



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

HOW AI WORKS

How AI Systems Actually Decide Which Businesses to Recommend

By Scott Tischler · July 28, 2026 · 10 min read

Ask ChatGPT, Google's Gemini, Perplexity, or Microsoft Copilot to recommend "the best commercial HVAC company in Denver" or "a good fractional CFO for a Series A startup," and you'll get an answer. Usually a confident one. Often just two or three names, with reasons attached.

For the businesses named, that moment is worth more than a page-one ranking used to be. For everyone else, it's invisible. There is no scroll, no page two, no "see more results." You are recommended, or you do not exist.

I've spent the last few years at Alrecommend.ai studying that decision — why one business surfaces and its equally qualified competitor doesn't. What follows is my best, honest account of how these systems actually decide. I'm going to be careful throughout to separate what is publicly documented from what is reasonable inference, because the fastest way to lose your trust is to pretend I have a schematic of ChatGPT's internals. Nobody outside those labs does. But the mechanics are far less mysterious than the marketing hype suggests, and they are more influenceable than most executives believe.

Two different machines are answering your question

The first thing to understand is that "AI recommends my business" is not one process. It's at least two, and they behave differently.

The first is the **model's trained parameters** — everything the system absorbed from its training data up to a cutoff date. When you ask a question and the AI answers purely from what it "already knows," it's drawing on statistical patterns formed across billions of documents. If your business was mentioned often, in credible contexts, alongside the right topics, the model has effectively learned an association: **this name goes with this kind of expertise**. You cannot edit this directly. You can only influence it over time by being present, consistently and credibly, in the kind of content these models train on.

The second machine is **retrieval** — often called RAG, retrieval-augmented generation. Here the system runs a live or near-live search, pulls documents, and generates an answer grounded in what it just fetched. Perplexity is built around this. ChatGPT with search, Gemini with grounding, and Copilot all do it too. Retrieval is where recency lives, where citations come from, and where a business the model has never "memorized" can still get recommended because it showed up in a fresh, relevant, trustworthy source at query time.

Most real answers blend both. The model uses its trained intuition to frame the response and retrieval to ground the specifics. This matters enormously for strategy: optimizing for the trained layer is a slow, reputation-driven game, while optimizing for the retrieval layer overlaps heavily with things you can affect this quarter — your content, your structured data, your presence in the sources these systems trust.

Step one: the system has to know what — and who — you are

Before any recommendation, the AI has to resolve your business into a distinct **entity**. Not a string of characters, an entity: a specific organization, separate from the dozen others with similar names, connected to a location, a category, people, and a body of work.

This is entity recognition and disambiguation, and it's the single most underrated step. If the system can't confidently tell that "Meridian Partners the boutique M&A advisory in Chicago" is different from "Meridian Partners the real estate group in Phoenix," it hedges — and hedging usually means omission. AI systems are tuned to avoid confidently stating things they're unsure about, so ambiguity gets you left out far more often than it gets you a wrong mention.

Strong entities share traits: a consistent name, category, and location across the web; presence in structured knowledge sources; unambiguous associations with their domain of expertise; and clear connections to real, identifiable people. When I audit why a qualified business isn't being recommended, a muddled or thin entity is the most common root cause I find. It's rarely a content problem first. It's an identity problem.

Step two: not all sources count equally

Once you're a recognizable entity, the question becomes whether the sources that mention you are ones the system trusts.

Every major AI system inherits, in some form, decades of information-retrieval research into source quality — the same lineage that gave search engines link-based authority and site-quality signals. The specific weightings are proprietary and I won't pretend to know them. But the direction is well established and consistent across systems: a claim carries more weight when it appears in an

independent, established, editorially credible source than when it appears on the business's own website.

This is why self-description has a ceiling. Your homepage saying you're "the leading provider" is a claim the system heavily discounts, because every competitor's homepage says the same thing. What moves the needle is **corroboration you don't control**: an industry publication, a reputable directory, a university page, a well-regarded review platform, a journalist's write-up, a professional association. In practice, authority is less about any single prestigious mention and more about a credible **web** of them pointing to the same entity.

Step three: consensus is the quiet kingmaker

If I had to name the most important and least understood signal, it's **consensus** — agreement across multiple independent sources.

Language models are, at their core, machines for detecting patterns that repeat. When many credible, unrelated sources describe your business the same way — same category, same specialty, same strengths — that repetition reads as consensus, and consensus reads as truth. When your positioning is described five different ways across five places, there's no pattern to lock onto, and the system defaults to whatever competitor **does** have a coherent story.

This has a blunt practical consequence. Message discipline is now a technical ranking factor, not just a branding nicety. If your site calls you a "growth consultancy," your LinkedIn says "marketing agency," a directory lists you under "advertising," and a podcast introduced you as "brand strategists," you have handed the AI four contradictory signals about who you are. The business that says one clear thing, everywhere, and gets echoed saying it, wins the consensus that AI systems are built to detect. In Alrecommend.ai's State of AI Search 2026 research, coherent, corroborated positioning was one of the clearest dividing lines between businesses that got cited and those that got overlooked.

Step four: structured data, so machines don't have to guess

Humans read prose fluently. Machines prefer explicit structure. Structured data — Schema.org markup like Organization, LocalBusiness, Product, Review, and FAQ — is how you hand an AI system unambiguous facts instead of making it infer them from paragraphs.

I want to be precise here, because there's a lot of overclaiming in this space. No AI lab has published a rule saying "Schema markup increases recommendation probability by X percent," and anyone who quotes you a number like that is guessing. What we can say responsibly is this: structured data reduces ambiguity, and reduced ambiguity is exactly what these systems reward at the entity and retrieval stages. When your name, category, location, hours, service area, and credentials are

machine-readable and internally consistent, you make it easy to be understood correctly — and easy is what gets you included. Structured data isn't a magic ranking lever. It's insurance against being misread.

Step five: reviews, reputation, and recency

Three human-facing signals round out the picture.

Reviews and reputation. AI systems lean on the same third-party reputation signals people do, because those signals are abundant, structured, and hard to fake at scale — volume of reviews, average sentiment, recency, and the substance of what people actually say. A cluster of recent, specific, positive reviews describing ***what you're good at*** does double duty: it feeds sentiment and it reinforces consensus about your specialty. Sentiment is increasingly legible to these models, so a wall of detailed praise about your onboarding or your emergency response tells the system something a star rating alone can't.

Recency. In the retrieval layer, freshness matters. A current, well-maintained source often outranks a stale one, and an entity that's visibly active — publishing, being written about, gathering new reviews — signals ongoing relevance. This is also why a business can rise in AI recommendations without touching the trained-model layer at all: consistent freshness keeps you in the retrieval set.

Consistency over time. The businesses that win aren't usually the ones with a single brilliant asset. They're the ones whose signals all agree and keep agreeing — entity, authority, consensus, structure, reputation, and recency pointing the same direction, month after month.

The signals that matter, ranked

If you take one thing from this piece, take this hierarchy — my practitioner's synthesis of the evidence, ordered by how much leverage each tends to give:

1. **Entity clarity** — Can the system unambiguously identify you and connect you to your expertise? Everything else is downstream of this.
2. **Consensus** — Do multiple independent, credible sources describe you the same way?
3. **Source authority** — Are the places that mention you trusted and independent, not just your own properties?
4. **Reputation and reviews** — Is there recent, specific, positive third-party sentiment?
5. **Structured data** — Have you removed ambiguity by making your facts machine-readable?
6. **Recency and activity** — Are your signals current, and are you visibly still in the game?

Notice what's **not** at the top: your website copy, your ad spend, your clever tagline. Those matter, but they're inputs to the signals above, not signals themselves.

What this means for how you operate

The strategic shift is uncomfortable for a lot of teams: you're no longer optimizing primarily for a crawler that ranks pages. You're building an evidence base that a reasoning system can use to justify recommending you over someone else. This discipline — Answer Engine Optimization — is what I spend my days on at Alrecommend.ai, and the pattern is consistent. In our State of AI Search 2026 research, businesses that executed a full AEO program saw 3.8 times higher AI citation rates and roughly 142% more AI-referred leads within 90 days versus their starting baseline. I'd caution against reading those as guarantees — they reflect businesses that did the whole job, not a single tactic — but the direction is unambiguous.

The urgency is real because the behavior has already shifted. Our research found 37% of people now start their searches with an AI system rather than a traditional search box, and 58 to 68% of those AI searches end without a single click to any website — the answer itself was the destination. When the answer is the destination, being **in** the answer is the entire game.

Here's my honest bottom line, offered as opinion rather than fact: I don't believe these systems are trying to be gatekeepers, and I don't think the mechanics are adversarial. They're trying to give a confident, correct, well-sourced answer. The businesses that win are the ones that make being recommended the **easy, safe, well-evidenced choice** — clear identity, consistent story, credible corroboration, current signals. You can't hack your way into an AI recommendation. But you can, deliberately, become the answer that the machine is most comfortable giving.

Key takeaways

- ② AI recommendations come from two different mechanisms — the model's trained knowledge and live retrieval (RAG). Retrieval is where most of your near-term leverage lives.
- ② Before anything else, the system must resolve you as a distinct entity. Ambiguous identity is the most common reason qualified businesses get left out.
- ② Consensus is the quiet kingmaker: multiple independent, credible sources describing you the same way reads as truth. Message discipline is now a technical ranking factor.
- ② Authority comes from corroboration you don't control — self-description has a hard ceiling because every competitor makes the same claims.

- ② Structured data doesn't have a proven "ranking boost" number; treat it as insurance against being misread, not a magic lever.
- ② Reviews, recency, and consistency over time compound. The winners aren't the flashiest; they're the ones whose signals all agree and keep agreeing.

Frequently asked questions

Can I pay an AI system to recommend my business?

Not through the recommendation mechanics themselves. The organic answer is generated from trained knowledge and retrieved sources, not from an ad auction. Some systems are testing separate ad placements, but those are labeled and distinct from the recommendation itself. Influence comes from strengthening your entity, authority, and consensus signals — not from payment.

How is this different from traditional SEO?

SEO optimizes a page to rank in a list of links a human then chooses from. AEO builds an evidence base a reasoning system uses to justify naming you in an answer where there's often no list and no click. They share roots — source authority, structured data, reputation — but AEO weights entity clarity and cross-source consensus far more heavily.

How long does it take to start showing up in AI recommendations?

The retrieval layer can respond within weeks as fresh, credible sources and structured data get picked up. The trained-knowledge layer is slower and reputation-driven, changing over model cycles and months of consistent presence. In our experience most businesses see meaningful movement inside a quarter when they address entity clarity and consensus first, rather than chasing individual mentions.

About the author. Scott Tischler is the Founder & Chairman of Alrecommend.ai and a practitioner-authority on AI search and Answer Engine Optimization. With 20+ years in marketing technology — including American Express, MetLife, and UBS — and executive study at Wharton, Harvard, Yale, and Oxford, he helps businesses become the ones AI recommends.