

AEO

AI



Why ChatGPT has probably never named your practice

By Nathan Woo

THE THESIS

Most independent practices are invisible to AI today. Most practice marketing budgets go to channels that do not shape what AI surfaces when a patient asks. A small number of independent practices show up consistently in AI answers, and they share a specific set of tactics — none of which appear on a typical practice marketing line item. This article walks through what those moves are and why most marketing dollars sit elsewhere.

THE STAKES

One in three U.S. adults has turned to AI chatbots for health information — the same share that uses social media for that purpose.¹ In Google's AI Overviews, brands cited in the answer receive about 35% more organic clicks than uncited brands; when AI Overviews appear and the brand is not cited, organic

click-through has fallen roughly 65% year-over-year.² A first-page ranking is no longer enough; being named in the AI overview is the new goal.

To measure where independent practices stand on that question, we sampled 200 practices at random from the NPES registry across 18 specialties and 20 metros, then ran a fixed battery of patient-realistic prompts (scoped to each practice's specialty and metro) through GPT-5.5 via the OpenAI API in May 2026 — about 4,950 prompts in total, with multiple runs per query to absorb variance. (We used the API rather than the consumer ChatGPT interface so prompts ran at consistent settings across thousands of queries.)

ChatGPT mentioned and cited zero of the 200 practices. (A *mention* is the practice's name appearing in the text a patient reads, as distinct from a *citation* — a URL ChatGPT links to as a source. Mentions are the metric that converts

TABLE 1. UPSTREAM SOURCE POOLS FOR DIFFERENT PATIENT QUESTION TYPES

Question type	Who dominates the named entities	Practice-org share
“Best [specialty] in [metro]”	Hospital-affiliated physicians via regional magazine “Top Doctors” lists	7%
“I have [symptom], who do I see”	Hospital systems’ find-a-doctor directories	12%
“Where can I get [procedure] in [metro]”	A mix – room for independent practices	17%

to bookings; the rest of this article anchors on mentions wherever the data supports it.)

Frankly, this surprised us – we expected the rate to be low, not zero. AI isn’t yet the dominant patient-acquisition channel for most practices, but the share of patient research flowing through it is growing fast enough that the gap is worth closing while it is still small.

AUDIT YOUR OWN PRACTICE IN 10 MINUTES

Before reading further, open ChatGPT and ask:

1. “What’s the best [your specialty] practice in [your metro]?”
2. “I have [a common condition you treat], who should I see in [your metro]?”
3. “Where can I get [a top procedure you do] in [your metro]?”
4. “Who are the top [your specialty]s in [your metro]?”

Count how many of these queries name your practice. If the count is zero, you share a baseline with most of the 200 practices we tested.

The mechanism: the questions a patient asks decide which source pool ChatGPT draws from

The 0/200 result becomes clearer when you flip the experiment: start from the practices ChatGPT does name, and work backwards. That corpus covered ~35,000 queries, ~18,000 distinct named entities, and ~41,000 mentions across nine recurring patient-question patterns – pure discovery, “best” reputation, condition-driven, procedure-driven, comparison, cost, symptom, and others. Patients ask in more shapes than the nine we tested; even within this slice ChatGPT drew on a noticeably different upstream source pool for each. [See the contrast in Table 1.]

The full nine-row breakdown shows similar shape across the other six patterns. Hospital-affiliated doctors are the largest bucket in

four of the nine cohorts (41%–55%). Hospital systems own symptom queries at 49%. Independent practice organizations sit at 7%–17% across every cohort and are never the largest.

For “best dermatologist in [metro]” queries, ChatGPT consistently cites the regional city magazine’s annual “Top Doctors” list and names the listed physicians. For symptom queries, it cites the local hospital system’s find-a-doctor directory or a Mayo Clinic patient-education page. For procedure queries the source pool fragments more – the only cohort where independent practices’ own procedure pages show up in volume.

There is no such thing as “AI visibility” in the singular: different patient questions feed different source pools, and most practice marketing spend – Google Ads, generic SEO, vendor-managed Healthgrades profiles – doesn’t change the pool for any specific cohort.

YOUR SPECIALTY AND MARKET DETERMINE YOUR CEILING

Practice-org mention share varies by about 2× across specialties and 3× across metros. Dermatology, ENT, and ophthalmology sit at 16%–17%. Plastic surgery, cardiology, and family medicine sit at 8%. The drivers are familiar: specialties dominated by academic medical centers (cardiology), or where individual physician brands eat the oxygen (plastic surgery), leave less space for practice organizations.

Markets vary even more. Charlotte’s practice-org share is 23%; Boston’s is 7%. Metros with high AMC density (Boston, San Francisco, Chicago) crowd out independent practices. Metros with fewer dominant AMCs leave room for practice brands to accumulate share. An administrator benchmarking against the cross-market average will mis-estimate their own ceiling – a Boston cardiology practice at 7% is doing well for its market; a Charlotte ENT practice at 7% is underperforming.



➤ 3 THINGS THE VISIBLE PRACTICES DO THAT MOST DO NOT

Lever 1: Page-per-procedure content depth.

Among practices in our sample that ChatGPT cited at least once, the share of citations landing on the practice’s own domain (rather than on Healthgrades or Birdeye) scales smoothly with how many distinct own-domain pages ChatGPT finds worth quoting. A practice with one services page listing 20 procedures has one page available to the model regardless of which procedure the patient asked about. Building a substantive page per procedure gives the model 20 distinct pages, and the data suggests the model uses them. As with any patient-facing material, clinical content on these pages should be drafted or reviewed by qualified clinicians and kept current with practice guidelines.

Even when ChatGPT names a practice, the description often comes entirely from third-party listings. One Phoenix dermatology practice accumulated 21 mentions in our corpus with zero own-domain citations – the description came from Healthgrades and Birdeye boilerplate, the same lines those sites publish for every practice in the metro. Three substantive own-domain pages (a practitioner profile, a top-procedure page, a condition landing page) give the model something specific to quote instead.

Lever 2: On-page structural signals. Of 516 cited practices, 66.1% had none of the five on-site features we tested (see Table 2).

Structured data (e.g., schema markup) give search systems explicit clues about a page – which physician, which procedure, which credentials. Pages without it may be harder for search and answer systems to interpret



as specific entities; pages with it offer clearer entity signals that can help match the content to specific patient-question patterns. The data show that practices that combine content depth with markup features appear in cohort-specific responses more often than either factor alone would predict, and adoption is rare enough that the technique still discriminates between sites that have it and sites that do not.

Lever 3: Top Doctors and city magazine inclusion. Half of “best specialty” mentions are hospital-affiliated physicians, and the mechanism that puts them there is the regional magazine “Top Doctors” pipeline – the city magazine publishes its annual list, ChatGPT cites it, the listed physicians become the named entities. This is the only counterplay independent practices have on reputation queries. The concrete move: nominate your physicians

TABLE 2. ON-SITE FEATURES ASSOCIATED WITH AI VISIBILITY

Feature	Incidence
Long-form content (≥500 words on entity pages)	29.7%
On-page Q&A blocks	7.2%
FAQPage schema	4.8%
Hero-intro paragraph with named physician + credentials	2.3%
Physician schema	0.6%
None of the five	66.1%



An independent practice isn't going to displace Mass General on "best dermatologist in Boston." The contestable ground is **procedure queries — patients searching for a specific treatment, which is also where commercial intent is highest and bookings actually come from.** To get named in those procedure-query responses, work all three levers: **content depth, structural markup, and editorial inclusion.**

annually for Castle Connolly and your regional city magazine's "Top Doctors" issue — both accept peer nominations from other physicians in your specialty. Calendar it the same way you calendar CE credits. A follow-up article will go deeper on the off-site terrain — which lists, what nomination cadences work, and how ChatGPT's preferred sources shift across specialties and metros.

Most practice marketing budgets currently allocate to Google Ads, Healthgrades premium profiles, and a generic SEO retainer. Relatively few practice marketing plans prioritize any of the three levers above.

WHAT HAPPENS WHEN THESE LEVERS ARE DEPLOYED

Disclosure: Halcy is my company, and what follows is observational customer data, not third-party research. We deploy the three levers as our core service.

Across our live customer base, recent-30-day mention rates run 25% above early-monitoring baselines after schema and content-depth deployments. AEO citation rates climb on a similar trajectory, typically reaching the mid-twenties or higher within the first 60–90 days of active deployment. *Caveats: ChatGPT model updates over the window are confounders, and the customer base is not a random sample. The trajectory is directionally consistent with the corpus findings.*

An independent practice isn't going to displace Mass General on "best dermatologist in Boston." The contestable ground is procedure queries — patients searching for a specific treatment, which is also where commercial intent is highest and bookings actually come from. To get named in those procedure-query responses, work all three levers: content depth, structural markup, and editorial inclusion.

METHODOLOGY AND CAVEATS

The random-sample experiment used 200 NPPEs-sampled practices and ~4,950 queries; the cohort analysis used ~35,000 queries across nine question patterns and ~41,000 mentions. All queries ran against GPT-5.5 via the OpenAI API in May 2026; behavior on Perplexity, Gemini, and Google AI Overviews may differ. Cohort data is survivor-biased. Lever analysis is correlational. ■



Nathan Woo is the founder of Halcy, an AI growth platform for healthcare. www.halcy.ai

NOTES

1. Montero A, Montalvo J, Kearney A, Valdes I, Kirzinger A, Hamel L. "KFF Tracking Poll on Health Information and Trust: Use of AI for Health Information and Advice." KFF. March 25, 2026. Available from: <https://www.kff.org/public-opinion/kff-tracking-poll-on-health-information-and-trust-use-of-ai-for-health-information-and-advice/>
2. McDonald T. "AIO Impact on Google CTR: September 2025 Update." Seer Interactive. Nov. 4, 2025. Available from: <https://www.seerinteractive.com/insights/aio-impact-on-google-ctr-september-2025-update>

Disclosure

Halcy is a commercial service; the practices they work with pay for these deployments. The research corpus was conducted by Halcy and is being shared publicly so administrators can audit their own AI presence regardless of vendor relationship.