

# Where ChatGPT gets its answers:

**THE OFF-SITE MAP THAT DECIDES WHICH INDEPENDENT PRACTICES GET CITED**

By Nathan Woo



## THE THESIS

There is no off-site playbook. There is your specialty’s off-site playbook in your metro, and most of what an outside vendor would sell you isn’t on it. The biggest off-site citation lanes are largely unmanaged, no or low cost, and almost never sold as a service. Once the shape of your cell is visible, three concrete moves emerge — and none of them is what most administrators expect.

## THE CHATGPT OFF-SITE MAP

Most agencies pitching independent practices haven’t updated their playbook for AI search at all — they sell the same local-SEO bundle they have sold for years: Google Business Profile optimization, Healthgrades premium, Yelp management, paid directory placements. Practice owners reasonably assume those investments translate to ChatGPT visibility. Our citation corpus shows they largely don’t.

Across 33,340 ChatGPT citation events, national medical directories account for 8.3% of the citation surface — the fourth-largest source bucket, not the first (See Table 1).

Figure 1 shows the distribution within the hospital-system slice. The 27.5% is the corpus-wide share across all 200 practices; each cell in Figure 1 is the share for a single (metro ×

A note on terms before the data. “On-site” means everything on a practice’s own website — the procedure pages, physician bios, location pages, and the structured data behind them. “Off-site” means everything ChatGPT cites about a practice that isn’t the practice’s own site — the hospital roster that lists a physician, the specialty-association directory, the state medical board profile, the city magazine’s “Top Doctors” issue.

The [previous article](#) mapped the on-site work that consistently-cited practices share, and drew the distinction between a mention — your name in the answer a patient reads — and a citation, the source ChatGPT actually pulls from. This article focuses on the off-site surfaces that govern citations. ChatGPT can only name a practice it has a source for, so that surface determines which practices can be mentioned at all. This is the map of that supply.

specialty), and that per-cell view hides a large range — below 10% (Phoenix × cardiology) to above 80% (St Louis and Detroit family medicine, Atlanta cardiology).

Three things stand out from this distribution. First, the biggest off-site lanes are the ones most independent practices do not actively manage. Hospital-system staff rosters and specialty-association credentialing pages together carry 49.6% of the cited surface; most practices have zero or one vendor managing these. National directories carry 8.3%,




**TABLE 1. SHARE OF CHATGPT CITATION EVENTS BY SOURCE BUCKET**

Source bucket	Share of citation events
Hospital systems + academic medical center staff rosters	27.5%
Practice-owned websites (independent solo + multi-specialty groups)	23.0%
Specialty associations + credentialing authorities (aad.org, acog.org, abms.org, heart.org, etc.)	22.1%
National medical directories (Healthgrades, Zocdoc, WebMD, RealSelf, Vitals, Doximity)	8.3%
State medical boards + state .gov	5.1%
Review aggregators (Birdeye, Tebra, Chamber of Commerce)	3.8%
Federal authority (NIH, CDC, FDA)	3.3%
Regional city magazines	2.7%
Ranking aggregators + magazine doctor-finder lookups (Newsweek rankings, magazine directory, subdomains, RealPatientRatings)	2.3%
Lab and imaging chain locators (LabCorp, Quest, RadNet, SimonMed)	1.3%
Other (Castle Connolly, local news, miscellaneous)	0.6%

**FIGURE 1. HEALTH-SYSTEM CITATION SHARE BY (METRO X SPECIALTY)**

Cells colored by share of the cited surface routing through the dominant local hospital systems and academic medical centers.

	cardiology	gastroenterology	dermatology	family-medicine	internal-medicine	ob-gyn	ophthalmology	ent	orthopedic-surgery	plastic-surgery
new-york	54%	43.9%	18.5%	33.4%	46.9%	26.7%	32%	47.8%	50%	30%
los-angeles	51%	45.1%	21.1%	36.3%	40.8%	37.5%	32.2%	26.6%	40.2%	37.3%
chicago	62.9%	53.4%	22.6%	40.3%	53.7%	41.1%	25%	47%	21%	14.1%
dallas-fort-worth	51%	34.3%	17.5%	30.7%	35.2%	26.3%	15.4%	24.8%	18.7%	4.8%
houston	48.3%	46.1%	12.1%	41%	53.6%	51%	22.4%	52.1%	39.5%	10.6%
washington-dc	59.4%	50.3%	4.2%	19.3%	35.4%	23.9%	21.9%	34%	28.4%	11.5%
philadelphia	64.9%	58.2%	11.1%	22.7%	58.8%	33.6%	21.2%	41.2%	39.8%	14%
atlanta	76.7%	35.9%	4%	41.4%	43.6%	17.9%	12.5%	30.4%	41.9%	1.5%
boston	55.7%	52.7%	26.9%	44.3%	49%	56.6%	42.4%	62.1%	42.9%	24.5%
phoenix	25.8%	24.6%	4.4%	35.8%	26.7%	12.9%	0.3%	9.4%	23%	0%
detroit	60.6%	53.7%	10.8%	67.6%	42.7%	37%	25.5%	41%	60%	13.9%
seattle	48.5%	39.7%	7.6%	36.4%	40.2%	29.9%	4.3%	32.5%	30%	2.9%
minneapolis-st-paul	65.3%	33.3%	14.8%	40%	36.8%	40.7%	8.5%	35.8%	43.1%	6%
san-francisco-oakland	45.2%	24.7%	24.4%	38.3%	51.9%	30.8%	25.5%	28.3%	38.9%	5.9%
tampa	64.2%	36.8%	12.1%	58.6%	51.4%	26.1%	13%	48.6%	29.9%	5.5%
san-diego	45%	36.5%	18%	36.4%	50.6%	35%	21.7%	30.8%	41.7%	2.9%
denver	41.6%	31.4%	12.8%	31.3%	43.3%	36.3%	21.6%	44%	68.1%	11.8%
baltimore	65.9%	33.3%	18.2%	31.3%	29.5%	25%	40%	40.3%	58.5%	17.1%
st-louis	47.1%	36.3%	11%	87.8%	32.1%	33.3%	12.2%	32.9%	26.2%	0%
charlotte	44%	28.6%	2.4%	16.7%	40.3%	47.4%	2.2%	28%	39.5%	3.1%



➤ and most practices have two or three vendors there. Second, review management is a minor citation lane, despite being the largest vendor category aimed at practices. Review aggregators are 3.8% combined; Yelp is 0.11%; GBP has zero direct citations because Google does not expose reviews as crawlable URLs. Reviews remain critical for local SEO and for persuasion once a practice is in a ChatGPT answer, but they are not where the AI-citation surface lives. Third, no single off-site source

dominates the way administrators expect — the largest buckets are hospital rosters and specialty associations, not directories or ranking sites, and which one leads varies sharply by specialty and metro.

An interactive companion at [halcy.ai/research/chatgpt-citation-map](https://halcy.ai/research/chatgpt-citation-map) lets you pivot the source-bucket map by specialty — including the four specialties this article doesn't enumerate.



**TABLE 2. OFF-SITE COMPOSITION VARIATION BY SPECIALTY**

Specialty	Practice-owned	Specialty association	Hospital system	Regional magazine	Federal authority
Plastic surgery	28.6%	30.1%	8.4%	1.7%	1.3%
Dermatology	31.8%	19.4%	11.2%	6.3%	5.3%
Cardiology	9.4%	22.6%	33.4%	5.1%	6.2%
Orthopedic surgery	15.6%	20.7%	28.6%	7.3%	7.0%
Family medicine	6.7%	19.2%	23.8%	6.0%	7.9%
OB-GYN	12.8%	24.1%	23.5%	6.4%	3.7%

**THE ON-SITE MULTIPLIER**

The off-site terrain decides part of the story. The on-site layer compounds it through three mechanisms.

First, on-site content is itself a citation surface. More entity-rich pages mean more URLs ChatGPT can ground on — the 23% practice-owned share above is a snapshot, not a ceiling. On-site investment expands the share itself.

Second, on-site content propagates outward. Hospital rosters, specialty associations, magazines, and directories describe practices based on what those practices publish. Rich on-site content shows up in the off-site descriptions ChatGPT then cites. On-site investment compounds into the off-site 77% — the share of the cited surface that isn't the practice's own site — indirectly.

Third, when a practice is cited, the description reflects what the source page provides. In one corpus query for “best mole removal in Phoenix,” ChatGPT cited five dermatology practices from their own websites. One was described as having “fellowship-trained Mohs surgeons” and a named physician “board-certified and fellowship-trained in Mohs/dermatologic oncology.” Another comparable practice in the same metro was described

as “a reasonable option if location is convenient” offering “board-certified care across Phoenix-area locations.” Same query, same metro, same specialty, same source type. The most visible difference between the two descriptions was what each website made available for ChatGPT to ground on.

Off-site moves earn baseline visibility. On-site work multiplies it across all three mechanisms. The rest of this article maps the off-site terrain because that's where administrators have least visibility — but the on-site layer is the multiplier.

**YOUR SPECIALTY DETERMINES YOUR PLAYBOOK**


The off-site composition shifts sharply by specialty. The same tactical advice works differently for two practices in the same metro.

The practice-owned column tracks citations to independent and multi-specialty group practice websites (Duly Health, Kelsey-Seybold, single-physician domains, doctor-owned multi-location groups). The variance by specialty is the story.

Plastic surgery and dermatology lead the practice-owned column (29-32%): most of the cited surface in these specialties is independent. Specialty-association pages carry a meaningful

**FIGURE 2. TOP SOURCE BUCKETS PER SPECIALTY (% OF CITED SURFACE)**

*The composition shifts so sharply that the same tactical advice works differently for two practices in the same metro.*



**TABLE 3. REGIONAL MAGAZINE CITATIONS BY SAMPLED METROPOLITAN AREAS**

Metro (sampled)	Dominant magazine	Citations
San Diego	sandiegomagazine.com	202
Seattle	seattlemag.com + seattlemet.com	211
Washington, D.C.	washingtonian.com	95
Boston	bostonmagazine.com	79
Philadelphia	phillymag.com	62
Dallas-Fort Worth	dmagazine.com + directory.dmagazine.com	50
Chicago	chicagomag.com	43
Atlanta	atlantamagazine.com	33

➤ slice for plastic surgery (30.1%), but the regional-magazine pipeline is thin (1.7%). The off-site moves available concentrate in specialty-association credentialing depth.

Cardiology and orthopedic surgery lean hospital-system. A third of cited URLs for cardiology and 29% for orthopedics are AMC or hospital-network staff rosters; the practice-owned slice drops to single digits for cardiology (9.4%) and the mid-teens for orthopedics (15.6%). Independent practices that show up share two characteristics: deeper specialty-association presence (SCAI for interventional cardiology, AAOS for orthopedics — roughly 21-23% of cited URLs), and partial AMC affiliation where available.

For an unaffiliated cardiology or orthopedic practice in a metro without a magazine pipeline, the realistic off-site lanes narrow to specialty-association credentialing and partial AMC affiliation. There is no off-site lever here at the scale of a hospital-system roster, so the answer runs through on-site authority.

Family medicine and OB-GYN sit between these poles. Family medicine has the lowest practice-owned share (6.7%) and the strongest federal-authority share (7.9%) — NIH, CDC, and government patient-education pages get cited as grounding for primary-care answers, rarely for subspecialty queries. State medical boards feature more heavily for primary care (16-20%) than for any subspecialty.

The off-site playbook is per-(specialty × metro). The version for Boston cardiology has almost nothing in common with Dallas dermatology, so figuring out which cell you’re in is the first question.

### THE MAGAZINE PIPELINE — WHERE IT EXISTS

The standard “Top Doctors” story runs through Castle Connolly. In our corpus, Castle Connolly and adjacent national “Top Doctor” sites account for 14 citations across 3 domains — 0.04% of the cited surface. The regional city

magazines (which source their lists partly from Castle Connolly and partly from peer nominations) account for roughly 60 times that: 886 citations across 29 publications.

The pipeline is metro-specific. Within our 20-metro sample, half the metros show negligible regional-magazine citations (see Table 3).

This is illustrative, not exhaustive. Many MGMA-member metros not in our 20-metro sample also publish Top Doctors issues (*Indianapolis Monthly*, *Cincinnati Magazine*, *Minneapolis/St. Paul Magazine*, *Houstonia*, etc.). The right inference is “verify directly for your own metro,” not “use this list as definitive.”

Within our sample, the pipeline is absent in NYC, LA, St Louis, and Baltimore. The NYC citation surface concentrates in NYU Langone, Mount Sinai, and Weill Cornell staff directories despite the city’s editorial density.

The pipeline runs like this: a practice nominates physicians, peer review filters the list, the magazine publishes, and ChatGPT cites it for years afterward. The lever is knowing the nomination deadline and submitting on time — the window typically closes six to nine months before the issue runs.

In metros with a strong pipeline, missing the window is a measurable miss in citation share. In metros without one, the off-site work falls back on specialty-association credentialing and state-board profiles.

### THREE OFF-SITE MOVES TO START WITH

Three concrete off-site moves emerge that most practices underinvest in.

The first is to audit your specialty-association credentialing pages. Confirm that ABMS, your specialty board, and any major society directory you qualify for (AAD, ACOG, ASPS, SCAI, AAOS, ASOPRS, and so on) reflect current credentials, current practice addresses, and subspecialty fellowships. The combined citation surface is 22.1% — more than every national





➤ directory put together, and the work boils down to practice operations, not vendor spend. The most common stale data point is an old practice address from a prior employer, which ChatGPT reproduces faithfully. When a physician leaves a practice, update each association profile to the new location and ask the previous employer to take down the physician’s bio page — otherwise the old page keeps feeding ChatGPT the wrong address.

The second is to find out whether your metro publishes a “Top Doctors” issue and, if it does, when nominations open. The deadlines are public but obscure, and most practices miss them not because the dates are hard to track but because nobody at the practice realizes the issue affects ChatGPT citations for years afterward. Set a recurring reminder months in advance; the window typically closes six to nine months before publication.

The third is to audit your state medical board profile. It is free to access, verifiable in minutes, and 5.1% of the citation surface — yet most practices have never checked what the board shows publicly. Whatever is there is what ChatGPT trusts over any vendor-managed directory entry.

These three moves cover roughly 30% of the entire citation surface that an independent practice can directly act on (22.1% specialty associations + 2.7% regional magazines + 5.1% state boards). Add the hospital-system roster lane (27.5%) for any physician with hospital privileges or an academic affiliation, and the addressable surface rises past half (57.4%). Almost none of it is sold as a vendor service, which is part of why most independent practices have nobody managing these lanes.

## METHODOLOGY AND CAVEATS

Exploratory secondary analysis of 33,340 ChatGPT citation events drawn from 4,950 patient-realistic queries across 200 NPI-stratified practices, 20 metros, and 10 specialties (May 2026, GPT-5.5 via OpenAI API). Citation events were grouped into source buckets (hospital systems, specialty associations, national directories, state medical boards, regional magazines, and so on) and analyzed in aggregate and by (specialty × metro) cell. The interactive companion at [halcy.ai/research/chatgpt-citation-map](https://halcy.ai/research/chatgpt-citation-map) carries the full per-cell breakdown for all 10 specialties.

Single-platform, cross-sectional, US-only. Behavior on Perplexity, Gemini, and Google AI Overviews may differ. Aggregate and per-specialty findings generalize across the sample; specific metro-level findings are illustrative, not exhaustive. The analysis is correlational — no causal claims. ■



Nathan Woo is the founder of Halcy, an AI growth platform for healthcare. [www.halcy.ai](https://www.halcy.ai)

### 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.