THE DECISION SCIENCE SERIES

The Decision Sequence

Designing Engagement Strategies Based on How Consumers Decide.

Artificial Intelligence Reads You Before the Buyer



DS 15

AI-Mediated Discovery

–A Strategic Resource by Joe Worden | More @ www.jworden.com

DECISION ARCHITECTURE RESEARCH LIBRARY

The AI-Era Research · Entry 15

AI-Mediated Discovery and AI Search

Forrester · Semrush · Gartner · NIST AI RMF · Aggarwal et al. (GEO)

Questions this paper answers

- How can we find out whether AI assistants describe our firm accurately?
- Why might an AI answer mention our brand without citing our website?
- Which public information should we check when an AI summary gets our firm wrong?
- How should we assess AI visibility without assuming one answer represents every buyer?

Definition. AI-mediated discovery is the quiet shift where a buyer's first encounter with a firm is often not a homepage, a pitch deck, or a salesperson, it's a machine. Search assistants and AI tools now handle the opening stretch of the decision, they surface options, condense them into summaries, line them up side by side, and translate "expert" material into something a buyer can act on. The implication is plain: the first steps in trust now happen in a layer the firm neither owns nor fully observes, and can't reliably steer.

The Introduction Nobody Wrote

A while back I started opening client engagements with a simple exercise. Before we talk about your marketing, let's read the introduction your buyers are getting. I ask a few AI assistants the questions a real prospect would ask: what does this firm do, who is it best for, how does it compare to the alternatives.

Then I read the answers to the room, out loud.

The first time I did this, the silence was something you could pour. The machine's introduction was confident, fluent, and about eighty percent right, and the missing twenty percent was the exact twenty the firm considered its edge. Their specialty was described a category too broad. Their proudest engagement was nowhere. One answer leaned, at length, on a review thread nobody in the room had ever seen.

Somebody finally asked the right question: "Who wrote that?" And the honest answer was: everyone and no one. The machine had assembled their introduction from whatever the public record offered, and the firm had never once read its own most-delivered pitch.

That's the new condition in one sentence. One version of your pitch may be reaching buyers without anyone at your firm hearing it. The research below says exactly how that happened, and how much of the decision now rides on it.

The Search Now Starts with a Machine

If you want the cleanest signal, Forrester provides it. The State of Business Buying, 2026 draws on a survey of nearly 18,000 buyers worldwide and lands on a simple finding: generative AI has become the front door for business research. In the underlying data, genAI tools show up as the most frequently named meaningful interaction used to research a purchase. And AI is no longer a niche assist, close to everyone, about 94 percent, uses it somewhere along the way. The first pass, the stage that used to be spread across a firm's website and clarified by a rep's explanation, is increasingly completed inside the tool.

But Buyers Still Validate With Humans

None of this means the machine has taken over the decision. The same Forrester research points to a persistent caution: buyers don't fully trust what AI tells them. Some finish the process feeling more assured because genAI helped them orient, but a meaningful slice, about one in five, walk away less confident, after running into gaps, errors, or information that doesn't hold up.

So they compensate in a very human way. They take the machine's output, then run it through people they trust: peers, domain specialists, and their own internal buying circle, in Forrester's survey, an average of 13 internal stakeholders and nine external participants influenced the decision. A useful shorthand is this: AI increasingly influences the consideration set, while human trust still decides the winner. The system narrows the field. People make the call.

Your Reputation Is Built From Sources You Don't Own

Once the machine becomes the early interpreter, the next question is obvious: what is it reading, and whose words is it repeating. Semrush's 2026 AI Visibility Index, built on an analysis of 126 million AI search prompts, argues that a firm's AI-facing story is no longer assembled primarily from its own site and polished, controlled materials. Reviews matter. Community threads matter.

Independent publishers matter. Industry sources, trade commentary, and third-party explainers all get pulled in, blended, and reissued as "the" description of a brand.

Semrush also separates two things many teams still treat as the same. Being mentioned in an AI answer is one outcome. Having your own site cited as the source is another, and on some platforms the overlap between those outcomes is surprisingly small.

Two practical implications follow. First, coherence starts to matter in a new way. When many independent sources describe a firm similarly, the machine tends to repeat that picture with confidence. When the signal is scattered, the brand can be distorted, diluted, or skipped altogether.

Second, visibility is still a blind spot for a lot of organizations. The measurement task needs its own discipline: record mentions, citations and accuracy separately rather than assuming they describe the same outcome. You are being characterized in spaces you can't walk into, and often without instruments that tell you what was said.

A New Discipline, and a Readiness Gap

As a result, a new body of work has started to gather around "being findable" to generative systems. The paper that effectively put a name on it, "GEO: Generative Engine Optimization" (Aggarwal et al., 2024), shows how generative engines assemble responses by synthesizing across sources rather than simply ranking pages, and it reports that intentional methods can materially increase how often, and how prominently, a particular source gets pulled into answers. Those were results in the researchers' evaluation, not a guarantee across today's services. For Google's AI search features, ordinary Search eligibility and people-first content guidance still apply.

In parallel, guardrails are being sketched out. NIST's AI Risk Management Framework offers a voluntary approach for trustworthy AI, organized around governing, mapping, measuring, and managing AI risk. Yet ambition is outrunning preparedness. Gartner's 2026 CMO Spend Survey captures the mismatch: around seventy percent of marketing leaders say becoming an AI leader is a critical goal, while only about thirty percent say they are ready to scale it. Adoption is moving quickly. Discipline is struggling to keep pace.

Why This Belongs in Decision Architecture

The older model of market entry assumed something comforting: a firm controlled its introduction. It chose the wording, set the timing, and shaped the first impression.

AI disrupts that premise. Prospects, and the machines they consult, interpret a firm's public signals before the firm shows up in the conversation, and in many cases they decide whether the firm gets into the conversation at all. That is why public expertise, plainspoken clarity, consistency across sources, source authority, and structured thought leadership are no longer just outputs for "marketing." They function as Decision Architecture, the raw inputs a machine uses to represent a firm at the opening of the trust sequence.

What hasn't moved is the endpoint. Across the research and inside the framework, the pattern holds: buyers still face uncertainty, still look for verification from trusted sources, and still need enough confidence to commit.

AI speeds up research. It doesn't speed up belief. The firm that is clear, consistent, and credible enough that a machine can portray it accurately earns the chance to be considered. A vague or contradictory public record can leave a firm poorly represented in one route to consideration. Other research and referrals may still bring it into the decision. The environment shifted. Trust did not.

The Marketing Read

Your firm now has two websites: the one you built, and the one the machines assemble from everything everyone else says about you. Either one may arrive first. You should know what both say.

So make the audit a ritual, and run it before spending anywhere else. Monthly, ask the major assistants the questions your buyers ask: best firms for this problem, compare us to the two competitors, what does this firm do and who is it for. Log the answers. Track what changed. Those answers are a sample of a discovery channel, not a count of what buyers saw. Log the service, date, prompt and any cited sources, and repeat the check before drawing a conclusion. The transcript is useful evidence, and until you're reading it too, your positioning strategy is a rumor about a document you've never seen.

Then feed the machine coherence, because Semrush's finding is a gift once you see it: the machine repeats with confidence whatever the record says consistently. One firm, one story, told the same way on your site, your bios, the directories, the press quotes, the conference listings. Every contradiction in the public record, the outdated specialty, the partner bio from two positioning strategies ago, is noise the machine averages into your introduction.

Invest in the sources the machine trusts more than it trusts you. Mentions come from the third-party record: trade press, reviews, community threads, independent explainers. This is Entry 08's borrowed trust, replayed for a machine audience, and it means an independently earned article in

a relevant publication may reach buyers through sources your own site does not. Citation, having your own material quoted, is a separate prize, and the GEO study tested ways of improving source visibility in its evaluation, not a universal recipe for citation. A sound editorial starting point remains: plain claims, real data, language a machine can lift whole.

Keep the human lane funded, because the research is blunt about where machines stop. AI assembles the consideration set; the buyer's dozen stakeholders and trusted voices still pick the winner. Machine visibility is the ticket to the room, and everything in Entries 06 through 14 still decides what happens inside it.

What backfires: sludge. Flooding the zone with AI-generated content to impress AI readers is the one move that fails with both audiences, machines weight authority and consistency over volume, and buyers who catch the slop discount everything else. The other failure is the comfortable one: deciding your buyers don't use AI. Forrester's survey is a reason to investigate AI use among your buyers, not to assume every prospect researches the same way.

The hard call: open an AI assistant right now and ask it about your firm the way a stranger with money would. Read every word. That's your first impression now, it's being delivered while you read this, and until today, nobody at your firm had heard the pitch.

Primary Sources

- Forrester, The State of Business Buying, 2026 (January 2026; based on Forrester's 2025 Buyers' Journey Survey of nearly 18,000 global buyers).
- Semrush, 2026 AI Visibility Index (June 2026; 126 million U.S. AI search prompts, January–April 2026). Publisher release: https://www.semrush.com/news/463141-semrush-releases-expanded-2026-ai-visibility-index-analyzing-126-million-ai-search-prompts/
- Business Insider and other business-press coverage of Semrush's AI search and visibility findings.
- Gartner, 2026 CMO Spend Survey (May 2026).
- National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0). NIST, 2023.
- Pranjal Aggarwal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan & Ameet Deshpande, "GEO: Generative Engine Optimization." Proceedings of ACM SIGKDD (KDD '24); arXiv:2311.09735, 2023.
- Research on AI Overviews, AI search visibility, source selection, and brand representation (including Semrush's AI Overviews study and related analyses of AI summaries and click behavior). Current eligibility guidance: Google Search Central, AI Features and Your Website, https://developers.google.com/search/docs/appearance/ai-features