



Scott Tischler
FOUNDER & CHAIRMAN · AIRECOMMEND.AI

ORIGINAL RESEARCH

The AI Recommendation Study: What AI Actually Recommends (2026 Pilot)

By Scott Tischler · July 30, 2026 · 8 min read

Everyone is talking about AI search. Very few people have actually **measured what AI recommends** when a real person asks it to name a business. So we ran the test ourselves.

This is the pilot round of the **AI Recommendation Study** — original research from Alrecommend.ai. We asked a live consumer AI engine a fixed set of natural "what's the best..." questions across products, services, franchises, and local businesses, and we recorded, verbatim, exactly what it returned: how many businesses it named, who they were, whether it cited sources, and — where we could see it — what appeared to decide the pick. Every number below is counted from a real answer. Nothing is estimated or invented. ***Prefer it visual? See the illustrated edition — [The Tischler AI Index](#).***

How we ran it

We ran **12 "recommend me" prompts** through **Perplexity** in a single U.S. session on **July 30, 2026**, and logged each result: the count of distinct businesses named, whether the answer surfaced citations, and the named entities. We then re-ran **five of them through Google's AI Overviews** to see how much the two engines agree. This is deliberately a **pilot** — a small sample, one location, one moment in time. It is enough to establish the method and show clear, directional patterns; the full study scales the sample, adds ChatGPT and Gemini, and tests multiple locations. We're publishing the pilot openly, limits and all, because the patterns are already striking.

Query	Businesses named	Cited sources?
best AI marketing agency for a small business	0 (criteria only)	minimal
best CRM for a small business	6	Yes (9)
best project management software	5	Yes
best email marketing platform	6	Yes

Query	Businesses named	Cited sources?
best franchise to buy in 2026	12	Yes (8)
best financial advisor for retirement	~4 (with a hedge)	Yes
best personal injury lawyer	5 (localized)	Yes (9)
best jeweler for a custom engagement ring	2 recommended (from a local list)	Yes
best accounting software for freelancers	5	Yes
best HVAC company near me	ratings-ranked local list	Yes
best website builder for a small business	2–3	Yes (9)
best password manager	6	Yes

Finding 1: AI names a shortlist, not ten blue links

The engine named businesses in **11 of the 12 queries**, and when it did, it named a **median of five** (ranging from two to twelve) — and for local services it narrowed to an explicit **top two or three**. This is the single most important shift for any business owner to absorb. Traditional search gives the user ten links and lets them choose. **AI gives a shortlist and does the choosing for them.** If you are not on the shortlist, you are not in the consideration set at all — there is no page two to be found on.

Finding 2: For local businesses, reviews literally decide it

Every one of our local-service queries — the lawyer, the jeweler, the HVAC company — returned a geo-personalized "Places" module, and in each case the businesses the AI **recommended were the ones with the highest star ratings and the most reviews**. For the custom-engagement-ring query, the two jewelers it singled out both held 5.0-star ratings with strong review counts; it said so explicitly. If you run a local business, your review profile is no longer just a trust signal for humans — it is the input the AI uses to decide whether to name you.

Finding 3: Third-party sources get you named — not your own website

Look at where the recommendations came from. The CRM, project-management, and password-manager answers were sourced almost entirely to **independent review and roundup sites** — PCMag, Zapier, the U.S. Chamber, Forbes — not to the vendors' own homepages. The AI didn't recommend the company with the best marketing copy; it recommended the companies that **other trusted sources** already talked about. Being written **about** elsewhere is what earns you a place in the answer. This is the whole game of entity authority in one data point.

Finding 4: If your category is ambiguous, AI names no one

The most revealing result was the query that named **zero** businesses: "best AI marketing agency for a small business." Instead of naming anyone, the engine hedged — "there isn't one universal best" — and returned a list of **criteria** for choosing. When a category is broad or ill-defined, AI declines to pick, which means it can't pick **you**. Clarity of category is a prerequisite for being recommended at all.

Finding 5: AI personalizes by location even when you don't give one

We never specified a city. The engine inferred location from the session and localized the lawyer, jeweler, and HVAC results anyway. For any business that serves a geographic area, AI is already deciding whether you're "near me" — based on how clearly your location and service area are established across the web.

Finding 6: In mature categories, the same names win across every engine

We re-ran five queries on Google's AI Overviews and compared them to Perplexity. For established product categories, the two engines **largely agreed on the top names**: for CRMs, both led with HubSpot, Pipedrive, and Zoho; for password managers, both named 1Password, Bitwarden, NordPass, and Proton Pass; for email platforms, both named MailerLite and Brevo; for project management, both led with Trello and Asana. For a fragmented category like franchises, the engines **diverged** — only Visiting Angels and Wingstop appeared on both lists. The lesson cuts two ways: in a mature category a handful of names win across **every** AI, so being one of them is everything; in a fragmented category you have to earn the recommendation on each engine separately. And both engines leaned on the same kind of evidence — independent review sites, roundups, forums, and video — not the brands' own pages. Off-site reputation drives the recommendation everywhere.

What this means for your business

Put the five findings together and the playbook writes itself. To be the business AI recommends, you need to be **unambiguously categorized, well-reviewed, clearly located**, and — above all — **talked about by independent, trusted sources** the models read. That is exactly the entity authority and structured-knowledge work that decides who makes the shortlist. Ranking on Google was about being on the list. Winning in AI search is about being **the recommendation** — and this data shows the recommendation is earned off your own website, in the sources and signals AI trusts.

Honest limits, and what's next

This is a pilot: twelve queries on Perplexity plus a five-category Google AI Overviews cross-check, one location, one day. It is not the final word, and we're not going to pretend it is. What it does show — clearly and reproducibly — is **how** to measure AI recommendation behavior and **what patterns** emerge when you do. The full study will scale the sample, run the same prompts across ChatGPT, Gemini, and Google AI Overviews to measure how much they agree, and test different markets. We'll publish that openly too. If you want to be included or have a category you'd like tested, reach out.

Key takeaways

- ② AI gives a shortlist and makes the choice for the user — it named businesses in 11 of 12 queries, a median of 5 each (and only 2–3 for local services).
- ② For local businesses, the AI recommended the highest-rated, most-reviewed options — reviews now decide recommendations, not just trust.
- ② Recommendations were sourced to independent review sites, not the businesses' own websites — being written about elsewhere is what gets you named.
- ② One ambiguous B2B query named zero businesses — if your category isn't clear, AI can't recommend you.
- ② AI geo-personalizes even without a stated city, so a clear location and service area matter.
- ② Across engines, mature product categories showed the same winners (CRMs, password managers matched on Perplexity and Google); fragmented categories like franchises diverged.
- ② This is an honest pilot (12 queries on Perplexity plus a 5-category Google AI Overviews cross-check, one session); the full study scales engines, sample, and locations.

Frequently asked questions

What is the AI Recommendation Study?

It's original research from Alrecommend.ai that measures what AI engines actually recommend when someone asks them to name a business, product, or service. We run real "recommend me..." prompts against live AI engines and record exactly what they return. This is the 2026 pilot round.

How many businesses does AI recommend on average?

In this pilot, the AI named businesses in 11 of 12 queries, and when it did it named a median of five (range two to twelve), narrowing to an explicit top two or three for local services. That's a dramatic compression from a traditional search page of ten links — AI hands the user a shortlist and does the choosing.

What decides which local business AI recommends?

In every local-service query we ran, the AI surfaced a "Places" module and recommended the businesses with the highest star ratings and the most reviews. Your review profile is now a direct input to whether AI names you, not just a signal for human shoppers.

Does AI recommend businesses based on their own websites?

Largely no. In this pilot, the recommendations were sourced to independent review and roundup sites — PCMag, Zapier, Forbes, the U.S. Chamber — rather than the businesses' own pages. Being covered and cited by trusted third parties is what earns a place in the answer.

Why did one query name zero businesses?

The broad query "best AI marketing agency for a small business" was too ambiguous, so the AI declined to name anyone and returned selection criteria instead. When a category is unclear, AI won't pick a winner — which means it can't pick you. Category clarity is a prerequisite for being recommended.

Is this study statistically definitive?

No, and we don't claim it is. It's a pilot — twelve queries on Perplexity plus a five-category Google AI Overviews cross-check, one location, one point in time. It establishes the method and shows clear directional patterns; the full study will scale the sample, add ChatGPT and Gemini, and test multiple markets.

Do different AI engines recommend the same businesses?

It depends on the category. In this pilot, Perplexity and Google's AI Overviews largely agreed on mature product categories — both led with HubSpot, Pipedrive, and Zoho for CRMs, and both named 1Password, Bitwarden, NordPass, and Proton Pass for password managers. For a fragmented category like franchises they diverged, sharing only a couple of names. In mature categories a few winners dominate every engine; in fragmented ones you must earn each engine separately.

How can a business get recommended by AI?

Based on this data: be unambiguously categorized, earn strong reviews, make your location and service area clear, and — most importantly — get written about by independent sources AI trusts. That entity-authority and structured-knowledge work is what puts you on the shortlist.

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.