



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

AI & SEARCH

Perplexity vs. ChatGPT vs. Google AI: Where Your Customers Actually Ask

By Scott Tischler · July 28, 2026 · 9 min read

The question I get most from founders isn't "how do I do AEO." It's "which one matters — Perplexity, ChatGPT, or Google's AI?" The honest answer is that they all matter, but they don't work the same way, they don't serve the same buyers, and optimizing for them identically leaves results on the table. Your customers are asking different kinds of questions in different places, and the engine they choose shapes whether — and how — you get recommended.

Here's the practitioner's frame. Each answer engine makes three decisions about your brand: what it **retrieves** (which sources it pulls), what it **cites** (which sources it shows the user), and what it **recommends** (which brands it actually names in the answer). The three engines weigh those decisions differently. Understanding how is the difference between spreading effort thin and putting it where it converts.

How the three engines actually differ

Let me describe each one the way I'd describe it to a client deciding where to focus.

Perplexity is a search-first answer engine. It's built around live retrieval — for most queries it runs a real web search, reads the top sources, and synthesizes an answer with visible, numbered citations right in the response. It behaves like a research analyst who shows their work. Because citations are front and center, ***being in the retrieved source set is the whole game***. Perplexity users tend to be researchers, analysts, and high-intent buyers who specifically want sources they can click and verify.

ChatGPT is model-first with retrieval layered on. A large share of what it says comes from its training — its internalized sense of who the trusted brands are — and it reaches out to live search when the query needs freshness or specifics. This has a big strategic implication: with ChatGPT, your reputation in the training data matters enormously, and that reputation was built from broad, consistent, corroborated presence across the web over time. ChatGPT has the largest and most general user

base of the three, spanning casual questions to deep buying research, which makes it the highest-volume surface for most brands.

Google AI — the AI Overviews and AI Mode sitting on top of Search — is retrieval-native and grounded in Google's index and its existing ranking systems. If Google already understands your site (structured data, entity clarity, established ranking), that understanding feeds the AI answer. It reaches the broadest audience by far, because it appears above ordinary search results for enormous query volume, often for users who didn't opt into an "AI tool" at all. Traditional SEO fundamentals carry more weight here than anywhere else.

The comparison at a glance

Dimension	Perplexity	ChatGPT	Google AI
Primary mechanism	Live retrieval, search-first	Model knowledge + retrieval	Index-grounded, ranking-native
Citation visibility	High — numbered, prominent	Moderate — links when it searches	Moderate — sourced, less prominent
What drives inclusion	Being in the retrieved source set	Corroborated reputation over time	Strong SEO + entity clarity
Freshness sensitivity	High	Moderate	High
Typical user	Researchers, analysts, verifiers	Broadest general + deep research	Broadest reach, mainstream search
Best-fit query	"Compare X vs Y, with sources"	"What should I use for..."	"Best [category] near/for me"
Biggest optimization lever	Rank + be extractable now	Be the consistent, corroborated brand	Classic SEO + structured data

The rows to internalize: Perplexity rewards **current retrievability**, ChatGPT rewards **accumulated reputation**, and Google AI rewards **search fundamentals**. Same brand, three different scoreboards.

Which buyers use which — and for what

The engines don't just work differently; they attract different people at different moments in a buying journey.

Someone doing **serious, verify-as-I-go research** — an analyst comparing vendors, a technical buyer vetting claims, a journalist — gravitates to Perplexity precisely because it shows sources. If you sell

into considered, high-stakes purchases where buyers want to check your claims, Perplexity punches above its user-count weight.

Someone doing **open-ended exploration** — "what's a good tool for onboarding a remote team," "how should I think about X" — is most often in ChatGPT. It's conversational, it's where people already are, and it's where the mainstream "what should I use" question gets asked. For most B2B and considered-consumer brands, this is the largest single pool of buying-intent conversations.

Someone doing **fast, mainstream lookups** — "best CRM for small business," "top med spas in Austin" — increasingly gets an answer from Google AI before they ever scroll to the links, simply because it's there by default. This is the broadest reach and the least self-selected audience. If your category gets searched on Google (nearly all do), you're being described in AI Overviews whether you've optimized for it or not.

The practical read: don't ask "which engine is biggest." Ask "where does **my** buyer ask **my** kind of question." A compliance-heavy enterprise tool and a local service business should not have the same engine priorities.

How optimization priorities shift across engines

Because the three weight retrieval, citation, and recommendation differently, the work is not identical. Here's how I'd sequence effort for each.

Optimizing for Perplexity

Perplexity's answer is only as good as what it retrieves in the moment, so the levers are the most "SEO-adjacent" and the most immediately actionable. **Be retrievable and extractable right now.** That means rank for the buying and comparison queries in your category, structure content so a clean claim can be lifted (clear headings, direct answers, comparison tables), and make sure your key facts live in crawlable text, not locked in images or PDFs. Because Perplexity favors freshness, recently updated, well-sourced pages get an edge. If you want a fast, measurable AEO win, Perplexity is usually where it shows up first.

Optimizing for ChatGPT

ChatGPT is the patience game, because much of what it "knows" about you is baked into its sense of the web. **You can't shortcut reputation — you build it.** The levers are broad, consistent, corroborated presence: the same positioning and key claims repeated across your site, third-party mentions, directories, reviews, and any place the model's training and retrieval both look. The goal is that whenever the web talks about your category, you're one of the consistently named, consistently

described brands. This compounds slowly and then pays off durably, because once you're part of the model's default answer, you tend to stay there.

Optimizing for Google AI

Google AI rewards the fundamentals you may already have partly in place. **Double down on entity clarity and structured data.** Make sure Google unambiguously understands what your business is, who it serves, and how it's categorized — through clean site structure, schema markup, consistent business information, and the ranking signals Google already trusts. If your traditional SEO is strong and your entity is clear, you're most of the way to being included in AI Overviews. If it's weak, no AI-specific trick compensates.

Where to focus if you can't do everything

Most teams can't optimize for all three at full effort at once, so here's the prioritization I actually recommend.

Start by matching to your buyer, as above. Then, as a general default for a brand starting from zero, I sequence it like this: **build the corroborated foundation first**, because it's the one investment that pays off across all three engines. Consistent positioning, corroborated claims, and clear entity data feed ChatGPT's reputation model, satisfy Perplexity's retrieval, and align with Google's index simultaneously. There is no wasted motion there.

From that foundation, chase the **fastest feedback loop next** — usually Perplexity — because you can see movement in weeks and use those wins to prove the approach internally. Layer in **Google AI's SEO fundamentals** in parallel if you have the resources, since much of that work overlaps with what you should be doing for classic search anyway. Treat **ChatGPT's reputation** as the long game running underneath all of it: never neglected, never expected to turn overnight.

Priority	Focus	Why first
1	Corroborated foundation	Pays off across all three engines at once
2	Perplexity retrievability	Fastest visible feedback loop
3	Google AI / SEO fundamentals	Broadest reach, overlaps with classic SEO
4	ChatGPT reputation	Highest volume, but compounds slowly

The mistake to avoid

The most common error I see is treating "AI search" as one monolithic thing and either optimizing for none of it or optimizing for all of it identically. Both waste effort. Perplexity, ChatGPT, and Google AI are three different scoreboards measuring three different things about your brand — current retrievability, accumulated reputation, and search fundamentals. Your customers have already sorted themselves across these engines by the kind of question they're asking. Your job is to figure out which of those questions are **your** buying questions, and to be the confident, consistent, corroborated answer on the surface where they're asked.

Key takeaways

- ② Each answer engine makes three distinct decisions — what it retrieves, cites, and recommends — and Perplexity, ChatGPT, and Google AI weight them differently.
- ② Perplexity is search-first and rewards current retrievability and extractability; it's usually where a fast AEO win shows up.
- ② ChatGPT is model-first and rewards accumulated, corroborated reputation built across the web over time — the long game that compounds durably.
- ② Google AI is index-grounded and rewards classic SEO fundamentals, entity clarity, and structured data, with the broadest mainstream reach.
- ② Match engine priority to where your specific buyer asks your specific question, not to which engine has the most users overall.
- ② Build the corroborated foundation first — it satisfies all three engines at once — then chase Perplexity for speed, Google AI for reach, and ChatGPT as the underlying long game.

Frequently asked questions

What's the core difference between Perplexity, ChatGPT, and Google AI for brand visibility?

Perplexity is search-first and rewards being in the live retrieved source set with extractable claims. ChatGPT is model-first and rewards accumulated, corroborated reputation across the web. Google AI is index-grounded and rewards classic SEO fundamentals and entity clarity.

Which AI engine should I optimize for first?

Start with the corroborated foundation — consistent positioning and clear entity data — because it pays off across all three at once. Then prioritize by where your specific buyer asks, but as a default, Perplexity gives the fastest visible feedback while ChatGPT reputation compounds over the long term.

Why does Perplexity show sources so prominently and ChatGPT often doesn't?

Perplexity is built as a research tool that runs live search and shows numbered citations by design, so users can verify. ChatGPT draws much of its answer from model knowledge and only surfaces links when it searches the web, so citations are less consistently prominent.

Does traditional SEO still matter for AI search?

Yes, especially for Google AI, which is grounded in Google's index and existing ranking systems. Strong SEO and clear entity data also help Perplexity retrieve you. ChatGPT is less directly tied to rankings, but the broad web presence good SEO creates still feeds its sense of your reputation.

Which engine do high-intent B2B buyers actually use?

It varies by buying stage, but verify-as-you-go researchers and technical evaluators gravitate to Perplexity because it shows checkable sources, while open-ended "what should I use for X" exploration happens heavily in ChatGPT. Google AI captures the broad, mainstream lookups above ordinary search results.

How fast can I see results from AEO on each engine?

Perplexity tends to move fastest because it retrieves live pages, so improving rankings and extractability can show up in weeks. Google AI moves with your SEO and index changes. ChatGPT is the slowest, since it depends on reputation accumulated across the web over time.

Should I optimize differently for each engine or use one approach?

Use one foundation and then differentiate. The corroborated foundation of consistent, verifiable claims serves all three, but from there Perplexity needs current retrievability, Google AI needs SEO and structured data, and ChatGPT needs sustained, broad corroboration. Treating them identically wastes effort.

Is Google AI worth targeting if fewer people use dedicated AI tools?

Yes, arguably the most, for reach. Google AI Overviews appear above ordinary search results for enormous query volume, reaching mainstream users who never opted into a separate AI tool. If your category gets searched on Google, you're already being described there whether or not you've optimized for it.

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.