

---

Digital Article / Brand Management

# Forget What You Know About Search. Optimize Your Brand for LLMs.

How to sell your product when generative AI becomes the consumer.  
by *David Dubois, John Dawson, and Akansh Jaiswal*

Published on HBR.org / June 4, 2025; updated June 5, 2025 / Reprint [H08QC6](#)



Grant Faint/Getty Images

**Over the past year, consumers have migrated** en masse from traditional search engines to Gen AI platforms including ChatGPT, Gemini, DeepSeek, and Perplexity. In a [survey](#) of 12,000 consumers, 58% (vs. only 25% in 2023) reported having turned to Gen AI tools for product/service recommendations. [Another study](#) reported a 1,300% surge in AI search referrals to U.S. retail sites during the 2024 holiday season.

Consumers who use Large Language Models (LLMs) to discover, plan and buy are on average younger, wealthier, and more educated. Their customer journey no longer begins with a search query or a visit to your website—it starts with a dialogue. Consumers are asking AI assistants questions like “What’s the best coffee machine under \$200?” or “Plan me a weekend getaway that won’t break the bank.”

For brand leaders, the implications cannot be overstated. Your digital strategy must now include optimizing for AI recommendation engines, not just search algorithms. In short, you must boost LLMs’ awareness of your brand.

### **The Rise of “Share of Model”**

To date, measuring awareness meant assessing consumers’ attention—either offline through recall surveys (e.g., “Which brands come to mind when you think of running shoes?”) or online, through search or social media volumes, manifesting private intent or popularity.

But the growing role of LLMs as an intermediary between consumers and brands demands that marketers consider another kind of awareness: how often, how prominently, and how favorably a brand is surfaced by LLMs to consumers. We call this awareness Share of Model (SOM). Think of it as the AI-era’s offshoot of share of search (“How much do people search for my brand via-a-vis competitors?”) and share of voice (“How much do people talk about my brand vis-a-vis competitors?”). SOM uniquely emulates LLMs’ perceptions and recommendations given a prompt, rather than reflecting human intent (SOS) or available content (SOV).

## The Differences Among “Share of” Concepts

A brand's share of model details how often, how prominently, and how favorably it is surfaced by large language models (LLMs), compared to other brands. It's like the AI version of share of search (How much do people search for my brand, compared to competitors?) and share of voice (How much do people talk about my brand, compared to competitors?).

	Share of model (SOM)	Share of voice (SOV)	Share of search (SOS)
	% of brand mentions in AI-generated content and responses	% of brand mentions in media and social platforms	% of brand-related search queries online
<b>Focus</b>	LLM recommendations for and perceptions about a brand	Public perceptions and social prominence for a brand	Private consumer interest for a brand
<b>Brand control</b>	Lower (dependent on AI algorithms and model-training data)	Medium (through earned and paid media mix)	Higher (through search engine optimization [SEO] and search engine marketing [SEM] strategy)
<b>Importance</b>	Growing (along with AI adoption for consumer purchases)	Stagnating (but increasingly AI-influenced)	Shrinking (shift to AI-driven discovery)
<b>Examples of channels</b>	ChatGPT, Gemini	Social media, news	Google, Bing
<b>Metrics</b>	Mention rate, average position, brand and category perception, model weight	Social and other media coverage, engagement rate, sentiment	Search volume, keyword ranking

Source: David Dubois, John Dawson, and Akansh Jaiswal



See more HBR charts in [Data & Visuals](#) on HBR.org.

Two of the coauthors' marketing agency, Jellyfish, has pioneered a methodology to measure SOM through prompting at scale. Building on this approach, we offer a new three-prong lens to unpack what and how AI "think" about brands: *mention rate*, which tracks how often a brand is mentioned by a specific LLM; *human-AI awareness gap*, which measures the disparity in brand awareness when surveying people vs surveying LLMs; and *brand and category sentiment*, which breaks down LLMs rationale for recommendations into associated strengths and weaknesses.

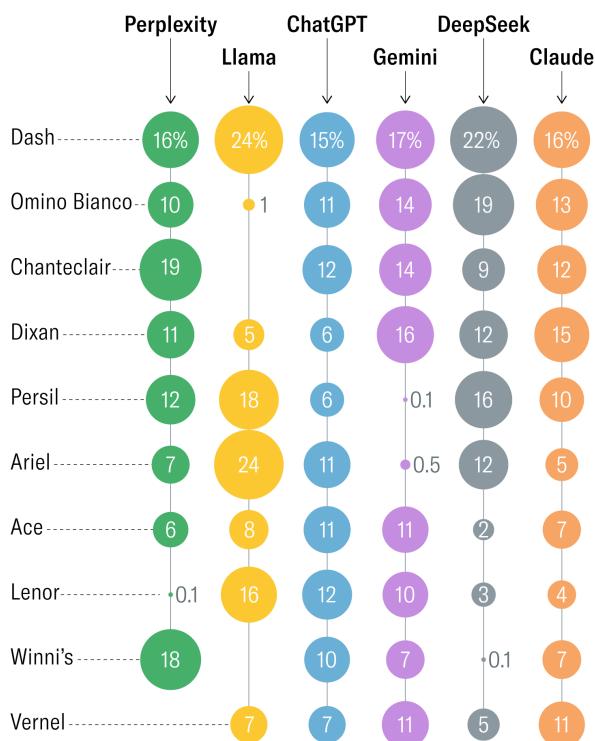
Take for example the laundry detergent market in Italy. We analyzed the top brands' mention rate among six LLMs using Jellyfish's proprietary Share of Model platform. Two observations stand out. First, brands' SOM varies significantly across the models, reflecting differences in how LLMs process brand information. For instance, Ariel's SOM ranges from almost 24% on Llama to less than 1% on Gemini.

Second, some brands are totally absent from at least one model. For instance, while Chanteclair enjoys a 19% SOM on Perplexity, it is missing from Llama. Clearly, LLMs either feature brands or not, unlike search engines or social media where brands that don't excite the algorithm are still represented, albeit less prominently. Failure to register on an LLM means a brand doesn't appear at all before consumers. On ChatGPT, unlike Google, there is no "page two."

## How Often LLMs Mentioned Different Laundry Brands in Italy

One key component of the “share of model” analysis is how often a brand is mentioned by large language models (LLMs), relative to other brands—in other words, their “share” of mentions.

Share of mentions of Italian laundry brands, by major LLM  
April 2025



Source: Jellyfish Share of Model



See more HBR charts in [Data & Visuals](#) on HBR.org.

### Probing the human-AI brand awareness gap

Importantly, a brand’s visibility on LLMs can differ significantly from its market share or other awareness metrics. Therefore, brand managers’ first task is to probe the link between human awareness (e.g., through SOS or SOM) and LLM awareness of their brands.

Quick note: Although a brand's SOM often varies across LLMs as we show above, the next examples in this article focus on brands' SOM across LLMs for ease of discussion. We'll outline the implications of SOM variability across LLMs later.

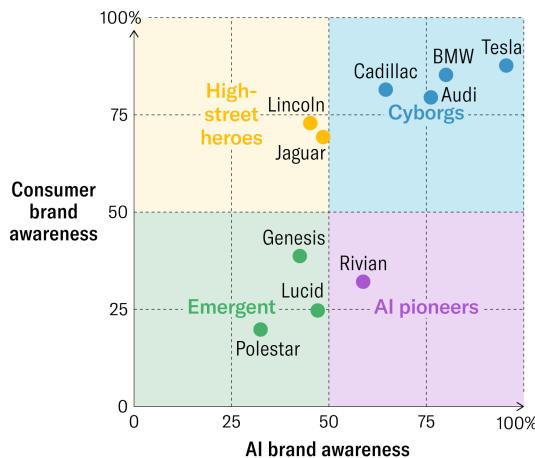
Consider our analysis of U.S. automobile brands' visibility in general and on LLMs during the first half of 2024. We constructed a Human-AI Awareness Matrix (Figure 2) that reflects brand awareness on LLMs, assessed through Jellyfish's tool, and in general, assessed by YouGov market research. Brands fall into 4 distinct categories:

## Human vs. AI Awareness of Popular Car Brands

High brand recognition among humans doesn't always translate to high AI awareness.

Quadrant analysis of brand positioning (consumer awareness) for select car brands in the U.S.

January–July 2024



### ● Cyborgs

Cyborgs rank highly both on traditional search metrics and with LLMs. Tesla, for example, has a strong media presence (a traditional search metric) and focuses on the cars' specific features, both of which elevate the brand with LLMs.

### ● AI pioneers

AI pioneers aren't well-known but are frequently surfaced by LLMs. Often, they are AI-native and fill a specific niche in the market. Rivian, for example, is favored by LLMs for being resolution-focused.

### ● High-street heroes

High-street heroes are well-known brands, but they aren't surfaced often by LLMs. Lincoln, for example, is likely not favored by LLMs because of its focus on intangible qualities such as elegance.

### ● Emergent

Emergent brands have low awareness in both the marketplace and LLMs. Polestar, for example, has a lack of digital footprint and appeal for LLMs.

Source: Jellyfish Share of Model, YouGov



See more HBR charts in [Data & Visuals](#) on HBR.org.

1. **Cyborgs:** These brands have top awareness in both traditional measures (e.g., surveys, search ranking, share of voice) as well as among LLMs. Take Tesla's position in this chart, for example. Elon Musk's ubiquity helps make consumers highly aware of the brand. Tesla also scores well among LLMs because of the brand's emphasis on its specific features. Its new digital advertising strategy attempts to rise the company's scores even higher among both people and large language models.
2. **AI Pioneers:** These brands are well-represented on LLMs but lack marketplace awareness. Often, they are AI-native brands or emerging digital players that are niche in broader digital spaces. Rivian's spot in this quadrant likely stems from its resolution-focused (which we'll touch on later) content strategy, which aligns with its positioning as a solution creator.
3. **High-Street Heroes:** These are established brands with high marketplace awareness but underrepresented or missing in AI-generated content. Case in point: Lincoln, which Frank Lloyd Wright famously said makes "the most beautiful car in the world." This is likely due to the brand's focus on intangible attributes such as elegance or heritage, which are less prized by LLMs.
4. **Emergent:** These brands struggle with low awareness in both the marketplace and among LLMs. They risk falling into digital irrelevance as AI-driven search becomes the norm. Despite its premium positioning, Polestar struggles in our analysis to achieve visibility across the spectrum, reflecting a lack of scaled digital footprint or lack of appeal for LLMs' processing style.

The main takeaway? Marketers need to come up with strategies designed to push their brands up the "consciousness" of LLMs. These strategies are likely to be very different from those designed to appeal to humans. For what we know about LLMs is this: LLMs are not optimizing for attention; they are optimizing for *resolution*. Identifying the "job to

be done" thus becomes the number one priority for brand leaders if they want to score big on SOM.

### **How to increase brand awareness on LLMs**

Our analyses across product categories reveals how models' perceptions of different categories presents specific opportunities to brands in those industries. This has implications for not only what content to produce (across text, image, and video), but also where brands may seek to distribute their messages (website, media, expert, or community contexts). LLMs are looking beyond keywords, focusing on concepts and relationships which create new ways to build brand awareness for LLMs.

Brands should create content that explains not just what the product is, but how it relates to broader contexts, use cases and user needs.

For example, instead of proclaiming "we sell superb running shoes", go for "our carbon-plated midsole design improves performance for long-distance runners."

Brands should also highlight proof of expertise. A skincare brand that references dermatologist-backed studies or links to PubMed research is likely to outshine competitors that don't. Brands that "narrowcast" about pain points—needs, questions and tasks—are more likely to be surfaced. Brands that simply broadcast may be left out.

This could explain why traditional car brands like Lincoln, which push aspirational and marketing-heavy content, are less salient to LLMs compared to Tesla or Rivian, which emphasize functions and features including battery life, tech stack, and software.

Similarly, although they dominate SOV, fast-fashion brands such as Shein lags in AI awareness due to an overwhelming volume of undifferentiated content and lack of trust signals such as reviews and certifications.

In contrast, the Ordinary brand of skincare products offers highly structured product pages with ingredient explanations, transparent science-backed content (explains the “how” and “why” of why a face cream works). Nike and its customer-generated content (runners’ blogs, Reddit, Strava), detailed product pages with clear use cases (e.g., “best shoes for marathon training”) and integrated app ecosystems (Nike Run Club, Nike Training Club). Both brands topped their respective category in our analyses.

Notably, legacy brands can also thrive in the age of AI—if they invest strategically in relevance, representation, and structured digital storytelling. Case in point: Cadillac. The century-old automobile brand scores highly in both human and AI brand awareness. Campaigns like “Audacity” and “The Daring 25” as well as international partnerships helped increase its AI visibility.

### **Gauging LLM sentiment**

Beyond looking into AI brand awareness and how it relates to other awareness metrics—marketers can also explore brand and category sentiment through sentiment (positivity) and semantics (associated terms). This helps them answer questions such as: What are my brand’s perceived strengths and weaknesses? How can I change how LLMs perceive my brand?

For example, our analysis of the travel industry in the U.S. shows that LLMs value characteristics such as convenience, variety, and space, with Booking taking the overall top spot across models. We also surfaced brands’ strengths and weaknesses relative to their competitors. Vrbo, for instance, scores much higher than Booking on privacy and uniqueness—strengths it could exploit to optimize AI awareness.

## What LLMs Value in the U.S. Travel Industry

Unlike traditional search engines, large language models (LLMs) tend to value utility characteristics—in other words, how a brand's products might solve a problem or get a job done.

**LLM relative importance of attributes when consumers were searching for products or services related to travel in the U.S.**

*January-April 2025*



Source: Jellyfish Share of Model



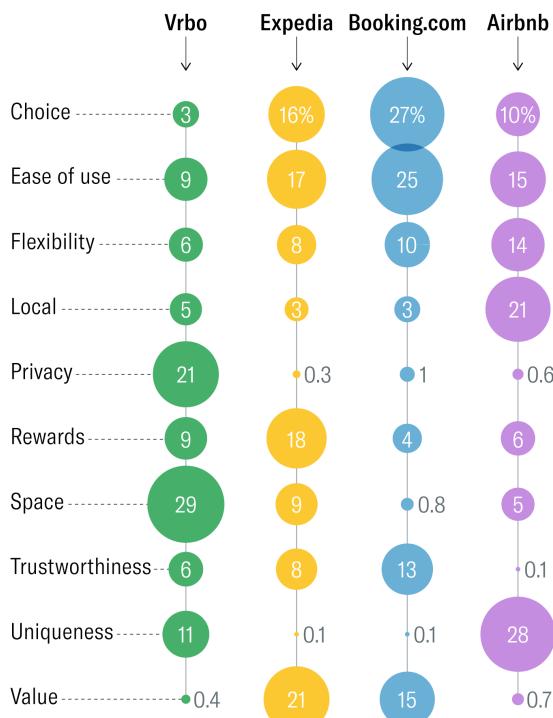
See more HBR charts in [Data & Visuals](#) on HBR.org.

## How Popular U.S. Travel Sites Rank on LLM Values

Knowing what large language models (LLMs) value most and how well your company conveys these values can help you successfully market to LLMs.

Share of each attribute, by travel site

January-April 2025



Source: Jellyfish Share of Model



See more HBR charts in [Data & Visuals](#) on HBR.org.

## How to Market to LLMs

Armed with insights on LLM sentiment, marketers may deploy several approaches to optimize their brand's AI visibility.

First, adopt a multi-pronged media strategy that covers text, images, videos, and structured data (e.g., tables, lists, reviews). Content that clearly links brands' offerings to broader contexts, use cases or consumer needs (e.g., "best EVs for winter driving" rather than just "electric SUV"), generates strong conceptual associations in LLMs.

Brands should also lead semantic niches—specific clusters of meaning where their products naturally fit (e.g., like the Ordinary brands with skincare science).

Importantly, just as each social media platform has its own "rules of engagement"—what works on TikTok probably won't fly on LinkedIn—each LLM applies its unique algorithmic lens to content.

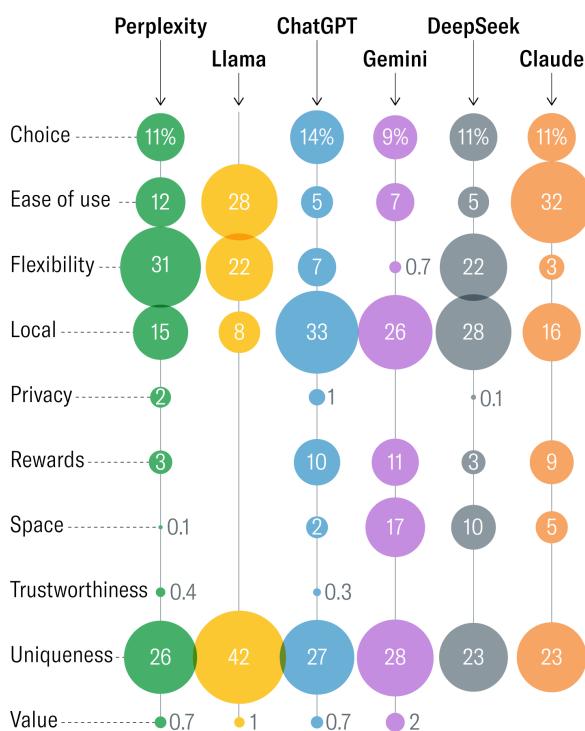
Take the U.S. travel industry again, focusing on LLM's perception of Airbnb. While Llama focuses on the uniqueness of a brand's offerings, ChatGPT focuses on the extent to which brands offer local options, whereas Perplexity seems to value flexibility most.

This ties in with our point earlier about brands' varying visibility across LLMs. We recommend that marketers tailor content to the LLMs whose processing style best amplifies their brand's content and narrative strengths, even as they apply overarching rules (e.g. solution-oriented messaging) across models. It is a fine balance: Tailoring content to the nuances of a dominant model can drive visibility but spreading efforts too thinly across all LLMs risks diluting impact.

## LLMs Weigh Positive Attributes Differently

While popular large language models (LLMs) each value many of the same attributes for a single brand (e.g., here, Airbnb), the LLMs differ in how much relative weight they give each attribute.

Share of each attribute for Airbnb, by major LLM  
January–April 2025



Source: Jellyfish Share of Model



See more HBR charts in [Data & Visuals](#) on HBR.org.

The shift away from traditional search engines is not just a technological evolution. It's a fundamental change in consumer behavior that demands corresponding shifts in marketing: from persuasion to precision, from keyword to advice, from market share to problem-share. Do it right, and brands can establish themselves as essential

participants in the algorithmic conversations that increasingly shape consumer decisions.

*Editors' note (June 5, 2025): A previous version of this article cited Meta instead of the company's AI tool, Llama, when discussing the laundry detergent Chanteclair.*

*This article was originally published online on June 4, 2025.*



**David Dubois** is an associate professor at INSEAD.



**John Dawson** is Vice President of Strategy at Jellyfish, where he advises clients on advertising strategy and AI marketing implementation. He has experience working with international and local brands across Australia, the Nordics, China, and the U.S. John holds an MBA from the University of Technology Sydney and a B.A. from the University of Sydney.



**Akansh Jaiswal** is Senior Manager, Investment Analytics at Jellyfish, where he architectures AI and analytics solutions that optimize global media campaigns and quantify brand perception in large language models. He drives data-led growth strategies through marketing analytics, automation, and predictive modeling. Akansh holds an M.S. in Data Analytics from Rutgers University and a B.Tech. in Computer Science from Symbiosis International University.