cites youtube.com 140 times, by a wide margin more than any other source. Google AI Mode is essentially a video-content reader for the AI coding niche. Written content matters far less than what reviewers and educators say on video.

Cursor, GitHub Copilot, Claude Code
Top 3 Brands

Claude Code

Distinctive Pick

4.7

Avg Brands/Response

Grok

The Tribal Operator

Mentions Cursor 81 times, more than any other model, and OpenAI's Codex zero times. Heavy on open-source incumbents (Continue 75, Aider 55). The corporate-affiliation pattern is unmistakable. xAI's product ally Cursor wins, OpenAI's product loses, older OSS gets a sympathy bump.

Cursor, Continue, GitHub Copilot
Top 3 Brands

Continue

Distinctive Pick

5.7

Avg Brands/Response

DeepSeek

The Open-Source Curator

Heavy reach into established tools, light on anything launched after 2024. Mentions Tabnine 58 times, Amazon Q 51 times, and Codex 23 times (one of its highest first-party mentions). The Chinese-AI window into Western developer tools, with a 12-month lag.

GitHub Copilot, Tabnine, Amazon Q Developer
Top 3 Brands

Tabnine

Distinctive Pick

3.4

Avg Brands/Response

MODEL INTELLIGENCE (CONTINUED)

What the models reveal *about themselves*



The Conservative European

Highest GitHub Copilot loyalty in the panel. 91 mentions, the model's most-recommended brand. Almost ignores Cursor (42), barely acknowledges Claude Code (1) or Codex (15). When Mistral recommends, it recommends the safe, established choice.


GitHub Copilot, Amazon Q Developer, Tabnine
Top 3 Brands

**Amazon Q
Developer**

Distinctive Pick

4.8

Avg Brands/Response



The Listicle Engine

Produces broad lists with mainstream consensus. GitHub Copilot 78, Cursor 64, Tabnine 35. Modest variance from the field. The closest thing to a global average AI surface, which is also what makes it forgettable as a distribution channel for any single brand.

GitHub Copilot, Cursor, Amazon Q Developer
Top 3 Brands

Continue

Distinctive Pick

4.7

Avg Brands/Response

CROSS-MODEL ANALYSIS

Where the models disagree: *and what it means*

Divergences between models are the most valuable insight in this report. When all models agree, the signal is obvious. When they disagree, it reveals structural biases, training data gaps, and exploitable asymmetries.

	CHATGPT	CLAUDE	GEMINI	PERPLEXITY	GOOGLE AI OVERVIEW	GOOGLE AI MODE	GROK	DEEPSEEK	MISTRAL	QWEN
GitHub Copilot	32.4	52.0	63.8	44.7	38.0	37.9	45.9	68.6	69.7	44.6
Cursor	43.7	42.4	23.3	38.9	43.3	50.0	49.8	15.3	27.7	37.6
Claude Code	36.8	33.8	0.0	28.0	27.4	38.3	13.6	0.0	5.4	16.2
Tabnine	15.4	26.3	42.9	24.0	24.8	29.0	28.3	39.8	38.5	22.6
Windsurf	21.9	23.1	27.7	19.1	26.9	29.9	23.4	18.8	12.0	23.9
Augment Code	21.6	24.7	22.3	24.1	28.0	21.9	8.7	5.8	15.0	7.1
Continue	15.1	22.3	18.9	22.5	16.0	25.6	42.9	16.2	15.7	24.0
Codex	18.5	17.5	31.6	16.2	19.2	12.3	0.0	33.5	36.7	6.6
Sourcegraph Cody	7.9	22.4	22.0	16.7	25.5	21.4	24.7	21.2	20.5	20.5
Pieces	41.7	4.0	23.6	5.4	9.7	0.0	0.0	0.0	0.0	0.0
Amazon Q Developer	12.4	20.4	35.1	17.1	9.8	23.0	24.2	38.6	48.6	25.3
Replit	13.8	16.8	8.3	17.4	21.5	29.9	12.2	21.1	28.8	11.4
Gemini Code Assist	16.4	21.7	39.5	17.6	11.6	18.5	7.4	26.6	0.0	6.4
Devin	16.9	10.8	25.4	25.8	14.7	25.0	15.5	43.3	17.0	18.3
Aider	13.4	14.2	0.0	18.5	20.6	20.8	32.7	9.5	14.2	18.1

CROSS-MODEL ANALYSIS (CONTINUED)

Where the models disagree: *and what it means*

	CHATGPT	CLAUDE	GEMINI	PERPLEXITY	GOOGLE AI OVERVIEW	GOOGLE AI MODE	GROK	DEEPSEEK	MISTRAL	QWEN
JetBrains AI Assistant	7.7	21.3	6.2	19.1	16.8	24.3	12.0	15.2	23.3	16.5
v0	27.9	9.4	0.0	11.7	7.2	8.5	7.5	34.5	9.5	8.1
Lovable	6.4	32.4	12.9	11.7	15.6	27.5	14.2	12.9	5.0	8.6
Cline	13.8	14.7	0.0	13.1	17.8	14.9	14.4	0.0	0.0	7.9
Qodo	11.7	11.6	0.0	16.1	16.6	25.8	14.4	14.3	4.4	11.8
Bolt.new	9.3	16.2	0.0	11.5	12.8	22.5	12.3	0.0	6.7	11.8
Phind	0.0	0.0	0.0	47.9	0.0	0.0	0.0	10.3	0.0	5.4
OpenHands	5.4	13.3	0.0	6.7	7.1	23.7	12.9	32.9	0.0	17.3
Roo Code	0.0	14.9	0.0	5.4	8.7	17.8	5.4	0.0	4.6	6.4
Kilo Code	10.8	4.1	0.0	5.3	4.6	32.9	0.0	0.0	0.0	0.0
Zed	6.2	10.9	25.4	8.7	3.5	23.4	8.5	0.0	0.0	3.9
Trae	0.0	0.0	0.0	0.0	0.0	6.7	0.0	0.0	0.0	0.0
Goose	7.4	4.6	0.0	0.0	0.0	5.3	0.0	0.0	0.0	0.0
Sweep	0.0	0.0	0.0	0.0	0.0	7.1	10.3	3.5	0.0	5.4
Plandex	0.0	0.0	0.0	0.0	0.0	21.1	0.0	0.0	0.0	0.0

CROSS-MODEL DIVERGENCE

Three asymmetries

worth acting on

The biggest stories in the data live in the gaps between models. Three asymmetries below shape how brands appear in AI search.

DIVERGENCE #1

Grok-4.3 mentions Cursor 81 times. Codex zero times. xAI's corporate feud is showing up in product recommendations.

Grok-4.3 is the only model in the panel that gives Cursor more visibility than GitHub Copilot. 81 Cursor mentions versus 73 for Copilot across 100 questions. It mentions OpenAI's Codex agent **zero times**, while every other model surfaces it between 6 and 26 times. xAI has been openly attempting to acquire Cursor, and Elon Musk's public conflict with Sam Altman is well-documented. The data shows that conflict expressed at the recommendation layer. For buyers researching AI coding tools, the model they ask is the answer they get.

DIVERGENCE #2

One small vendor controls the citation game. Augment Code is the #1 cited domain by four of five AI search models.

augmentcode.com appears 328 times across the citation footprints of Claude, Perplexity, Google AI Overview, and Google AI Mode. That is more total citations than any media property, more than any developer tool's own site, more than Wikipedia or GitHub. Augment Code itself ranks #11 on the brand leaderboard, a tier-2 challenger by visibility. But by media weight inside AI training and search, they sit at tier-1. The playbook is direct. Own the canonical comparison content for your category, and AI search will treat you as the reference.

DIVERGENCE #3

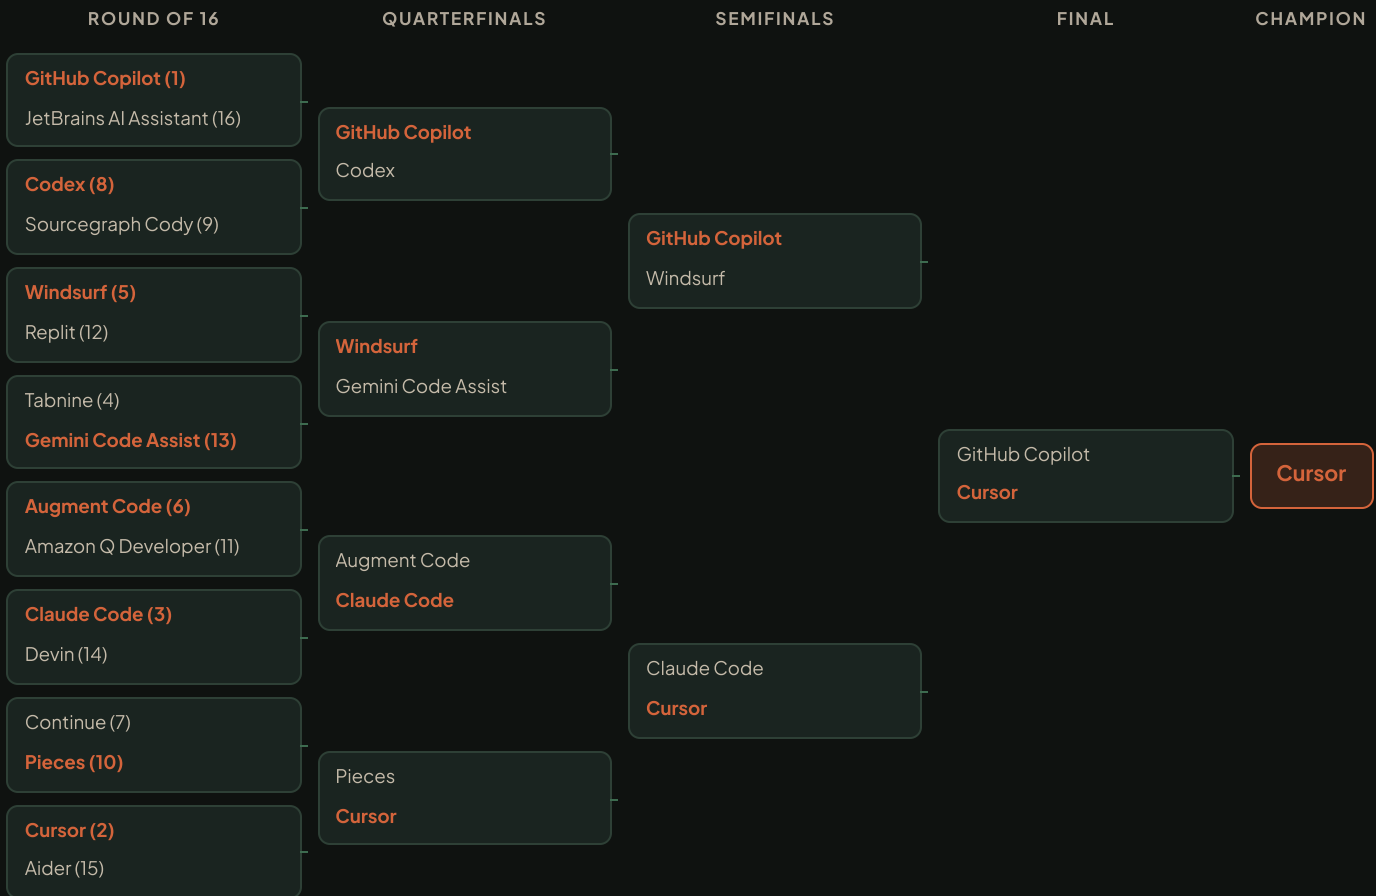
Self-promotion bias is reversed. ChatGPT mentions Claude Code 53 times. Claude itself mentions it 42.

The intuitive theory, that model providers will systematically push their own coding tools, does not survive the data. ChatGPT promotes OpenAI's Codex 2.3 times the cross-model average rate, but it under-promotes GitHub Copilot by 30%. Claude mentions Claude Code only marginally above the cross-model average. Gemini mentions Gemini Code Assist three times less often than other models do. Whatever is driving model-specific brand preferences, it isn't first-party loyalty.

BRACKET WARS

Head-to-head: *who wins when AI decides*

Top 16 brands seeded by visibility, matched head-to-head using ChatGPT. In each round, the model is asked which of the two tools it would recommend, and the winner advances.




TOURNAMENT RESULT

Cursor (seed #2) wins the ChatGPT bracket. The brand AI reaches for first in every head-to-head. When the same bracket runs on **Google AIO**, **Cursor** (seed #2) takes the title, same result, strong consensus.

QUADRANT ANALYSIS

The visibility landscape: *mapped*

AI Visibility measures how often and how prominently a brand appears across AI responses. Recommendation Strength reflects how positively AI models describe the brand when they mention it. Dot size indicates mention frequency. Color shows visibility tier: ● strong (60%+) ● good (40%+) ● moderate (20%+) ● low (10%+) ● minimal (<10%).



CHAMPIONS
High visibility and positive sentiment. The AI favorites.

CULT FAVORITES
Well-regarded but underexposed. Opportunity to grow visibility.

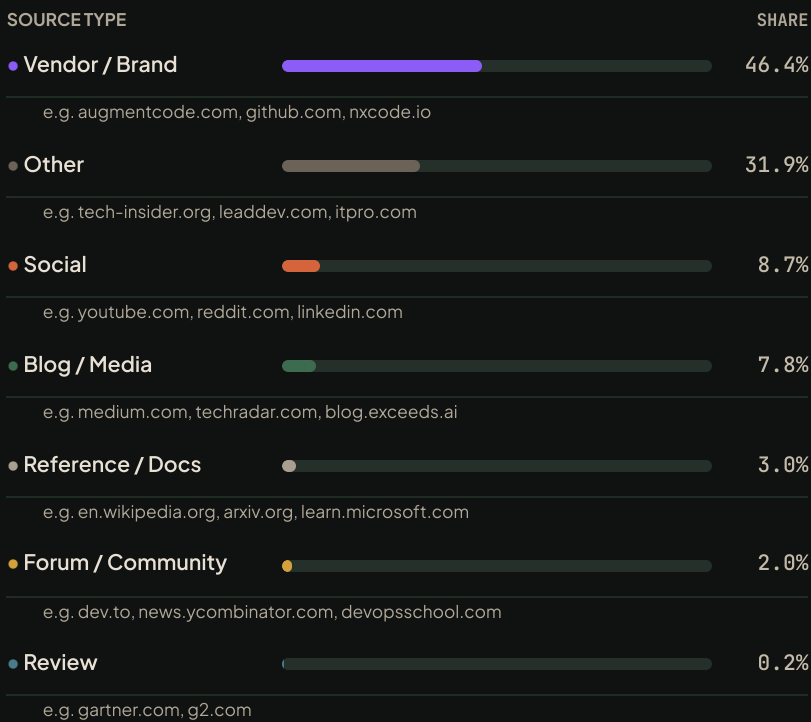
ALSO-RANS
Visible but with mixed perception. Reputation risk in AI.

GHOSTS
Low visibility, low sentiment. The danger zone.

CITATION ANATOMY

The sources AI trusts: *and the ones it ignores*

We analyzed thousands of citations across the AI platforms in our study that provide source attribution. Sources are grouped by type to reveal what categories of content AI models rely on most.



CONTENT FORMAT

"Listicle" is the most-cited content format at 43%

AI models overwhelmingly cite listicle content when answering questions in this category. Creating content in this format gives you the highest probability of being referenced.

KEY FINDING

AI citations are highly fragmented across hundreds of sources

No single domain captures more than 6% of citations. AI models draw from a wide range of sources, making it difficult for any one site to dominate. Winning requires presence across many content types and platforms.

OWNED CONTENT

Vendor-owned sites account for 46% of all citations

AI models frequently cite brands' own websites, documentation, and blogs as authoritative sources. This is the most controllable lever for AI visibility: invest in your own content and AI will reference it.

LONG TAIL

32% of citations come from niche and specialized sites

The "Other" category represents hundreds of smaller sites: niche comparison pages, regional tech blogs, and specialized tools. AI models actively discover and cite content beyond the obvious sources.

CONTENT STRATEGY


Blog and media content is cited 43x more than analyst reviews

Independent blogs and tech media dominate over traditional analyst reports (Gartner, G2). AI models favor accessible, detailed content over gated reviews. Publishing in open, indexable formats pays off.

CONTENT ANALYSIS

What kind of content *AI actually cites*

Content type classification based on all citation URLs and titles collected in our study. Understanding which formats AI prefers reveals where to invest content efforts.



EXAMPLE CITATIONS BY TYPE

- Listicle:** "The Ultimate AI Coding Tools Guide: When to Use What (F..."
- Other:** "Day 14: When AI Starts Hallucinating | 31 Days of Vibe..."
- Comparison:** "VS Code AI Extensions - Agentic Design | Agentic Design..."
- Product Page:** "ABOV3 - Enterprise AI Development Platform for Air-Gapp..."
- How-to / Guide:** "Speed up code reviews: Let AI handle the feedback imple..."




TAKEAWAY

Listicle content dominates at **42.6%**. Brands that produce comparison and listicle content on their own domains have a structural advantage in AI visibility.

VENDOR BIAS

Do models push their parent's product?

For each major foundation-model vendor that also ships a coding tool, we measured how often the vendor's own model mentions that tool vs. the average across peer models. The pattern is mostly truth-seeking with a small home-team lean, not aggressive self-promotion.



SYNTHESIS


Two of three pairs show a small home-team lean (~2x peer average), but peer averages are dragged down by training-only models. **The cleaner signal is absolute:** ChatGPT mentions Claude Code (53%) more than it mentions Codex (26%). Even OpenAI's own surface ranks a competitor's tool above its parent's. Models are largely truth-seeking on which tool is best. The strongest bias in this study lives elsewhere (see Grok-Cursor): corporate feud, not in-house loyalty.

MODEL LENS

How different model categories see the market

Not all AI models think alike. Grouping them by architecture, origin, and licensing reveals systematic biases in how they recommend tools. The gaps between groups are often larger than the gaps between individual models.

GOOGLE ECOSYSTEM



GOOGLE VARIANCE

Same model, different interface = different recommendations

All three Google engines are powered by Gemini, yet they produce very different results. Gemini direct mentions 3.0 brands/question vs AI Overview at 3.9. The interface layer reshapes the model's output.

OPEN VS CLOSED




OPEN-WEIGHTS SIGNAL

Open-weights models are more opinionated in their recommendations

Open-weights models recommend at 46.4% vs closed models at 48.3%. They cover 25 brands vs 27. Fewer guardrails appear to produce stronger opinions.

CHINESE VS WESTERN



CHINESE MODEL BIAS

DeepSeek is the only model that consistently puts open-source first

Chinese models show 37.4% recommendation rate vs Western at 50.5%. DeepSeek uniquely favors open-source tools, reflecting different training data and cultural priorities.

MODEL DIVERGENCE MAP

Where models disagree most

Each row shows one brand. Colored dots mark each model's mention rate. When dots cluster together, models agree. When they spread apart, there's an exploitable asymmetry: some models see value others miss entirely.



HIGH DIVERGENCE

Brands at the top have the widest spread across models. These represent the biggest asymmetries: a brand may be invisible on one model but dominant on another. Target the models where you're underperforming.

CONSENSUS

Brands near the bottom have tight clustering across all models. These positions are "locked in," either universally visible or universally absent. Changing consensus brands' AI visibility requires broader signal changes.

STRATEGIC RECOMMENDATIONS

The playbook

Actionable strategies derived from this report's data, ranked by impact and effort.

01 Audit the AI surface that matches your buyer.

URGENT WEEK 1

Your developer audience uses ChatGPT, Cursor's chat, and Claude as their primary AI search. Your enterprise buyer reaches the same questions through Google AI Overview embedded in normal search. Each surface ranks brands differently. Pick the one that matches your buyer's actual workflow and benchmark visibility against it specifically, not against a generic SEO dashboard.

02 Owned editorial is the cheapest AI-visibility lever you control.

HIGH IMPACT QUARTER

Augment Code's blog ranks above Wikipedia for AI citations on this niche. The playbook is direct: write the definitive comparison content for your category, structure it for machine readability, link it well, and AI search will pull from you. AI surfaces treat well-structured vendor content as canonical when third-party coverage is thin. That window is open right now.

03 If you launched after late 2024, current web presence is your only path.

EMERGING BRANDS ONGOING

Training-only models (Gemini, DeepSeek, Grok, Mistral, Qwen) literally don't know about Claude Code, Codex, Devin, Lovable, or Bolt yet. Each accounts for 0 to 3 mentions across 500 responses. Web-search-enabled surfaces are the only place these brands exist. Investment in current-year content, third-party comparisons, and AEO-optimized brand pages is the only lever newer brands have.

04 Build a sentiment moat. Listicle ubiquity is not enough.

MEDIUM IMPACT QUARTER

Tabnine appears in 402 responses with 18% positive sentiment. Cursor appears in 505 responses with 29% positive sentiment. Mention count is the floor. Framing is the ceiling. Brands that show up alongside praise compound visibility into qualified consideration. Brands that show up in neutral 'list of options' framing rarely convert.

05 Watch the model-provider feuds. They are now product feuds.

STRATEGIC ONGOING

xAI's public conflict with OpenAI is now visible in Grok's recommendation patterns. Anthropic's relationship with the Claude Code ecosystem will likely shape what Claude promotes. Google's Gemini surprisingly under-promotes Gemini Code Assist, possibly a deliberate neutrality choice. As more model providers ship product, expect more recommendation-layer politics. Choose your AI surface accordingly.

Track this in real-time.

Get continuous AI visibility monitoring, alerts on ranking changes, and model-specific optimization recommendations.

Start free at tryrenown.com →

METHODOLOGY

How we measure this



Prompt Design. Prompts simulate real-world purchase research in AI Coding Tools: comparison requests, recommendation asks, feature deep-dives, pricing inquiries, and use-case matching. Each is designed to elicit brand recommendations without biasing toward any vendor.

Model Execution. Every prompt was executed across 10 AI models spanning conversational AI, search-augmented AI, and specialized providers. Responses were collected with full citation metadata.

Scoring. AI Visibility reflects how likely a real user is to encounter a brand through AI-powered tools. Not all AI platforms have equal reach: a recommendation from ChatGPT or a Google AI Overview carries more weight than one from a niche model, because more people see it. We assign each model a weight proportional to its estimated real-world usage, then compute a brand's visibility as its weighted mention rate across all models. This ensures the rankings reflect actual user exposure rather than treating every model as equally influential.

Sentiment & Recommendation. Beyond visibility, we assess how each brand is talked about. Each mention is classified by sentiment and recommendation strength, identifying whether a brand is being actively recommended, passively listed, or criticized.

Citation & Content Analysis. URLs cited in AI responses are extracted, classified by source type (brand-owned, social, review, blog) and content format (listicle, comparison, docs). This reveals which platforms and formats AI models trust.

Hallucination Detection. For pricing queries, AI-stated prices are compared against known pricing. Hallucinations are categorized by type and severity.

Limitations. AI outputs are non-deterministic. This report uses single-run data collection and reflects model behavior at the time of collection. Model updates and training data changes can shift results between editions.

About Renown. Renown provides AI visibility intelligence for brands. We monitor how AI models perceive, rank, and recommend products, giving marketing teams the data they need to compete in AI-mediated discovery.
 Contact: hello@tryrenown.com · tryrenown.com



RENOWN RESEARCH
tryrenown.com