

THE LUCIE AI VISIBILITY INDEX

GEORGIA SENIOR LIVING STUDY

2026 EDITION | A COMPANION STUDY TO THE 2026 INDEX

When Families Ask AI

How AI Systems Find, Describe, and Recommend
Georgia's Senior Living Communities

By Craig Lucie, Akshita Gorti, and Prajwal Mohan Kumar

Edited by: Aaron Diamant

Lucie Research | Atlanta, Georgia | September 2026

Study data collected through August 31, 2026. External statistics and sources reviewed through September 21, 2026.

A Note on Naming

This study reports results for individual communities and operator-level brands in aggregate.

Communities that showed limited or inconsistent visibility signals are not identified, either by name or by description specific enough to identify them. A small number of consistently strong performers are named as verified, illustrative examples.

That choice is deliberate.

The purpose of this research is to describe an industry-wide pattern in how AI systems represent senior living, not to rank providers or single out any operator for criticism.

Low visibility in an AI answer is not evidence of poor care, poor management, or poor stewardship. Several of the communities with the weakest signals in this sample are long-established, well-regarded organizations. What this study measures is how an automated system represents them, which is a different question from how well they serve residents, and it is a condition that is common, recent, and correctable.

All 57 evaluated entities were assessed using the same measurement framework and scenario design. Individual communities and operator-level brands were evaluated using their respective web presences, as described in the Methodology section.

Any community included in this study may request its own complete results from Lucie Research at no cost and with no obligation. Details are in the Methodology section.

On the research. Lucie Research is the research division of Lucie Content. Lucie Content sells communications and AI visibility services, including to organizations in this sector. The measurement framework, scenario design, and scoring were applied uniformly across the sample and were not adjusted for any community, client or otherwise. No community in this sample paid to be included, excluded, or ranked. The methodology summary in this report is published so that the findings can be examined on their merits.

Contents

Executive Summary	1
1 Why Senior Living Is Different	2
2 What This Study Measured	4
3 Findings	6
4 Consistently Strong Performers	13
5 What This Means for Operators	13
6 How This Study Relates to the 2026 Index	14
7 Methodology	15
8 About the Authors	18
9 About Lucie Research	18
10 References	19

Executive Summary

Fifty-five of the 57 Georgia senior living entities in this study fell below the Strong readiness tier, and 14 did not appear at all when AI was asked the kinds of questions families ask.

In August 2026, Lucie Research measured how AI systems find, describe, classify, and recommend senior living communities operating in Georgia.

The 2026 Lucie AI Visibility Index examined this pattern across 94 companies in mixed sectors and found that AI-assisted discovery can shape a customer's shortlist before the first click [10]. This companion study applies the same measurement framework to a single sector in a single state, where the consequences of an incomplete or inaccurate answer are more immediate.

Each evaluated entity was tested through three brand-blind family-decision scenarios, each repeated three times, and evaluated against 31 signals grouped into eight measured pillars.

The sector matters. Adoption of AI for local recommendations has moved quickly: BrightLocal's 2026 survey found 45 percent of consumers using generative AI tools for local business recommendations, up from 6 percent a year earlier, making AI the third most-used source behind Google and Facebook [1].

Key findings

- **One in four evaluated entities did not appear at all.** Fourteen of the 57 evaluated entities, or 24.6 percent, were not retrieved in any tested scenario.
- **Being present was not the same as being recommended.** Eight evaluated entities were retrieved at least once and still carried recommendation scores below 50. One appeared in every tested scenario and still scored 28.2 on recommendation strength.
- **Accuracy lagged behind presence.** Twenty-two evaluated entities, or 38.6 percent, scored below 50 on knowledge credibility, the measure of whether the information AI holds about a community is accurate and well supported. Five scored zero.
- **Care type was frequently unclear to AI.** Twenty evaluated entities, or 35.1 percent, scored below 70 on category association, the measure of whether AI associates a community with the kind of care it actually provides. In senior living, that is a clinical distinction, not a marketing one.
- **Technical readiness did not carry through to performance.** Foundation signals averaged 75.9 against performance signals of 67.7. Seventeen evaluated entities, or 29.8 percent, showed a gap larger than 15 points between the two.
- **Complexity tracked with invisibility.** Independent living communities averaged 85.6 with

a 100 percent retrieval rate. Full-continuum communities, offering independent living, assisted living, and memory care together, averaged 67.5 with a 60.4 percent retrieval rate.

- **Location mattered more than brand.** Within one operator's seven Georgia locations, technical readiness was identical and visibility scores ranged from 46.0 to 85.4. Four locations were retrieved in every scenario; three were retrieved in none.

Taken together, the findings describe a sector that is largely technically competent and largely absent from the answer layer. The gap is not concentrated among small or under-resourced providers. It runs through nonprofits, independents, and multi-state operators alike.

This study did not measure inquiries, tours, move-ins, occupancy, or revenue. It measures representation, not outcomes.

Research at a Glance

- 57 Georgia senior living entities (51 individual communities and six operator-level brands)
- Three brand-blind family-decision scenarios, each repeated three times
- 31 signals across eight measured pillars
- Data collected through August 31, 2026
- Single-state, single-sector sample

Research boundary. Selected, nonrandom sample. Findings are directional and describe the entities evaluated during the collection window. They are not census-level estimates for the Georgia senior living industry or for every AI platform.

1. Why Senior Living Is Different

The 2026 Index argued that AI-assisted discovery changes business risk because absence is no longer the same as ranking lower [10]. Senior living concentrates that risk, for three reasons that are specific to the sector.

1.1 The search is short and it happens under pressure

Senior living decisions are rarely leisurely. In a survey of 1,104 family caregivers published by A Place for Mom, two-thirds of those who found care secured it in 60 days or less, and only 10 percent reported a search lasting longer than a year [2].

The same survey found that 54 percent of caregivers wished they had started planning sooner, a figure that rose to 77 percent among those still searching [2].

A shortlist assembled in a few weeks, under emotional strain, leaves little room for a family to discover an option that the first several answers omitted.

1.2 The searcher often cannot name what they are looking for

The same research found that 41 percent of caregivers were unsure what they were looking for in a care experience, and 34 percent had difficulty identifying the appropriate level of care [2].

This is the condition under which an answer system matters most. A searcher who can already name the category can correct a bad answer. A searcher who cannot is more dependent on whatever the system supplies, including its categorization.

1.3 Care type is a clinical distinction, not a label

Independent living, assisted living, memory care, personal care, and skilled nursing serve materially different populations with different needs and different regulatory frameworks. Because the distinctions carry through the rest of this report, they are set out here in the plainest terms.

The care types referred to in this report.

Independent living. Private residences for older adults who do not need help with daily personal care. Services are residential and social: dining, housekeeping, transportation, activities.

Assisted living. Housing with scheduled personal care and supervision, including help with bathing, dressing, and medication management, for residents who do not require continuous skilled nursing.

Memory care. A secured setting with staff trained in dementia care, designed for residents living with Alzheimer's disease or another dementia. It may be a dedicated community or a distinct wing within one.

Personal care home. Georgia's licensed category for a residence providing housing, meals, and personal services to adults who need help with daily living. The state also licenses assisted living communities separately, under a broader scope of permitted care.

Full continuum. A single community offering more than one of the above, most often independent living, assisted living, and memory care together, so a resident can move between levels without moving out.

These are descriptions of what each care type means to a family reading an AI answer, not a summary of Georgia licensure. Licensing categories, permitted services, and staffing requirements are set by the Georgia Department of Community Health.

The scale is substantial. Industry figures put the United States at roughly 41,465 assisted living communities with close to 1.4 million licensed beds and more than one million residents, of whom roughly four in ten are living with Alzheimer's disease or another dementia [4].

In Georgia specifically, an estimated 188,300 people are living with Alzheimer's, supported by roughly 391,000 unpaid family caregivers [5]. National prevalence and caregiving figures are reported annually by the Alzheimer's Association [6].

When an AI answer associates a community with the wrong care type, the error is not cosmetic. It can send a family toward a community that cannot serve their relative's needs, or away from one that can.

1.4 Capacity is tight, so a misdirected family may not be redirected

Senior housing occupancy across the 31 NIC MAP Primary Markets reached 89.5 percent in the first quarter of 2026, with independent living surpassing 91 percent for the first time since 2016. New supply is at historic lows, with roughly 16,400 units under construction, about 2.3 percent of existing inventory [3].

In a market with limited slack, the practical cost of being omitted from the first set of answers is higher than it would be in a sector with abundant alternatives.

What this study does not claim. This research did not measure whether families used AI to choose a community, and it did not track any family's search. It measured what AI systems returned when asked the kinds of questions families ask. The consumer-behavior figures above come from independent third-party sources and are provided as context for why representation in that layer is worth measuring.

2. What This Study Measured

Lucie Research evaluated 57 Georgia senior living entities, 51 individual communities and six multi-location operator brands, using audit runs from the Lucie production database completed through August 31, 2026.

2.1 Sample construction

The 51 individual communities were included where the organization maintained a distinct, independently reachable web presence for its Georgia operations.

Six multi-location operator brands were evaluated at the brand level, using the operator's primary web presence with location set to Georgia, because their individual Georgia locations were not independently reachable as standalone pages.

Each operator-level brand was counted as one evaluated entity. Brand-level results should not be interpreted as individual visibility measurements for every community operated by that

brand.

Where a community had more than one eligible audit run, the most recent run was used. One community appearing under two records was consolidated to a single entry.

2.2 Scenario design

Each evaluated entity was tested through three brand-blind family-decision scenarios, each repeated three times.

Scenarios were written as the questions a family would plausibly ask, not as branded queries. No scenario named the community being evaluated.

Repetition matters because AI answers vary across sessions and prompts. A single answer is an observation; repeated answers describe a pattern. Recent measurement research makes the same argument for observing AI visibility across repeated runs rather than single-run snapshots [8].

2.3 Measured pillars

The framework evaluated 31 signals grouped into eight measured pillars:

- Technical Readiness
- Eligibility
- Extractability
- Authority
- Knowledge Credibility
- Retrieval Consistency
- Category Association
- Recommendation Strength

Two composites are reported alongside the pillars. **Foundation** combines eligibility and extractability, the signals that determine whether a site can be read and used at all. **Performance** combines retrieval, category association, and recommendation strength, the signals that describe what AI actually does with a community once it can read the site.

2.4 Readiness tiers

Each evaluated entity was classified as Strong, At Risk, Weak, or Critical.

Readiness tiers are weighted classifications, not a ranking of composite scores. A tier reflects consistency across repeated runs and the balance of signals a community passed, not its composite score alone. An evaluated entity can carry a higher composite score than another and still sit in a lower tier, because its results were less consistent across repetitions. Tiers should be read as descriptions of condition, not as a league table.

The complete measurement framework, including the general-market findings this study builds on, is described in the 2026 Lucie AI Visibility Index [10]. The framework itself, including signal logic, weighting, and evaluation criteria, is proprietary to Lucie Research and is not fully disclosed in either report.

3. Findings

How to read these findings

These findings describe 57 Georgia senior living entities, comprising 51 individual communities and six operator-level brands, evaluated through August 31, 2026. They are directional. They describe representation in AI answers and do not measure inquiries, move-ins, occupancy, or revenue. Where a single community is used to illustrate a pattern, it is described without identifying detail, consistent with the naming policy.

3.1 Finding 1: Fewer than one in twenty evaluated entities reached the Strong tier

Two of the 57 evaluated entities, or 3.5 percent, reached the Strong tier. Thirty-two, or 56.1 percent, were classified At Risk. Twenty, or 35.1 percent, were Weak, and three, or 5.3 percent, were Critical.

The sample averaged 70.9 on the composite visibility score, with a median of 77.3 and a spread of 57.8 points between the highest and lowest scoring communities.

Why it matters

The distribution does not describe a handful of laggards. Fifty-five of 57 evaluated entities fell below the Strong tier, which makes limited AI visibility the ordinary condition in this sample rather than the exception.




Figure 1. Readiness tiers across the sample of 57 evaluated entities. Tiers are weighted classifications reflecting consistency and signal balance, not a ranking of composite scores.

3.2 Finding 2: One in four evaluated entities did not appear at all

Fourteen evaluated entities, or 24.6 percent, were not retrieved in any tested scenario.

These were not uniformly small or obscure organizations. The group included well-established single-site communities, locations belonging to regional chains, and operator-level brands running communities across multiple Georgia markets.

Why it matters

A community that is never retrieved has no opportunity to be compared, evaluated, or recommended. Within this sample, retrieval functioned less like a ranking and more like a threshold: communities tended to appear consistently or not at all. The 2026 Index observed the same clustering across its broader sample, where 84 percent of companies fell at one of the two retrieval extremes [10].

3.3 Finding 3: Being present was not the same as being recommended

Eight evaluated entities, or 14.0 percent, were retrieved in at least one tested scenario and still carried recommendation scores below 50.

The pattern was sometimes stark. One community in the sample was retrieved in every tested scenario and still scored 28.2 on recommendation strength. Another was retrieved in every scenario and scored 44.7.

Across the sample as a whole, 22 evaluated entities, or 38.6 percent, scored below 50 on recommendation strength. Eight of those scored zero on recommendation strength, and all eight were among the 14 that were never retrieved. The eight described above are a separate

group: they were retrieved and still not recommended.

Why it matters

Recommendation strength averaged 55.2 across the sample, the weakest of the eight measured pillars. Appearing in an answer is a necessary condition for being considered, but on this evidence it is not a sufficient one. The eight evaluated entities described above were visible and still went unrecommended, which is a different problem from invisibility and calls for a different fix.

3.4 Finding 4: What AI knew was often weaker than whether it knew

Knowledge credibility, the measure of whether the information available to the model about a community is accurate, consistent, and well supported, averaged 61.4 across the sample.

Twenty-two evaluated entities, or 38.6 percent, scored below 50. Thirty-nine, or 68.4 percent, scored below 70. Five scored zero, meaning the framework found no credible supporting information about the care those communities provide.

Why it matters

Knowledge credibility and recommendation strength were the two weakest pillars in the study. That combination points to a specific problem: for a meaningful share of this sample, the issue is not only that AI does not surface the community, but that what it holds about the community is thin or unsupported.

3.5 Finding 5: Care type was frequently unclear

Category association, the measure of whether AI associates a community with the kind of care it actually provides, averaged 75.3.

Thirty-three evaluated entities, or 57.9 percent, scored below 80. Twenty, or 35.1 percent, scored below 70. Eight, or 14.0 percent, scored below 60, which the framework treats as a severe classification problem.

Among the communities with the weakest category signals, the study observed providers offering independent and assisted living described in terms that did not reliably distinguish them from providers offering dedicated memory care. The reverse also occurred.




Figure 2. Average score by measured pillar across all 57 evaluated entities. The signals describing what AI knows and whether it recommends scored well below the signals describing whether a site can be read.

Why it matters

This is the finding most specific to senior living. In most sectors, a category error is a marketing inconvenience. Here, assisted living, memory care, and personal care serve different populations under different licensure. An answer that blurs those distinctions can cost a family time they do not have, particularly given that a third of caregivers already report difficulty identifying the appropriate level of care [2].

3.6 Finding 6: Technical readiness did not carry through to performance

Foundation signals, which describe whether a site can be read and used, averaged 75.9 across the sample. Performance signals, which describe what AI does with the community once it can read the site, averaged 67.7. The industry-wide gap was 8.2 points.

Seventeen evaluated entities, or 29.8 percent, showed a gap larger than 15 points. The widest gap observed in the sample was 68.7 points, at a community whose foundation score was 87.2 and whose performance score was 18.5.

Why it matters

A readable, well-structured, technically accessible website did not reliably produce AI visibility in this sample. The 2026 Index reported the same disconnect in its broader sample, where the correlation between eligibility and retrieval was effectively zero [10].

The practical implication is that conventional technical search work, while necessary, did not by itself determine whether a community appeared in these answers.

3.7 Finding 7: Complexity tracked with invisibility

Visibility declined as the care model became more complex.

Independent living communities averaged 85.6 with a 100 percent retrieval rate and 90.5 on recommendation strength. Retirement communities averaged 77.6 and memory-care-only providers 76.8. Assisted-living-only providers averaged 72.4.

At the other end, communities offering assisted living together with memory care averaged 68.0, with a 66.7 percent retrieval rate. Full-continuum communities, offering independent living, assisted living, and memory care together, averaged 67.5 with a 60.4 percent retrieval rate. Full-continuum was the largest single group in the sample, at 16 evaluated entities.

Why it matters

A community offering one clearly defined service appears easier for an AI system to categorize than one offering a blend of care levels. The consequence is uncomfortable: within this sample, the communities serving the most complex needs, including memory care and multi-level support, were the ones least reliably represented.

Visibility Declined as Care Complexity Increased

Ordered from single-service to multi-level care. Both measures use a 0–100 scale.




Figure 3. Average visibility score and retrieval rate by care type. Group sizes vary and the smallest groups contain very few communities, so these averages are directional.

3.8 Finding 8: Location mattered more than brand

The clearest evidence that AI visibility is a location-level condition came from multi-location operators in the sample.

One operator running seven Georgia communities recorded an identical technical readiness score of 86.7 at every location, on the same website framework. Visibility scores across those seven locations ranged from 46.0 to 85.4, a spread of 39.4 points. Four locations were retrieved in every tested scenario. Three were retrieved in none.

Two other multi-location operators showed the same pattern at smaller scale, with within-brand ranges of 38.3 points across three locations and 19.7 points across two.

Why it matters

A regional or national operator cannot assume that brand-level work will cascade to every community it runs. On this evidence, each location is represented, or not represented, largely on its own signals.




Figure 4. Seven Georgia locations operated by a single company. Technical readiness was identical across all seven; visibility outcomes were not. The operator is not named, consistent with the naming policy.

3.9 Finding 9: The most common technical failure was AI-specific

Across the 15 technical signals evaluated on the full sample, the most commonly failed signal was LLM readiness, which describes whether a site provides the machine-readable guidance AI agents use to interpret its content. Thirty-five evaluated entities, or 61 percent, failed it.

The second most commonly failed signal was entity schema presence, the structured data that tells an AI system what an organization is and what services it offers, failed by 19 evaluated entities, or 33 percent.

By contrast, conventional signals passed widely. Every community in the sample passed page latency and static content visibility. Indexing permissions and crawler access passed at 95 percent or better.

Why it matters

The signals this sample failed most often are the ones added specifically for AI systems. The signals it passed most often are the ones conventional search work has addressed for years. That is consistent with Finding 6, and it points to where the remaining work sits.




Figure 5. Pass rates for the 15 technical signals evaluated across the full sample. The two most commonly failed signals are those designed for AI interpretation rather than conventional search.

3.10 Finding 10: Budget and scale did not explain the results

Grouped by ownership, nonprofit communities averaged 73.5, independent and other single-operator communities averaged 72.2, and chain-operated communities averaged 67.0.

Within the chain group, results varied widely, from an operator-level brand scoring 87.0 to one scoring 53.1.

Why it matters

The communities that performed best in this sample were not consistently the largest or the best resourced. On this evidence, content quality, community-specific detail, and local signal

consistency tracked more closely with visibility than operational scale did. This is a directional observation from grouped averages, not a controlled comparison.

4. Consistently Strong Performers

Two communities in the sample reached the Strong readiness tier:

Brandon Wilde

Azalea Estates of Fayetteville

Each was retrieved consistently across repeated runs and showed comparatively strong signals on knowledge credibility and recommendation strength, the two pillars on which the sample as a whole scored lowest.

These are illustrative examples, not a complete ranking. Results reflect the August 2026 collection window and may change as platforms, available information, and model behavior evolve.

Both communities were evaluated through the same research process as every other community in the sample. Neither was required to participate in, review, or endorse the study, and inclusion here does not imply that either did.

5. What This Means for Operators

The findings point to four situations, each with a different implication.

5.1 If you operate multiple communities

Brand-level work did not cascade to every location in this sample. Identical technical readiness produced a 39.4-point visibility spread within one operator's seven Georgia communities.

Each location needs its own page-level detail, its own local signals, and its own community-specific content. A single corporate site describing the brand does not appear to carry individual communities into these answers.

5.2 If you offer complex or multi-level care

Full-continuum and assisted-living-plus-memory-care providers scored lowest on both visibility and retrieval in this sample.

The implication is not to simplify the care model. It is that a complex care model requires

more explicit description, not less: plain-language explanation of each care level offered, who it serves, and how the levels relate, rather than an umbrella term that leaves the distinctions to be inferred.

5.3 If your website is technically sound and you are still not appearing

This was the single most common profile in the sample. Foundation exceeded performance by more than 15 points at 17 communities.

Conventional technical search work was necessary but did not determine AI visibility here. The signals that separated communities were the AI-specific ones: machine-readable guidance, entity-level structured data, accurate and well-supported public information, and clear category description.

5.4 If you are a nonprofit or a single-site operator

The data does not suggest an inherent disadvantage. Nonprofit and independent communities averaged slightly higher than chain-operated ones.

Scale did not decide the outcome in this sample, which suggests the work is available to organizations without large marketing budgets.

The practical sequence. The 2026 Lucie AI Visibility Index sets out an eight-step framework for measuring and improving AI visibility, beginning with measuring how AI currently answers the questions customers actually ask and ending with re-measuring against a documented baseline [10]. That framework applies directly to the findings in this study and is not restated here.

6. How This Study Relates to the 2026 Index

The 2026 Lucie AI Visibility Index examined 94 companies across mixed sectors and reported five findings: retrieval clustered at the extremes, AI frequently struggled to identify what a business does, a competitor was often recommended when the subject business was not, readiness for conventional search did not predict AI retrieval, and a business had to be retrieved before it could be recommended [10].

This study applies the same framework to one sector in one state. Three of those patterns reproduced directionally here: retrieval clustered at the extremes, conventional readiness did not predict retrieval, and presence preceded recommendation.

Two things are specific to this study.

The first is **care-type classification**. Category association is measured in the general framework, but its consequences are sector-specific. In senior living, the categories correspond to different regulated service types serving different populations.

The second is **the location-level finding**. The Index sample was largely single-entity. This sample contained several multi-location operators, which made it possible to observe that visibility varied widely within a single brand on a single website framework.

Neither study measured conversions or revenue, and neither established causation. Both describe representation, tested repeatedly, within a defined collection window.

On the research context. Academic work has shown that content structure and evidence can influence prominence in generated answers [7], and independent industry research has found AI recommendation visibility to be more selective than conventional local search visibility [9]. The findings here are consistent with that direction of travel without depending on it.

7. Methodology

7.1 Dataset scope

The study analyzed 57 Georgia senior living entities, comprising 51 individual communities and six multi-location operator brands, drawn from full-report audit runs in the Lucie production database and current as of August 31, 2026.

Where a community had multiple eligible runs, the most recent was used. One community stored under two records was consolidated to a single entry. One community's name was stored with a typographical error in the database and is corrected in the reporting.

7.2 Sample composition

Communities were included where the organization maintained a distinct, independently reachable web presence for its Georgia operations. This includes individual locations belonging to a larger brand where the location maintained its own reachable page.

Six operator-level brands were evaluated at the brand level, using the operator's primary web presence with location set to Georgia, because those operators run communities in multiple Georgia cities without independently reachable location pages. These six are included in all reported statistics.

7.3 Signal coverage

Forty-nine of the 57 evaluated entities were assessed against the full set of 31 signals. Eight were evaluated against a reduced set of 20 to 28 signals, because some signals could not be assessed for those sites.

All 57 evaluated entities are included in every statistic reported here, including composite scores, readiness tiers, retrieval, recommendation, and category association. The reduced signal coverage affects only the denominator used for the signals-passed count at those eight evaluated entities. It does not exclude them from any reported measure.

7.4 Scenario and query design

Each evaluated entity was tested through three brand-blind family-decision scenarios, each repeated three times.

The scenarios were designed around recommendation, comparison, and category intent, written in the form of questions a family would plausibly ask. No scenario named the community under evaluation.

Community identification for the sample was carried out separately from the visibility evaluation. The list-building step used a retrieval prompt for independently reachable Georgia senior living providers. That step was not a consumer recommendation query and its results should not be read as such.

7.5 Measurement framework

Eight pillars were scored: Technical Readiness, Eligibility, Extractability, Authority, Knowledge Credibility, Retrieval Consistency, Category Association, and Recommendation Strength.

Foundation combines eligibility and extractability. Performance combines retrieval, category association, and recommendation strength.

Proprietary methodology statement

The Lucie AI Visibility measurement framework and evaluation methodology are proprietary to Lucie Research. This report describes the research approach but does not disclose the complete signal logic, weighting methodology, evaluation criteria, or implementation processes.

7.6 Naming policy

Findings are reported primarily in aggregate.

Communities that showed limited or inconsistent visibility signals are not identified, by name or by description specific enough to identify them. Where an individual result is used to illustrate a pattern, identifying detail is withheld.

A limited set of consistently strong performers is named to provide verified, illustrative examples. Each named result was confirmed against dated collection records.

Inclusion in this study does not imply participation in it or endorsement of it.

7.7 Data availability for included communities

Any community included in this study may request its own complete results, including its signal-level detail and the scenarios used, from Lucie Research at no cost and with no obligation. Requests can be sent to info@luciecontent.com.

7.8 Limitations

AI outputs are probabilistic and change over time. These results reflect a single collection window ending August 31, 2026.

The sample was selected rather than randomly drawn, and it covers one sector in one state. Findings are directional and describe the evaluated entities. They are not census-level estimates for Georgia senior living as a whole.

Care-type groups vary substantially in size. The smallest groups contain very few communities, and their averages should be read as indicative rather than precise.

Six of the 57 evaluated entities were assessed at the operator level rather than the individual-community level. These operator-level results may understate variation across individual communities within those organizations' Georgia portfolios and should not be interpreted as measurements of each individual location.

The study did not measure inquiries, tours, move-ins, occupancy, census, revenue, or any other business outcome. AI visibility should not be treated as a financial performance metric without additional evidence.

Correlation does not establish causation. The study describes measured relationships among signals, not the mechanisms behind them.

8. About the Authors

Craig Lucie is the founder and CEO of Lucie Content, an Atlanta-based strategic communications and media production firm.

He is a former television news anchor and reporter who worked at news stations across the United States, including the nation's top-rated local television station. He received a Southeast Emmy Award for Best News Anchor and was named Best On-Air Personality by the Georgia Association of Broadcasters. Under his leadership, Lucie Content has earned two Emmy Awards and 13 Telly Awards.

Craig leads the communications and research direction behind the Lucie AI Visibility Index and is the author of the children's book *Hold You*.

Akshita Gorti is the Founding Head of AI (GEO and Strategy) at Lucie Content, where she leads development of the Lucie AI Visibility measurement framework.

She earned a B.S. from Cornell University, where she studied Electrical and Computer Engineering and Robotics and was a Rawlings Cornell Presidential Research Scholar. She is also a FIDE Master and Woman International Master in chess.

Prajwal Mohan Kumar is the Founding Head of AI (Platform and System Architecture) at Lucie Content.

He holds an M.S. in Quantitative and Computational Finance and an M.S. in Computational Science and Engineering from Georgia Tech, as well as a B.S. in Computer Science and Economics. His research background includes machine learning and quantum computing.

9. About Lucie Research

Lucie Research is the research division of Lucie Content and the publisher of the Lucie AI Visibility Index.

Lucie Research studies how organizations are discovered, understood, and recommended in AI-generated answers, through original research at the intersection of strategic communications, content, and AI visibility.

This Georgia senior living study is the first sector study published under the Index.

The full study, the 2026 Index, and a free check of how AI currently describes an organization are available at lucievisibility.com.

Learn more at luciecontent.com. Media inquiries: info@luciecontent.com.

10. References

- [1] BrightLocal. “Local Consumer Review Survey 2026.” 2026.
<https://www.brightlocal.com/research/local-consumer-review-survey/>
- [2] A Place for Mom. “Senior Care Search Trends: Navigating Options in the U.S.” Survey of 1,104 family caregivers, published November 20, 2025.
<https://www.aplaceformom.com/senior-living-data/senior-care-search-trends-united-states>
- [3] National Investment Center for Seniors Housing & Care (NIC). “Senior Housing Occupancy Climbs to 89.5% as Record Demand Meets Record-Low New Supply in First Quarter 2026.” May 19, 2026.
<https://www.nic.org/blog/senior-housing-occupancy-climbs-to-89-5-as-record-demand-meets-record-low-new-supply-in-first-quarter-2026/>
- [4] American Health Care Association / National Center for Assisted Living (AHCA/NCAL). “Assisted Living Facts and Figures.”
<https://www.ahcancal.org/Assisted-Living/Facts-and-Figures/Pages/default.aspx>
- [5] Alzheimer’s Impact Movement. “Alzheimer’s Association Facts and Figures: Georgia.”
<https://alzimpact.org/Georgia>
- [6] Alzheimer’s Association. “2026 Alzheimer’s Disease Facts and Figures.”
<https://www.alz.org/alzheimers-dementia/facts-figures>
- [7] Aggarwal, Pranjal, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, Karthik Narasimhan, and Ameet Deshpande. “GEO: Generative Engine Optimization.” *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024.
<https://dl.acm.org/doi/10.1145/3637528.3671900>
- [8] Schulte, Julius, Malte Bleeker, and Philipp Kaufmann. “Don’t Measure Once: Measuring Visibility in AI Search (GEO).” arXiv preprint, April 2026.
<https://arxiv.org/abs/2604.07585>
- [9] SOCi. “The Most Visible Local Brands of 2026” and related 2026 Local Visibility Index analysis. 2026.
<https://www.soci.ai/insights/lvi/>
- [10] Lucie Research. *The Lucie AI Visibility Index, 2026 Edition: The Invisible Business, How AI Answers Shape Who Gets Considered*. Atlanta, Georgia, August 2026.
<https://lucievisibility.com>

© 2026 Lucie Research, a division of Lucie Content, LLC. All rights reserved.

Short excerpts may be quoted with attribution to the Lucie AI Visibility Index, Georgia Senior Living Study.

No part of this publication may be reproduced, distributed, or republished without prior written permission, except as otherwise permitted by law.