

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

A review of how ChatGPT and Google's AI pick an assisted living or memory care community, an original study of what they cite, and an honest list of what actually moves the answer.

Streben Marketing · July 2026 · For assisted living and memory care operators

00 THE SHORT VERSION

The adult daughter researching care for a parent increasingly starts by asking ChatGPT or Google's AI which community to consider. So the question for an operator is direct: how do you become the one it recommends? The common answer, optimize your website, is even more wrong here than in most industries.

We ran the measurement. For the query a family actually uses, **"best assisted living in [city]," a community's own website was cited as a source in 0 of 63 runs.** On "memory care in [city]," the highest-demand query in the category, the community website was cited in only **24% of runs**, while third-party directories, A Place for Mom, Caring.com, Seniorly, held **73%** of the grounding. A community's own pages only win the answer on narrow, hyper-specific queries ("small residential memory care home," "secured care for late-stage dementia"), where they were cited **91 to 92%** of the time.

This is a sharper version of a pattern that holds across local industries, and it is *more* extreme in senior living than in another healthcare market we measured the same way. The middle of the funnel, the category query where most families land, is owned by the referral directories, and those directories are paid a fee that can approach a full month's rent every time they place a resident. They also buy ads inside the AI answer: A Place for Mom ran a paid

placement inside **81% of the generic answers** we sampled. The levers that move an AI recommendation for a community sit almost entirely off its website: its Google Business Profile, its review rating and volume, its listing and reputation on the directories the AI trusts, and whether the AI can even reach its site. Its website earns citations on the specific service questions, and it has to publish the facts (pricing, care level, each location) plainly, or a stale directory answer will define them for you.

One caveat before the evidence, so none of this reads as "SEO is dead." This paper is about the AI *recommendation* layer. Classic search, ordinary Google and Bing results, still sends a large share of families to a community, and it rewards exactly the on-page work the AI layer shrugs off. The posture is additive: keep standard SEO for organic search, and add the off-site work described here for the AI layer. Two systems choose differently, and you want to show up in both.

01 WHY THIS IS WORTH AN OPERATOR'S ATTENTION

You do not need a statistic to tell you what a move-in is worth, so we will make a different point: the AI layer now sits upstream of the tour, the directory, and often the first phone call, and it is the part of the funnel most operators 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% at the end of 2025.¹ A separate representative survey found **45% of consumers had used AI for local business recommendations, up from 6% a year earlier.**² Your buyer is overwhelmingly the adult child doing the research, most often a daughter, and she behaves like every other AI user: she asks the assistant, then verifies on review sites before she calls.² The referral directories know precisely what that placement is worth to them, which is why they have moved so aggressively to own this layer, in the organic answer and, as our study shows, in paid placements inside it.

02 HOW AN AI ASSISTANT ACTUALLY PICKS A COMMUNITY

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 "memory care near me" typed into Google, the surface you compete on is still Google Maps and the Business Profile, not a generative answer. 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 directories, not your marketing pages. 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. In senior living there is an additional, category-specific layer on top of the place data: the referral directories, which is what our own study measures next.

03 OUR STUDY: WHAT CHATGPT ACTUALLY CITES, BY HOW SPECIFIC THE QUERY IS

The measurement that settles the "does my community's website get cited" question does not exist publicly for senior living. So we ran it, using the same method we used in a healthcare vertical we studied earlier, which lets us compare the two directly.

METHOD

We queried ChatGPT through the DataForSEO scraper (natural web-search behavior, US) across **12 markets**, nine national metros plus the three markets where our senior-living clients operate, and **four query tiers**, running each query **7 times** to measure run-to-run variance: **462 responses, zero failures**. For each answer we separated the *cited web sources* (the pages ChatGPT grounded on) from the *named businesses* (communities it listed from place data), and classified every source as a community's own site, a referral directory (sub-coded by name), a forum, editorial content, a medical authority, or news. Source domains were re-classified against a reviewed list, so aggregator sites that superficially look like a community's own page were coded as directories. The four tiers sample a broad-to-specific spectrum: generic ("best assisted living in [city]"), category ("memory care in [city]"), residential specialty ("small residential memory care home in [city]"), and acuity specialty ("secured memory care for late-stage dementia in [city]"). Full market list, coding rubric, and per-run data are in the study files.

RESULT: THE SPECIFICITY GRADIENT, WITH A HIGH DIRECTORY FLOOR

Citation of a community's *own website* rises with the specificity of the query, but the floor is far higher than in the healthcare vertical we studied earlier. On the generic query, **0 of 63 national runs cited any community website**; the answer grounded on directories (70%) and news (23%), even though ChatGPT still named five to nine local communities per answer from place data. That news share is not current-events reporting: nearly nine in ten of it came from just two publishers. The larger, Forbes Health, contributed its city-by-city "Top Assisted Living Facilities in [city]" pages, a ranked roundup that behaves like a directory; the rest was largely an evergreen Associated Press explainer on choosing a facility. In effect the generic answer is built almost entirely from ranked lists and place data, with no community's own website anywhere in it. At the category tier, "memory care in [city]," community websites were cited in only **24% of runs**; directories held 73% of the sources, with Caring.com alone appearing in 87% of those runs. Only on the two specialty tiers did a community's own pages win the answer: **91% and 92% of runs**.

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

QUERY TIER	COMMUNITY SITE	DIRECTORY	FORUM	OTHER	RUNS CITING ANY COMMUNITY SITE
"best assisted living in [city]"	0%	70%	6%	23% news	0% (0/63)
"memory care in [city]"	19%	73%	4%	3% content	24% (15/63)
"small residential memory care home in [city]"	95%	3%	1%	2% authority	91% (57/63)
"secured memory care for late-stage dementia in [city]"	73%	6%	10%	11% mixed	92% (58/63)

The senior-living result is the same shape as our earlier healthcare study, but the winnable band is much narrower. The comparison is the most useful thing in this paper:

Table 2. Community/practice website cited (share of runs), senior living vs the same measurement in another healthcare vertical.

QUERY TIER	HEALTHCARE VERTICAL	SENIOR LIVING
Generic ("best [category] in [city]")	0%	0%
Category / condition	89%	24%
Specialty	100%	91 to 92%

Read the category row. "Memory care in [city]" is the query with the most demand behind it, and in senior living it is **directory-owned** (24% community-site) where the equivalent query in the other healthcare vertical was **website-owned** (89%). The difference is the referral directories: A Place for Mom, Caring.com, Seniorly, and SeniorLiving.org run city-specific and care-level-specific landing pages that a single community's page cannot out-rank at the category level. The practical instruction is the same as in that vertical but weighted differently: **compete for your reviews and directory presence on the generic and category terms; compete with your website only on the specific service pages.** Editing service pages moves AI citation less here than in most industries, and only on the narrow residential and acuity band.

THE DIRECTORIES ADVERTISE INSIDE THE ANSWER

Senior living has a paid layer sitting on top of the organic one. In our sample, a paid placement appeared inside **81% of the generic answers and 32% of the category answers**, and A Place for Mom was the advertiser on the large majority of them (it was also cited organically in 40% of generic runs, so it appears twice). Tellingly, that ad presence **fell to 0% at the acuity tier**, "secured care for late-stage dementia," where the query is too specific for a directory to monetize. The paid layer mirrors the organic one exactly: the directories buy the top of the funnel where the volume and the referral fees are, and cede the clinically specific tail to the communities that actually provide that care.

THE "BEST" EFFECT (CONTROLLED TEST)

A superlative is a request for a *ranking*, and the model answers a ranking request from ranking and directory sources when they exist. Toggling only the word "best" against a fixed base (three metros, 21 runs per cell) shows it costs roughly one tier of citation, and at the specialty tier the drop is severe:

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

QUERY LEVEL	DESCRIPTIVE BASE	+ "BEST"	CHANGE
"assisted living in [city]"	10%	0%	-10
"memory care in [city]"	43%	0%	-43
"small residential memory care home in [city]"	95%	24%	-71

The lesson for on-page work: write your headings descriptively ("who provides memory care for late-stage dementia in [city]"), never as a superlative. That framing hands the answer straight to the directories.

One piece of good news for a specific community. In the three tighter client markets we also tested (versus the large national metros), the category tier was materially more winnable: a community's own site was cited in **62% of "memory care in [city]" runs**, against 24% nationally. Within a defined local market, a real, well-optimized community can be the site ChatGPT cites for the category query, where a large metro answer defers to national directories. The prize at the category level is reachable locally; it is just not reachable by editing pages alone.

Three honest limits. This is one model (ChatGPT), one instrument, US-only, one week; the scraper targets the search-enabled surface, so it under-counts answers given with no web search at all. Run-to-run variance is real: the top-named community averaged 21 different winners across 63 generic runs. And the client-market cells are small (21 runs), so treat the 62% as directional, strong enough to say the category tier is winnable locally, not as a precise rate. The national cells (63 runs each) are the solid gradient estimate.

04 REVIEWS AND DIRECTORY PRESENCE ARE THE GATE

Two things stand between a community and an AI recommendation on the high-volume queries, and neither lives on its website: its review rating, and its presence on the directories the AI cites.

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 And it is *volume*, not only rating, that separates winners: in a corrected study of 1,120 query executions, businesses that appeared in Google's AI Mode had **49% more reviews at the median** than those that appeared only in the traditional local pack.¹⁰ Review generation is therefore not a "nice to have" for AI visibility; on the current evidence it is one of the strongest observable correlates of inclusion, and it is off the website. (Correlate, not proven cause: review count also tracks a community's age and prominence.)

THE ECONOMICS BEHIND THE DIRECTORY WALL

The reason A Place for Mom and Caring.com hold the category tier is that being placed there is worth a great deal to them. These services are free to families and charge the community a placement fee when a resident moves in, reported across trade coverage at roughly half to a full month's rent, against a national median assisted-living rate of \$6,200 a month, with memory-care placements higher.³¹¹ The directory collects that fee once, at move-in, whatever the resident's length of stay (a median of 22 to 28 months),⁴ so a single placement is highly profitable to it. That is the incentive behind the wall: the directories that dominate the AI answer have a direct financial reason to keep ranking themselves above the communities they list. The counter-move is not to out-rank them on the category term, which our data says a single community usually cannot do; it is to (a) be listed and well-reviewed on the directories the AI actually cites, Caring.com first, and (b) own the specific and local queries where your own site does win, so a share of families reach you directly.

05 **HOW NARROW THE AI FUNNEL REALLY IS**

Before an operator spends 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.) The honest promise to a community is not "we will make you the AI answer." It is "we will make you the community the AI keeps finding, get you through the review gate, and make sure your own facts, not a directory's stale ones, define you when your site does get cited." That is measurable and reasonable; a guaranteed AI answer is not.

06 **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. The most important structural fact: **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%.¹² A benchmark of 1.9k queries with proper significance testing found only **3 of 54 tactics significantly helped, and none helped for question-answering**.¹³ 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."**¹⁴

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.¹⁵ 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. For a regulated, health-adjacent business like senior living, building your AI strategy on a fabricated statistic is a real risk, not a rhetorical one.

07 PUBLISH YOUR FACTS, OR A DIRECTORY WILL FILL THEM IN

On the specific and local queries where a community's own site *does* get cited, a second factor decides whether the AI states your facts or someone else's: if a key fact is not published in plain, crawlable text on your site, the AI fills the gap from a directory, and directory data is routinely stale or wrong. Price is where this matters most, because many communities keep it behind a "contact us" form.

THE QUICK WIN

When a community keeps its pricing behind a contact form, an AI asked about cost does not stay silent. It fills the gap from a third-party directory, and directory prices are routinely stale or wrong, sometimes well below a community's real rate. So a family can be shown a number the community never set and cannot see. Publishing your real price ranges as plain on-page text, and in structured data, is a fast way to take that number back. It runs against the old instinct to hide price and force a phone call, but in the AI layer, hiding it simply hands the number to whoever guesses.

The same logic applies to care levels, licensing, what is included in the base rate, and payment accepted (private pay, VA, long-term-care insurance, Medicaid waiver). Every fact a family asks about that is missing from your pages is a fact the AI will source from a directory. Publish them plainly.

08 THE MULTI-LOCATION ADVANTAGE, AND THE CRAWLER NO ONE CHECKS

Two more points, one an opportunity and one a quiet risk, are worth calling out for senior living.

Multiple locations are an advantage, when each is its own entity. An operator with several homes has something a single-location community does not: the chance to be the recommended answer in every town and county it serves. That works when each home is published as a distinct place, its own name, full address, service area, and details. Set up instead as one generic group, where every location shares the same name and differs only by phone number, the homes look identical to an AI, so it reaches for the brand rather than the specific home a county search wants.

THE OPPORTUNITY

For a group running homes across a region, this is one of the largest winnable footprints in AI search. Each home, set up as its own distinct entity, unique name, full address and geo, the towns it serves, its own price range and a one-line description, becomes eligible to be the recommended answer for its own town and county, not just the brand overall. Most operators have not set their locations up this way yet, which makes it both a clear advantage and an early-mover one for the groups that do.

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 shows up nowhere in your analytics or your robots.txt file. 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.¹⁶ Blocking a *training* crawler (GPTBot, Google-Extended) costs little in AI-search visibility; blocking a *search* crawler removes your current pages from cited answers. Every community should have this checked once at onboarding and after any host change. It is the cheapest high-stakes item on the list.

09 WHAT A COMMUNITY SHOULD ACTUALLY DO

Stripped to what survives the evidence, the list is short and unglamorous, and mostly not about your website:

1. **Get and keep reviews.** Rating clears the bar (recommended businesses average 4.3 / 4.1 / 3.9), and review volume is the strongest visible correlate of inclusion. This is the highest-impact action, and it is off your website.
2. **Be listed and well-reviewed on the directories the AI cites, Caring.com first,** then A Place for Mom, plus Seniorly and SeniorLiving.org. You will not out-rank them on the category term; you can make sure the version of you they carry is complete, accurate, and highly rated.
3. **Complete and maintain your Google Business Profile** for every location. This is the structured record the AI reads first.

4. **Publish your facts in plain text: price ranges, care levels, what is included, payment accepted, licensing.** Anything you hide, the AI sources from a directory, often wrong.
5. **Give every location a distinct identity** in your content and structured data (unique name, address, geo, service area, price, description). This is what lets a specific home get cited for a specific town or county.
6. **Write genuine, plain-language answers to the specific care questions you serve** (memory care, late-stage dementia, respite), phrased descriptively, not as "best." These are the queries where your own site earns citations. Skip schema added solely to chase AI citations; the visible answer does the work.
7. **Verify the AI search crawlers can reach your site.** Once at onboarding, and after any host or CDN change.
8. **Measure honestly.** A single AI query is noise; the top result changes run to run. Track appearance rates across many runs over weeks.

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

10 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.¹⁷ Our own study meets that seven-run bar for a point-in-time snapshot (462 runs, zero failures); it is not a multi-day stability study, and it varies reruns but not 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 actually measured. External entries carried over from Streben's earlier healthcare AEO evidence review were fetched and confirmed there; the senior-living-specific entries were confirmed against the sources noted (the referral-fee figure is cited as consistently reported across trade press, since A Place for Mom does not publish it).

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 recs; 42% always verify on review platforms.

3. CareScout (Genworth), *2025 Cost of Care Survey*, fielded July to November 2025. carescout.com/cost-of-care National median assisted living \$6,200/mo (\$74,400/yr), up 5% YoY; private nursing-home room \$129,575/yr. Memory care typically adds a premium above the assisted-living base.
4. AHCA/NCAL (National Center for Assisted Living), *Assisted Living Facts & Figures*. ahcancal.org Average assisted-living length of stay reported at 22 to 28 months (median ~22 months).
5. Whitespark, *Prevalence of AI Overviews in local search*, 540 queries, 6 industries. whitespark.ca Local-intent AIO 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 + 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%; avg 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 + corrections. prweb.com 49% median review-count gap between AI-Mode and local-pack-only businesses.
11. Referral-fee model of A Place for Mom / Caring.com. Seniors Housing Business and industry trade press. Free to families; the community pays a placement fee reported across trade coverage at roughly 50 to 100% of the first month's rent (memory-care placements higher). A Place for Mom does not publish the figure, so it is cited as consistently reported, not as an official rate.
12. Kim et al., *SAGEO Arena*, 2026, arXiv 2602.12187, 171,003 docs, retrieval reinstated. arxiv.org Body optimization reduced presence ~9%, citation ~6%.
13. Puerto et al., *C-SEO Bench*, NeurIPS 2025, arXiv 2506.11097. arxiv.org 3 of 54 tactics helped; none in QA.
14. Bagga et al. (MIT/Columbia), *E-GEO*, arXiv 2511.20867, 13,747 queries. arxiv.org 10 of 15 heuristics yield negligible or negative ranking changes.
15. 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.
16. 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.
17. Schulte et al. (Univ. St. Gallen), *Don't Measure Once*, Apr 2026, arXiv 2604.07585. arxiv.org Single run "essentially uninformative"; need 7+ runs; ChatGPT skipped web search on 57.8% of runs.
18. Streben Marketing, *Senior-Living AEO study*, Jul 2026. DataForSEO ChatGPT scraper, 12 markets × 4 query tiers × 7 runs = 462 responses, 0 failures. Primary source for the gradient (Table 1), the cross-vertical comparison (Table 2), the "best" isolation (Table 3), the ads-inside-the-answer finding, and the client-market result. Method and per-run data in the study files.

Method and sources. The original study (Section 3) is complete: 462 runs across 12 US markets, zero failures, figures in Tables 1 to 3. External references carried over from Streben's earlier healthcare AEO evidence review were fetched and confirmed there. The senior-living-specific figures were confirmed against their sources: the assisted-living cost figure is from CareScout (Genworth), the length-of-stay figure from AHCA/NCAL, and the directory referral-fee figure is cited as consistently reported across trade press (A Place for Mom does not publish it). The Section 7 and 8 examples are anonymized. Prepared by Streben Marketing, July 2026.