

AI BRAND ACCURACY · GCA · B2B STRATEGY · AI SEARCH

What to do when AI gets your brand wrong

You've found an AI error about your brand. This guide covers what to do about it: which sources to fix, in what order, with what realistic expectations across different AI platforms.

AIVIARA RESEARCH · MAY 2026

B2B SaaS brands run into AI brand errors in predictable categories: wrong pricing that persists after a plan restructure, discontinued features still appearing in vendor comparison recommendations, competitor capabilities misattributed to your product, or a company description frozen at a previous positioning. The companion guide on [when AI gets your brand wrong](#) covers detection and why these errors happen.

This article starts after detection. You have a specific error. What do you do?

The central point is structural: correction works by changing what AI reads, not by asking AI to update itself. Everything that follows flows from that. There is no correction form that routes to a training pipeline. There is no database entry to edit. The only available path is changing the signals that AI will encounter (owned properties, third-party profiles, and independent press) and then waiting.

You can't tell the model to fix it

No correction portal exists at any major AI platform for brand-level factual errors. For ChatGPT's base model and Claude, there is no documented direct correction pathway. For those platforms, feedback submissions document that you are unhappy; they do not change what the model says. One practitioner reported to the Am I Cited community that four months of feedback submissions to OpenAI produced only form letter responses and no change in outputs.

Perplexity is a partial exception. Perplexity provides a feedback mechanism, and because it retrieves from the live web on most queries, a correction to source material can propagate in days once you have submitted a report alongside the source fix. That speed reflects its

live-retrieval architecture.

The error type tells you which source to fix first

Before changing anything, identify which sources the AI response is actually drawing from for the specific wrong claim. Run the query that produces the error across ChatGPT, Perplexity, Google AI Overviews, and Gemini. Check what each platform cites.

The error type is usually a reliable indicator of the source category. Wrong pricing that persists across platforms often traces to a stale G2 or Capterra profile, or a comparison article that hasn't been updated since your last pricing change. A wrong founding date or "acquired by" attribution is more likely to trace to Wikidata or Wikipedia, where structured entity data feeds directly into Google's Knowledge Graph and into many model training corpora. A discontinued product still appearing in recommendations (RankScience found exactly this in January 2026, over a year after discontinuing multiple services) typically has a long tail of historical press and blog coverage that new signals haven't yet outweighed.

Knowing the source category tells you which fix is highest priority. Fixing your About page when the error originates in a G2 listing is busy work.

Start with what you control

Owned properties are not the most powerful signal. Third-party sources with high authority outweigh a brand's own website when both are present and say different things. But owned properties are the fastest to update and cost nothing, so they are the right starting point.

The most useful owned-property action is a structured company facts page. It should state the correct information explicitly: founding date, current product offering, current pricing, what your company does and does not do. Add JSON-LD schema markup using the Organisation and Product types. This makes the structured data machine-readable in a format that AI extraction pipelines can parse without ambiguity. For a brand where AI is getting the founding year wrong, a schema-tagged fact block on an owned domain creates a higher-confidence signal than a narrative paragraph in a blog post.

Search Engine Land puts it plainly: "Schema won't guarantee citations, but it helps AI understand entities." That is the right frame. Schema improves how clearly AI can extract your brand's specific claims when it retrieves your page. It does not override errors on high-authority third-party sites, and for platforms like ChatGPT's base model that answer from training data rather than live retrieval, changes are not reflected until the next training

cycle.

A common failure mode to avoid: a structured facts page with schema markup but thin surrounding content. Schema on a page with little substantive material is parsed but carries low credibility. The facts page needs to be a real resource, not a schema container.

Fix the third-party record

This is where fixing the actual source of the error happens. AI systems give substantial weight to established third-party sources: review platforms, structured business data, and comparison sites. The reason Ishtiaque Ahmed at ZipTie.dev puts it plainly: "LLMs don't know what's official — they know what's frequently mentioned." Fixing the sources that AI actually weighs changes what appears most often in its context. Updating your own website when the error lives elsewhere does not.

For each platform where you have an editable profile, make the correction now:

Crunchbase and LinkedIn Company Pages accept direct brand edits. Updates take effect immediately. These are the fastest third-party fixes available and worth doing in the first hour.

G2, Capterra, and TrustRadius allow brands to claim listings and update official product descriptions, pricing ranges, and feature lists. User-submitted reviews cannot be edited by the brand, but the official listing fields can be corrected. Some fields require platform verification, which adds a few days.

Wikidata feeds Google's Knowledge Graph and carries that data forward into multiple AI training corpora. Brands cannot directly edit Wikidata records about themselves, but corrections via talk pages are actioned, and Wikidata updates propagate relatively quickly compared with Wikipedia editorial cycles, often within a few weeks.

Wikipedia is where many teams underestimate the constraints. Wikipedia's editorial policy prohibits organisations from directly editing articles about themselves. Doing so risks having the edit reverted and being flagged for conflict-of-interest editing, which makes future correction requests harder to process. The compliant path is slower: document the factual error with a reliable third-party source, raise the correction on the article's talk page, and request that an independent editor review it. English Wikipedia's 2026 policy also bans LLM-generated article content, which eliminates the obvious shortcut.

One documented community case showed a company incorrectly labelled "acquired by [large company]" across AI responses. The correction required building a Wikipedia page from scratch, developing a Wikidata entry, distributing a press release, and updating the website FAQ. AI responses corrected approximately six weeks later. The Wikipedia and Wikidata work dominated the timeline.

Data inconsistency across platforms is itself a problem that teams often miss. If Wikidata records a founding date of 2018 and Crunchbase says 2019, AI systems assign lower confidence to both data points. Fixing one while leaving the inconsistency elsewhere can leave the AI producing hedged or averaged wrong answers. Update all editable profiles to be consistent with each other, not just internally correct.

If AI is describing your pricing as per-seat when you moved to usage-based billing six months ago, the sequence is G2 and Capterra first, then Crunchbase and LinkedIn, then a pricing page update with FAQ schema, then a direct-statement pricing explainer on a well-indexed domain. Platform timelines run from when those source corrections are live and indexed, not from when you started.

Entrenched errors require signal in sources AI actually reaches

Owned-property updates and profile corrections address the sources you control. When an error is entrenched across many independent sources (old TechCrunch articles, comparison blogs, cached pages from a previous product positioning), those changes are insufficient on their own.

AI weights frequency. According to 5W Public Relations' AI Platform Citation Source Index 2026, which analysed 680 million citations across ChatGPT, Google AI Overviews, Perplexity, Gemini, and Claude, Reddit is the top citation source across all major platforms at approximately 40% of citations. Wikipedia accounts for 26–48% of ChatGPT's top-10 cited domains. The top 15 domains capture 68% of consolidated AI citation share. Your company blog is not in that group.

Reddit accounts for approximately **40%** of citations across all major AI platforms. Wikipedia accounts for **26–48%** of ChatGPT's top-10 cited domains. The top **15 domains** capture **68%** of consolidated AI citation share.

5W PUBLIC RELATIONS, AI PLATFORM CITATION SOURCE INDEX 2026 — 680 MILLION CITATIONS

Correcting deeply entrenched errors requires building accurate signal in sources that AI actually reaches. For a B2B SaaS company, that means getting the correct facts into industry publications, analyst coverage, and high-authority category sites. Sources where the right information is stated explicitly, not buried in a forward-looking quote or implied by context. A press release that buries the pricing correction in a forward-looking quote is not the same as an article that states the current pricing directly.

One practitioner documented this approach in the Am I Cited community: securing mentions in local business journals, industry podcasts, and guest posts produced AI accuracy improvements in approximately six weeks, at a cost of around \$2,000 in distribution.

Another case showed that TechCrunch coverage was specifically what shifted ChatGPT's description of a company. Neither result is reproducible in a predictable way, but both indicate that credible third-party coverage stating correct facts is the most durable correction path for errors that have spread widely.

For errors entrenched across many independent sources, this is not work a content team can handle through publishing alone. It requires active press engagement or community participation in the platforms AI cites most.

The timeline is not linear, and the correction can reverse

Platform timelines for corrections differ dramatically. A common error is to treat one platform's improvement as general progress.

Perplexity is the fastest. Practitioners in the Am I Cited community reported corrections appearing within 24–48 hours for clear factual errors once sources were updated. That speed reflects its live-retrieval architecture.

Perplexity can reflect source corrections in **24–48 hours**. ChatGPT's base model may take **18–24 months** or longer for the same correction to propagate.

PRACTITIONER-DOCUMENTED ESTIMATES, AM I CITED COMMUNITY AND RANKSCIENCE

ChatGPT Search and Google AI Overviews sit in the middle. These platforms combine training data with selective web retrieval. OptimizeGEO.ai documents typical timelines of one to four weeks for ChatGPT's web-browsing version and two to six weeks for Google AI Overviews after source indexing.

ChatGPT's base model is where corrections stall. Dana Davis at RankScience documented this directly: after discontinuing several services in 2024 and updating all owned properties, ChatGPT, Gemini, and Perplexity were still recommending those discontinued services in January 2026, more than a year later. (The RankScience case is covered in more detail in the [companion article on detecting AI brand errors](#).) RankScience estimates 18–24 months or longer for corrections to propagate across all platforms for brands with that kind of historical footprint.

For B2B SaaS specifically, the categories where this bites hardest are pricing model changes and product deprecations. If your platform ran a per-seat model for three years and then moved to usage-based pricing, three years of comparison articles, review platform write-ups, and press coverage have encoded the old pricing. New sources have to outweigh all of that. A pricing page update and a schema block do not move the needle until the balance of new-to-old coverage shifts. For widely covered products, that ratio takes time.

The more disorienting timeline problem is reversion. Lily Ray documented this mechanism in April 2026: she published a deliberately fictional article about a January 2026 Google core update that never happened, and Google's AI Overviews cited it as fact within 24 hours. The system does not reliably distinguish accurate reporting from repeated AI-generated misinformation. Lily Ray named this the "AI slop loop": one AI-generated article hallucinates a detail, content pipelines scrape and republish it, and more AI-generated sites compound the misinformation. Repetition creates apparent corroboration.

A deliberately fictional article about a Google core update that never happened was cited by Google's AI Overviews as fact [within 24 hours](#) of publication.

LILY RAY, APRIL 2026

Practitioners in the Am I Cited community have reported corrections that appeared to hold reverting weeks later, introduced by AI-generated content that scraped and republished the original error before the fix fully propagated. (The GPTZero finding on AI-generated source encounter rates is covered in the [companion article on detecting AI brand errors](#).) Corrections can disappear again before teams realise the underlying source balance shifted back.

What this approach will not achieve

The ceiling matters.

You cannot retrain a model. The only path to changing what a static-training model says is influencing the sources that go into the next training cycle. There is no guarantee that updated sources will be included, weighted appropriately, or processed within any specific window.

You cannot guarantee a timeline. The ranges cited above are practitioner-documented estimates, not service-level commitments. For companies with large historical digital footprints, even an aggressive multi-channel correction effort can take well over a year to show consistent results across all platforms.

Reversion is a risk even after corrections appear to have propagated. The slop loop dynamic means errors can re-enter AI retrieval through AI-generated content that scraped the original error before the fix spread. Treating brand accuracy as a one-time project rather than a monitoring function leaves the problem unattended.

You cannot directly correct some third-party sources. Some comparison articles are not actively maintained. Some cached pages are not crawled with enough frequency for freshness to matter. Some coverage that established your old positioning simply will not be updated or taken down.

For brands with years of coverage behind a previous positioning (a rebrand, a pricing model change, a significant product pivot), the correction is harder than it is for a brand correcting a single misattributed detail. The accumulated old coverage is not neutral. It actively competes with new content, and for a long time it will win.

The reverse is also true. Brands without deep historical coverage sometimes correct faster than larger, more documented competitors. There is simply less old content competing with the update.

Citation volatility cuts in both directions. According to 5W Public Relations' data, AI citation patterns can shift within weeks, not years. That makes platform-specific fast paths (Perplexity, Google AI Overviews for recently indexed content) viable for corrections that would otherwise wait for a training cycle. It also means errors introduced by a single well-distributed AI-generated article can spread across live-retrieval platforms just as quickly. The same properties that make the problem fast to acquire make it fast to address, for the platforms where fast is possible at all.

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