RESEARCH PAPER · JUNE 2026
PRIMARY-SOURCE VERIFIED · INDEPENDENT ANALYSIS

CITATION ADVISORY GROUP · AI VISIBILITY INTELLIGENCE

Why Law Firms Are Invisible to AI — and What the Data Shows

A primary-source analysis of AI citation infrastructure in the legal market, including verified findings on schema gaps, platform fragmentation, citation instability, and the growing threat of AI hallucination, slop, and bias.

Rob Poncin
Founder, Citation Advisory Group LLC

16+ years in legal market intelligence
Former national sales leadership at major legal recognition and directory platforms

EXECUTIVE SUMMARY

Five Findings. One Conclusion.

> Generative AI has become a referral channel. Most law firms do not know where they stand on it. This paper presents primary-source findings from verified AI citation audits across the legal market — and explains why what firms assume about their AI visibility is frequently wrong.

FINDING ONE — THE SCHEMA GAP

Major legal recognition platforms carry some structured data on attorney profile pages, but not the type that matters most. Tested across seven platforms, none implement Attorney or Person schema — the markup that identifies a named, credentialed individual to AI systems rather than simply marking that a business exists.

FINDING TWO — DIRECTORY SPEND

Directory spend and AI citation are not correlated. Firms with comprehensive legal directory packages routinely fail to appear in generative AI recommendations for practice area queries. Directory presence is a different signal from generative citation, and conflating them is one of the most common and costly mistakes in legal marketing today.

FINDING THREE — FRAGMENTATION

No single competitor dominates AI citation across all platforms in any market audited. Different platforms surface different firms. This fragmentation represents a real opportunity — but it is closing as infrastructure-forward firms consolidate their positions. The window is measurable in quarters, not years.

FINDING FOUR — INSTABILITY

AI citations are unstable without structural infrastructure. In one verified case, a firm dropped from citation on five AI platforms to two in a single quarter — with no changes made by the firm or its vendors. Checking once provides a snapshot, not a position. Ongoing monitoring is the only methodology that tracks change over time.

FINDING FIVE — THE RELIABILITY PROBLEM

AI platforms hallucinate, produce slop, and reflect systemic bias — all of which directly affect law firms. A firm may be cited with incorrect practice areas, wrong verdicts, or inaccurate locations. A competitor may be elevated above more qualified firms based on training data bias. And firms that check their AI visibility while authenticated to a platform see a fundamentally different result than prospective clients who search incognito. The gap between what a firm thinks AI says about it and what AI actually says can be significant — and costly.

SECTION ONE

The New Referral Layer

> Legal clients have always used the internet to find attorneys. What has changed is that they now ask questions and receive recommendations — not just search results. This is a structural shift, not an incremental one.

For two decades, law firm visibility online was primarily a search engine optimization problem. Rank well for "personal injury attorney Los Angeles" and the phone rings. The mechanics were understood, the vendors were established, and the market had largely figured out how to compete.

Generative AI has introduced a different dynamic. When a prospective client opens ChatGPT and asks "who are the best employment attorneys in Chicago," they receive a recommendation — a synthesized response that names specific firms, describes their strengths, and sometimes explains why they are worth contacting. The client may never click through to a traditional search result at all.

ChatGPT
Conversational recommendations. Draws from training data and live web search. Strong correlation with structured data and entity authority.

Google Gemini
Integrated with Google's index. Schema-responsive. Favors firms with established Google entity presence and Knowledge Panel verification.

Perplexity AI
Web-crawl dependent. Surfaces firms with strong content depth and editorial authority. Frequently cites directory sources as partial mentions only.

Google AI Overviews
Appears in Google SERP. Highly schema-dependent. Firms without LegalService or Attorney structured data are structurally excluded from this surface.

Microsoft Copilot
Draws from Bing index. Responds to Bing Places optimization and editorial backlink authority. Shares underlying infrastructure with ChatGPT.

Voice / Siri
Routes through Apple Maps and web signals. NAP consistency and Apple Business Connect optimization are the primary citation drivers here.

The platforms handling these queries are not niche tools. ChatGPT has over 200 million weekly active users. Google AI Overviews appears in search results for hundreds of millions of queries per day. These are the interfaces through which a substantial and growing portion of legal consumer research now flows.

> *"A firm that is absent from AI recommendations does not receive a notification that this is happening. It simply does not get mentioned."*

FINDING ONE

The Attorney Schema Gap at Major Legal Recognition Platforms

> Law firms depend on major legal recognition platforms to establish credibility. What most do not know is that none of these platforms implement the specific schema type that tells an AI system a named individual is a credentialed legal professional — leaving even well-recognized attorneys structurally invisible as people, not just as listings.

Structured data — specifically JSON-LD schema markup — is one of the primary mechanisms by which AI platforms identify, categorize, and cite entities on the web. When a page contains Organization or Local Business schema, it signals that a business exists at a location. When Attorney or Person schema is present, it tells these systems that a specific, named individual is a credentialed professional. That second signal is the one generative AI platforms rely on most directly when recommending a specific lawyer rather than a generic firm. Its absence is one of the most reliable predictors of AI invisibility for an individual attorney.

Audits conducted in July 2026 using Google's Rich Results Test on individual attorney profile pages across seven major national peer recognition platforms found valid schema item counts ranging from zero to over one hundred — often inflated by generic markup unrelated to the attorney profiled. One result held across every platform and page type: not a single Attorney or Person schema item was found.

✓ Verified July 2026 via Google Rich Results Test (search.google.com/test/rich-results) on individual attorney profile pages across seven major national recognition platforms. Results are reproducible and publicly verifiable by any analyst with access to the tool.

SCHEMA AUDIT — INDIVIDUAL ATTORNEY PROFILE PAGES, SEVEN PLATFORMS

CATEGORY	PLATFORM	VALID SCHEMA ITEMS	ATTORNEY/PERSON SCHEMA
Firm-Based Directories	Platform A	4	**No**
	Platform B	5	**No**
	Platform C	5	**No**
Attorney-Based Directories	Platform D	1	**No**
	Platform E	5	**No**
	Platform F	33	**No**

Firm-based platforms center pages on the practice; attorney-based platforms exist specifically to spotlight a named professional. Yet even on attorney-based platforms, no Attorney or Person schema was found — the category built to highlight individuals tells AI systems nothing about them as individuals.

The implication is significant. A law firm that has invested heavily in building and maintaining recognition platform profiles — that has earned peer recognition, obtained tier rankings, and paid for enhanced placement — receives essentially no AI citation benefit from those investments. Recognition is not the same as AI visibility. The two do not reliably correlate. This is the most consistently surprising finding for managing partners and marketing directors who encounter it for the first time.

FINDING TWO

Directory Spend Does Not Drive AI Citation

> The assumption that more directory presence equals more AI visibility is one of the most widespread and costly misconceptions in legal marketing. The data does not support it.

The major legal directory networks represent a multi-billion-dollar segment of the legal marketing industry. Firms routinely spend $15,000 to $50,000 or more annually on directory placement, enhanced profiles, and premium positioning. The pitch has always been presence: be where clients are looking.

There is a version of that argument that holds up. Directory profiles do provide NAP corroboration — Name, Address, Phone data indexed across multiple high-authority domains that helps AI platforms confirm a firm exists and where it is located. Major legal directories carry genuine domain authority, and links from those domains contribute modestly to a firm's backlink profile. These are real, if indirect, benefits.

What directory spend does not do is produce generative AI citations. When ChatGPT recommends a personal injury attorney in Detroit, it is not pulling from a directory listing. When Google AI Overviews surfaces an employment attorney in Los Angeles, it is not responding to directory review counts. The signals that drive generative citation are fundamentally different: JSON-LD schema on the firm's own website, entity verification through Wikipedia and Google Knowledge Panel, editorial backlink authority from non-directory sources, and deep practice area content.

KEY DISTINCTION

NAP corroboration — what directories provide — helps AI platforms confirm a firm exists. Generative citation — what firms want — requires AI platforms to actively recommend a firm. These are different outcomes driven by entirely different signals.

$0

GENERATIVE AI CITATION VALUE OF A COMPREHENSIVE DIRECTORY PACKAGE

Verified audits show that firms with comprehensive legal directory packages — including premium placement across major attorney directory platforms — routinely appear in zero generative AI recommendations for practice area queries. Directory presence at any spend level does not produce the structured data, entity authority, or content depth that generative AI platforms require to cite a firm by name.

Most legal marketing conversations conflate these two outcomes. When a marketing director reports that "we're on all the major directories," they are describing one type of digital presence. When a managing partner asks "does AI recommend us to prospective clients," they are asking about a completely different thing. The directory investment addresses the first question. It does not answer the second.

FINDING THREE

Platform Fragmentation — Every Market Is Different, and the Window Is Open

> In every market audited, no single competitor has locked up AI citation across all platforms. Different platforms surface different firms. This fragmentation represents a real opportunity — but it is closing, and it closes differently in every market.

In the Los Angeles sexual harassment attorney market, verified audits found five different platforms surfacing five different groups of firms. ChatGPT cited one set. Google Gemini cited a different set. Perplexity surfaced different firms still. Google AI Overviews, Microsoft Copilot, and Voice/Siri each returned their own results with minimal overlap. No single firm achieved consistent citation across all six platforms.

The Boston personal injury market tells a different story — and the contrast is instructive. In that market, one firm has achieved cross-platform dominance that no competitor in the Los Angeles market has yet built, appearing consistently across all six platforms for practice area queries. The difference is not market size or firm reputation. It is infrastructure. That firm built deep schema implementation, verified entity presence, and editorial authority before its competitors did. The window in that market has largely closed around a single incumbent.

0

FIRMS DOMINATING AI CITATIONS ACROSS ALL 6 PLATFORMS — IN MOST MARKETS

In market after market, no single firm achieves consistent citation across all six primary AI platforms for practice area queries. Where one firm does dominate, the infrastructure explanation is consistent: deep schema, verified entity presence, editorial authority. The space remains contestable in most markets — but the window is narrowing every quarter.

The practical implication is that every market must be evaluated on its own terms. A firm cannot assume that because a competitor appears in ChatGPT results, that competitor is also dominating Perplexity or Copilot. Nor can a firm assume its strong position on one platform reflects its standing across all six. What is consistent across every market audited is the infrastructure explanation for dominance. The market varies. The infrastructure requirements do not.

> *"The firms that move first are not just getting citations now. They are building structural foundations that make future citation more likely and competitive disruption more difficult."*

FINDING FOUR

AI Citations Are Unstable — The Monitoring Imperative

> A firm that checks its AI citation status once and finds favorable results has a data point, not a position. AI platforms update continuously, and citation status changes without warning and without notice to the firms affected.

In a verified case documented in the course of primary research, a law firm was cited on five of six major AI platforms for its primary practice area queries in one quarter. In the following quarter — with no changes made by the firm, its vendors, or its marketing team — that citation had dropped to two platforms. The firm did not know this had happened. There was no notification, no alert, no change in its own digital presence that would have suggested a shift. From the firm's perspective, nothing had changed. From the AI platforms' perspective, something had.

5→2

VERIFIED CITATION DROP — ONE QUARTER, NO FIRM ACTION

A firm went from being cited on 5 of 6 AI platforms to being cited on 2 — in a single quarter, with no changes to its website, directory presence, or marketing activities. AI platforms updated their signals in ways that affected this firm's citation status without any notification to the firm. This finding illustrates the central argument for ongoing, primary-source monitoring: citation status is not a permanent condition.

AI platforms are not static. ChatGPT regularly updates its training data and the weighting it applies to authority signals. Google AI Overviews adjusts in response to algorithm changes that occur hundreds of times per year. Perplexity's web crawl assigns freshness and authority scores that shift as the broader web changes. A firm cited because it had the most schema-rich pages in its market may find itself displaced when a competitor adds more comprehensive schema. A firm cited because of an editorial mention may see that signal decay as the content ages.

◉ A single AI citation audit answers "where do we stand today?" Quarterly audits using a consistent methodology answer "is our position improving, holding, or eroding — and why?" The difference between a data point and intelligence is the pattern over time.

There is an additional risk that compounds this instability: newly gained citations are often fragile. When a firm adds schema and begins appearing on a platform where it was previously absent, that citation is contingent on the continued presence and completeness of the schema that triggered it. If schema is removed, modified, or broken by a website update, the citation typically disappears — sometimes immediately. The firms with the most stable citation histories are those whose infrastructure is most complete, most redundant, and most regularly maintained.

FINDING FIVE

The Authenticated vs. Incognito Gap

> The most common method law firms use to check their AI visibility produces results that are meaningfully different from what prospective clients actually see. In documented cases, the difference is between appearing and not appearing at all.

When a firm's marketing director wants to know whether the firm appears in AI recommendations, the natural instinct is to open ChatGPT and type the relevant query. The problem is that they are not doing what a prospective client does. They are logged in. Their accounts carry months or years of browsing history, preference signals, and prior search behavior. The AI platform has been personalized based on that history.

A prospective client searching for an attorney is almost always doing so in a neutral context — a fresh browser session, incognito mode, or a device with no prior legal search history. What that client sees is the unfiltered, unpersonalized response of the AI platform to the query alone. These two results can diverge significantly.

8→1

DOCUMENTED AUTHENTICATED VS. INCOGNITO GAP

A major marketing technology platform claimed its product produced citations on 8 out of 10 ChatGPT queries for a test brand. Independent incognito testing of the same queries produced citations on 1 out of 10. The discrepancy — 8 versus 1 — illustrates the scale of the authenticated session problem. Firms that rely on authenticated testing to evaluate their AI visibility may be measuring something entirely different from what prospective clients experience.

This finding has direct implications for how law firms assess the claims of vendors who sell AI optimization products. Any tool that produces AI citation metrics from authenticated sessions is producing data that reflects a personalized experience, not the experience of a prospective client. The only methodology that answers the question law firms are actually asking — "does AI recommend us to someone who doesn't already know us?" — is incognito, primary-source testing with no session history.

The authenticated session problem also affects competitive research. A firm that believes a competitor is not appearing in AI recommendations, based on authenticated testing, may be wrong. The competitor may be appearing consistently in the incognito sessions that prospective clients actually use. Accurate intelligence requires incognito, primary-source, date-stamped methodology — consistently applied to self-assessment and competitive research alike.

SECTION SEVEN

Slop, Bias, and Hallucination — The Reliability Problem

> AI visibility is not just a question of whether a firm appears. It is also a question of what AI says when it does appear. The reliability problem in AI is real, consequential, and largely unaddressed in the legal market.

A

Slop — Confident, Fluent, and Wrong

AI-generated content that is factually incorrect but grammatically fluent and stylistically confident is colloquially referred to as slop. For law firms, slop typically manifests as incorrect verdict amounts, wrong practice area descriptions, inaccurate attorney credentials, or outdated firm information presented as current fact. A firm that recovered $8.5 million may be described as having recovered $3.2 million. An attorney who practices employment law may be described as a personal injury specialist. The AI platform presents these errors with the same authority it uses to present accurate information — and prospective clients have no indication that any of it may be wrong.

B

Bias — Systemic Underrepresentation of Regional and Boutique Firms

AI platforms are trained on data that reflects the existing distribution of content on the web — and that distribution is not neutral. Large firms in major metropolitan markets with heavy media coverage are overrepresented in training data. Regional firms, boutique practices, and newer firms are systematically underrepresented regardless of their quality, results, or client satisfaction. This creates a self-reinforcing cycle: visibility begets visibility, and absence begets absence. A small firm that has achieved extraordinary results but has limited media coverage may be structurally invisible to AI platforms simply because it has not been written about at scale.

C

Hallucination — The Fabrication Risk

AI hallucination — the generation of plausible but fabricated information — is a documented, persistent characteristic of large language models. For law firms, hallucination can take several forms: a firm credited with verdicts it did not achieve, an attorney described as having credentials they do not hold, entirely fabricated cases attributed to real attorneys, or a firm described as practicing in markets where it has no presence. This is a two-sided reputation risk: a firm may be falsely elevated by AI content that sets expectations it cannot meet, or falsely diminished by content that cost it clients based on information that is simply untrue.

The defense against all three failure modes is the same: robust entity infrastructure. Firms that establish verified Wikipedia articles, confirmed Google Knowledge Panels, deep structured data, and consistent information across all indexed sources give AI platforms the authoritative signals needed to generate accurate content — and less room to fabricate or misrepresent.

SECTION EIGHT

What the Data Suggests Law Firms Should Understand

> Across every market and practice area audited, the firms with the strongest AI citation profiles share a consistent set of infrastructure characteristics. These are structural attributes, not vendor-specific solutions.

1

Schema on the Firm's Own Domain

LegalService, Attorney, Organization, and FAQ schema implemented in JSON-LD format on the firm's primary website — not on directory profiles or third-party platforms. The schema must be on the firm's own domain to provide the entity signal AI platforms require. Firms with 20 or more valid schema items consistently outperform firms with fewer items, across all platforms tested.

2

Verified Entity Presence

A verified Google Knowledge Panel, a Wikipedia article for the lead attorney or firm where notability criteria are met, and a Wikidata Q-item for the entity. These establish what AI platforms call entity legitimacy — confirmation that the firm is a real, documented, credentialed organization in verifiable third-party sources. Entity verification is the single factor most correlated with ChatGPT citation in the legal market.

3

Content Depth That AI Can Draw On

Practice area pages with sufficient depth (1,000+ words with structured FAQ content), attorney bios with specific credentials and documented achievements, and case result pages with concrete outcomes. AI platforms generate recommendations by synthesizing available content — firms that provide more accurate, specific, and well-structured content give AI platforms more to work with and less room to fabricate.

4

Editorial Authority Beyond Directories

Backlinks and mentions from bar association publications, legal media, regional business press, and educational institutions. These editorial sources carry a different signal than directory links — they suggest the firm is recognized by credible third-party voices, not just listed in a paid directory. AI platforms weight editorial mentions more heavily than directory listings.

5

NAP Consistency Across All Indexed Sources

Exact consistency in firm name, address, and phone number across the firm's website, Google Business Profile, Apple Maps, Bing Places, and all directory profiles. AI platforms use NAP consistency as an entity verification signal. Inconsistencies — old addresses, variant firm name spellings, phone number discrepancies — weaken the entity signal and can suppress citations, particularly on Voice/Siri.

SECTION NINE

The Case for Verified, Independent Intelligence

Law firms face three methodological problems when assessing their own AI visibility. The authenticated session problem produces personalized results that do not reflect what prospective clients see. The self-assessment problem means a firm's own team has an inherent interest in finding favorable results. The vendor-reporting problem means vendors selling AI optimization products have an obvious incentive to report metrics that support the value of their service.

3

REQUIREMENTS FOR RELIABLE AI CITATION INTELLIGENCE

Primary source — the analyst queries the platform directly, not through a tool or proxy. Incognito — no session history, no personalization. Date-stamped — AI platform results change over time; a finding without a date cannot be compared to future findings. Any methodology missing any of these three elements produces data that cannot be trusted as a reliable representation of what prospective clients see.

A single audit answers "where do we stand today?" Quarterly audits using a consistent methodology answer "is our position improving, holding, or eroding — and why?" The competitive dimension of AI citation intelligence is often more valuable than self-assessment — because it tells a firm not just where it stands, but how fast the landscape around it is moving.

CLOSING

What This Means for Law Firms

Generative AI is not coming to the legal referral market. It is already there. The firms that appear in AI results — consistently, accurately, across multiple platforms — are already receiving referrals that competing firms are not. The firms that are absent are already paying a cost they cannot easily measure because they do not have the methodology to see it.

The good news is that the AI citation landscape in most legal markets remains highly contestable. The infrastructure requirements for meaningful AI citation are understood and achievable. The window to build a durable competitive position is open. The data suggests it is narrowing.

ABOUT THE AUTHOR

Rob Poncin

Rob Poncin has spent 16 years in legal market intelligence, including national sales leadership roles at major legal recognition and directory platforms. He founded Citation Advisory Group to bring primary-source, verified AI citation intelligence to elite law firms — an independent practice not affiliated with any directory, vendor, agency, or AI platform.

ABOUT CITATION ADVISORY GROUP

The Lighthouse. Not the Ship.

Citation Advisory Group LLC is an independent advisory practice providing AI visibility intelligence, competitive market analysis, and digital authority assessment for elite law firms. CAG does not build websites, manage directories, or execute marketing campaigns. We diagnose, verify, and advise. Implementation is at the firm's direction. citationadvisorygroup.com