

How AI assistants actually recommend dentists, and what the evidence says you should do about it

A review of how ChatGPT and Google's AI pick a dental practice, an original 546-run study of what they cite, and an honest list of what moves the answer.

Streben Marketing · September 2026 · For dental practice owners

00 THE SHORT VERSION

Patients increasingly start by asking ChatGPT or Google's AI which dentist to see. The question for a practice owner is direct: how do you become the one it recommends? The usual advice is to stop worrying about your website, because the AI ignores it. For dentistry that advice is mostly wrong.

We ran the measurement across 12 markets and 546 responses. On **"best dentist in [city]," ChatGPT cited a practice's own website in 1 of 63 runs.** That tier belongs to Healthgrades, WebMD Care, hosted review pages, and city magazine "Top Dentists" lists. But on **"dental implants in [city]," it cited the practice's own site in 90% of runs,** and on a full-arch specialty query in **100%.** Once a patient's question gets specific, your service pages win the answer outright.

Two results were new, and they are the two most useful things in this paper. **"Emergency dentist for a cracked tooth in [city]" cited a practice's own website in 63 of 63 runs, with zero directory citations.** It cited no emergency-dentist lead-generation site in any national run, and almost nobody has built a real page for this query. In the other direction, **"dentist that takes Delta Dental in [city]" cited a practice website in 0 of 63 runs.** The carrier's own find-a-dentist tool answered it in 62 of 63. That tier is not a marketing problem; it is a network-participation and records-accuracy problem, with one exception we will come to that matters a lot for smaller markets.

We also found that **no advertiser bought a placement inside any of our 546 answers.** The directories and booking marketplaces that hold the generic tier have not started buying their way into the answer itself. Treat that as a window rather than a permanent condition.

One caveat before the evidence, so none of this reads as "SEO is dead." This paper is about the AI recommendation layer. Ordinary Google and Bing results still send a large share of patients to a practice, and they reward much of the same on-page work. The posture is additive: keep standard SEO, and add what follows for the AI layer.

01 WHY THIS IS WORTH A PRACTICE OWNER'S ATTENTION

The AI layer now sits upstream of the phone call, the insurance lookup, and often the first search result, and it is the part of the funnel most practices are not yet watching.

Adoption is climbing fast. In a survey of nearly 1,000 US patients, **47% said they had used AI to find a new provider**, up from 31% nine months earlier.¹ A separate representative survey found **45% of consumers had used AI for local business recommendations, up from 6% a year earlier.**² Those patients behave consistently. They ask the assistant, then verify on review sites before they call.²

Dentistry has two demand patterns that most local businesses do not, and the AI layer treats them as opposites. The first is **urgency**. There are roughly **2 million emergency-department visits a year in the US for dental pain**, which the ADA notes could largely be redirected to dental offices.⁴ Someone with a cracked tooth at 8pm is not browsing; they are asking one question and calling the first credible answer. The second is **coverage**. Roughly **284 million Americans, 83% of the population, carry some form of dental benefit**, though enrollment fell 2.3% in the most recent year measured.³ A large share of dental search is not "who is good" but "who takes my plan." We measured both, and they are the two tiers where the usual advice is most wrong.

02 HOW AN AI ASSISTANT PICKS A DENTIST

The mechanism is more legible than the hype suggests, and it is different from classic SEO. Two facts drive everything below.

First, for genuinely local queries, Google usually does not serve an "AI Overview" at all. It serves the local pack. In a 540-query study across six local industries, local-intent queries returned an AI Overview only **15% of the time, while the local 3-pack appeared 93% of the time.**⁵ So for "dentist near me" typed into Google, the surface you compete on is still Google Maps and the Business Profile. Asking ChatGPT or Perplexity directly is the separate case Section 3 measures.

Second, when AI does answer, it grounds heavily on structured place data and a small set of trusted third parties. Across 1.3 million Google AI Mode citations, **google.com was the single most-cited domain at 17.4%, more than the next six domains combined**, resolving to Business Profiles and knowledge panels.⁶ An independent analysis of 32 million observations measured an **8.4x surge** in Google citing its own place cards this way, concentrated in local verticals including healthcare.⁷ And as of July 2026 ChatGPT reads a licensed Yelp feed of hundreds of millions of reviews directly.⁸

None of those inputs is your website copy. That is the part most people already know. What almost nobody has measured is where the line falls. Which dental questions does the place-data-and-directory layer answer, and which ones do your own pages answer? That is what we went and measured.

03 OUR STUDY: WHAT CHATGPT CITES, BY QUERY TYPE

No independent measurement of this existed for dentistry. What exists online is agency marketing content, much of it repeating figures with no traceable source, including a widely circulated claim that structured FAQ content raises AI citation by 37 to 40%, which Section 8 shows is fabricated. So we ran the study ourselves.

METHOD

We queried ChatGPT through the DataForSEO scraper (natural web-search behavior, US) across **12 markets**, nine national metros plus three smaller towns, and **five query tiers**, running each query **7 times** to measure run-to-run variance. That is **546 responses, zero failures**. For each answer we separated the *cited web sources* (the pages ChatGPT grounded on) from the *named businesses* (practices it listed from place data), and classified every source as a practice's own site, a directory, an insurance carrier, a forum, editorial content, a medical authority, or news. Source domains were then re-classified against a reviewed list, so that aggregator sites which superficially look like a practice's own page were coded as directories, and every state Delta Dental member company was coded as an insurer. That validity step is load-bearing. Before it, the generic tier appeared to cite practice sites 11% of the time and the insurance tier 5%. After it, the true figures are 2% and 0%. The five tiers sample the real spread of dental demand: generic ("best dentist in [city]"), category ("dental implants in [city]"), specialty ("All-on-4 full arch dental implants in [city]"), urgency ("emergency dentist for a cracked tooth in [city]"), and insurance ("dentist that takes Delta Dental in [city]"). Full market list, coding rubric, override log, and per-run data are in the study files.²⁰

Table 1. Cited-source distribution by query tier, 9 national metros, 63 runs per tier. The middle columns are the share of cited web sources; the last column is the share of runs that cited any practice's own site.

QUERY TIER	PRACTICE SITE	DIRECTORY	INSURER	OTHER	RUNS CITING ANY PRACTICE SITE
"best dentist in [city]"	1%	87%	1%	9% news, 2% authority	2% (1/63)
"dental implants in [city]"	87%	9%	0%	5% authority	90% (57/63)
"All-on-4 full arch dental implants in [city]"	97%	3%	0%	<1% news	100% (63/63)
"emergency dentist for a cracked tooth in [city]"	97%	0%	0%	3% authority	100% (63/63)

QUERY TIER	PRACTICE SITE	DIRECTORY	INSURER	OTHER	RUNS CITING ANY PRACTICE SITE
"dentist that takes Delta Dental in [city]"	0%	0%	100%	0%	0% (0/63)

Read the second row first, because it is the one that decides how you spend your money. **"Dental implants in [city]" is website-owned.** Your own service page is the source of the answer in 9 out of 10 runs. Nothing sits between you and that answer. The consumer directories that do exist, Healthgrades, WebMD Care, hosted review pages, and a long tail of city-and-procedure landing pages, hold the generic tier and almost nothing else. Once the question names a procedure, they disappear from the answer.

Which makes the practical map simple. Each tier is won somewhere specific, and only two of the five are won anywhere other than your own pages:

Table 2. Who supplies the answer at each tier, and where the work sits.

PATIENT'S QUESTION	WHO THE AI CITES	YOUR LEVER
"best dentist"	Review directories and "Top Dentists" lists	Reviews, Google Business Profile, directory records
"dental implants"	Your service page	A specific, descriptive procedure page
"All-on-4 full arch"	Your service page	A dedicated page per procedure, not one combined page
"emergency dentist, cracked tooth"	Your emergency page	Hours, walk-in policy, same-day availability in plain text
"takes Delta Dental"	The insurance carrier	Network participation and carrier record accuracy

Three of the five tiers are decided on pages you control. The other two are decided by things you also control, just not through your website. For a dental practice, service pages are the main lever rather than a consolation prize.

NAMED IS NOT CITED

On the generic query ChatGPT still named practices in 62 of 63 runs, averaging **11.8 named practices per answer**, while citing a practice's own site in one. Across those 63 runs there were **28 different top-named practices**, so the "winner" of "best dentist in [city]" is close to a coin flip from one run to the next.

Appearing in the answer and being the source of the answer are two different games, and the generic tier is the former.

THE "BEST" EFFECT (CONTROLLED TEST)

A superlative is a request for a ranking, and models answer ranking requests from ranking sources when those exist. Toggling only the word "best" against a fixed base (three metros, 21 runs per cell):

Table 3. Practice-site citation, descriptive base vs adding "best."

QUERY LEVEL	DESCRIPTIVE BASE	+ "BEST"	CHANGE
"dentist in [city]"	0%	0%	0
"dental implants in [city]"	95%	62%	-33
"All-on-4 full arch dental implants in [city]"	100%	86%	-14

The effect is real but survivable. The likely reason it is not worse is supply. Practices have published so much implant content that it outweighs the thin "best implants in [city]" listicle corpus, so the model bends toward ranking sources without abandoning practice pages. Write your headings descriptively anyway. A superlative heading costs you about a third of your citation rate on the query type you most want to win.

Three honest limits. This is one model, one instrument, US-only, sampled on one day; the scraper targets the search-enabled surface, so it under-counts answers given with no web search at all (3 of our 546 runs returned no web sources). Run-to-run variance is real, which is why every figure here is a rate across seven runs rather than a single observation. And we tested one insurance carrier and one emergency phrasing, so read Sections 4 and 5 as strong signals about a mechanism rather than as precise rates for every plan and every emergency wording.

04 THE EMERGENCY QUERY IS THE LARGEST UNCONTESTED OPPORTUNITY WE FOUND

"Emergency dentist for a cracked tooth in [city]" cited a practice's own website in **63 of 63 national runs**. Practice sites made up 97% of all cited sources. **It cited no directory at all in any national run**, including the emergency-dentist lead-generation sites that dominate paid search for this query.

100%

of runs cited a practice's own site on the emergency query²⁰

0

directory citations across 63 national emergency runs²⁰

2M

US emergency-department visits a year for dental pain⁴

The mechanism is straightforward. The query is urgent and factual. To answer it the model needs hours, walk-in policy, and whether someone can be seen today, and those facts exist only on practice websites. A directory listing cannot answer "can I be seen tonight."

The practices winning it are the ones that built a real page for it. The most-cited domains in this tier were purpose-built emergency pages, alongside general practices with a dedicated emergency section that states hours and same-day availability plainly. This is the clearest instruction in the paper. **If you accept emergency patients and do not have a page that says so, with your hours, your walk-in policy, and what to do for a cracked tooth or a knocked-out tooth, you are leaving a 100%-first-party query type on the table.** It is also the query attached to the highest-intent patient you will get all week.

05 THE INSURANCE QUERY IS A SHUTOUT, AND WHAT TO DO INSTEAD

"Dentist that takes Delta Dental in [city]" cited a practice's own website in **0 of 63 national runs**. Every cited source was a payer property. **deltadental.com appeared in 62 of the 63 runs**, with the remainder its state member companies. The model averaged 1.65 cited sources on this tier, the fewest of any. It finds the carrier's find-a-dentist tool and stops.

This is a different problem from the generic tier. There is no listicle to out-rank and no review score to raise. The levers are **being in the network** and **having the carrier's provider record correct**, which is an operations task rather than a marketing one. If your address, phone, or accepting-new-patients status is stale in a carrier directory, that is the record the AI reads.

THE SMALL-MARKET EXCEPTION

In the three smaller towns we tested, practice websites were cited in **10 of 21 insurance runs, 48%**, against 0% across the nine large metros. In one town it was 6 of 7 runs. The reason is supply. Where a carrier's directory returns only a handful of local providers, the model keeps looking, and it finds practice pages that state which plans they accept in plain, readable text. **Publishing your accepted plans as on-page text is close to worthless in Denver and close to decisive in a town of 15,000.** If you are a single-location practice outside a major metro, this is a page you should write this month.

06 REVIEWS AND PLACE DATA ARE THE GATE ON THE GENERIC TIER

The generic tier is the one your website cannot win. Two things win it instead, and neither is on your site: your review rating, and your presence in the place data and directories the AI reads.

4.3★

Average rating of businesses
ChatGPT recommends⁹

4.1★

Average rating Perplexity
recommends⁹

3.9★

Average rating Gemini recommends⁹

A 350,000-location study across industries found reviews function **"as a filter rather than a ranking signal"**: the AI gates on rating, then chooses from what is left.⁹ Volume matters too. In a corrected study of 1,120 query executions, businesses appearing in Google's AI Mode had **49% more reviews at the median** than those appearing only in the traditional local pack.¹⁰ Review generation is therefore not optional for AI visibility, and it is off your website. (Correlate, not proven cause: review count also tracks a practice's age and prominence.)

In our generic-tier runs the cited sources were led by a small, stable set: Healthgrades, WebMD Care, hosted review pages, aggregated "find a dentist" platforms, and, in one metro, the city magazine's annual "Top Dentists" feature, which appeared in all seven runs. Those are the properties worth keeping accurate and well-reviewed. You will not out-rank them with a service page, and you do not need to, because they only hold this one tier.

A NOTE ON PAID MARKETPLACES

Booking marketplaces charge a practice a one-time fee each time a new patient books, applied at the moment of booking rather than at the visit, because it is a marketing fee rather than a fee for care.¹² Public rate cards do not exist, but per-new-patient fees are consistently reported across third-party coverage in the range of roughly \$35 to \$110, with general dentistry in large markets at the upper end.¹³ That is the market rate someone else is willing to pay for the patient who found you through a directory. Our study is the argument for spending some of that budget on the tiers you can own outright instead, because those tiers are most of the map.

07 HOW NARROW THE AI FUNNEL REALLY IS

Before spending heavily to "win AI," it helps to see the size of the prize. The same 350,000-location study measured how often each platform recommends any given business:

Table 4. Share of locations recommended, by surface.⁹

SURFACE	LOCATIONS RECOMMENDED
Google local 3-pack	35.9%
Gemini	11%
Perplexity	7.4%
ChatGPT	1.2%

The study frames AI visibility as **3 to 30 times harder** than ranking in traditional local search.⁹ (These percentages come from one vendor's dataset, the largest of its kind; read them as directional and the pattern as solid.) That selectivity is why the tier-by-tier picture matters. Competing for the generic recommendation means

fighting for a 1.2% slot against every practice in your city. Competing for "emergency dentist for a cracked tooth in your town" means writing one page, in a tier where every cited source is somebody's practice page and no directory holds any of it.

08 WHAT THE SCIENCE SAYS ABOUT "OPTIMIZING" YOUR CONTENT

There is a genuine academic literature on generative engine optimization, and it is far more skeptical than the agency version. One structural fact explains most of the gap. **Every study that reported a positive optimization effect on a real commercial engine did so by skipping the retrieval step**, handing the engine the document and then measuring what it did with it. One large study that put retrieval back in, across 171,003 documents, found that rewriting a page to "optimize" it **reduced** its presence in results by roughly 9% and its citation rate by 6%.¹⁴ A benchmark of 1.9k queries with proper significance testing found only **3 of 54 tactics significantly helped, and none helped for question answering**.¹⁵ An MIT and Columbia study of 13,747 queries found **10 of 15 hand-written optimization heuristics "yield negligible or even negative changes in ranking."**¹⁶

THE PAPER EVERYONE CITES DOES NOT SAY WHAT THEY CLAIM

The foundational "GEO" paper (Princeton and collaborators, KDD 2024) is real and honest, but it was a *simulated* pipeline scored by an older model, and its famous "+40%" is a source's share of a five-source answer it was *already in*, not a 40% increase in getting found.¹⁷ It tested nine content tactics; **none was FAQ or Q&A formatting**. The widely repeated "structured FAQ content increases AI citation 37 to 40%" is fabricated. Those numbers belong to a different tactic in the paper, reattached to something it never tested. It circulates heavily in dental marketing content, so you have probably been quoted it. For a regulated, health-adjacent business, that is worth knowing before you build a budget on it.

None of that means content does not work, and our own results show why. What works is *being the page that answers a specific question*, which is why implant, full-arch, and emergency pages get cited at 90 to 100%. What does not work is rewriting an existing page with optimization tricks to chase a citation. Write the page a patient needs, then leave it alone.

09 PUBLISH YOUR PRICES, OR SOMEONE ELSE DEFINES THEM

On the tiers where your site does get cited, a second factor decides whether the AI states your facts or a stranger's. If a fact is not published in plain, crawlable text on your site, the model fills the gap from wherever it can.

We watched this happen live. In several implant-tier runs, ChatGPT built its cost answer directly out of the handful of practices in that city that publish price ranges on their own pages, quoting them by name for single-implant and full-arch figures. Every other practice in that market was named in the answer but had no part in the number the patient was shown. **Whoever publishes the number becomes the source of the number.**

Independent price figures for implants and full-arch cases are unreliable. We checked: the "national averages" circulating online come from lead-generation and cost-guide pages, and they disagree with each other by a factor of two or more. That is the material the AI builds a patient's answer from when you stay silent. Publishing your real ranges, along with financing options, runs against the instinct to hold price for the phone call, but in the AI layer, withholding it hands the number to whoever guesses. The same logic applies to accepted insurance plans, emergency hours, and which procedures you perform in-house.

10 MULTIPLE LOCATIONS, AND THE CRAWLER NO ONE CHECKS

Multiple locations are an advantage, when each is its own entity. The single most-cited practice domain in our entire study was one national group practice, cited in 25 of the 63 implant runs and 27 of the 63 emergency runs, more than any independent practice anywhere. It wins for an unglamorous reason. It publishes a real, distinct page for each city and each service, so it is eligible to be the cited answer in every town it serves. With **13% of US dentists now affiliated with a dental support organization, and 27% of those who graduated within the last five years,**¹¹ that is the competitor an independent practice is up against in this layer.

Size alone does not do it, which is the encouraging half. A second national chain, one of the best-known full-arch implant brands in the country, was *named* in 174 of our 546 answers and had its own site cited as the source in 11. Patients heard the brand and read somebody else's page about it. Brand recognition buys you the mention; specific published pages buy you the citation.

The counter is available to any practice with more than one location, and to any single location with more than one service. Give each its own page with its own address, its own service description, and its own details, rather than one generic page that differs only by phone number. Locations that look identical to an AI get collapsed into a brand. Locations published as distinct places get recommended by name.

One market in our study showed the other side of this. In Portland, a university health system was cited in all seven implant-tier runs, taking share no independent practice recovered. In academically dominant markets, expect to lose some of the category tier to the university system, and to win it back on the specialty and emergency tiers, where we measured 100%.

Can the AI even reach your site? An AI search assistant cannot cite a page its crawler is blocked from fetching, and this fails silently: it appears nowhere in your analytics. Modern security edges and some managed hosts block the AI *search* crawlers (OAI-SearchBot, Claude-SearchBot, PerplexityBot) by default or by misconfiguration, while letting browsers and Google through.¹⁸ Blocking a *training* crawler (GPTBot, Google-Extended) costs little in AI-search visibility; blocking a *search* crawler removes your current pages from cited answers. Have this checked once, and again after any host or security change. It is the cheapest high-stakes item on the list.

11 WHAT A DENTAL PRACTICE SHOULD DO

Stripped to what survives the evidence, in rough order of return:

1. **Build a real emergency page.** Hours, walk-in and same-day policy, what to do for a cracked or knocked-out tooth, and how to reach you after hours. This is a 100%-first-party query type with no

directory competition and the highest-intent patient in the funnel.

2. **Give every significant procedure its own specific page**, written descriptively rather than as "best." Implants, full-arch, sedation, root canal, and the rest. These pages are cited in 90 to 100% of runs, which is not true in every industry, and it is true in yours.
3. **Publish your prices, or at least real ranges, and your financing**. Whoever publishes the number becomes the source of the number.
4. **Fix your insurance operations, not your insurance copy**. Be in the networks you claim, and make sure each carrier's provider record for you is current. Then, if you are outside a large metro, add a plain-text accepted-plans page, which our data says wins roughly half the runs in smaller markets and none in big ones.
5. **Get and keep reviews**. Rating clears the bar (recommended businesses average 4.3, 4.1, and 3.9 across the three assistants), and review volume is the strongest visible correlate of inclusion. This is what wins the generic tier, and it is off your website.
6. **Complete and maintain your Google Business Profile**, and keep your Healthgrades and WebMD Care records accurate. This is the structured record the AI reads first, and those directories hold the generic tier.
7. **Give every location a distinct identity** in your content and structured data. This is what lets one office get cited for its own town.
8. **Verify the AI search crawlers can reach your site**. Once at setup, and after any host or security change.
9. **Measure honestly**. A single AI query is noise. Track appearance rates across many runs over weeks.

Everything on that list is measurable, and none of it requires believing a fabricated statistic.

12 A NOTE ON MEASUREMENT, AND WHY SINGLE RUNS MISLEAD

AI answers are unstable in a way that breaks most "AI visibility" reporting. A 45-day study across four engines found that two runs of the *same* prompt minutes apart shared only about a quarter to a third of their cited sources, and that **a single run is "essentially uninformative."** Reliable monitoring needs at least 7 runs per prompt.¹⁹ Our own study meets that seven-run bar (546 runs, zero failures), and our generic tier illustrates exactly why it is needed: 28 different practices took the top-named position across 63 runs of the same question. It is not a multi-day stability study, and it varies reruns rather than phrasings, so it does not capture how much a reworded prompt would move a given result. Any vendor selling a single-number "AI visibility score" from one weekly query is selling you noise.

REFERENCES

Each entry notes what the source measured, as distinct from what it is often quoted as saying. Every external source here was fetched and confirmed, and where a figure is reported rather than published, it is labelled as such.

1. rater8, *2026 Patient Choice Report*. SurveyMonkey, Apr 2026, ~1,000 US patients. rater8.com 47% used AI to research a provider, up from 31%.
2. BrightLocal, *Local Consumer Review Survey: AI Trust*, 2026, n=1,002. brightlocal.com 45% used AI for local recommendations; 42% always verify on review platforms.
3. National Association of Dental Plans, *2025 Dental Benefits Report: Enrollment*, published Dec 2025 (2024 data). nadp.org Nearly 284 million Americans, 83% of the population, have some form of dental benefit; enrollment down 2.3% year over year.
4. American Dental Association, *Emergency Department Referrals* (Action for Dental Health). ada.org An estimated 2 million US emergency-department visits a year for dental pain; redirecting them to dental settings is estimated to save \$1.7 billion.
5. Whitespark, *Prevalence of AI Overviews in local search*, 540 queries, 6 industries. whitespark.ca Local-intent AI Overview 15%, local pack 93%.
6. SE Ranking, *Google links in AI Mode*, Feb 2026, 68,313 keywords, 1.32M citations. seranking.com google.com = 17.4% of all AI Mode citations.
7. Profound, *google.com is now AI Mode's #2 cited domain*, 32M+ observations, 2026. tryprofound.com 8.4x surge in Google self-citation via Business Profile cards; healthcare among top-affected verticals.
8. OpenAI and Yelp licensing deal, announced Jul 2026. via Search Engine Land ChatGPT reads a licensed Yelp feed (~330M reviews / 8M listings).
9. SOCi, *2026 Local Visibility Index*, 350,000+ locations / 3.2M AI queries. soci.ai · per-platform detail via Search Engine Land Recommendation rates 35.9 / 11 / 7.4 / 1.2%; average recommended ratings 4.3 / 4.1 / 3.9; reviews as a filter not a ranking signal; 3 to 30x selectivity vs local search.
10. Local SEO Data, 1,120 query executions, May 2026, bootstrap CIs and corrections. prweb.com 49% median review-count gap between AI-Mode and local-pack-only businesses.
11. ADA Health Policy Institute data, reported in ADA News, *Practice modality by the numbers*, 2024 (analysis by Marko Vujicic; n=191,669 practicing US dentists). adanews.ada.org 13% of US dentists DSO-affiliated in 2022, up from 8.8% in 2017; 27% of dentists who graduated within five years; 24% of dentists 10 years or less out are in practices with 10 or more locations.
12. Zocdoc, *Pay-Per-Booking for Providers: Fees and Cancellations*. zocdoc.com The marketplace's own description of the model: no subscription, a one-time fee charged when a new patient books, applied at the moment of booking regardless of attendance, because it is "a marketing fee, not based on the visit outcome"; amount varies by specialty and location.
13. Third-party coverage of booking-marketplace per-patient fees (Emitrr 2026 pricing review; DentalVitals 2026 alternatives guide). Per-new-patient booking fees reported at roughly \$35 to \$110 depending on specialty and region, with general dentistry in large markets reported at \$50 to \$110. Reported, not published: no public rate card exists, so this is cited as consistently reported rather than as an official rate.
14. Kim et al., *SAGEO Arena*, 2026, arXiv 2602.12187, 171,003 docs, retrieval reinstated. arxiv.org Body optimization reduced presence ~9%, citation ~6%.

15. Puerto et al., *C-SEO Bench*, NeurIPS 2025, arXiv 2506.11097. [arxiv.org](#) 3 of 54 tactics helped; none in question answering.
 16. Bagga et al. (MIT and Columbia), *E-GEO*, arXiv 2511.20867, 13,747 queries. [arxiv.org](#) 10 of 15 heuristics yield negligible or negative ranking changes.
 17. Aggarwal et al., *GEO: Generative Engine Optimization*, KDD 2024, arXiv 2311.09735. [arxiv.org](#) Simulated pipeline; "+40%" is share-of-answer; no FAQ condition tested.
 18. OpenAI, Anthropic, and Perplexity crawler documentation. OpenAI bots Search crawlers (OAI-SearchBot, Claude-SearchBot, PerplexityBot) are distinct from training crawlers; blocking the search crawlers removes current pages from cited AI answers.
 19. Schulte et al. (Univ. St. Gallen), *Don't Measure Once*, Apr 2026, arXiv 2604.07585. [arxiv.org](#) Single run "essentially uninformative"; need 7 or more runs.
 20. Streben Marketing, *Dental AEO study*, Sept 2026. DataForSEO ChatGPT scraper, 12 markets × 5 query tiers × 7 runs = 420 responses, plus a 126-run controlled "best" isolation. 546 responses, 0 failures. Primary source for the tier gradient (Table 1), the cross-market comparison (Table 2), the "best" isolation (Table 3), the emergency and insurance tier findings, the zero-ad result, and the naming instability figures. Method, coding rubric, source-reclassification log, and per-run data are in the study files.
-

Method and sources. The original study (Sections 3 to 5) is complete: 546 runs across 12 US markets, zero failures, figures in Tables 1 to 3. Every source domain was re-classified against a reviewed override map before the headline figures were computed; each override was checked against the actual cited URL or by fetching the site, and the full list with reasons is in the study files. Every external figure was fetched and confirmed at its source: the dental benefits figure from the National Association of Dental Plans, the emergency-department figure from the American Dental Association, the DSO affiliation figures from ADA Health Policy Institute data as reported in ADA News, and the booking-marketplace fee model from the marketplace's own documentation, with the dollar range labelled as reported rather than published. Practice names in Sections 9 and 10 are described but not identified. Prepared by Streben Marketing, September 2026.