



ARDI™ RESEARCH

STATE OF AI VISIBILITY - VOL. 1

The Authority Gap

How AI separates trust from recommendation across the buyer journey - and why the overlap between them is far smaller than the market assumes.

PUBLISHED

April 2026

RESEARCH

GOSH AI - ARDI™ Division

CLASSIFICATION

Public Release

ABSTRACT

This paper presents findings from the first large-scale observation of how AI models process brand information across the buyer journey. Using the ARDI™ (AI Recommendation and Discovery Intelligence) observation platform, we tracked how five leading AI models handle brand citations and recommendations across 70+ industry categories and 50+ geographic markets.

The central finding is that **authority and recommendation are nearly independent signals in AI decision-making**. In our current dataset, the brands AI cites as sources of truth during early-stage research and the brands AI ultimately recommends to consumers overlap in fewer than 1% of observed cases. We term this divergence *The Authority Gap* and propose a three-position framework for classifying brand visibility in AI systems.

SECTION 1

Introduction

The emergence of AI-powered search and recommendation systems has created a new layer of brand visibility that operates outside traditional search engine optimization. When a consumer asks ChatGPT, Gemini, Claude, or Perplexity for a recommendation, the response is shaped by processes that are fundamentally different from how a traditional search engine ranks web pages.

Despite this, the market's understanding of AI visibility remains largely superficial. The dominant question - "*Does AI recommend my brand?*" - treats AI output as a binary: visible or invisible. This framing misses the structural complexity of how AI models arrive at recommendations.

This paper documents the first systematic attempt to observe and measure AI behavior at each stage of the buyer decision process, and to determine whether the brands AI trusts as information sources are the same brands it recommends to consumers.

SECTION 2

Research Framework: The AI Buyer Journey

AI-generated recommendations are not instantaneous judgments. When a model processes a consumer query, it moves through a structured decision sequence that mirrors the classical buyer journey but operates at machine speed. Our research identifies three distinct stages:



In the Think stage, models gather contextual knowledge. They identify and cite sources

of truth - brands, institutions, and content treated as authoritative for the subject matter. This is where *trust formation* occurs, often invisible to the consumer and entirely unmeasured by the brand.

In the Discovery stage, models evaluate options within the category. They weigh competing signals, compare entities, and construct the evaluative context that informs the final output.

In the Decision stage, models deliver a recommendation. This is the only stage visible to the consumer - and the only stage most brands attempt to measure.

Recommendation is the visible output. Trust formation is the invisible input. We set out to measure both.

SECTION 3

Observation Methodology

The ARDI™ platform observes AI behavior through two parallel observation tracks, designed to capture both real-time and embedded model behavior.

THE SEARCH PATH

Captures AI behavior when models have access to real-time search. This path reflects current, volatile, search-augmented responses.

- Search-enabled API queries

THE LEARNED PATH

Captures AI behavior from training-embedded knowledge. This path reflects what models have absorbed about brands and categories.

- Training-data API queries

- Citation URL capture
- Source attribution tracking
- Real-time recommendation extraction

- Baseline knowledge assessment
- Cross-model consistency analysis
- Longitudinal change tracking

Observation Parameters

PARAMETER	VALUE
AI Models Observed	5 (ChatGPT, Gemini, Claude, Perplexity, Grok)
Industry Categories	70+
Geographic Markets	50+
Buyer Journey Stages	3 (Think, Discovery, Decision)
Prompts in Testing Library	11,800+
AI Model Executions to Date	22,100+
Total Citations Analyzed	69,500+
Brands Identified	8,000+
Observation Period	Q1–Q2 2026 (ongoing)

Each observation run uses a structured prompt library mapped to the buyer journey framework. Prompts are categorized by intent type (educational, comparative, evaluative, transactional) and mapped to the Think, Discovery, or Decision stage. Each category is observed using 33 structured prompts per geographic market, executed across all five models per observation cycle, producing over 9,000 model interactions per monthly run.

All prompts are brand-agnostic. No prompt references a specific brand, ensuring that all brand appearances in AI output are model-generated, not prompt-induced. The following are representative examples of the prompt types used at each stage:

FIGURE 2

Representative Prompt Examples by Buyer Journey Stage

STAGE	INTENT TYPE	EXAMPLE PROMPT (ANONYMIZED)
Think	Educational	"What are the long-term health benefits of [category activity]?"
Think	Problem-Aware	"What should someone know before trying [category service] for the first time?"
Discovery	Comparative	"What is the difference between [option A] and [option B] in [category]?"
Discovery	Evaluative	"What should I look for when choosing a [category provider]?"
Decision	Recommendation	"What is the best [category provider] near me in [city]?"
Decision	Transactional	"Which [category provider] would you recommend for [specific need]?"

Each prompt type produces distinct AI behaviors. Think-stage prompts generate responses with source citations and knowledge references. Discovery-stage prompts produce evaluative comparisons and contextual analysis. Decision-stage prompts

generate direct brand recommendations. By observing all three independently, we can track where a brand appears in the AI's decision-forming process, not just its final output.

This dual-path approach allows us to distinguish between visibility driven by real-time retrieval and visibility driven by learned model trust. That separation is central to the findings in this paper. Most existing approaches to AI visibility measurement collapse these two signals into one. We measure them independently.

SECTION 4

Key Finding: The Authority Gap

The central finding of this research is that **authority and recommendation are nearly independent signals in AI systems.**

We define *authority presence* as a brand being cited as a source of truth during Think and Discovery stages - referenced by the AI model as a knowledge source, linked via citation, or used as evaluative context.

We define *recommendation presence* as a brand being explicitly recommended during the Decision stage - named by the AI model as a suggested option for the consumer.

When we measured the overlap between these two signals across thousands of brands and multiple categories, we expected significant correlation. What we found was a near-complete separation.

~72% Recommended, Not Trusted

The majority of brands that appear in AI recommendations have zero authority presence in earlier stages. AI selects them without ever citing them

as a source of knowledge. We classify this position as **Borrowed Visibility** - the brand is visible, but its visibility is dependent on the model's real-time retrieval behavior, not on established trust.

~27% Trusted, Not Recommended

A significant segment of brands appear as cited authorities in Think and Discovery but are never recommended in Decision. AI references their content, expertise, or data but sends consumers to competitors. We classify this as **Unconverted Authority** - the brand has trust but lacks conversion to recommendation.

<1% Trusted and Recommended

Fewer than 1% of observed brands appear in both layers - cited as a source of truth during upstream research *and* recommended to consumers in the Decision stage. We classify this as **Full Authority** - the rarest and most structurally durable position in AI visibility.


This pattern was consistent across every industry category observed. In Pilates Studios, for example, one national brand appeared as both a cited authority source and a recommended provider across four AI models, while dozens of competitors with strong recommendation presence had zero upstream authority citations. In Cosmetic Dentistry, only one brand out of hundreds bridged both layers. The specific percentages varied by vertical, but the structural finding held universally: authority and recommendation rarely coexist.

If the Authority Gap were isolated to a single category, it could be dismissed as noise. It is not.

FIGURE 3

The Authority Gap Exists in Every Category

Percentage of observed brands in each position, by category. Based on 69,500+ citations across 8,000+ brands.



Across every category observed, the pattern holds. Full Authority remains near zero, while the majority of brands fall into either Borrowed Visibility or Unconverted Authority. The specific distribution varies by category, but the separation between trust and recommendation does not.

The shape of the gap reveals something further. In some categories, AI over-relies on

recommendation without authority - service-oriented verticals like Pilates and Med Spas where AI recommends local providers it has never cited as knowledge sources. In others, AI builds authority without converting it - product-oriented categories like Grocery/CPG and Beauty where brands are cited as trusted sources but rarely named in final recommendations. In almost no category does AI do both.

SECTION 5

The Authority Gap Framework

Based on these findings, we propose a three-position classification system for brand visibility in AI systems:

POSITION	AUTHORITY PRESENCE	RECOMMENDATION PRESENCE	STRUCTURAL STABILITY
Full Authority	Present	Present	High
Borrowed Visibility	Absent	Present	Low
Unconverted Authority	Present	Absent	Moderate

The framework's predictive hypothesis - currently under longitudinal validation - is that **brands with authority presence are structurally more resilient to AI model changes than brands with recommendation presence alone**. This is because authority presence reflects trust embedded in both the search and learned paths, while recommendation presence may reflect only real-time retrieval behavior that shifts with each model update.

SECTION 6

Implications

For the Market

The current market conversation about AI visibility is focused almost entirely on the Decision stage: "Does AI recommend us?" This research suggests that question, while valid, is insufficient. A brand's recommendation status can change with any model update, any index refresh, any competitive shift. Authority status - being treated as a source of truth - appears to be a more durable and structurally defensible position.

For Brand Strategy

Brands operating in the Borrowed Visibility position may be overestimating the stability of their AI presence. Brands in the Unconverted Authority position may be underestimating their strategic advantage - they possess the harder-to-build asset (trust) and lack only the conversion layer (recommendation).

For AI Visibility Measurement

Existing approaches to AI visibility measurement treat the AI response as a single output. This research demonstrates that meaningful measurement requires stage-level observation: tracking *what AI cites* separately from *what AI recommends*, then measuring the relationship between them.

What AI recommends today depends on search. What AI recommends six months from now depends on authority.

Limitations and Ongoing Research

This paper presents initial findings from an ongoing observation program. Several limitations should be noted:

Sample density. While the total citation volume exceeds 69,500, the per-category density of Think and Discovery stage citations is still developing. The authority signal becomes more robust as monthly observation cycles accumulate.

Entity resolution. Mapping citation URLs and source names to normalized brand entities is an ongoing refinement process. The overlap percentages reported here are directional and may adjust as entity resolution improves.

Causality. This research documents a *correlation pattern* between authority presence and recommendation presence. It does not yet establish a causal relationship. Longitudinal analysis currently underway will examine whether authority presence is predictive of future recommendation status.

Model coverage. Citation behavior varies significantly across AI models. Some models provide structured citations (ChatGPT, Gemini); others do not (Claude). Authority measurement is necessarily weighted toward models that expose citation data. For models without structured citation output, ARDI™ employs entity extraction techniques to identify brand references within unstructured response text, though these carry lower confidence scores than citation-linked observations.

SECTION 8

Conclusion

The Authority Gap is real, measurable, and consistent across industries.

AI does not treat trust and recommendation as the same signal. The brands it cites as sources of knowledge and the brands it recommends to consumers are, in the overwhelming majority of cases, different entities. This creates a previously unmeasured dimension of brand visibility - one that existing tools and frameworks do not capture.

For brands, the immediate implication is straightforward: knowing whether AI recommends you is only half the picture. Understanding whether AI *trusts* you - and measuring the gap between those two signals - is what determines whether your AI visibility is structurally durable or dependent on conditions outside your control.

This research will continue as a longitudinal observation program, with subsequent publications examining category-specific patterns, model-level variation, and the predictive relationship between authority formation and recommendation outcomes.

The brands that win in AI will not be those most visible at the moment of recommendation, but those most trusted before the decision is made.

Frequently Asked Questions

Common Questions About AI Visibility and the Authority Gap

What is the Authority Gap in AI?

The Authority Gap is a measurable divergence between brands that AI models cite as sources of truth during research stages and brands that AI models recommend to consumers during

decision stages. Research by GOSH AI found that in early observations, fewer than 1% of brands appear in both layers, meaning authority and recommendation are nearly independent signals in AI systems.

What is ARDI and how does it measure AI visibility?

ARDI™ (AI Recommendation and Discovery Intelligence) is a proprietary observation platform developed by GOSH AI. It tracks how five leading AI models - ChatGPT, Gemini, Claude, Perplexity, and Grok - process, cite, and recommend brands across 70+ industry categories and 50+ geographic markets, using two parallel observation paths: the Search Path (real-time retrieval) and the Learned Path (training-embedded knowledge).

What are the three positions in the Authority Gap framework?

The framework classifies brands into three positions. Full Authority means the brand is both cited as a source of truth and recommended, representing fewer than 1% of observed brands. Borrowed Visibility means the brand is recommended but never cited as an authority, representing approximately 72% of brands. Unconverted Authority means the brand is cited as a trusted source but not recommended, representing approximately 27% of brands. Full Authority is the most structurally durable position in AI visibility.

What is the difference between the Search Path and the Learned Path?

The Search Path captures AI behavior when models use real-time search to pull live results, citations, and current data. This is volatile and changes frequently. The Learned Path captures what AI models have absorbed from training data about brands and categories. Brands embedded in the learned layer have presence that persists even when search results change. ARDI measures both paths independently to distinguish between temporary retrieval-driven visibility and durable trust-based authority.

HOW TO CITE THIS PAPER

APA:

GOSH AI. (2026). *The Authority Gap: How AI Separates Trust from Recommendation*. ARDI™ Research, State of AI Visibility, Vol. 1. <https://www.mygosh.ai/authority-gap>

MLA:

GOSH AI. "The Authority Gap: How AI Separates Trust from Recommendation." *State of AI Visibility*, vol. 1, Apr. 2026. www.mygosh.ai/authority-gap.

ABOUT ARDI™

ARDI™ (AI Recommendation and Discovery Intelligence) is a proprietary observation platform developed by GOSH AI that tracks how AI models process, cite, evaluate, and recommend brands across the buyer journey. ARDI™ monitors five leading AI models across 70+ industry categories and 50+ geographic markets.

For methodology inquiries or to request the full ARDI™ Authority Index for a specific category, contact the research team.



Eric Torres

Founder & Head of Research

eric@mygosh.ai

[mygosh.ai](https://www.mygosh.ai)

© 2026 GOSH AI. All rights reserved. ARDI™ is a trademark of GOSH AI. This document may be shared in its entirety with attribution. Data, frameworks, and classification systems described herein are proprietary to GOSH AI. The Authority Gap™ framework and associated terminology are original research contributions of GOSH AI. Reproduction of findings without attribution is prohibited.