NOTES FROM THE BOT LAB · EDITION OF NOVEMBER 25, 2026

Don't Poke Wet Concrete With a Stick: Why AI Search Has a Curing Window

By Will Tygart · Tygart Media · Canonical edition at https://tygartmedia.com/dont-poke-wet-concrete-ai-search-curing-window/

Pour a concrete driveway on a Tuesday morning, trowel it smooth, and then walk out of the garage every twenty minutes to poke the middle with a broomstick to see if it is hard yet. By suppertime you will not have a faster driveway. You will have a slab full of holes that has to be jackhammered out on Friday.

Watch smart people work on AI search visibility right now and you will see the broomstick come out before the coffee gets cold. Somebody fixes a pricing table or rewrites an opening paragraph at 9:00 a.m. At 9:05 a.m. they open ChatGPT or Claude, type in their favorite test question, see the exact same old answer from last week, decide the edit failed, and rip the page apart again before lunch.

Today we tested that habit with two clocks running at the exact same second — one on the chat window, and one on our raw origin server logs.

The "Searching the Web" Illusion

When modern AI assistants spin a little indicator that says Searching the web, it is completely natural to picture a courier sprinting out to your website, reading your latest HTML, and carrying it back to the screen. Sometimes — when a human pastes a specific URL into a chat box and asks a model to summarize it — a live user-triggered agent really does knock on your front door. We documented thousands of those live fetches in our 25-to-1 crawler-versus-agent server log study.

Open-ended discovery prompts work differently. Today we ran time-sensitive product and technical prompts across ChatGPT, Claude Sonnet with live web search enabled, Grok, Perplexity, and Microsoft Copilot while tailing our Nginx access logs down to the second. During those exact test windows, 154 automated requests crossed our origin server — routine background sweeps from bingbot, OAI-SearchBot, GPTBot, PerplexityBot, and Applebot on older URLs. Know how many synchronous requests hit the exact pages those five engines were actively answering questions about? Zero.

Not one of them visited the origin server during the prompt. Even when Claude logged three explicit search tool calls and pulled seventeen candidate web pages into its working memory, it pulled them out of a pre-built search and snippet cache — not off our live disk.

Why Index Digestion Beats Noon Panic

Once you see that gap in the logs, two mysterious things suddenly make plain sense. First, traditional search indexation still sits upstream of almost everything an AI recommends. Across 295 settled multi-engine test runs in our lab this week, pages that Google Search Console marked as indexed were cited 21.29 percent of the time. Pages still sitting in an unindexed neutral queue were cited 4.88 percent of the time — more than a fourfold drop before anyone even judged the writing.

Second, crawl speed and index refresh run on different clocks across engines, something we first measured in our IndexNow crawl-speed comparison. A bot visiting your server tonight does not mean its downstream answer cache rebuilt tonight. Between crawl queues, snippet extraction, and ranking caches, most structural upgrades need a seven-to-fourteen-day curing window before a recheck tells you the truth.

Leave the Broomstick in the Garage

We wrote recently about knowing when lab mode ends and letting crab pots soak. Search and AI visibility punish fidgeting the exact same way. When you rewrite a page three times in forty-eight hours because ChatGPT has not flinched yet, you are grading wet concrete.

Do the structural work cleanly once: make the opening answer clear, put real comparison numbers in a plain table, verify your sitemap is clean, and step away from the prompt window for two weeks. As we noted in Billboards, Bots, and Heading Toward the Mountain, you are never talking to the exact same model twice anyway — and next door in You Made the Shortlist on the Desk, you will see why even a #1 background search rank needs the right opening paragraph once the slab finally sets.

Stage in the Chain	What People Assume Happens	What the Logs Actually Show	Operator Rule
1. You publish or edit a page	AI models can see the new text immediately	Nothing changes upstream until Googlebot, bingbot, or OAI-SearchBot re-crawls	Submit your XML sitemap and record the timestamp
2. Five minutes later, you test a prompt	Searching the web means the bot is visiting your URL right now	0 synchronous origin hits across ChatGPT, Claude, Grok, Perplexity, and Copilot tests	Do not judge a fresh edit with a same-day prompt
3. Search index digestion	Every engine keeps its own independent live crawler	Indexed PASS pages earn citations at 21.29% versus 4.88% on unindexed pages (4.36x lift)	Fix search indexation before rewriting prose twice
4. Days 7 through 14	If the answer hasn't moved by noon, the copy failed	Background search snapshots refresh asynchronously across crawlers over one to two weeks	Let the concrete cure before running your recheck

Key Questions Answered: Why didn't ChatGPT or Claude show my page edit five minutes after I published it? Most open-ended AI search prompts pull from pre-built search and snippet caches rather than fetching your origin server live on every query, so edits stay invisible until search crawlers re-index the page. **How long should I wait before re-testing an updated page across AI engines?** Give structural upgrades a 7-to-14-day curing window (or confirm fresh Googlebot, bingbot, and OAI-SearchBot hits in your origin logs) before judging citation movement. **Does Google Search Console indexation still affect AI citations?** Yes. In our 295-run multi-engine benchmark, pages indexed in Google Search Console earned citations at 21.29 percent compared with 4.88 percent for unindexed pages — a 4.36x difference.