Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

HOW AI WORKS

Knowledge Cutoffs: Why AI Doesn't Know Your Business Exists Yet — and How to Fix It

By Scott Tischler · October 11, 2026 · 9 min read

If you launched, rebranded, or relocated your business in the last year and AI assistants act like you don't exist, you're not imagining it — and it usually isn't a quality problem. It's a timing problem. Every AI model is trained on a snapshot of the world that ends on a fixed date, its **knowledge cutoff**. Anything that happened after that date isn't in the model's baked-in memory at all. So when ChatGPT, Gemini, Claude or Perplexity answers purely from what it "knows," a recent business can be simply invisible. The good news: there's a second path to being known, it's faster than the first, and it's the one you can actually control.

I run into this constantly in the work I do at AIrecommend.ai. A founder is frustrated that the AI "has never heard of us," takes it as a verdict on their relevance, and starts rewriting their homepage. But the model isn't rejecting them. It just froze its memory before they showed up. Understanding that distinction changes everything about what you do next.

What a knowledge cutoff actually is

A large language model learns by training on an enormous collection of text gathered up to a certain point in time. When training ends, the model's internal knowledge is effectively frozen. It doesn't keep learning as the world moves on. That frozen snapshot is the knowledge cutoff, and it's why a model can describe events and companies that existed before its cutoff in detail, yet draw a complete blank on something that launched a month later.

This isn't a bug, and it isn't personal. It's just how trained memory works. Think of the model's built-in knowledge as a very well-read person who stopped reading on a specific day. Ask them about anything up to that day and they're brilliant. Ask them about something that happened the week after and they genuinely don't know — not because the thing is unimportant, but because it's outside the window they ever saw.

For an established business, the cutoff is mostly harmless: you were already part of the snapshot. For a new, rebranded, or relocated business, the cutoff is the whole problem. Your current reality lives on the far side of the date the model's memory stops.

The two ways an AI can know about you

Here's the part that most people miss, and it's the part that sets you free. There are two completely different ways an AI assistant can come to know about your business, and they have nothing to do with each other.

The first is **trained memory** — what the model absorbed during training, frozen at the cutoff. The second is **live retrieval** — the engine going out to search the web, or a connected index, at the exact moment it's answering your question, reading current pages, and building its response from what it finds right now. Most modern AI answers are some blend of the two, and the blend is the key to your whole strategy.

When an engine answers from trained memory alone, the cutoff rules. A new business loses by default, because it isn't in the snapshot. But when an engine retrieves — when it searches live to ground its answer — the cutoff stops mattering. The retrieval layer doesn't know or care when you launched. It cares whether it can find you, read you, and trust you *today*. That's a question about your current web presence, not about any training date. And that is a game a new business can win immediately.

So the strategic reframing is simple: stop trying to get into the model's memory and start trying to get into its live retrieval. One of those is slow, uncertain, and out of your hands. The other is fast and largely within your control.

Why retrieval is the lever you can actually pull

Getting included in a future training run is not something you can schedule. You'd be waiting on a model provider to build a new version, hoping their training data happened to include enough high-quality references to you, and hoping you're represented accurately when it does. There's no application form, no timeline, and no guarantee of how you'll come out the other side. Treat it as a bonus that may arrive someday, not a plan.

Retrieval is the opposite. When a customer asks an engine a question and the engine decides to search the live web to answer it, your fate in that answer is decided by things you can change this quarter: whether your pages are crawlable and readable, whether your facts are clear and consistent, and whether the wider web corroborates what you say about yourself. I've watched recently launched businesses get named in AI answers quickly — not because any model was retrained, but because

when the retrieval layer went looking, they were the cleanest, most findable, best-corroborated answer available. The cutoff was irrelevant the moment the engine chose to search.

My read, not a guarantee: for any business younger than the current crop of models, retrieval is roughly the entire game, and effort spent worrying about training inclusion is effort spent on the one variable you can't move.

How to become the business retrieval surfaces

If the goal is to be what the live retrieval layer finds and trusts, the work is concrete. Start with your own house. Your key pages need to be genuinely readable by crawlers — not locked behind logins, not buried in JavaScript the bot never executes, not blocked in robots.txt. The answer a customer needs should be stated plainly and early on the page, in clean text, not implied across a brochure. If an engine can't extract a direct answer from you in seconds, it'll extract one from someone else.

Then get your facts consistent everywhere you appear. The retrieval layer builds confidence through corroboration — it trusts what multiple independent sources agree on. If your name, location, focus, and founding details read one way on your site, another way on your directory listings, and a third way on an old profile you forgot about, you've handed the engine a reason to hesitate. For a new or rebranded business this is especially urgent, because the stale version of you may be the only version some sources still show. Hunt down the contradictions and fix them.

Next, get mentioned in the places engines actually pull from. Retrieval doesn't just read your website — it reads the third-party sources that discuss your category: industry write-ups, legitimate directories, reputable coverage, the communities where your customers compare options. Being present and accurately described in those sources is often what tips an engine from "this business exists" to "this business is the one to name." You're not just publishing; you're making sure the web around you tells a consistent, credible story.

Finally, make yourself current and crawlable on an ongoing basis. Fresh, clearly dated content and a clean technical setup signal that you're an active, maintained entity worth surfacing — which matters a great deal when the whole point is to be found in the moment, by the live layer, rather than remembered from a frozen past.

Don't argue with the model — fix the web it reads

One mistake I see often: a founder discovers the AI has stale or wrong information about them and tries to correct the model directly, in the chat, as if they're updating a record. You can't. Trained memory doesn't update because you told it to in a conversation, and that conversation won't change

what the model says to the next person. The lever is never the chat window. The lever is the web the retrieval layer reads.

So when the memory is stale, don't debate it — go make the live web unambiguous, then test whether retrieval picks up the change. Ask the engines the questions your customers ask, and watch. When you start getting named, it's almost never because the model's memory finally caught up. It's because the retrieval layer found your cleaned-up, corroborated, current presence and chose it. That's the signal you're looking for, and it's the one that arrives on a timeline you can actually affect.

The knowledge cutoff feels like a wall when you first hit it. It isn't. It's just the boundary of one of the two doors into the model's knowledge — and it's the slow one. The fast door, retrieval, is standing open right now, and walking through it is almost entirely about doing the unglamorous, controllable work of being findable, readable, and consistent on the live web. New businesses don't need the model to remember them. They need to be the answer when it goes looking.

Key takeaways

- Every AI model has a knowledge cutoff — a date after which nothing is baked into its trained memory — so anything you launched recently may simply not exist to the model itself.
- There are two ways an AI can know about you: from training (frozen at the cutoff) and from live retrieval (search/grounding at answer time). New businesses win through the second path.
- If an engine answers from memory alone, a recent launch, rebrand or move can be invisible; if it retrieves, a strong, crawlable, corroborated web presence can get you named the same day.
- Your job is to be the business the retrieval layer surfaces — clear pages, consistent facts across the web, and presence in the third-party sources engines pull from.
- Getting into the next training run is a slow, uncertain bonus; getting into live retrieval is the fast, controllable lever — prioritize it.
- Don't try to correct a model's stale memory by arguing with it — fix the web it reads from, then test whether retrieval picks up the change.

Frequently asked questions

How do I find out a model's knowledge cutoff?

You can usually just ask the assistant directly — most will tell you their training cutoff date — and the providers publish it in their documentation. But the more useful question for your business isn't the exact date; it's whether

the engine is answering you from trained memory or from live retrieval in that moment. A cutoff only hurts you when the engine is relying purely on memory. The second it searches the live web to answer, the cutoff stops being the constraint and your current web presence takes over.

My business launched this year. Does that mean AI can't recommend me at all?

Not at all — it means the trained-memory path is probably closed to you for now, so you have to win through retrieval. When an engine searches the web to answer a question, it doesn't care when you launched; it cares whether you're findable, readable, and corroborated right now. I've seen recently launched businesses get named in AI answers quickly, not because they made it into any training set, but because their live web footprint was clean and strong enough for the retrieval layer to surface. That path is open to you today.

If I get everything right, how fast can AI start naming me?

Through retrieval, it can be fast — sometimes days, once the engines' crawlers have fetched your pages and the corroborating sources are in place. Through training, it's slow and genuinely uncertain: you're waiting on a future model to be built, and there's no guarantee you'll be included or how you'll be represented. My read, not a guarantee: treat training inclusion as a bonus you can't schedule, and put your real effort into the retrieval path, because that's the one you can actually influence on a timeline that matters.

About the author. Scott Tischler is the Founder & Chairman of AIrecommend.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 and professional study at Wharton, Harvard, and Oxford, he helps businesses become the ones AI recommends.

© 2026 Scott Tischler · scotttischler.com
scott@airecommend.ai · AIrecommend.ai