

 RENOWN RESEARCH

AI Visibility Report

Outbound Sales Stack

May 2026

THE THESIS

Apollo has overtaken Outreach and Salesloft inside AI search. The category your buyers see no longer matches the one you compete on at G2.

A data-driven analysis of how AI models perceive, rank, and recommend brands in the Outbound Sales Stack space.

For two decades, "outbound sales tools" meant Outreach versus Salesloft. The only real question was which sequencer your SDR team would standardize on. Apollo changed that by bundling data, dialer, sequencing, and analytics into one platform at a price point that scales down to startups. AI search engines noticed before most RevOps leaders did.

We asked ten AI search engines a hundred questions about modern outbound. Apollo leads on every dimension we measured. First-pick rate: 96 wins. Total mentions: 645. Sentiment: 0.41, the highest in the niche. Recommendation rate: 21% across those 645 mentions. The old duopoly framing no longer matches what buyers actually see.

That is the surface story. Three findings underneath it reshape the playbook.

One. ZoomInfo's editorial property pipeline.zoominfo.com is the most-cited domain on Claude (75 citations) and Perplexity (28) for outbound queries. Across four AI search engines its owned content racks up 188 citations, more than any media outlet and more than any other vendor. Long-form, comparison-rich content built with care is getting pulled into answers as canonical reference material.

Two. Training-only models (Mistral, Qwen, Grok, DeepSeek) account for between 35% and 65% of any given brand's visibility share. Brands launched after the 2023 training cutoff get most of their AI exposure through web-search-enabled models, because the training-only ones barely know they exist. For newer entrants, current web presence is the only lever that works.

Three. Lavender is mentioned only 59 times, but 32% of those mentions are recommendations, the highest rate in the niche. Mention count and recommendation rate measure different things. Some brands sit at the floor on visibility and the ceiling on intent. Others do the reverse.

If you sell into outbound, this report tells you where you sit on AI's map of your category and what your tier needs to do next.
































Apollo leads on every metric we measured. ZoomInfo's blog is cited more than Wikipedia. Lavender shows up rarely but gets recommended a third of the time it does. None of this matches what your G2 dashboard tracks.

SCOPE











What we measured

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

BRANDS EVALUATED

 Apollo	 ZoomInfo	 Outreach	 Salesloft
 Clay	 Gong	 LinkedIn Sales Navigator	 HubSpot Sales Hub
 Cognism	 Lusha	 Reply.io	 Instantly
 Smartlead	 Lemlist	 Chorus	 Chili Piper
 6sense	 Demandbase	 Bombora	 Common Room
 Avoma	 Attio	 11x	 Artisan
 Regie.ai	 Lavender	 UserGems	 Reo.dev
 Nooks		 Orum	 Clearbit


AI MODELS EVALUATED

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

61% #1 Apollo	9% Median Visibility	31 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












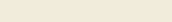







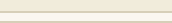










•• MODEL BREADTH

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




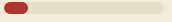



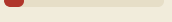



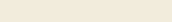







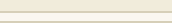










LEADERBOARD

AI Visibility Rankings

#	BRAND	AI VISIBILITY	AVG POSITION	MODELS
1	 Apollo	 60.5	#2-3	10 / 10
2	 Salesloft	 40.8	#3-4	10 / 10
3	 ZoomInfo	 39.9	#3-4	10 / 10
4	 Outreach	 37.7	#1-2	10 / 10
5	 HubSpot Sales Hub	 23.3	#4-5	10 / 10
6	 Clay	 21.1	#4-5	9 / 10
7	 Lemlist	 19.5	#4-5	10 / 10
8	 Gong	 19.1	#3-4	10 / 10
9	 Instantly	 19.0	#3-4	9 / 10
10	 LinkedIn Sales Navigator	 17.4	#4-5	10 / 10
11	 Cognism	 17.3	#4-5	10 / 10
12	 Smartlead	 16.3	#4-5	10 / 10
13	 6sense	 14.6	#4-5	10 / 10
14	 Reply.io	 14.1	#4-5	10 / 10
15	 Lusha	 10.3	#4-5	10 / 10

LEADERBOARD

AI Visibility Rankings (cont.)

#	BRAND	AI VISIBILITY	AVG POSITION	MODELS
16	 Bombora	 9.5	#4-5	10 / 10
17	 Chorus	 9.2	#4-5	10 / 10
18	 Nooks	 8.8	#4-5	6 / 10
19	 Demandbase	 7.8	#4-5	10 / 10
20	 Clearbit	 7.8	#5-6	10 / 10
21	 Avoma	 6.7	#4-5	10 / 10
22	 Orum	 5.1	#3-4	9 / 10
23	 Lavender	 4.3	#3-4	10 / 10
24	 Chili Piper	 3.6	#5-6	10 / 10
25	 Artisan	 3.3	#4-5	6 / 10
26	 UserGems	 2.8	#4-5	7 / 10
27	 11x	 2.7	#3-4	7 / 10
28	 Common Room	 1.9	#3-4	7 / 10
29	 Regie.ai	 1.9	#5-6	10 / 10
30	 Attio	 1.1	#6-7	6 / 10

KEY OBSERVATION

Apollo leads with **60.5%** 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 HubSpot Sales Hub at **23.3%**.

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 Consumer Researcher

Cites Reddit 77 times, by far the most of any source. This is the default surface for SDRs and founders doing informal vendor research. Apollo, Outreach, and ZoomInfo dominate, and ChatGPT mirrors the broad consensus of the B2B sales community.

Apollo, Outreach, Salesloft

Top 3 Brands

Orum

Distinctive Pick

3.4

Avg Brands/Response

Claude

The Vendor-Trusting Analyst

Cites pipeline.zoominfo.com 75 times. A single vendor's editorial property is its top source for outbound queries. Treats well-structured vendor content as authoritative, which makes it the most likely surface to expose mid-market challengers when their content is dense and credible.

Apollo, ZoomInfo, Salesloft

Top 3 Brands

ZoomInfo

Distinctive Pick

4.4

Avg Brands/Response

Gemini

The Generalist

Modest variance from the cross-model average. Top picks: Apollo, Outreach, Salesloft. Doesn't favor any specific vendor property or community channel. Useful as a baseline for what mass-market users see.

Salesloft, Outreach, Apollo

Top 3 Brands

Outreach

Distinctive Pick

4.7

Avg Brands/Response

Perplexity

The Technical Evaluator

Top citations: pipeline.zoominfo.com (28), apollo.io (24), salesforce.com (20), unifygtm.com (17), trellus.ai (15). Skews heavily toward vendor-owned and high-quality educational content. The most likely surface to feature your product when your owned content is technically thorough and citation-rich.

Apollo, Salesloft, ZoomInfo

Top 3 Brands

Apollo

Distinctive Pick

4.4

Avg Brands/Response

MODEL INTELLIGENCE (CONTINUED)

What the models reveal *about themselves*

Google AI Overview

The Mass-Market Filter

Embedded directly in Google search. Cites Reddit, YouTube, ZoomInfo's pipeline blog, Apollo's blog, and Unify GTM. The single largest visibility lever for any B2B sales brand. If you're not in Google AIO results, you are invisible to the broadest segment of researching buyers.

Apollo, ZoomInfo, Salesloft
Top 3 Brands

Nooks

Distinctive Pick

4.6

Avg Brands/Response

Google AI Mode

The Multi-Source Synthesizer

Pulls from YouTube (59), Reddit (59), ZoomInfo pipeline (50), Apollo (44), and Unify GTM (36). The widest mix of citation types of any surface we measured. To land here you need video, community, owned, and third-party content all working together.

Apollo, ZoomInfo, Outreach
Top 3 Brands

Clay

Distinctive Pick

5.7

Avg Brands/Response

Grok

The Loud Outsider

High mention rates for popular tools (Apollo 73, Salesloft 70, Outreach 50) but unusual emphasis on lesser-known names (Lemlist 47, Smartlead 37, Reply.io 27). X/Twitter as a training source skews toward what the founder community is actively talking about, which isn't always what enterprise sales teams use.

Apollo, Salesloft, ZoomInfo
Top 3 Brands

Lemlist

Distinctive Pick

7.2

Avg Brands/Response

DeepSeek

The Conservative Recommender

Mentions established tools heavily and ignores most newcomers. Apollo (48), Salesloft (60), Outreach (62). Knows Clay only 3 times and Instantly only 4. Strongly biased toward pre-2024 training data, which leaves recent category leaders absent from its recommendations.

Outreach, Salesloft, HubSpot Sales Hub
Top 3 Brands

**HubSpot
Sales Hub**

Distinctive Pick

5.1

Avg Brands/Response

MODEL INTELLIGENCE (CONTINUED)

What the models reveal *about themselves*

 **Mistral**

The Old-Guard Pragmatist

Favors the legacy sequencer pair: Outreach 67, Salesloft 61. Mentions Reply.io 56 times (the highest of any model) and Clay only 7. Reflects what enterprise European buyers see, with a conservative bias toward established brands.

Outreach, Apollo, Salesloft
Top 3 Brands

Reply.io

Distinctive Pick

6.8

Avg Brands/Response

 **Qwen**

The Comprehensive Lister

Highest Apollo mention rate of any model (84). Produces broad lists with mainstream consensus. The closest thing to a global B2B sales view, which is also what makes it forgettable as a distribution channel for any single brand. To show up in Qwen's lists, you need to be on every vendor list.

Apollo, Salesloft, Outreach
Top 3 Brands

Salesloft

Distinctive Pick

6.7

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
Apollo	39.6	43.6	36.6	40.5	35.5	47.9	48.5	32.8	36.0	47.0
Outreach	36.1	35.2	46.4	29.2	27.6	37.4	34.9	49.8	49.6	37.2
Salesloft	25.0	30.7	40.7	28.0	30.0	38.8	39.2	40.1	36.7	40.5
ZoomInfo	21.4	33.3	28.0	29.9	30.4	40.2	37.6	33.7	26.8	29.8
HubSpot Sales Hub	15.8	28.9	31.6	28.7	20.0	36.0	26.5	42.3	36.4	21.2
Instantly	19.3	25.4	0.0	18.8	26.6	31.0	22.9	31.6	17.8	24.1
Clay	23.5	18.8	0.0	18.9	23.3	30.2	22.6	28.9	9.2	24.1
Gong	20.1	15.5	17.5	19.9	19.1	30.3	23.5	26.4	28.5	23.2
LinkedIn Sales Navigator	21.5	13.7	25.6	13.0	15.5	25.6	23.5	30.2	23.3	11.1
Smartlead	15.4	22.8	13.3	16.4	17.4	29.8	25.6	32.4	18.5	19.7
Lemlist	16.8	22.1	18.4	13.6	16.0	18.0	31.3	31.2	30.3	14.0
Cognism	18.6	13.1	14.4	15.7	23.5	28.7	14.0	5.9	7.5	9.7
Chili Piper	21.9	24.6	6.7	25.4	15.1	31.4	22.3	19.7	7.2	13.0
Lavender	12.9	15.9	17.2	16.3	22.0	23.5	6.4	35.2	15.9	12.9
6sense	12.3	16.7	15.0	17.3	15.2	26.7	23.8	16.8	15.1	18.3

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
Reply.io	9.9	16.3	12.8	14.0	16.8	20.7	19.4	20.6	33.6	10.8
Orum	7.5	24.6	0.0	18.9	20.9	32.6	15.4	38.2	5.4	12.6
Chorus	14.2	10.9	11.6	6.9	12.7	27.3	22.5	20.7	22.5	17.3
Lusha	14.0	13.3	12.0	7.4	10.5	23.0	14.9	21.3	16.6	12.8
Nooks	6.5	22.5	0.0	13.8	15.8	20.9	38.2	0.0	0.0	0.0
Bombora	9.1	8.5	15.0	9.6	15.6	20.4	22.5	20.6	15.1	14.5
11x	25.4	27.9	0.0	7.9	7.8	19.5	9.1	0.0	0.0	17.5
UserGems	5.8	43.3	0.0	25.8	13.3	21.7	2.1	0.0	0.0	16.2
Artisan	6.4	20.6	0.0	9.2	17.2	14.1	0.0	0.0	0.0	7.9
Demandbase	7.1	9.0	12.9	11.4	15.1	17.6	11.2	18.7	11.1	10.1
Avoma	7.9	7.1	5.9	14.5	15.4	20.9	7.0	5.4	7.3	11.6
Clearbit	5.0	5.9	12.7	13.4	10.0	17.5	26.2	22.2	16.4	17.4
CommonRoom	0.0	25.4	0.0	12.9	14.6	21.4	4.1	3.5	0.0	12.9
Attio	2.9	60.4	0.0	5.4	4.6	4.0	8.6	0.0	0.0	0.0
Regie.ai	4.6	8.7	12.9	14.6	5.3	20.8	19.6	3.6	6.3	8.3

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

Apollo wins every metric we measured. The outbound duopoly era is over.

Apollo is mentioned in 645 of 1,000 AI responses. It earns 96 first-pick designations, more than any other brand. Sentiment runs at 0.41, also the niche high. Among brands with this much visibility it holds a 21% recommendation rate, meaning **134 of those mentions are active endorsements**. By raw count it is already ahead of Salesloft (528) and Outreach (477). By weighted visibility (factoring in which AI surfaces buyers actually use) Apollo's lead is decisive: 205.7 points to Salesloft's 138.7. The cause is structural. Apollo bundled data, sequencing, and analytics at startup-friendly pricing while the legacy sequencers stayed enterprise-priced. AI search reflects modern buying patterns more than installed-base inertia.

DIVERGENCE #2

ZoomInfo's owned content is cited more than any media property.

pipeline.zoominfo.com is the #1 most-cited source on both Claude (75 citations) and Perplexity (28) for outbound sales queries. Google AI Mode cites it 50 times, Google AI Overview 35. Total: 188 citations across four AI search engines, more than TechRadar, more than Wikipedia, more than Reddit, more than any third-party outlet. ZoomInfo's editorial property has become reference material for AI search in this category. For B2B marketing teams the takeaway is direct. Long-form, comparison-rich, well-structured content gets pulled into AI answers the same way any other authoritative source does. Owned media that helps the buyer reason about the category gets used.

DIVERGENCE #3

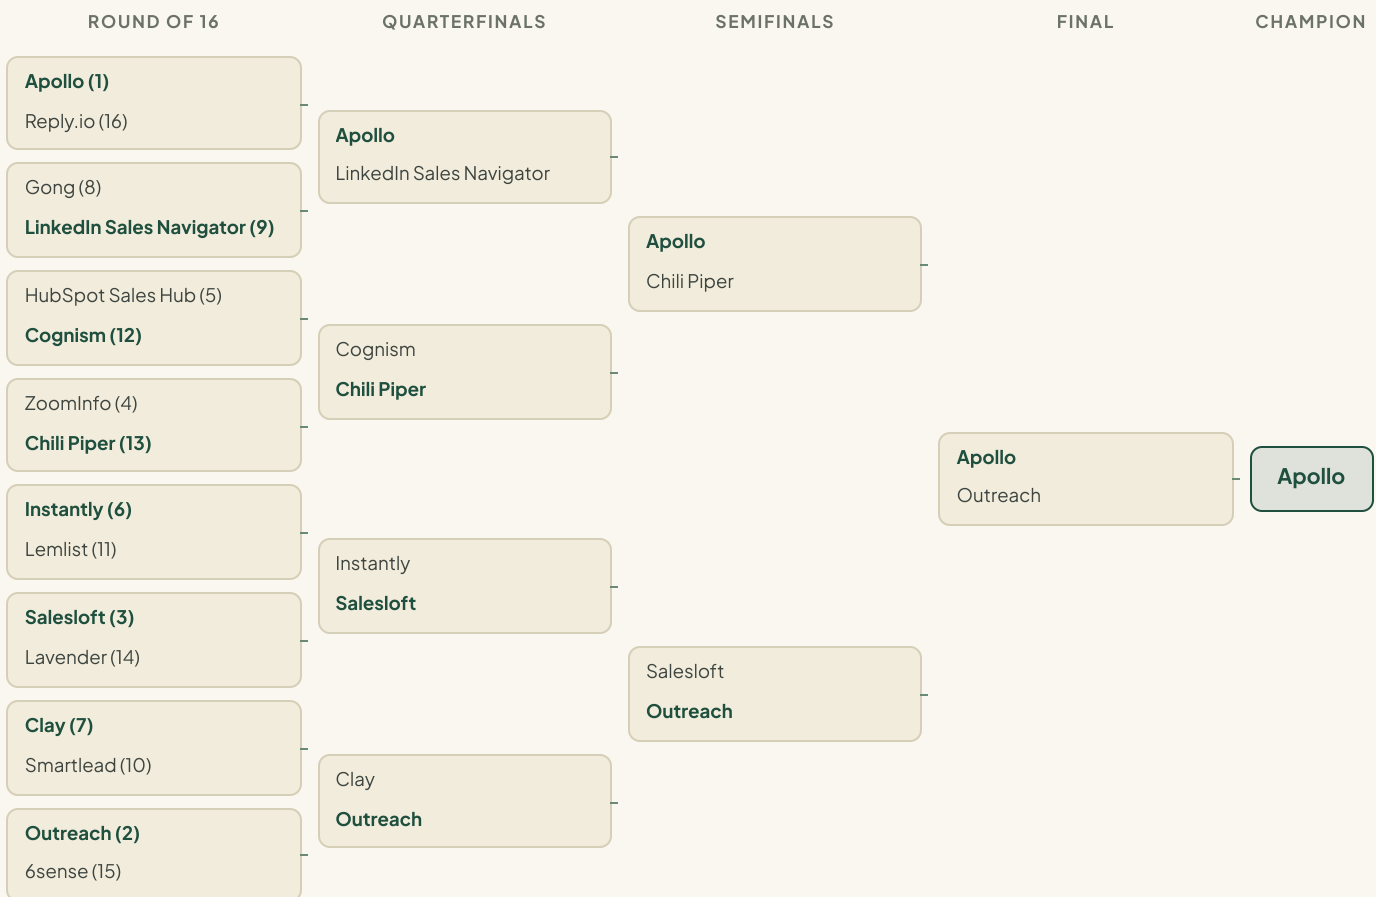
Lavender hits 32% recommendation rate. Mention count and intent are not the same metric.

When Lavender (an AI email assistant) is mentioned, 32% of those mentions are recommendations, the highest rate of any brand in the niche. It earns fewer mentions than the leaders but a far larger share of endorsement when it does appear. Apollo wins on volume (645 mentions at 21% rec rate). Lavender wins on intent density. Reply.io is the inverse case: 194 mentions at 25% rec rate, plenty of listings but weaker endorsement signal. For brands sitting under the leaders, the right metric depends on the goal. Mentions for top-of-funnel awareness. Recommendation rate for whether the AI surface actively closes the consideration loop.

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

Apollo (seed #1) wins the ChatGPT bracket. The brand AI reaches for first in every head-to-head. When the same bracket runs on **Google AIO**, **Apollo** (seed #1) 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.

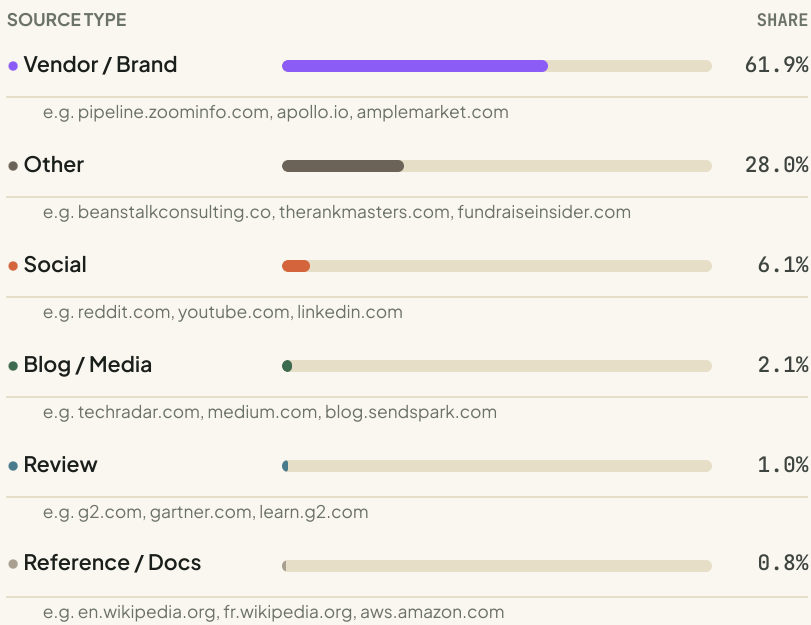
ALSO-RANS
Visible but with mixed perception. Reputation risk in AI.

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 57%

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 3% 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 62% 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

28% 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 2x 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:** "10 Best Outbound Agencies for IT Companies (2026) - 100..."
- Other:** "6sense Salesloft Integration | Partner Network"
- How-to / Guide:** "When to Use Outbound Marketing for SaaS Companies: A St..."
- Comparison:** "6sense vs. Demandbase for Account Based Marketing"
- Product Page:** "The Top B2B Contact Databases According to"

TAKEAWAY


Listicle content dominates at **57.5%**. Brands that produce comparison and listicle content on their own domains have a structural advantage in AI visibility.

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 4.7 brands/question vs AI Overview at 4.6. 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 53.2% vs closed models at 46.7%. They cover 30 brands vs 30. 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 44.6% recommendation rate vs Western at 52.0%. DeepSeek uniquely favors open-source tools, reflecting different training data and cultural priorities.

MODEL DIVERGENCE MAP

Where models disagree most

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 your AI surface visibility against your buyer's actual journey.

URGENT WEEK 1

B2B sales buyers research vendors across ChatGPT (consumer-grade research), Perplexity (technical evaluation), and Google AI Overview (comparison shopping). Each surface ranks brands differently. ZoomInfo's pipeline blog dominates Claude and Perplexity, while Reddit dominates ChatGPT. Map your buyer's actual research workflow and benchmark visibility on those specific surfaces instead of a generic SEO dashboard.

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

HIGH IMPACT QUARTER

ZoomInfo's pipeline.zoominfo.com is the single most-cited source in this niche, ahead of Wikipedia, TechRadar, and Reddit. They got there through sustained editorial investment in long-form, comparison-rich content tied to clear use cases. AI search treats authoritative vendor content as canonical when third-party coverage is thin. The window is open. Produce the definitive content for your category, structure it for machine readability, and AI search will pull from you.

03 If your tool launched after 2023, current web presence is your only path to AI visibility.

EMERGING BRANDS ONGOING

Training-only models (Mistral, Qwen, Grok, DeepSeek) account for between 35% and 65% of any brand's visibility share. Newer brands like Nooks, Cognism, and Clay get most of their visibility from web-search-enabled models because the training-only ones barely know they exist. If you launched recently, every visibility point comes from current web content: third-party comparisons, owned editorial, listicle inclusion, structured product pages. There is no historical corpus to fall back on.

04 Recommendation rate matters more than mention count for category fit.

MEDIUM IMPACT QUARTER

Lavender (59 mentions, 32% rec rate) is better positioned to convert than Reply.io (194 mentions, 25%). Apollo (645 mentions, 21%) wins on raw volume. The metric that matters depends on your goal: visibility for awareness, recommendation rate for consideration. The bottom eight by sentiment (Avoma, Lusha, Regie.ai, Clearbit, Reply.io, Chili Piper, Cognism, Bombora) get listed in AI answers but rarely endorsed. Being mentioned without being recommended does not move pipeline.

05 Inverse recency bias keeps newer entrants invisible. Plan around it.

STRATEGIC QUARTER+

Reo.dev is a fast-growing, funded signal-based-selling platform with active customers and industry presence. It is mentioned zero times across the 1,000 AI responses we collected. The cause is inverse recency bias. LLMs over-index on content that existed and got linked to during their training window, and even retrieval-augmented surfaces lean on those older anchor sources. The fix for newer entrants is coordinated content density: third-party comparisons, owned editorial that benchmarks against established names, listicle inclusion, and structured product pages crawlers can parse. Track citations from each major AI surface monthly and treat invisibility as a measurable, fixable problem with a quarterly cadence.

Track this in real-time.

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

METHODOLOGY

How we measure this



Prompt Design. Prompts simulate real-world purchase research in Outbound Sales Stack: 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