AI VISIBILITY · MEASUREMENT · GBMF · B2B STRATEGY

The three measures that decide whether your brand makes the AI shortlist

Visibility, alignment, and sentiment are three different problems. The GBMF taxonomy turns them into a shortlist diagnosis: AI Champion, Contender, Wildcard, Pariah, Absent.

AIVIARA RESEARCH · MAY 2026

When a B2B buyer asks an AI assistant for shortlist candidates in their category, three things determine whether your brand appears on it. Whether you are mentioned at all. Whether what the AI says about you is accurate. And whether the response speaks favourably about you. These are different conditions with different remediation paths, and treating them as one number tells you very little about which actual problem to fix.

The companion explainer What is AI visibility? covers why AI-mediated discovery matters for B2B buying. This article goes one layer deeper. It introduces the open measurement specification Aiviara Research publishes for brand representation in AI answers, and the shortlist diagnosis that comes out the other side.

Three measures, not one number

The Generative Brand Mention Framework (Aiviara Research, GBMF Working Paper, 2026) names three separable measures:

Brand-mention visibility (MV). The probability that the brand is mentioned by name when an AI engine answers a category-relevant prompt. Reported on a 0–100 scale, per engine and as a headline figure across the declared engine set. MV is unconditional. Every prompt in the declared set contributes whether or not the brand appears.

Brand-mention alignment (MA). Given a mention, how accurately the description matches the brand's own declared facts. Reported on 0–100. Computed only over the prompt-engine cells in which the brand was actually mentioned, against a version-controlled Brand Profile the brand publishes about itself.

Brand-mention sentiment (MS). Given a mention, whether the response speaks positively, neutrally, or negatively about the brand. Reported as three percentage shares plus a net figure on a −100 to +100 scale.

The conditioning carries weight in practice. A brand mentioned in 8% of category answers has thin data on whether AI describes it accurately or speaks of it favourably; there isn't enough mention volume to read the conditional measures. MV stands on its own; MA and MS depend on it.

How the measures combine into a shortlist diagnosis

Take a mid-market B2B analytics platform. A buyer types a prompt close to "best analytics software for B2B mid-market companies" into an AI assistant. The AI response names six vendors and characterises each. Each vendor ends up somewhere on a sixteen-state grid plus two off-grid positions. The state determines how the buyer reads each candidate.

The framework names four reception roots, set by the sentiment band:

Reception root	Sentiment	What it means for the buyer's shortlist
AI Champion	Positive	On the shortlist with momentum. The AI advocates for the brand by name
AI Contender	Neutral	On the shortlist as a candidate. Present in the consideration set without advocacy
AI Wildcard	Mixed (both positive and negative shares ≥ 25%)	On the shortlist as the bold-but-divisive option. Buyers consider it knowing the tradeoffs
AI Pariah	Negative	Unlikely to make the shortlist. Characterised unfavourably enough that the buyer skips ahead

The reception root is then modified by a position prefix, set by where the brand sits on the visibility-alignment plane. *Misrepresented* (visible but inaccurately described) and *Undiscovered* (accurate but rarely surfaced) carry their own commercial meaning. A *Misrepresented AI Pariah* is visible, described inaccurately, and characterised unfavourably; the negative sentiment often tracks the misrepresentation itself. An *Undiscovered AI Contender* is accurately and neutrally described on the few occasions

the AI surfaces it, but doesn't surface often enough to register.

Two off-grid states sit outside the sixteen on-grid combinations:

AI Absent. MV = 0. No mentioning cells exist. The buyer's AI session never names the brand. Not a candidate at all in the AI answer space.

Insufficient-data case. Mentions exist but too few clear the framework's reportability floor to read MA or MS reliably. The brand has presence but no measurable standing on the conditional axes. Reported by readable metrics only, with the shortfall flagged.

The taxonomy is a presentation convention derived from the three scores, not a separately validated construct. What gives it commercial meaning is the alignment between the state diagnosis and the shortlist outcome.

Why three measures must be separable

The strongest argument for measuring visibility, alignment, and sentiment as separate axes is that they describe different commercial problems and call for different responses.

A brand in *Visible & Misaligned* (high MV, low MA) has the inverse of an exposure problem. The AI surfaces it often and gets the facts wrong. Adding more visibility against this baseline amplifies the misrepresentation, because AI is now wrong about the brand at a higher volume. The remediation is alignment work first.

A brand in *Aligned but Unseen* (high MA, low MV) has the foundation right but isn't reaching the answer space. Investing in brand-perception work or fact correction here solves the wrong problem; the few mentions the AI does produce are already accurate. The remediation is reach.

A brand in *Visible & Aligned* with a negative MS net (an AI Pariah in the visible-and-aligned quadrant) is being described correctly and unfavourably. Generic positive PR does not move the sentiment signal; the remediation requires understanding which sources AI is drawing on for the negative framing and addressing those specifically.

A single composite "AI visibility score" collapses these into one number, and that number tells a brand almost nothing about which intervention to invest in. The framework's central commitment is the separation; the state taxonomy is the visible diagnosis that comes out of it.

What this framework does not measure

GBMF measures the answer space, not the buyer's downstream decision. The state tells you what the AI says about your brand to category-relevant queries; what an individual buyer does with that answer is a separate research question. The framework explicitly disclaims downstream-outcome validation and does not attempt to predict purchase intent, click-through, or revenue.

The measurement is conditional on a declared prompt set and engine set. Two reports with different prompts or different engines are measuring different things, even when both report a number labelled MV. Reports without a conditions block are not interpretable or comparable to other reports.

The framework measures mention behaviour. A brand cited only as a source URL footnote in an AI response, without being named in the response text, is not counted as a mention. That is a deliberate boundary; the framework asks whether the brand exists in the answer space the buyer reads, not whether its content was used as a source.

GBMF is presented as a proposed open specification. Empirical validation against a held-out brand sample is part of a separate proposed validation study, and the paper sets out the falsifiability thresholds in advance.

What changes if you measure this way

The diagnostic payoff is that a brand can identify which of the eighteen states it actually holds, per engine, and let that diagnosis tell it which remediation to invest in.

For a B2B SaaS team running their first audit, the practical effect is to stop debating whether the issue is "AI visibility" in general and start working on the specific problem the diagnosis surfaces. An *Undiscovered AI Contender* needs reach. A *Misrepresented AI Pariah* needs facts fixed and the resulting sentiment investigated. An *AI Absent* needs the corroboration density that AI engines weight when assembling answers in the first place.

The how-to article in this series walks through running a first GBMF measurement: prompt set, engine set, Brand Profile, evaluator approach, and reading the output. The companion article on remediation by state covers what each off-target position calls for and routes to the existing tactical guides on engine-specific exposure, alignment correction, and sentiment work.

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