

AI BRAND SENTIMENT · BRAND ACCURACY · AI VISIBILITY · B2B STRATEGY

AI brand sentiment: how it forms and what shapes it

AI can describe your brand with every fact correct and still frame it as the fallback option. Brand sentiment and brand accuracy are separate problems with different causes and different fixes. This article explains how AI forms a view of your brand's tone, what drives cross-platform variation, and what interventions are actually plausible.

AIVIARA RESEARCH · MAY 2026

AI can describe your brand with every fact correct and still make buyers hesitate.

That's the version of the problem most monitoring checklists miss. A brand can have accurate pricing, the right product description, and current feature information in an AI response and still receive language that positions it as the fallback option, the "powerful but has a learning curve" choice, the vendor you'd consider if your first preference doesn't work out.

AI brand sentiment and AI brand accuracy are separate dimensions. They share some causes and they interact, but they require different diagnostics and different responses. Getting them confused leads to investing in the wrong work.

How AI forms a view of your brand

When an AI model generates a response mentioning your brand, it isn't retrieving a stored opinion. It's synthesising a pattern from the sources it encountered during training and, for platforms with live retrieval, from the current web.

Snezzi's Gautham Seshadri describes the core process in terms of how LLMs retrieve candidate pages, compare sources for consensus, and weight domain trustworthiness. Where traditional search matches keywords to pages, an LLM "examines patterns and seeks consensus across multiple trusted sources to form an opinion." The facts it surfaces and the framing it applies both reflect that same aggregation process. If most sources discussing your brand use cautious language ("popular but expensive," "powerful but has a learning curve"), the model learns to apply that framing. Not because it decided the framing is deserved, but because the majority signal is cautious.

This is the same mechanism that drives factual errors (covered in the [companion article on AI brand accuracy](#)), but applied to tone rather than facts. A consistent signal about your brand's complexity or pricing concerns in the sources AI draws from will show up in AI responses about your brand, regardless of what your own website says.

Visible, an AI brand monitoring service, categorises the signal sources that drive this framing into four types: training data accumulated before deployment, real-time retrieval via RAG for platforms that use it, structured data signals, and third-party mentions. For many brands, third-party community content is one of the most visible shaping forces: review platforms, forums, and discussion threads where buyers and former customers describe their experiences in their own words.

A separate mechanism operates alongside the citation-source effects. A 2025 study in PNAS Nexus (Wang, Eshghi, Ding, and Gopal) found that when AI tools rephrase human-written content, they consistently produce more neutral emotional tone regardless of model or instruction. The researchers analysed 50,000 tweets and Amazon product reviews, finding that AI-rephrased versions showed significantly reduced emotional intensity compared to the originals. This study is about how LLMs rephrase source content generally, not about brand sentiment formation specifically, but it suggests that even when your brand has strongly positive source material, the synthesis process may moderate the warmth. Whether that creates a structural ceiling on AI enthusiasm for a brand is an inference the data doesn't directly confirm.

Reddit, G2, and Trustpilot are doing more than you think

Three platform categories shape AI brand sentiment more than almost anything else a brand can influence directly.

Reddit's citation rate across major AI platforms is 40.1%, ahead of Wikipedia at 26.3%, according to Liam Dunne's analysis at Discovered Labs (February 2026). Reddit's upvote-based quality signals, long-form community discussions, and Q&A; structure match how LLMs synthesise information. And critically, AI doesn't filter for positive Reddit discussions.

Positive and negative brand mentions are cited at nearly the same rate: **5%** for positive, **6.1%** for negative. AI does not filter out the criticism.

LIAM DUNNE, DISCOVERED LABS · FEBRUARY 2026

The average cited Reddit post is approximately a year old, and 4% of cited posts date from 2019 or earlier, which means legacy community sentiment (good or bad) persists in AI responses long after circumstances change.

For B2B software brands, G2 can carry significant weight. G2's visibility in AI-generated answers grew from a 6.3% visibility score in August 2025 to 14.9% by October 2025, ahead of Microsoft, HubSpot, Salesforce, and Google in that measure. Research cited by G2 found the platform accounts for roughly 33% of review-site citations in ChatGPT and Google AI Overviews, and approximately 75% in Perplexity. Kevin Indig's analysis of 30,000 citations found that a 10% increase in G2 reviews correlates with a 2% increase in citations. The relationship is real but not dominant. Reviews explain only around 2% of variance, while brand authority and content quality carry more weight. Still, G2 listings are one of the clearest review-based signals AI systems may draw on when characterising customer sentiment for B2B software brands.

Trustpilot follows similar logic. Peec AI's analysis found that HubSpot's Trustpilot profile had 46% one-star ratings against 33% five-star, and documented that this bimodal distribution directly influenced how AI platforms characterised HubSpot's customer support quality. For Revolut, Peec AI scored the AI-generated characterisations of customer support at 52 out of 100, with pricing scoring 38 and safety scoring 45, each reflecting the negative community sentiment on review platforms about those specific dimensions.

Peec AI found HubSpot's Trustpilot profile had **46%** one-star ratings against **33%** five-star. That distribution directly shaped how AI platforms characterised its customer support quality.

PEEC AI / TOMEK RUDZKI · JANUARY 2026

These platforms are among the most visible third-party reference pools AI systems may draw from when forming a view of your brand. Your own website's messaging doesn't override them. It competes with them, usually from a position of lower aggregate authority.

The same brand reads differently across platforms

AI brand sentiment is not consistent across platforms, and the differences can be substantial.

Michael Brito's case study comparing how ChatGPT, Perplexity, Claude, and Google AI Mode describe the athleisure brand Vuori is the clearest published single-brand cross-platform comparison available. Brito found ChatGPT and Google AI Mode applied "strongly positive framing" with language including "premium," "versatile," and "sustainable" without qualification. Perplexity and Claude were more restrained. Claude, specifically, structured its response into "Strong Points" and "Common Criticisms" with roughly equal weight, explicitly listing quality decline, sizing issues, and customer service problems. Same brand, same facts, measurably different tone depending on which system generated the response. (This is a single-brand case study, not a controlled experiment across categories.)

Onely's Bartosz Góralewicz notes platform-level tendencies from practitioner observation. GPT-4 tends toward neutral-to-negative framing; Claude toward cautious and nuanced characterisation; Perplexity toward citation-led, fact-checking framing; GPT-3.5 toward a positive bias. These are practitioner observations, not findings from a controlled study, and model versions evolve. Onely's practitioners observe this variation consistently across their client work, though the cause at the model level remains unconfirmed.

A DerivateX study of 50 B2B SaaS companies (reported in Demand Gen Report, April 2026) points in a different direction. The researchers ran 1,400 buyer-intent prompts across ChatGPT, Perplexity, Claude, and Gemini. According to the Demand Gen Report coverage, the large majority of the companies scored at or near the top of the study's sentiment scale. Their conclusion was that the visibility gap for that sample was driven by mention frequency and platform breadth, not brand perception. When B2B SaaS brands appeared in AI responses at all, they appeared positively.

These findings aren't contradictory. Vuori is a consumer brand with substantial Reddit, Trustpilot, and review coverage, which gives AI enough signal to form a nuanced, differentiated view. The B2B SaaS brands in the DerivateX study likely have less third-party community discussion. Where third-party coverage is sparse, AI has fewer signals to form differentiated sentiment and the response defaults toward mild endorsement or neutral description. More community discussion means more surface area for positive and negative framing alike.

For many B2B SaaS brands, the sentiment problem may not be that AI says something negative. It may be that AI says very little, and what it does say lacks the specificity to differentiate from competitors.

The bias underneath the framing

Independent academic research establishes that AI brand representation isn't neutral at the model level, before any community signal shapes things.

Kamruzzaman, Nguyen, and Kim published "Global is Good, Local is Bad? Understanding Brand Bias in LLMs" (accepted to EMNLP 2024). Across four brand categories, the study found LLMs disproportionately associate multinational corporations with favourable characteristics compared to local competitors. Country-of-origin effects appeared in the data; while global brands received systematic favouritism overall, local brands showed boosted preference in specific contexts. A separate arXiv preprint (not yet peer-reviewed), Rienecker et al.'s "Auditing Preferences for Brands and Cultures in LLMs" (submitted March 2026), analysed over 2,000 questions across 10 topics using three models and found that "U.S.-developed models Gemini and GPT show marked favouritism toward American entities," while DeepSeek, a China-developed model, showed more balanced geographic preferences. The paper connects these patterns explicitly to "real-world economic outcomes" and "market fairness."

Both papers measure preference bias in recommendations (which brand a model favours when asked to compare or recommend), not the descriptive tone in open-ended responses. The inference that systematic recommendation preferences translate into warmer or cooler descriptive language is reasonable, but it's one step beyond what either paper directly shows. The more precise claim is that LLMs have learned systematic preferences from training data that persist across query types, and those preferences weren't designed, they emerged from what was in the corpus.

For brands headquartered outside the U.S., this structural bias is worth knowing about, even if the practical response is the same as for any other AI visibility problem. Build more authoritative coverage in the kinds of sources AI systems frequently retrieve and cite.

Sentiment vs. accuracy: why the distinction changes what you do next

Accuracy problems and sentiment problems look similar in monitoring dashboards but require different responses.

Accuracy problems (wrong pricing, discontinued services, misattributed features) involve AI stating something that can be checked against objective facts and corrected. The fix, as the [guide on correcting AI brand errors](#) covers, is changing the underlying sources where the

wrong facts live: G2 listings, Wikipedia, comparison articles. The question "is this claim true?" has a yes/no answer and a direct remediation path.

Sentiment problems are different. Esteve Castells at LLM Pulse (March 2026) defines brand sentiment in AI as "the qualitative tone and context surrounding brand mentions within responses generated by large language models." A brand with an accurately described product can still receive language that positions it as the complex option, the one that "may be suitable for enterprise teams comfortable with onboarding time." That framing is not wrong. It's just not helpful. And you can't fix it by correcting a fact.

Góralewicz at Onely puts the practical separation plainly. Factual accuracy asks whether the information is correct; sentiment asks what the tone and recommendation framing are. Many companies find their brand gets technically accurate mentions with lukewarm language that positions competitors as the superior choice.

The interventions for sentiment are harder than for accuracy. Changing the qualitative impression AI has formed of your brand, through aggregation of community discussions, review patterns, and press coverage, requires changing the weight and tone of that coverage over time. For most brands, that's months of work, not weeks. And unlike factual corrections, there's no controlled study confirming that specific actions produce specific sentiment shifts in AI outputs. The logical chain (better coverage leads to better source signals, which leads to better AI framing) is sound; the measured effect size isn't established.

Where accuracy failures and sentiment problems co-exist, one compounds the other. Góralewicz notes that incorrect pricing generates a downstream "appears expensive" framing even after the factual error is corrected, because the historical community discussion about pricing has already shaped the pattern AI learned. In those cases, fixing the fact is necessary but not sufficient.

Running branded queries to monitor what AI says about your company surfaces two types of information: whether the facts are right, and whether the framing is working for you. Most teams are doing the first check. Fewer are systematically assessing the second.

What brands can realistically influence

The most plausible intervention points for AI brand sentiment are the same sources that drive AI coverage generally, with a focus on qualitative framing rather than factual presence.

The most defensible intervention is addressing the underlying sources of negative community sentiment. If Trustpilot one-star reviews about customer support are driving "lukewarm" AI characterisations of your support quality, no amount of owned content changes that pattern. The reviews are part of the signal. Addressing the product problems that generate those reviews is the most durable approach.

"Most negative experiences aren't inherent flaws but **fixable product problems**."

TOMEK RUDZKI, PEEC AI · JANUARY 2026

For earned media coverage, the direction of the evidence is consistent even where specific percentages aren't reliable. Authoritative third-party sources carry more weight than owned content in shaping AI's reference base. Industry publication coverage, analyst commentary, and high-authority category sites are the places where AI looks when forming a view of a B2B brand. Building that coverage with language that describes your brand accurately and specifically creates better signals than publishing on your own domain.

Structured content (clear FAQ sections, well-organised headers, explicit answer-first paragraphs) makes it easier for AI to extract accurate factual information about your brand. Whether that also shifts qualitative framing is less clear; the mechanism affects extraction quality, not the emotional valence of what gets extracted.

For platforms with live retrieval (Perplexity, Google AI Overviews, and others), changes in source material can propagate faster than for models relying on fixed training data. For training-data-dependent responses, sentiment changes require changes in what goes into the next training cycle. That's a timeline measured in months, determined by the platform's retraining schedule, not your publication calendar. Any specific timeline claim shorter than that isn't well-evidenced.

AI brand sentiment is more manageable as a long-term discipline than as a remediation project. Brands that monitor consistently, address real product weaknesses that generate negative community discussion, and build a broad body of accurate, authoritative third-party coverage are in a structurally better position. Brands that treat it as a one-time optimisation project are likely to find the underlying signals reassert themselves.

Aiviara is building infrastructure for monitoring AI brand citations and factual accuracy across LLM platforms. Early access information is available at aiviara.com.