THE LUCIE

AI VISIBILITY INDEX

2026 EDITION | INAUGURAL EDITION

The Invisible Business

How AI Answers Shape Who Gets Considered

By Craig Lucie, Akshita Gorti, and Prajwal Mohan Kumar

Edited by: Aaron Diamant

Lucie Research | Atlanta, Georgia | August 2026

Index data collected May–July 2026. External statistics and sources reviewed through August 5, 2026.

Contents

Executive Summary

> **One in four businesses in the 2026 Lucie AI Visibility Index was completely invisible to AI.**

That finding highlights one of the central challenges of AI-assisted discovery: a business can have an established website and a credible market presence yet still be absent when an AI system answers a customer's question.

AI-assisted discovery can shape the shortlist before a customer visits a website. A person can ask a full question and receive a synthesized answer, comparison, or short set of recommendations. Traditional search still matters, but part of the consideration now happens before the first click.

Adoption is already substantial. BrightLocal found that 45% of consumers use generative AI for local business recommendations, up from 6% a year earlier [1]. Cox Automotive found that 25% of new-vehicle buyers used AI while shopping [2].

To examine what this shift means for businesses, Lucie Research began with 133 requested business names and produced a clean sample of 94 companies after matching, deduplication, and quality exclusions.

Each company was tested through three brand-blind scenarios, repeated three times, producing 846 Gemini queries. We also examined 564 recommendation-response objects using confidence intervals, correlation testing, distribution summaries, and company-clustered analysis.

Key findings

- **AI often found businesses everywhere—or missed them entirely.** Twenty-five of 94 companies were never retrieved, while 54 appeared in every tested scenario. Overall, 84% landed at one retrieval endpoint or the other. The 95% confidence interval for complete invisibility was 17.7%–35.5%.
- **AI often struggled to understand what businesses do.** Thirteen companies (13.8%) showed limited knowledge-credibility signals; 26 (27.7%) showed limited category-association signals.
- **When the subject business was missed, a competitor often took its place.** Competitors were recommended in 229 of 564 response objects, or 40.6%.
- **Being ready for Google did not guarantee AI visibility.** Of 90 companies meeting the high-readiness threshold, 24 were never retrieved. Readiness and retrieval were essentially uncorrelated ($r = -0.009$, $p = 0.93$).
- **A business had to be found before it could be recommended.** Retrieval and recommendation were strongly associated ($r = 0.90$), with retrieval accounting for 81% of observed variance in recommendation strength. The measures came from separate modules.

Taken together, the findings show that AI visibility is not simply a technical-search problem. It reflects whether an organization is present, clearly understood, and credibly represented when AI answers a customer's question.

The Index does not prove that visibility causes revenue or that every platform behaves the same way.

Research at a Glance

- 94-company clean sample from 133 requested names
- 846 Gemini queries: three scenarios × three runs per company
- 564 recommendation-response objects
- May–July 2026 collection period
- Statistical testing and confidence intervals
- One-platform dataset: Gemini consistency mode

Research boundary. Selected, nonrandom sample; findings are directional, not census-level estimates for all businesses or platforms.

Why This Matters

> **AI can shape who makes the shortlist before the first click.**

In traditional search, businesses compete on a visible results page. Even a lower-ranked company can still be seen, compared, and considered.

AI-assisted discovery can narrow the field earlier. A synthesized answer may name only a few organizations, attach specific strengths or weaknesses to each, and leave others out entirely. The customer may never know which businesses were considered, which were omitted, or why.

That changes the business risk.

Being absent is no longer the same as ranking lower; it can mean never entering the consideration set.

Being included with unclear or inconsistent information can be nearly as damaging because the first impression forms before the customer reaches the organization's own channels.

For leaders, AI visibility connects communications, reputation, and discoverability. Websites, media coverage, reviews, executive commentary, case evidence, video, imagery, awards, and basic facts all contribute to the public record from which AI systems may build an answer.

The leadership question is no longer just, "Where do we rank?" It is also, "Are we present, accurately understood, and credibly represented when AI answers the questions our customers ask?"

1. From Results Pages to Answer Systems

Search still matters. What has changed is how people increasingly discover businesses. Search engines and standalone AI platforms now answer more questions directly, often combining multiple sources into one response.

A customer looking for a dealership, law firm, healthcare provider, financial adviser, contractor, or senior living community may ask a conversational question instead of entering keywords. The answer may compare tradeoffs and name only a few options. Part of discovery has shifted from ranking to recommendation.

BrightLocal found that generative AI had become the third most-used source for local business recommendations in its 2026 survey, behind Google and Facebook [1].

Cox Automotive also found that 25% of new-vehicle buyers used AI tools while shopping, and 83% of surveyed consumers expected AI to reshape car buying within ten years [2]. The finding is industry-specific, but it shows AI-assisted research entering a high-consideration purchase.

SOCi's 2026 Local Visibility Index, covering 2,751 multi-location businesses and more than 350,000 locations, also found AI recommendations to be more selective than Google's local three-pack [3]. The exact rates depend on SOCi's sample and prompts, but the direction matches the Lucie Index: AI recommendation visibility is concentrated.

AI has not replaced search. Organizations now face two related questions:

1. Can a customer find the business in traditional search?
2. Does an AI system retrieve, understand, and recommend the business when answering a decision-oriented question?

Those questions overlap, but they are not identical.

2. What the 2026 Index Studied

Lucie Research began with 133 requested business names. Matching and deduplication produced 132 unique companies and 134 company-domain pairs; two businesses appeared under both www and non-www domains. Both were later excluded, so the clean denominator was unaffected.

Thirty-eight companies were excluded because access was blocked, the site returned an unusable JavaScript-only shell, or the system relied on invalid category detection. The final clean sample contained 94 companies evaluated from May through July 2026.

Each company was tested through three brand-blind customer scenarios, repeated three times. That produced nine queries per company and 846 total Gemini queries. Competitive substitution was evaluated across 564 recommendation-response objects.

The Index measured seven production pillars:

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

The Fidelity pillar was excluded because that module was not operational in production during the collection window.

The findings are based on Gemini consistency mode and should not be interpreted as representative of every AI platform. Because AI answers vary across prompts, sessions, users, and time, each scenario was repeated three times.

Recent measurement research also supports repeated observations rather than single-run precision [4].

What the study demonstrates

The Index describes the visibility outcomes and signals observed in the clean sample, tests relationships among measured variables, and quantifies how often the subject business or a competitor appeared in the tested responses.

Study Limitations

The study did not measure conversions, revenue, market share, customer acquisition, or long-term causal effects.

It did not isolate every content, communications, or rich-media feature, test every platform, or use a random sample. Its findings are directional rather than universal.

Methodology and Data-Quality Flow

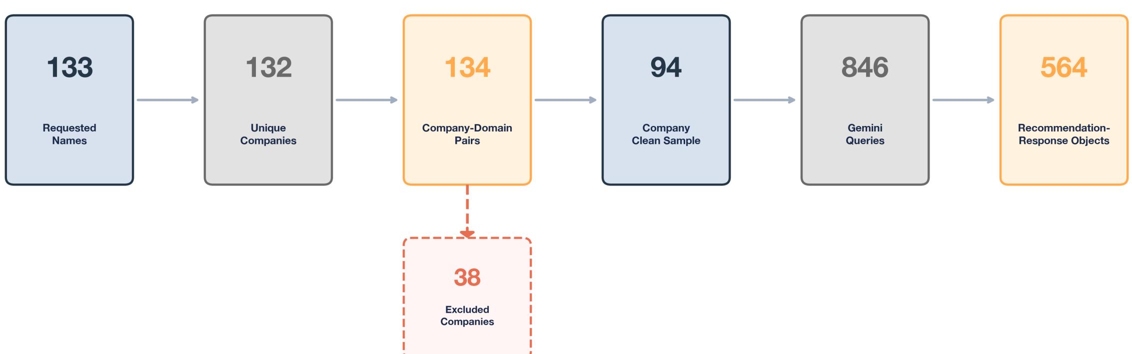

Figure 1: *The final Index dataset was developed through matching, deduplication, and quality exclusions.*

3. The 2026 Index Findings

How to read these findings

The findings describe a selected 94-company sample tested in Gemini consistency mode from May through July 2026.
They are directional, not universal, and do not measure revenue or conversion. Finding-specific limits are noted only where necessary.

Finding 1: AI either found businesses—or missed them entirely

Twenty-five of 94 companies, or 26.6%, did not appear in any of the tested scenarios. The 95% confidence interval was 17.7% to 35.5%.

At the other extreme, 54 companies, or 57.4%, appeared in every tested scenario. Only 15 companies were retrieved in some, but not all, scenarios.

In total, 79 of 94 companies, or 84.0%, fell at one of the two retrieval endpoints.

Why it matters

Within this sample and testing design, AI visibility behaved less like a ranking system and more like an inclusion threshold. Many companies were either consistently retrieved or not retrieved at all.

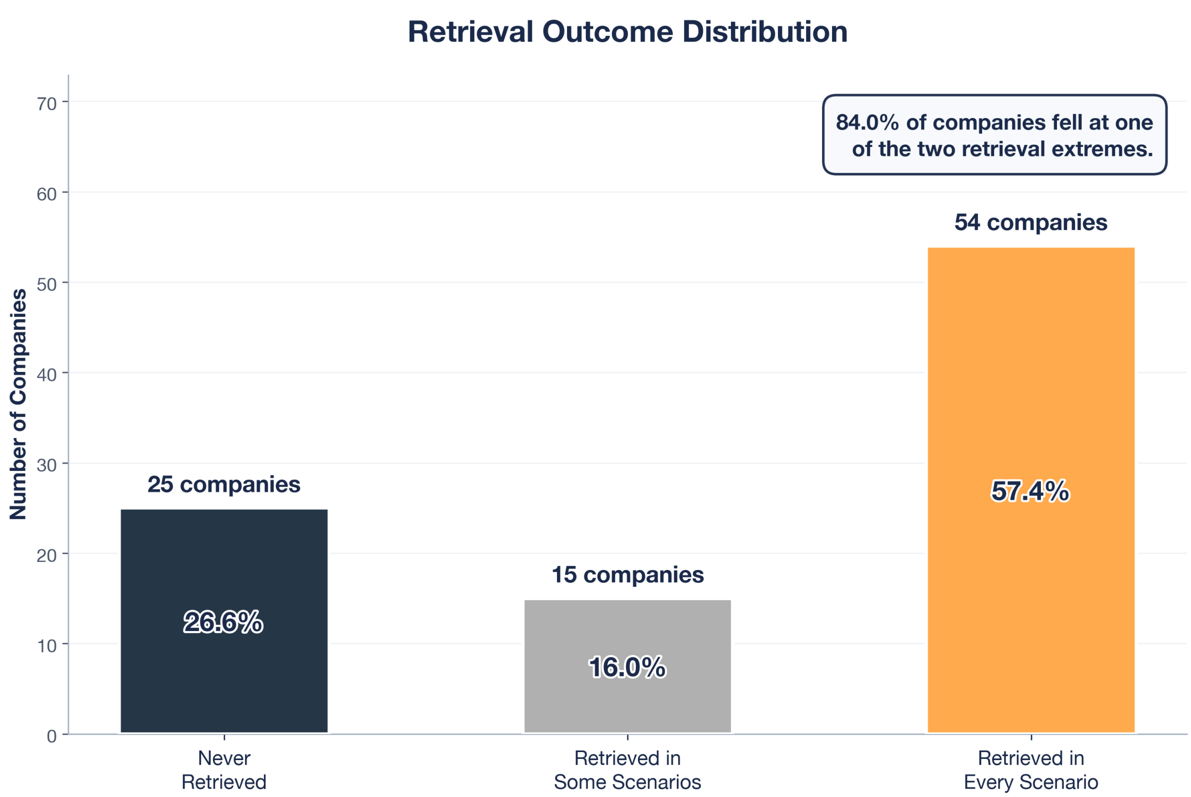

Figure 2: *Retrieval clustered at the endpoints: 25 companies were never retrieved, while 54 were retrieved in every tested scenario.*

Finding 2: AI often struggled to understand what businesses do

Thirteen companies, or 13.8%, showed limited knowledge-credibility signals under the framework.

In those cases, the information available to the model was not sufficiently complete, consistent, or well-supported to produce stronger credibility signals.

Twenty-six companies, or 27.7%, showed limited category-association signals. In those cases, the model had difficulty identifying or consistently associating the business with what it does.

Why it matters

A business can be present online without being represented clearly enough for an AI system to identify its category or evaluate its claims.

The dataset did not identify which specific fact caused the limited signal for each company.

Finding 3: When AI didn't recommend the subject business, it often recommended a competitor instead.

Across 564 recommendation-response objects, a competitor was recommended instead of the subject business 229 times, or 40.6% of the time.

A 10,000-resample company-clustered bootstrap produced a 95% interval of 31.7% to 49.6% for the competitor-substitution rate.

The wider interval reflects the fact that multiple responses were grouped within each company.

Why it matters

When AI didn't recommend the subject business, it often recommended a competitor instead.

That creates competitive exposure before the customer reaches a website or speaks with a representative. The study did not measure clicks, leads, purchases, or revenue.

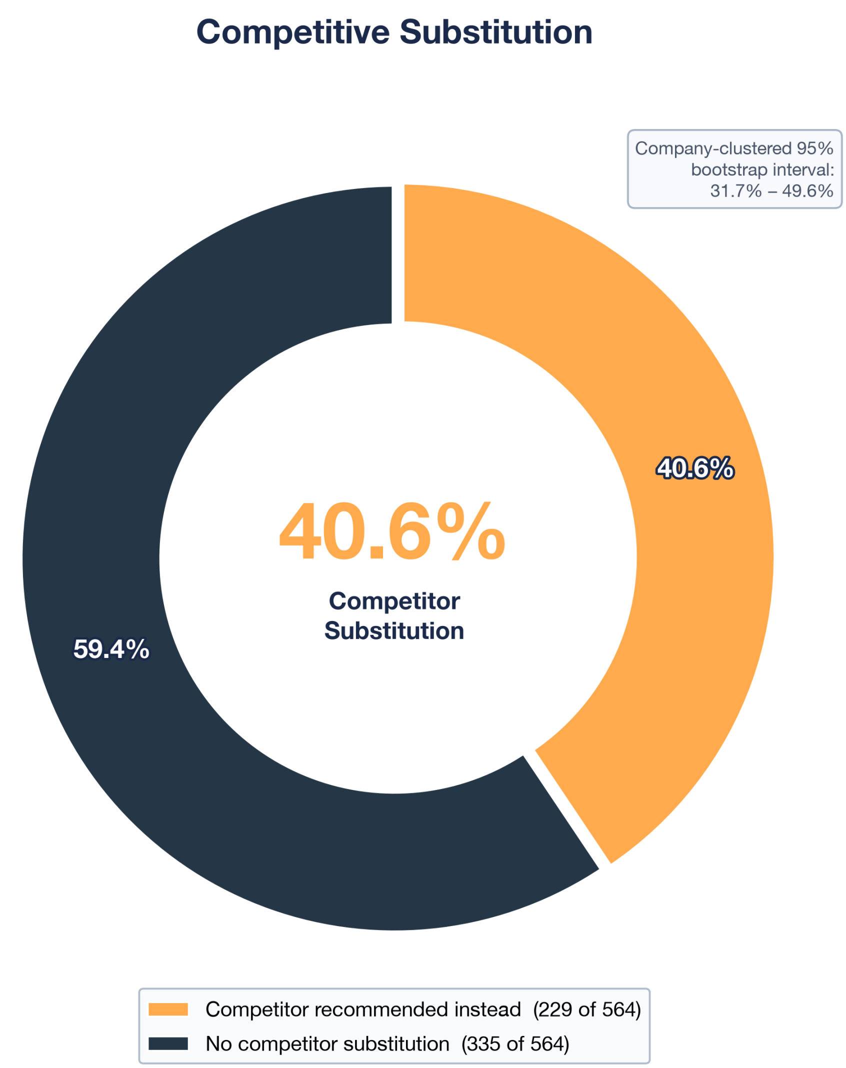

Figure 3: *In 40.6% of recommendation-response objects, the answer recommended a competitor instead of the subject business.*

Finding 4: Being ready for Google did not guarantee AI visibility

Ninety of the 94 companies met the framework's high Google-readiness threshold. Of those 90 companies, 24, or 26.7%, were still not retrieved in any tested scenario.

The correlation between eligibility and retrieval was effectively zero:

$$r = -0.009, \qquad p = 0.93.$$

Why it matters

Technical readiness for conventional search did not explain whether a company appeared in the tested AI answers.

The Lucie eligibility measure reflects readiness, not page-one ranking, organic traffic, or search performance.

SOCi's separate 2026 study also found limited overlap between leaders in traditional local search and AI recommendations [3].

Finding 5: A business had to be found before it could be recommended

The retrieval and recommendation measures were strongly associated:

$$r = 0.90, \qquad p = 5.67 \times 10^{-35}.$$

Retrieval accounted for 81% of the observed variance in recommendation strength:

$$r^2 = 0.81.$$

The relationship was not mechanically built into the measurement framework. Retrieval and recommendation were calculated by separate modules.

Some companies showed measurable recommendation strength despite not being retrieved by the retrieval module, while another company was retrieved in some scenarios but showed no recommendation strength.

Why it matters

In practical terms, being found came first. A company that was not retrieved had little opportunity to be recommended.

The relationship is strong, but it does not prove that retrieval alone causes recommendation.

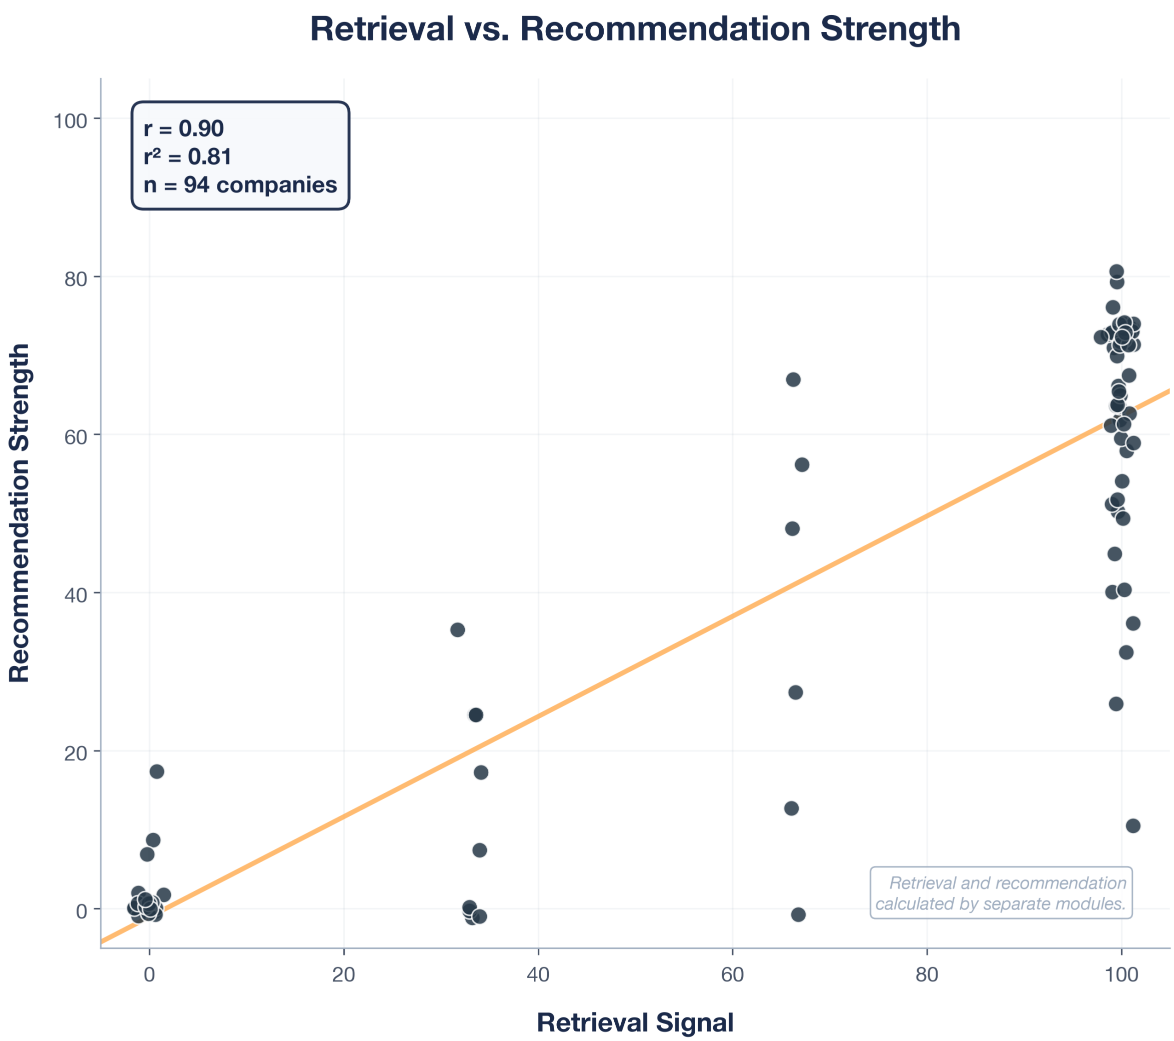

Figure 4: *Companies with stronger retrieval signals generally showed stronger recommendation signals.*

The executive takeaway

Across the five findings, the pattern is clear: AI visibility was uneven, competitors frequently filled the gap, and retrieval mattered far more to recommendation than Google readiness did to retrieval.

In this sample, being present was the first requirement for being considered.

Selected Strong Performers in the 2026 Sample

Seven organizations demonstrated consistently strong AI visibility signals across the tested customer scenarios:

- Olansky Dermatology
- Homrich Berg
- Atlanta Dental Spa

- Georgia Aquarium
- Georgia Power
- OneTrust
- Piedmont Healthcare

Each was retrieved in every scenario and showed comparatively strong recommendation signals within the sample.

The examples span healthcare, financial services, technology, utilities, and consumer-facing organizations, showing that strong signals were not limited to one sector.

The examples are illustrative, not a complete ranking. Results reflect the May–July 2026 collection period and may change as platforms, available information, and model behavior evolve.

The organizations were evaluated through the same research process and were not required to participate in or endorse the study.

4. Interpreting the Findings

The Index measured outcomes, not every possible cause. The discussion below combines findings from the 2026 Index with peer-reviewed generative-search research, independent industry studies, and Lucie's clearly labeled applied observations to help explain those patterns.

Evidence and clarity matter

In 2024, researchers from Princeton University, IIT Delhi, Georgia Tech, and the Allen Institute for AI published "GEO: Generative Engine Optimization" at the ACM SIGKDD conference [5].

In a 10,000-query benchmark, credible citations, relevant quotations, and statistics improved source visibility in the experimental setting, with gains of up to 40%.

Results varied by domain and do not establish a universal formula, but they show that content structure and evidence can influence prominence in generated answers.

Independent research supports the findings

Independent industry studies show that AI recommendation visibility is selective and can differ from conventional local-search visibility [3].

Lucie's applied work also considers video and imagery because they can provide concrete evidence about people, facilities, work, products, and outcomes.

The 2026 Index did not isolate a rich-media effect, so this remains an applied communications recommendation rather than a quantified finding.

Repeated testing matters

AI answers can vary across sessions, prompts, and platforms.

Lucie observes contradictions in live diagnostics, but the 2026 dataset did not measure their frequency.

Repeating each scenario three times reduced the risk of treating a single answer as stable; recent measurement research similarly argues that AI visibility should be observed across repeated runs [4].

5. Why AI Visibility Is a Communications Issue

AI-generated answers are built from a broader public information environment than a website alone. That makes strategic communications part of discoverability.

AI systems may encounter websites, business profiles, reviews, media coverage, executive interviews, reports, videos, images, directories, awards, and other third-party sources.

Platforms use different methods, but organizations are represented through the information available across that environment.

This creates five communications responsibilities.

1. Owned information

Service pages, location pages, executive biographies, FAQs, newsroom content, reports, case studies, and product information should state what the organization does in clear, specific language.

Important facts should be visible in the page content rather than existing only in hidden markup.

2. Earned validation

Journalism, trade coverage, independent reviews, community recognition, awards, analyst references, and credible third-party citations can provide evidence beyond the organization's own claims.

The foundational GEO research found that citations and relevant external support improved visibility within its experimental setting [5].

3. Executive and expert authority

Interviews, speeches, testimony, authored commentary, research, and subject-matter explanations can connect an organization to the expertise and categories it wants to be known for.

These materials should be factual, attributable, and consistent with the organization's public record.

4. Visual evidence

Useful video and imagery can show people, facilities, work, processes, products, and outcomes.

They can strengthen the evidence available to customers and make important claims easier to understand and verify.

The Index did not independently quantify a rich-media effect, so visual content should be treated as a communications and evidence-building opportunity rather than a guaranteed AI-visibility tactic.

5. Factual consistency

Names, locations, services, leadership, credentials, categories, hours, and contact information should agree across the organization's website and major external profiles.

The Index's knowledge-credibility and category-association gaps show that clear representation cannot be assumed merely because a business has a website.

For communications teams, the shift is practical: reputation, discoverability, and factual governance now meet in the same information environment.

AI visibility is therefore not only a web or search task; it is an organizational communications responsibility.

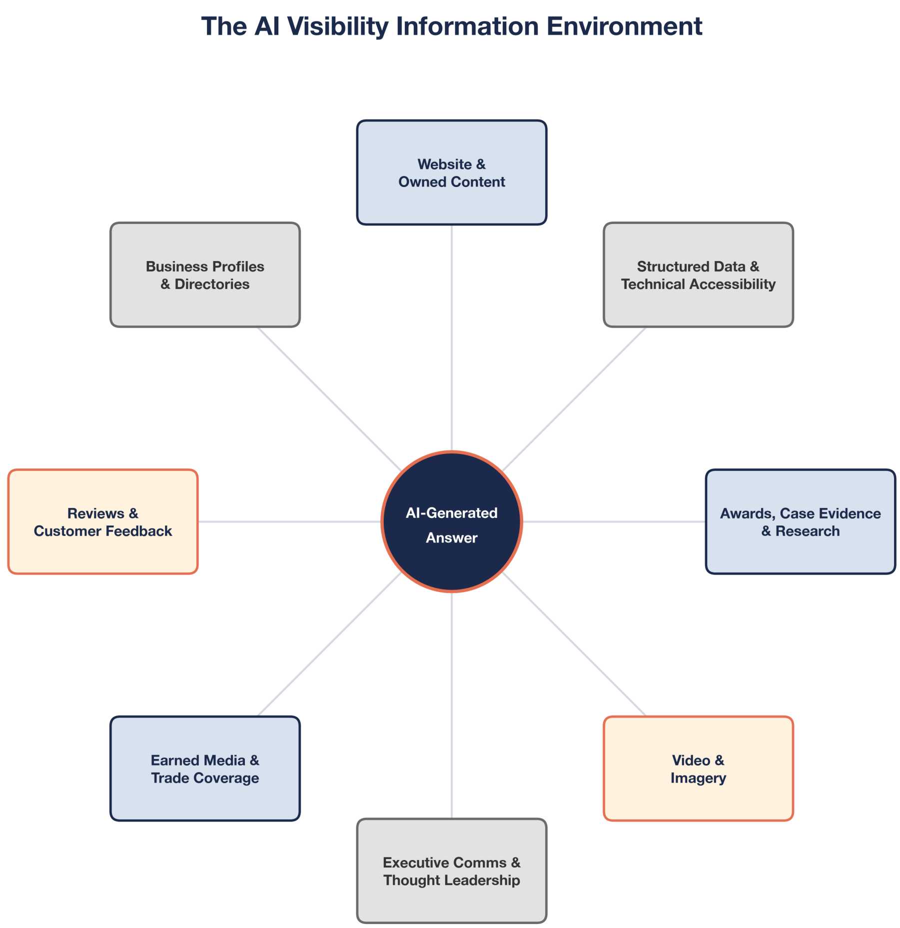

Figure 5: *AI visibility is shaped by an organization's broader public information environment, not by schema or website copy alone.*

6. A Practical Framework

The research points to a practical sequence: measure the current answer, correct the information environment, strengthen evidence, and test again.

1. **Measure how AI answers the questions customers actually ask**

   Create a defined set of brand-blind recommendation, comparison, trust, and category prompts. Run them repeatedly and record whether the business appears, how it is described, which sources are cited, and which competitors are recommended.

   One answer is an observation; repeated answers provide a more credible visibility pattern.

2. **Correct the core facts**

   Audit names, locations, services, hours, leadership, credentials, and contact information across the website, business profiles, directories, and review platforms.

   Resolve conflicts and remove obsolete information.

3. **Clarify categories and service relationships**

   Use plain language to explain what the organization does, who it serves, where it operates, and how its services relate to customer needs.

   Avoid relying on slogans or internal terminology that does not identify the category clearly.

4. **Answer real questions on owned channels**

   Publish useful FAQs, guides, comparison pages, service explanations, research, and newsroom content based on the questions customers and stakeholders ask.

   Use descriptive headings and structured data only when it matches visible page content.

5. **Support claims with evidence**

   Use named awards, dated statistics, credible citations, detailed case evidence, independent reviews, and attributable expert commentary.

   The goal is not to add claims; it is to make important claims verifiable.

6. **Strengthen third-party authority**

   Pursue relevant journalism, local coverage, trade publications, interviews, community recognition, and independent references.

   Earned validation expands the evidence environment beyond brand-owned statements.

7. **Use video and imagery where they add proof**

   Show the people, work, facilities, products, processes, and results that matter to customer decisions.

   Accompany media with clear titles, descriptions, transcripts or supporting text, and appropriate markup.

8. **Re-measure and document change**

AI platforms evolve, and individual outputs vary. Re-run the same prompt set at defined intervals, compare results with the prior baseline, and document both improvements and regressions.

Visibility should be treated as an ongoing measurement program rather than a one-time snapshot.

These actions do not guarantee recommendation. They improve the likelihood that an organization will be accurately understood, retrieved, and considered by both people and AI systems.

7. Framework Validation

Before applying the framework externally, Lucie tested it against its own digital presence.

The initial assessment identified areas of limited and inconsistent visibility. After changes to question-based content, factual consistency, evidence-rich pages, and supporting media, later assessments showed stronger overall visibility signals.

This example is illustrative. It is not part of the 94-company Index analysis, and the before-and-after observation does not isolate the effect of any single change or independently establish causation.

Its purpose is to document that Lucie used the framework on itself before applying the approach in client work.

Frequently Asked Questions

Is GEO a replacement for SEO?

No. Traditional search remains a major discovery channel, and many search fundamentals also support AI retrieval.

The Index found that Google readiness did not guarantee AI retrieval, not that SEO or search ranking no longer matters. Businesses should measure both traditional and AI-assisted discovery.

What is AI visibility?

AI visibility describes whether and how an organization appears in AI-generated answers.

It includes retrieval, factual representation, category association, cited support, comparison, and recommendation.

What is Generative Engine Optimization?

Generative Engine Optimization, or GEO, is the practice of improving the visibility of sources and organizations within responses generated by AI systems.

The term was formalized in research published at ACM SIGKDD in 2024 [5].

In practice, AI visibility also involves communications, reputation, technical accessibility, and information consistency.

Does AI visibility guarantee leads or revenue?

No. This Index did not measure conversions, revenue, or customer acquisition.

Visibility creates an opportunity to enter the consideration set; it does not guarantee that a customer will choose the business.

How is AI visibility measured?

A measurement program begins with the questions customers ask, repeats those prompts across defined scenarios and platforms, and records whether the business is retrieved, how it is described, which sources are used, and who is recommended instead.

Because outputs vary, repeated testing is more reliable than a single spot check.

How quickly can visibility change?

There is no universal timeline.

Results depend on the platform, crawl and retrieval behavior, market competition, the nature of the changes, and the consistency of public information.

Progress should be evaluated through repeated measurements against a documented baseline.

Who should own AI visibility inside an organization?

AI visibility is inherently cross-functional.

Communications, marketing, web, SEO, reputation, public relations, subject-matter experts, and leadership may all influence the information environment.

Clear ownership is more important than which department ultimately leads the effort.

Methodology

Dataset Scope

Across its broader research and system-validation work, Lucie Research has evaluated approximately 150 businesses.

The formal 2026 Index dataset began with 133 requested names after early, incomplete, or otherwise unreliable records were set aside before statistical analysis.

The counts below refer only to the Index dataset, not to the full body of evaluations.

Sample and data-quality trace

The initial analysis list contained 133 requested names. After matching and deduplication, it represented 132 unique companies and 134 company-domain pairs.

Two businesses appeared under both www and non-www versions of their domains.

Thirty-eight companies were excluded because of firewall or web application firewall blocking, JavaScript-only application shells without usable server-rendered content, or invalid system-side category detection.

The two businesses with duplicate domain records were among the excluded companies, so the duplicate entries did not affect the final clean denominator.

The final dataset trace was:

> **133 requested names → 132 unique companies → 134 company-domain pairs → 38 excluded companies → 94 companies in the clean sample.**

Four names required alias matching, one required a partial-name match, and 14 companies had multiple runs.

When a company had more than one eligible run, the most recent eligible result was used.

Query design

Each company was tested through three brand-blind customer decision scenarios, with three runs per scenario.

This produced nine AI queries per company and 846 total queries. Scenario design covered recommendation, comparison, trust, and category intent.

Platform

Retrieval and recommendation were run in Gemini consistency mode.

Because the proprietary statistics come from one platform, they should not be read as estimates of how every AI platform behaves.

Measurement framework

The production pipeline generated research outputs for:

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

The Fidelity pillar was excluded because that module was not operational in production during the collection period.

Proprietary methodology statement

The Lucie AI Visibility Index 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.

Competitive substitution was evaluated across 564 recommendation-response objects.

Statistical analysis

The analysis used Pearson correlation, a z-test for proportions, Wilson confidence intervals, distribution summaries, and a company-clustered bootstrap implemented with `scipy.stats` and `statsmodels`.

Unless otherwise noted, the significance threshold was:

$$\alpha = 0.05.$$

For the competitor-recommendation result, the report presents the observed figure of 229 out of 564 responses, or 40.6%.

A 10,000-resample company-clustered bootstrap produced a 95% interval of 31.7% to 49.6%.

Research records and reproducibility

Lucie Research maintains internal archives of the supporting datasets, prompt records, screenshots, calculations, analysis files, methodology versions, and final publication materials used for this edition.

These records are retained to support internal verification and the reconstruction of reported findings, subject to confidentiality, privacy, contractual, and proprietary-methodology restrictions.

Naming policy

The Index findings are reported primarily in aggregate.

Businesses that showed limited or inconsistent visibility signals are not identified.

A limited set of strong performers is named to provide verified, illustrative examples; each named result was confirmed against dated collection records and supporting screenshots.

Inclusion does not imply participation in or endorsement of the study.

Limitations

AI outputs are probabilistic and change over time. The results reflect the May–July 2026 collection window.

The sample was selected rather than random and supports directional findings, not census-level claims.

Because blocked records were excluded, the clean-sample analysis may understate the access barriers businesses experience in practice.

The proprietary statistics come from one AI platform. Correlation does not prove causation.

Retrieval and recommendation were calculated by separate modules, but unmeasured variables may influence both.

The database did not independently isolate every content, communications, technical, or rich-media feature.

Recommendations concerning content evidence, rich media, and communications practice therefore draw on the Index findings together with peer-reviewed GEO research, independent industry research, and Lucie's applied observations.

The study did not measure clicks, leads, purchases, revenue, market share, or conversion outcomes.

AI visibility should not be treated as a direct financial-performance metric without additional evidence.

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 & 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 & 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.

About Lucie Research

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

The 2026 report is the inaugural edition of the 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.

Learn more at luciecontent.com.

Media inquiries: info@luciecontent.com.

References

[1] BrightLocal. " Local Consumer Review Survey 2026: Star Ratings Keep Rising, but Trust Is Getting Harder to Earn. " 2026.

https://www.brightlocal.com/research/local-consumer-review-survey/

[2] Cox Automotive. " Cox Automotive Car Buyer Journey Study Finds Efficiency, Digital Tools and AI Drive Record Satisfaction. " January 13, 2026.

https://www.coxautoinc.com/insights/cox-automotive-car-buyer-journey-study-finds-efficiency-digital-tools-and-ai-drive-record-satisfaction/

[3] SOCi. " The Most Visible Local Brands of 2026 " and related 2026 Local Visibility Index analysis. 2026.

https://www.soci.ai/insights/lvi/

https://www.soci.ai/blog/how-to-rank-in-chatgpt-perplexity-and-google-ai-overview/

[4] 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

[5] 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

https://arxiv.org/abs/2311.09735