EVIDENCE REVIEW

Why AI Visibility Is Not Just SEO

A Multi-Paper Analysis of the Evidence Behind the V.O.I.C.E. Methodology

Compiled by Dan Cartwright

Founder, ScopeSite Digital Studios

April 2026

Contents

1	**Foreword**
2	**The Core Claim**
3	**Evidence Layer 1** — Structured Data Directly Influences AI Visibility
4	**Evidence Layer 2** — Generative Engines Require Different Optimisation
5	**Evidence Layer 3** — Structured Content Frameworks Increase AI Citations
6	**Evidence Layer 4** — AI Engines Have Measurable Content Preferences
7	**Evidence Layer 5** — Entity Presence and Trust Signals Determine Recommendations
8	**Evidence Layer 6** — Real-World Proof (H4TLT Case Study)
9	**How V.O.I.C.E. Maps to the Evidence**
10	**References**

Foreword

A claim keeps showing up in the comments under posts about AI visibility. It runs something like this: if your site ranks on Google, you are already visible to ChatGPT, Claude, Gemini and Perplexity, so anyone selling AI visibility is just repackaging old SEO and charging a premium for it.

This paper was put together to answer that claim, not with opinion, but with published research. Every finding below comes from a peer-reviewed journal, a top-tier computer science conference, or a controlled empirical study. The sources are listed in full at the end. Anyone who wants to check the work can do so.

> ***AI platforms do not work like Google. They use different crawlers, different ranking signals, different content processing logic, and different citation behaviour. Optimising for Google does not automatically optimise for AI.***

That is the core claim. The six evidence layers that follow each test a specific part of it. At the end, the V.O.I.C.E. methodology is mapped against the findings, so you can see exactly which piece of research validates which part of the work.

If you read this and still want to argue that AI visibility is repackaged SEO, fine. The evidence has been laid out, and the door is open. But it is an argument with the researchers, not with ScopeSite.

The Core Claim

Put plainly:

> ***AI platforms do not work like Google. They use different crawlers, different ranking signals, different content processing mechanisms, and different citation logic. Optimising for Google does not automatically optimise for AI. The two require distinct, complementary strategies.***

The six evidence layers that follow each stand on their own. Taken together, they make this claim very difficult to argue with.

Evidence Layer 1

Structured Data Directly Influences AI Visibility

RESEARCH SOURCE

Schanbacher, P. (2026)

The Impact of JSON-LD Metadata on ChatGPT Visibility

Journal of Advance Research in Business, Management and Accounting, 12(1), 16-23
DOI: 10.61841/xt3he524

This is the largest empirical study of its kind to date. Schanbacher analysed 1,508 real estate agencies in Germany to work out which website attributes correlated with a business being known or referenced by ChatGPT.

Key Findings

FAQPage schema was the strongest single predictor of ChatGPT visibility. Businesses with FAQPage JSON-LD were 13 times more likely to be visible to ChatGPT than those without (odds ratio 13, $p < 0.001$). Only 6.2% of ChatGPT-visible agencies had FAQ schema, against 0.8% of non-visible ones, but the predictive power was overwhelming.

Product schema quadrupled the odds. Agencies with Product schema were 4 times more likely to appear in ChatGPT responses (OR 4, $p < 0.001$). 17.2% of visible agencies had it, against 1.8% of non-visible ones.

Technical foundations mattered, but were not enough on their own. Mobile-optimised sites were 5.2 times more likely to be visible. A robots.txt file increased odds by 3.4 times. Proper heading structure (H2 and H3) increased odds by 2.3 to 3.3 times. Faster page load correlated with visibility (418ms average against 623ms).

75% of web pages still lack any Schema.org metadata. This figure comes from Iliadis et al. (2025). It means the competitive opportunity remains wide open for any firm that implements schema properly.

Supporting Dataset

A dataset of 500 UK personal injury law firm websites has been compiled (SchemaMarkupAdoptionDataset500PersonalInjuryLawFirmWebsites2026.csv) with 32 variables

tracking schema adoption, type coverage, property completeness, entity disambiguation, hierarchical nesting, and validation status. The numbers reinforce Schanbacher: schema adoption across professional services remains extremely low, which creates a measurable first-mover advantage for firms that invest in it properly.

What This Means for V.O.I.C.E.

V.O.I.C.E. mandates JSON-LD schema implementation on every page: Organisation, Service, FAQPage, BreadcrumbList, and SpeakableSpecification. This is not a nice-to-have. Schanbacher's data shows it is the single most influential factor in determining whether an AI assistant knows your business exists.

Evidence Layer 2

Generative Engines Require Different Optimisation

RESEARCH SOURCE

Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K. and Deshpande, A. (2024)

GEO: Generative Engine Optimization

Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24), Barcelona

DOI: 10.1145/3637528.3671900

This is the foundational academic paper on GEO, published at one of the top computer science conferences in the world. The Princeton team built GEO-bench (10,000 queries across 25 domains) and tested 9 different optimisation methods.

Key Findings

Traditional SEO methods actively harmed visibility in generative engines. Keyword stuffing, the bread and butter of old-school SEO, performed 10% worse than the unoptimised baseline on Perplexity.ai. That is not diminishing returns, that is active damage.

GEO methods boosted visibility by up to 40%. Adding citations from relevant sources, incorporating statistics, and including quotations from authoritative sources each improved visibility by 30-40% on the Position-Adjusted Word Count metric.

Combining GEO strategies amplified the results. The best-performing combination (Fluency Optimisation plus Statistics Addition) outperformed any single strategy by more than 5.5%.

Lower-ranked websites benefited most from GEO. Traditional search engines rely heavily on backlinks and domain authority, which small businesses struggle to build. Generative engines prioritise content quality and structure, which levels the playing field.

The methods generalised across platforms. Testing on Perplexity.ai, a commercially deployed generative engine with millions of users, confirmed real-world visibility improvements of up to 37%.

What This Means for V.O.I.C.E.

V.O.I.C.E. does not rely on keyword density or backlink volume as primary levers. It engineers content for AI extraction: structured FAQ formats, statistical evidence, cited sources, and authoritative tone. The Princeton research validates that this approach produces measurably better outcomes than traditional SEO tactics when targeting generative AI platforms.

Evidence Layer 3

Structured Content Frameworks Measurably Increase AI Citations

RESEARCH SOURCE

A Strategic Outlook on LLM SEO (2025-2026)

Quasi-experimental study across ChatGPT, Perplexity, and Google AI Overviews

100 prompts tested across 5 geographic locations (India, USA, UK, Australia, Germany)

This study introduced the LLM SEO Files framework: six structured file formats (llms.txt, YAML, Markdown, Schema JSON-LD, evidence-of-facts.json, Open Graph JSON) designed specifically for machine parsing by AI platforms.

Key Findings

> **Structured optimisation produced a 64% increase in brand mentions and domain citations.** Non-optimised content averaged AIM Visibility Scores of 0.8-1.8. Optimised content using the structured framework scored 3.3-5.0 consistently.

> **Results were consistent across every major AI platform.** ChatGPT brand mentions increased from 100 to 154. Perplexity mentions rose from 69 to 112. Google AI Overviews showed the largest proportional gain, from 45 to 102 mentions.

> **High-intent queries saw the biggest gains.** Purchase-stage queries (bottom of funnel, where revenue actually happens) showed the most significant improvements. Brand mentions went from 13 to 55, citations from 8 to 39. The framework works hardest where it matters most commercially.

> **The awareness stage also improved dramatically.** Brand mentions jumped from 20 to 59, and citations from 10 to 56, confirming that the framework operates across the entire marketing funnel.

What This Means for V.O.I.C.E.

V.O.I.C.E. incorporates multiple layers of machine-readable content: JSON-LD schema, llms.txt files, ai-context.json, speakable schema, and FAQ-first content architecture. The research demonstrates that this multi-layer approach produces consistent, measurable improvements across every major AI platform and across every stage of the buyer journey.

Evidence Layer 4

AI Engines Have Measurable Content Preferences

RESEARCH SOURCE

Ma, L., Qin, J., Xu, X. and Tan, Y. (2025)

The Style and Semantic Effects of Generative Search Engine

Submitted to Information Systems Research

arXiv: 2509.14436

RESEARCH SOURCE

Chen, M., Wang, X., Chen, K. and Koudas, N. (2025)

Generative Engine Optimization: How to Dominate AI Search

University of Toronto

arXiv: 2509.08919

The first paper analysed roughly 10,000 websites collected from Google's conventional and generative search platforms, looking at how AI Overviews select and cite content. The second ran extensive experiments across ChatGPT, Claude, Gemini and Perplexity, covering brand visibility, source selection, and media-type preferences across multiple verticals, languages, and query intents.

Key Findings

> **Generative engines prefer content with lower perplexity.** In plain English: AI platforms cite content that is clearly structured, semantically consistent, and easy for language models to parse. Content cluttered with marketing jargon, complex sentence structures, or inconsistent terminology gets deprioritised.

> **Semantic similarity between cited sources matters.** AI Overviews tend to cite sources that are semantically consistent with each other. Your content needs to align with how authoritative sources in your sector describe the same concepts.

> **AI engines overwhelmingly favour earned media over brand-owned content.** Claude and ChatGPT are described as "extremely earned-heavy." Third-party validation (reviews, directory listings, industry mentions) carries more weight than content you publish on your own website.

> **Schema markup is described as non-negotiable.** The research states directly: "The primary action is to treat your website as an API for AI. This requires rigorous implementation of technical SEO fundamentals combined with detailed schema markup (Schema.org) for all entities."

> **Different AI engines have different source preferences.** Google maintains a balanced mix of brand, social, and earned content. ChatGPT suppresses social content and emphasises earned media. Perplexity takes a blended approach. Optimisation has to account for platform-specific behaviour, not treat all AI as one monolithic system.

> **Language localisation has a larger impact than content optimisation alone.** For businesses targeting multiple markets, local-language authority building is critical, particularly for ChatGPT and GPT-based platforms which heavily localise their source selection.

What This Means for V.O.I.C.E.

V.O.I.C.E. tests across four AI platforms (ChatGPT, Claude, Gemini, Perplexity) specifically because each has different citation behaviour. The methodology includes entity building through third-party directories and platforms (Wikidata, Yell.com, TechBehemoths, industry-specific directories) because the research confirms earned media signals are weighted more heavily than brand-owned content.

Evidence Layer 5

Entity Presence and Trust Signals Determine Recommendations

RESEARCH SOURCE

de Rosen, T. (2025)

Methodology for AI Visibility Optimization (AIVO Standard) Version 3.0

Available at SSRN

SSRN: 5431877

The AIVO Standard is the first comprehensive framework designed to help products, services, and brands become discoverable and recommendable by LLMs. It defines a 9-stage methodology with certification criteria.

Key Findings

AI assistants rely on entirely different signals than search engines. The framework identifies the key signals: natural language prompt-response training, semantic content mapping, trusted ecosystem metadata, and public data from schema.org, Wikidata, GitHub, and indexed publishing channels.

Foundational presence is a prerequisite, not a bonus. Entities without structured data in knowledge graphs (Wikidata), semantic web elements (Schema.org JSON-LD), and public repositories "risk becoming invisible or misrepresented in AI recommendations."

Trust signals serve as credibility amplifiers. Without them, AI responses may omit a business in favour of better-cited alternatives. The framework requires at least 4 distinct trust signal types for compliance: Organisation schema, verified directory listings, third-party endorsements, and cross-linking.

Consistency across platforms is critical. Name, Address, Phone (NAP) consistency, and matching descriptions across Wikidata, Schema.org, LinkedIn, Crunchbase, and directory listings all feed into whether an AI system treats a business as a trustworthy, citable entity.

Regulated industries require enhanced implementation. Solicitors, accountants, healthcare providers, and financial services firms need additional trust signals and compliance markers to be reliably recommended by AI assistants.

What This Means for V.O.I.C.E.

V.O.I.C.E. scoring includes entity disambiguation, sameAs signals, and cross-platform consistency. Every ScopeSite client build includes Wikidata entity creation, directory submissions with consistent NAP data, and schema markup that explicitly defines what the business is, what it does, and where it operates. The AIVO Standard validates that this entity-building work is not supplementary, it is foundational.

Evidence Layer 6

Real-World Proof — The H4TLT Case Study

The Client

Hear 4 The Long Term (H4TLT) is a workplace hearing health surveillance business founded by Mark Ashmore, an HCPC-registered audiologist based in Somerset. H4TLT provides self-service hearing tests and earplug fit testing as an alternative to traditional mobile audiometry vans.

The Starting Point

When ScopeSite began working with H4TLT in late 2025, the business had minimal online presence. It was not appearing in any AI assistant recommendations for workplace hearing tests, HSE hearing surveillance, or any related queries.

What Was Implemented

The V.O.I.C.E. methodology was applied systematically:

- **Server-side rendered website** built on Wix Studio, with clean HTML that AI crawlers can read without executing JavaScript.
- **Comprehensive JSON-LD schema** on every page: Organisation, Service, FAQPage, BreadcrumbList, with SpeakableSpecification on FAQ headings.
- **Entity building** through directory submissions, consistent NAP data, and structured sameAs references.
- **Content architecture** designed for AI extraction: FAQ-first structure, clear service descriptions, and regulatory citations linked to authoritative sources (HSE, legislation.gov.uk).
- **AI crawler configuration** ensuring the site was accessible to ChatGPT-User, ClaudeBot, PerplexityBot, and Googlebot.

The Results

> ***The H4TLT website went live on 1 January 2026. By April 2026, Google AI Overview was recommending H4TLT nationally as a leading HSE industrial hearing test provider.***

H4TLT now appears as the first listed provider, with pricing included, in AI Overview results for queries like "best HSE industrial hearing tests UK." It shows up alongside firms that have been established for decades, with substantially larger marketing budgets and higher domain authority.

This was achieved for a small, single-operator business in under 4 months. No paid advertising. No link-building campaigns. No domain authority manipulation. Just structured data, server-side rendering, entity building, and content architecture, the exact pillars of V.O.I.C.E.

How V.O.I.C.E. Maps to the Evidence

The V.O.I.C.E. methodology was not designed in a vacuum. Each component corresponds directly to findings from the research cited in this paper. The mapping is set out below.

V

Visibility — *AI Platform Testing*

Tests across ChatGPT, Claude, Gemini, and Perplexity because each platform has different citation behaviour.

Source: *Chen et al., 2025; Ma et al., 2025*

O

Optimisation — *Content Architecture for AI Extraction*

Content is structured for machine parsing: FAQ formats, statistical evidence, cited sources, clear entity descriptions. GEO methods outperform traditional SEO by 30-40%. Keyword stuffing actively harms AI visibility.

Source: *Aggarwal et al., 2024*

I

Intelligent — *Schema Engineering*

JSON-LD schema is the strongest single predictor of ChatGPT visibility. FAQPage schema provides a 13x increase in odds. V.O.I.C.E. mandates WebPage, Service, FAQPage, BreadcrumbList, and SpeakableSpecification on every page.

Source: *Schanbacher, 2026*

C

Crawler — *Server-Side Rendering and AI Crawler Access*

A significant share of AI crawlers cannot execute client-side JavaScript. If a website relies on JavaScript rendering (React SPAs, Angular, Vue without SSR), the AI crawler sees an empty page. V.O.I.C.E. requires SSR HTML and proper robots.txt configuration.

Source: *Chen et al., 2025*

E

Engines — *Entity Building and Trust Signals*

AI assistants rely on knowledge graphs, directory listings, and cross-platform consistency to determine entity trustworthiness. Earned media signals outweigh brand-owned content. V.O.I.C.E. includes Wikidata entity creation, directory submissions, sameAs signals, and NAP consistency.

Source: *de Rosen, 2025; Chen et al., 2025*

References

1	Schanbacher, P. (2026). *The Impact of JSON-LD Metadata on ChatGPT Visibility.* Journal of Advance Research in Business, Management and Accounting (ISSN: 2456-3544), 12(1), 16-23. https://doi.org/10.61841/xt3he524
2	Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K. and Deshpande, A. (2024). *GEO: Generative Engine Optimization.* Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24), Barcelona. https://doi.org/10.1145/3637528.3671900
3	A Strategic Outlook on LLM SEO (2025-2026). *Quasi-experimental study across ChatGPT, Perplexity, and Google AI Overviews.* 100 prompts tested from 5 geographic locations.
4	de Rosen, T. (2025). *Methodology for AI Visibility Optimization (AIVO Standard) Version 3.0.* Available at SSRN 5431877. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5431877
5	Ma, L., Qin, J., Xu, X. and Tan, Y. (2025). *When Content is Goliath and Algorithm is David: The Style and Semantic Effects of Generative Search Engine.* Submitted to Information Systems Research. https://arxiv.org/abs/2509.14436
6	Chen, M., Wang, X., Chen, K. and Koudas, N. (2025). *Generative Engine Optimization: How to Dominate AI Search.* University of Toronto. https://arxiv.org/abs/2509.08919
7	Testbed for Generative Engine Optimization in E-Commerce (2025). *E-GEO benchmark, 7,000+ Reddit product queries matched with Amazon listings.*
8	Iliadis, A., Acker, A., Stevens, W. and Kavakli, S. B. (2025). *One Schema to Rule Them All: How Schema.org Models the World of Search.* Journal of the Association for Information Science and Technology, 76(2), 460-523.

9	SchemaMarkupAdoptionDataset500PersonalInjuryLawFirmWebsites2026.csv. *Dataset of 500 UK personal injury law firm websites analysing schema markup adoption across 32 variables.* Internal ScopeSite dataset, 2026.

About ScopeSite Digital Studios

ScopeSite Digital Studios is a veteran-owned web design and AI visibility agency based in Beckington, Frome, Somerset. Founded by Dan Cartwright (British Army, REME), ScopeSite builds AI-visible websites for professional services firms using the V.O.I.C.E. methodology.

V.O.I.C.E. stands for Visibility, Optimisation, for Intelligent, Crawler, Engines.

Website	scopesite.co.uk
AI Scanner	voice.scopesite.co.uk
Company	ScopeSite Ltd — Company Number 16130355
Location	Beckington, Frome, Somerset BA11

This document is intended for distribution to prospective clients, industry professionals, and anyone who claims that AI visibility optimisation is "just SEO." The evidence says otherwise.