

• BENCHMARK STUDY • FIRST EDITION

THE 2026

Triangle AI Visibility Study

How 52 Triangle-area small businesses show up
when people ask AI about them.

AUTHOR

Lenise Kenney
Founder

PUBLISHED

July 22, 2026
Data collected July 20 and 21

• KEY FINDINGS

Not one of 52 came through clearly.

In July 2026 we asked four major AI platforms, ChatGPT, Claude, Gemini, and Perplexity, about 52 small businesses across the Triangle. We scored every answer against each business's own published information. The results were consistent, and they were poor.

0 of 52

Businesses that reached the top visibility tier on any platform

79%

Landed in the two lowest visibility tiers (41 of 52 businesses)

2.48/10

Mean accuracy of the best-performing platform, Perplexity

45 of 52

Claude responses that included a fabricated fact

Not one of the 52 businesses came through clearly. Zero reached the top visibility tier. Across four AI platforms, every business had gaps.

Ten businesses (19%) were absent entirely. No platform could identify them at all, or the platforms described a different business.

The platforms failed in two different ways. ChatGPT and Gemini more often said nothing. Claude and Perplexity more often filled the gap with invented details.

The businesses that did best had done ordinary web work. Those with structured data and fuller website content scored higher, even though none had set out to prepare for AI.

“

Not one of the 52 Triangle small businesses we tested came through clearly to AI. The best-performing platform still got most of them wrong.

• THE TEST

Four platforms, three questions, 208 answers.

In July 2026, Signal & Structure AI checked how 52 small businesses across the Triangle appear to four AI platforms: ChatGPT, Claude, Gemini, and Perplexity. The businesses span six communities and eleven industries. Durham accounted for 24 of them, Cary 11, Raleigh 9, Chapel Hill 5, Morrisville 2, and Apex 1.

For each business, we asked the four platforms the kinds of questions a customer would ask before choosing a local service: what the business is, what it offers, and whether it would be recommended. We then measured two things. First, accuracy: how correctly each platform described the business against verified facts, on a 0 to 10 scale. Second, hallucination: how often a platform stated something false as if it were true. That produced 208 platform responses in total, four for each business.

The result is a snapshot of one moment in time for one region. It is not a forecast, and it does not name any business. Every number in this study is aggregate.

Finding 1: Nobody reached the top

We sort each business into one of four visibility tiers based on how well the four platforms find and describe it: Strong, Weak, Low, or No Signal. Figure 1 shows where the 52 businesses fell.




FIGURE 1 Most Triangle businesses tested landed in the two lowest visibility tiers. None reached the top. Strong Signal: 0 (0%). Weak Signal: 11 (21%). Low Signal: 31 (60%). No Signal: 10 (19%).

Zero businesses reached Strong. Eleven reached Weak: found, but described with gaps or errors. The other 41 were barely visible or not visible at all, and the ten in the No Signal tier are the starkest case: the platforms could not identify them, or described a different business entirely. This is the headline of the study. A customer who asks an AI platform about businesses in these categories is unlikely to hear any of these 52 named correctly, and for ten of them, unlikely to hear them named at all.

• FINDING 2

Even the best platform got most things wrong.

Accuracy was low across all four platforms. Figure 2 shows the average score each platform earned against verified facts.




FIGURE 2 Every platform averaged under 2.5 out of 10 on accuracy. Perplexity led, ChatGPT trailed.

PLATFORM	MEAN ACCURACY (0 TO 10)	MEDIAN
Perplexity	2.48	2.0
Claude	2.29	2.0
Gemini	0.90	0.0
ChatGPT	0.65	0.0

Perplexity scored highest and still averaged under a quarter of the available points. ChatGPT and Gemini both had a median of zero, which means at least half of their responses earned no accuracy credit at all.

The hallucination numbers, shown in Figure 3, tell the other half of the story and look contradictory at first.




FIGURE 3 The platforms with higher accuracy also invented the most. Claude produced a false statement in 45 of 52 responses (87%). Perplexity: 37 of 52 (71%). Gemini: 19 of 52 (37%). ChatGPT: 8 of 52 (15%).

Claude and Perplexity earned the best accuracy scores, yet they also invented the most. ChatGPT and Gemini invented the least, yet scored lowest on accuracy. The two sets of numbers fit together once you see that the platforms failed in two different ways.

When ChatGPT and Gemini did not have solid information about a business, they more often returned little or nothing. That keeps their hallucination rate low, but it leaves the business invisible. Claude and Perplexity more often attempted a full answer. Sometimes that produced a correct detail, which lifted their accuracy score. More often it produced a confident description of the wrong business, wrong services, or wrong location.

For a business owner, neither failure mode is good. One means a customer hears nothing about you. The other means a customer hears something false about you and has no way to know it is false.

“**The platforms that tried hardest to answer were also the ones most likely to make things up.**

• FINDING 3

Dental practices led, restaurants trailed.

Visibility varied by industry. Each business in the study receives a single combined visibility score from 0 to 10, which blends how accurately the four platforms describe it with the quality of the machine-readable information the business publishes about itself. The appendix defines this score and the tier bands built on it. The table below shows the average across the categories with enough businesses to report.

INDUSTRY	BUSINESSES	MEAN VISIBILITY SCORE (0 TO 10)
Dental	14	4.77
Home Services	14	3.96
Beauty	3	3.70
Financial Services	3	3.33
Restaurant	7	1.49

Dental practices came out ahead of every other category. They tend to keep detailed, consistent listings and structured websites, which gives AI platforms more to work with. Home services businesses, the other large group in the sample, sat in the middle.

Restaurants scored worst by a wide margin, at 1.49. Many rely on third-party platforms and social media rather than their own detailed websites, so the information an AI platform can read directly about them is thin. Several other industries appeared in the sample in smaller numbers, including legal, fitness, healthcare, real estate, and consulting. Those groups were too small to report on their own, so we have left them out of the industry comparison rather than draw conclusions from two or three businesses.

• FINDING 4

Ordinary web work carried over.

None of the 52 businesses had set out to prepare for AI. Yet some were easier for AI platforms to read than others, and the reason was ordinary website quality.

Businesses with structured data on their sites scored higher on accuracy than those without. Businesses with fuller, more complete website content also scored higher than those with thin content. Neither group had done anything AI-specific. They had built solid websites for regular search reasons, and that work carried over to how AI platforms understood them.

This is a correlation, not proof of cause, and the sample is small. But the direction is consistent with what we see in individual audits. AI platforms read the open web. A business that has published clear, structured, complete information about itself gives those platforms something accurate to repeat. A business that has not leaves the platforms to guess, and the guesses are often wrong.

“

The businesses that did best had never done AI optimization. They had built solid websites for regular search, and that work carried over.

• WHAT THIS MEANS

Three observations for Triangle businesses.

More people now ask AI platforms for local recommendations before they open a search engine or a map. This study measured what the platforms say when asked directly about 52 Triangle businesses, and the answer is: very little, and much of it wrong.

01 **Being invisible to AI is not a sign of a failing business.**

Several of the businesses in the lowest tiers are established and well run. They simply have not published their information in a way AI platforms can read cleanly, and most owners do not yet know the gap exists.

02 **A good review profile is not enough on its own.**

AI platforms build their answers from what they can read across the open web, not from a single strong listing. When the underlying website information is thin or inconsistent, even a business with strong reviews can come through as No Signal.

03 **The work that correlated with better results is not exotic.**

Clear website content and structured, consistent business information are the same things that help with ordinary search. The businesses that scored best in this study got there without any AI strategy at all.

The businesses that scored best had no AI strategy at all. They published clear, complete, consistent information, and the platforms repeated it.

• APPENDIX

How this study was done.

WHO RAN IT

This study was conducted by Signal & Structure AI, a Durham company that works on AI visibility for businesses. We are naming that up front because it matters when reading a vendor-produced benchmark. The aggregated dataset behind every figure in this study is available to journalists and researchers on request.

WHEN

All platform queries were run July 20 and 21, 2026. Every business was tested in the same window, so no business was measured against a different version of any platform.

THE SAMPLE

52 small businesses across Durham (24), Cary (11), Raleigh (9), Chapel Hill (5), Morrisville (2), and Apex (1), spanning eleven industry categories. Businesses were identified through public business directory searches, not drawn from our client list, and none is a client. Each business needed a working public website to be included. A small number of businesses whose checks failed to complete for technical reasons were excluded before analysis rather than scored incomplete.

THE PLATFORMS

Four AI platforms were queried through their standard interfaces: ChatGPT (GPT-4o), Claude (Claude Sonnet 4), Gemini (Gemini 2.0 Flash), and Perplexity (Sonar). Model versions are listed so the snapshot can be compared fairly against future editions.

THE QUESTIONS

Each platform was asked the same three kinds of questions about every business: what the business is, what it offers, and whether the platform would recommend it. These mirror what a prospective customer asks before choosing a local service. Illustrative phrasings, not the exact wording used: "Tell me about [business name]," "What services does [business name] offer?," "Would you recommend [business name]?" We do not publish the exact wording, for one reason: this is a recurring benchmark, and keeping the question set fixed and unpublished is what lets future editions be compared honestly against this one.

SCORING

Each platform response was scored 0 to 10 for factual accuracy against the business's own published information: its website and its public listings. The same rubric was applied to every response, every platform, every business. One boundary case matters for reading the numbers: a statement the published record could not confirm earned no accuracy credit, but it was only counted as a fabrication when it

conflicted with that record or described a different business. So the accuracy scores measure agreement with what the business has published, and the hallucination flags mark active contradiction, not mere absence. A response was flagged for hallucination when it stated something false as if it were fact, such as a wrong location, a service the business does not offer, or a description of a different business entirely.

WHY SILENCE SCORES LOW

A platform that returns nothing about a business, or declines to answer, earns a low accuracy score. This is deliberate, and it is why the median accuracy for ChatGPT and Gemini is zero. The study measures what a customer actually learns when they ask, and a customer who gets no answer learns nothing. A skeptic could argue this is unfair to cautious platforms, and we would agree it penalizes caution. From the business’s point of view, though, the outcome is identical: the customer moves on. We report hallucination rates separately for exactly this reason, so silence and fabrication are visible as different failures.

THE COMBINED SCORE AND THE TIERS

Beyond per-response accuracy, each business receives one combined visibility score from 0 to 10. It blends two things: how accurately the four platforms describe the business, and the quality and completeness of the machine-readable information the business publishes about itself, which is what the platforms have to work from. The exact weighting inside that score is proprietary. The tier bands applied to it are not, and they were fixed before any Triangle business was tested, not drawn afterward.

STRONG SIGNAL	7.5 and above	Platforms find the business and describe it accurately.
WEAK SIGNAL	5.0 to 7.4	Platforms find the business but with gaps or errors.
LOW SIGNAL	2.5 to 4.9	Platforms have little correct to say about the business.
NO SIGNAL	Below 2.5	Platforms cannot identify the business, or describe a different one.

LIMITATIONS

This is a one-time snapshot of one region, not a trend line. AI platform answers vary between runs, so any individual response may differ on a later date. Several industries appear in numbers too small to report separately. One class of business-listing signal was set aside for this edition because of a data-collection issue on our side, so this study reports on website content and platform-response signals only. Business names are not disclosed, and all figures are aggregate.

ABOUT SIGNAL & STRUCTURE AI

Signal & Structure AI is a Durham, NC company that measures how AI platforms find and describe businesses. Founded by Lenise Kenney and Julian Bass. signalstructure.ai.

The Triangle AI Visibility Study is the first edition of a recurring benchmark. Future editions will track how these numbers change over time and extend to other cities.

MEDIA AND DATA REQUESTS

Journalists and researchers can request the aggregated dataset behind every figure in this study, industry-level breakdowns, or a walkthrough of the methodology. No individual business results are shared with anyone, including the press.

Contact: Lenise Kenney, Founder · hello@signalstructure.ai · signalstructure.ai

The 2026 Triangle AI Visibility Study. Signal & Structure AI. Data collected July 20 and 21, 2026. n=52 businesses, 208 platform responses across ChatGPT, Claude, Gemini, and Perplexity. Published July 22, 2026.