

VISIBILITYINDEX.AI

August 2026

Local AI Visibility Index Report

A measurement study of how consistently AI platforms agree on local business recommendations, across 12 US cities, 12 service industries, and 5,760 individual searches.

First edition, Version 1.0 • Published August 2026 • VisibilityIndex.ai Research

Table of Contents

1. Executive Summary
2. Methodology
3. Finding 1: Most Businesses Are Visible on Only One Platform
4. Finding 2: True Consensus on a Market Leader Is Rare
5. Finding 3: A Concrete Example of Total Disagreement
6. Finding 4: Local, Independent Businesses Dominate Recommendations
7. Finding 5: Platform Pairs Do Not Agree Equally
8. Limitations and Scope
9. Appendix: Detailed Methodology
10. About VisibilityIndex.ai

Executive Summary

Consumers increasingly ask AI assistants directly for local business recommendations rather than browsing a search results page. This report provides a structured, cross platform measurement of how often major AI assistants agree on local business recommendations, an area where clear, data backed comparisons have been limited.

VisibilityIndex.ai tested ChatGPT, Gemini, Claude, and Perplexity against 1,440 unique local business recommendation questions, each run identically across all four platforms for 5,760 total individual searches, spanning 12 US cities and 12 service industries including HVAC, plumbing, roofing, electrical, moving, pest control, and home cleaning. Each platform was queried using its default, automatic search behavior to reflect real world usage rather than an artificially forced test condition.

The headline findings:

- 74 percent of businesses recommended by any platform were recommended by exactly one platform, and did not appear on the other three.
- Only 4.6 percent of recommended businesses were named by all four platforms tested.
- Across 144 city and industry markets studied, only 3.5 percent produced a single business that all four platforms agreed was the top recommendation.
- 84 percent of all recommendations went to businesses operating in a single city, confirming that AI recommendations are dominated by genuinely local businesses rather than national chains.

Together, these findings indicate that a business's visibility to one AI platform provides little information about its visibility elsewhere. This report is the first release in an ongoing, recurring index intended to track how AI recommendation behavior evolves over time.

Methodology

This study was designed to measure real world AI recommendation behavior under conditions that reflect how an ordinary consumer would use these tools, rather than a synthetic or artificially constrained test.

Platforms tested	ChatGPT, Gemini, Claude, and Perplexity
Cities	12 major US metropolitan markets
Industries	12 local home service categories
Search prompts per market	10 standardized buyer intent questions
Unique search queries	1,440
Total individual searches	5,760 (across all platforms)
Search behavior	Default automatic search mode for each platform, not forced or constrained
Data collection window	Late July through early August 2026

For each search, the top ranked businesses returned by each platform were recorded and matched against one another by business name and web domain. Agreement between platforms was measured using Jaccard similarity, a standard method for comparing the overlap between two sets, applied here to the sets of businesses each platform recommended for the same search.

All data referenced in this report reflects genuine platform output collected during the study window. No recommendations were generated, simulated, or estimated.

Finding 1

Most Businesses Are Visible on Only One Platform

Across every business recommended anywhere in the study, 74 percent were recommended by exactly one of the four platforms tested and never appeared on the other three. Only 4.6 percent of businesses achieved recommendation from all four platforms.




Figure 1. Share of recommended businesses by number of platforms naming them.

One representative example: a regional pest control company appeared 24 times across 6 different cities through a single platform, and did not appear once across the other three, despite a consistent, multi market presence on the one platform that did surface it.

Finding 2

True Consensus on a Market Leader Is Rare

Looking at each of the 144 city and industry markets studied as a whole, rather than individual searches, only 5 markets, 3.5 percent, had all four platforms converge on the exact same top recommended business. In 42 markets, 29.2 percent, all four platforms each named a different top business entirely.




Figure 2. Outcome distribution across all 144 markets studied.

This measure is a stricter test than platform level overlap, since it asks whether consensus exists at the level of an entire market rather than a single search. The rarity of full agreement at this level suggests that even well established, multi platform visible businesses rarely hold an undisputed top position across every AI platform simultaneously.

Finding 3

A Concrete Example of Total Disagreement

The clearest illustration in the dataset comes from a single search: which HVAC contractors are good for a smaller job in Denver, Colorado. All four platforms were asked this identical question. Each returned three different businesses. Across all twelve names given between them, not one business was named by more than one platform.

Platform	Top recommendation
ChatGPT	Business 1
Gemini	Business 2
Claude	Business 3
Perplexity	Business 4

Table 1. First ranked recommendation per platform for an identical search. Business names shown as generic labels; each represents a distinct, real business named in the platform's actual response.

A consumer asking this exact question would receive an entirely different answer depending solely on which AI assistant they happened to open. Every industry and city combination in this study received an identical number of searches, so this result is not an artifact of one category being tested more heavily than another.

Finding 4

Local, Independent Businesses Dominate Recommendations

Measuring the geographic footprint of every recommended business, meaning how many different cities in the study each business appeared in, 84.3 percent of all recommendation appearances went to businesses operating in a single city. The remaining 15.7 percent was spread across roughly 260 businesses with a presence in more than one city, out of more than ten thousand distinct businesses recommended overall. A small subset of about three dozen businesses operating in seven or more cities alone accounted for nearly 8 percent of all recommendations.




Figure 3. Share of recommendation appearances by business geographic footprint.

This finding runs counter to a common assumption that AI platforms default to well known national brands. The large majority of what these platforms actually recommend is local and independent, meaning genuine visibility, and the risk of invisibility, is concentrated among small businesses rather than national chains.

Finding 5

Platform Pairs Do Not Agree Equally

Agreement between platforms is not uniform. Measuring pairwise overlap between every combination of the four platforms tested, the most aligned pair shared roughly 20 percent overlap in their recommendations for the same searches, while the least aligned pair shared under 8 percent, and returned zero overlapping businesses for the same search roughly half the time.




Figure 4. Mean recommendation overlap between each pair of platforms tested.

This means the size of the AI visibility gap for a given business depends heavily on which specific pair of platforms are being compared, not just how many platforms are considered in total.

Limitations and Scope

This report reflects a single measurement window in August 2026 and describes what was observed during that window rather than a claim about permanent or guaranteed platform behavior. AI platforms update their underlying models and search behavior on an ongoing basis, and results from a repeated study at a later date may differ.

Each platform was tested using its own default automatic search setting rather than a forced, identical search configuration across all four. This was a deliberate choice, intended to reflect how each platform actually behaves for a typical user, but it does mean the platforms were not tested under perfectly identical technical conditions.

The 12 cities and 12 industries studied were selected to represent a broad, high volume cross section of local home services, not an exhaustive survey of every US market or industry category. Findings should be read as representative of the categories studied, not as universal across all local business categories.

This is the first release in an ongoing, recurring index. A single measurement establishes a baseline. The value of continued tracking comes from observing how these patterns shift, or hold steady, across future editions.

Appendix: Detailed Methodology

This appendix documents the technical configuration and calculation methods used in this study, for readers seeking to independently evaluate or reproduce the approach.

Study design

- 12 cities, 12 industries, 10 standardized buyer intent prompts per city and industry combination, yielding 144 markets and 1,440 unique search questions.
- Each of the four platforms was queried once per search, for a total of 5,760 platform level observations.
- Sample prompt patterns used across markets included variations such as "Who are the best [industry] in [city]?", "Which [industry] are good for a smaller job in [city]?", "Which [industry] have the best customer service in [city]?", "Which [industry] offer the best value in [city]?", and "Which [industry] can handle an urgent service request in [city]?", adapted per industry and city.

Platform configuration

Platform	Configuration
ChatGPT	Model gpt-5, low reasoning effort, low search context size, batch submission mode
Claude	Model claude-sonnet-5, automatic tool choice, up to 5 search uses per query
Gemini	Native automatic search grounding
Perplexity	Native search integration, one search executed per query by design

Entity matching and deduplication

- Businesses were matched across platforms primarily by web domain, normalized to strip protocol prefixes and "www" subdomains, with business name used as a secondary signal where a domain was not available.
- Where a platform returned a business without a resolvable official domain, that item was retained and analyzed by name but excluded from domain level cross platform matching.
- Duplicate results within a single platform's response to a single search were removed prior to analysis.

Handling of incomplete or malformed responses

- Across all 5,760 observations collected, 1 failed to produce a usable result and is recorded as a zero result observation.
- That zero result observation was retained as an empty recommendation set rather than excluded from analysis: it was included in pairwise overlap calculations involving that platform for that market, where an empty set against a non-empty set yields a similarity score of zero, and it prevented that market from qualifying for full four platform consensus, since a market leader could not be established for that platform.

- A small number of responses across platforms required automated structural repair, for example correcting malformed output formatting, before they could be included; the content of the business recommendations themselves was not altered by this process.
- No responses were regenerated, resubmitted, or replaced based on the content of what a platform recommended.

Calculation methods

Platform overlap was measured using Jaccard similarity: the number of businesses two platforms both recommended for the same search, divided by the total number of distinct businesses either platform recommended for that search. A score of 1.0 indicates identical recommendation sets; a score of 0 indicates no shared businesses.

Market level consensus was evaluated by comparing each platform's top ranked business for every market, rather than the full recommendation list, to test agreement on a single leading recommendation rather than partial list overlap.

Geographic footprint was calculated by counting the number of distinct cities in which each business appeared across any platform's recommendations over the course of the full study.

The complete underlying dataset, including the full prompt set, per market results, and platform response records, is available on request at support@visibilityindex.ai.

About VisibilityIndex.ai

VisibilityIndex.ai measures how local businesses are represented across AI powered search and recommendation platforms. The Local AI Visibility Index is a recurring research program tracking AI recommendation consistency across US markets and industries over time.

For the full underlying dataset, additional market and industry level detail, or press inquiries, contact support@visibilityindex.ai.

© 2026 VisibilityIndex.ai. This report may be cited with attribution. Full methodology detailed in the appendix; complete underlying dataset available on request.