

GEO READINESS AUDIT AS A BRAND IMPERATIVE: RETHINKING DIGITAL VISIBILITY IN THE ERA OF AI- MEDIATED CONSUMER DISCOVERY

Liviu CRĂCIUN

ORCID: <https://orcid.org/0000-0002-7661-7002>

University of Craiova, Craiova, Romania

Email: liviu.craciun@edu.ucv.ro

Andreea-Maria TRAISTARU

National University of Science and Technology Politehnica Bucharest, Romania

Email: andreea.maria.traistaru@gmail.com

Pavel-Cristian CRĂCIUN

ORCID: <https://orcid.org/0009-0000-7077-7913>

Bucharest University of Economic Studies, Bucharest, Romania

Email: craciunpavel18@stud.ase.ro

DOI: 10.52846/MNMK.24.1.09

Abstract :

This study examines how brands remain visible and discoverable as consumers shift toward AI-powered search tools, with a focus on the strategic implications of Generative Engine Optimization (GEO) for brand management. The growing adoption of large language models (LLMs) such as ChatGPT, Gemini, and Perplexity as everyday information sources is fundamentally changing how consumers find and evaluate brands, in ways that traditional search engine optimization (SEO) frameworks fail to address. The study employed a conceptual research design, combining purposive literature synthesis with secondary data analysis of industry behavioural analytics encompassing over 1.96 million LLM-mediated consumer sessions. The analysis reveals that LLM citation is driven by a distinct set of content and authority signals that diverge substantially from classical SEO ranking factors, including content evidentiality, earned media breadth, and semantic entity consistency. Key findings indicate that LLM-referred visitors convert at rates up to nine times higher than organic search visitors, yet only 16% of brands systematically monitor their AI search presence. Based on the synthesised evidence, this study proposes a five-dimension GEO Readiness Audit framework and discusses implications for brand strategy, content governance, and marketing channel allocation.

Keywords: digital marketing, brand visibility, marketing, generative engine optimization, LLM, AI.

1. Introduction

The mechanisms through which consumers discover brands, products, and services have changed fundamentally in recent years (Swanson, 2026). For the past two decades, search engine optimization (SEO) has constituted the primary strategic

instrument through which brands secured visibility in the digital environment, with organic ranking on search engine results pages (SERPs) functioning as a widely accepted proxy for relevance, authority, and consumer reach (Erdmann, Arilla et al., 2022). Substantial academic and practitioner literature has been developed around the strategic management of SEO as a brand asset, examining its relationship with consumer attraction, brand equity, and digital marketing performance (Berman and Katona, 2013; Mladenovic et al., 2023).

The accelerating proliferation of large language model (LLM)-powered generative engines, including ChatGPT, Google Gemini, and Perplexity, presents a structural challenge to this established paradigm. These systems do not return ranked lists of hyperlinks; instead, they synthesize information drawn from multiple retrieved sources and deliver direct, conversational responses tailored to user queries (Aggarwal et al., 2024; Swanson, 2026). The commercial and strategic implications for brand managers are considerable: visibility is no longer a function of ranking position in a paginated results list, but of citation, recommendation, or named reference within a synthesized AI-generated answer. A brand absent from that answer is, in practical terms, invisible to the consumer conducting that query. The urgency of this challenge is reflected in emerging industry forecasts. Gartner (2024) projected that traditional search engine volume would decline by 25% by 2026 as generative AI solutions function increasingly as substitute answer engines. Behavioural data from large-scale session analysis indicates that LLM-mediated traffic exhibits conversion rates substantially higher than those associated with traditional organic search, compounding the strategic importance of AI visibility (Bell, 2025). Despite these signals, the academic marketing literature has not yet produced a structured, evidence-based approach to auditing brand readiness for the generative search environment. This study addresses that gap. The research examines how generative engines determine brand visibility, identifies the signal dimensions that differentiate GEO from conventional SEO, and proposes a structured GEO Readiness Audit framework applicable to brand management practice. The study is informed by the following research questions:

(RQ1) What mechanisms govern brand visibility in generative engine responses?

(RQ2) How do the determinants of LLM-based brand citation differ from those of traditional SEO?

(RQ3) What constitutes a viable GEO readiness audit framework for brand management?

The paper covers the following: a literature review on SEO and LLM-based search in Section 2, the methodology in Section 3, findings and discussion in Section 4, and conclusions in Section 5.

2. Literature Review

2.1 Search engine optimization as a brand visibility mechanism

Search engine optimization has been studied extensively in the marketing literature as both a technical discipline and a strategic brand management instrument. Berman and Katona (2013) provided a foundational formal treatment demonstrating that SEO shapes the competitive dynamics between content producers and advertisers, with higher-quality websites securing systematic organic positioning advantages through a self-reinforcing quality signal mechanism. Their analysis showed that although SEO reduces search engine advertising revenues by

displacing paid placements, it produces net welfare gains for consumers through improved content quality on organic results pages.

Subsequent research has examined SEO from a strategic resource allocation perspective. Erdmann et al (2022) analysed the long-term keyword investment strategies of leading fashion e-commerce firms, demonstrating convergence toward equilibrium keyword expenditure paths and the differential value of branded versus generic keyword positioning. Their findings established that organic keyword strategy constitutes a durable strategic asset, with brand-keyword convergence generating compounding visibility advantages over time. Mladenovic et al. (2023) extended this inquiry to a service sector context, identifying domain authority, content relevance, and technical site health as primary determinants of search visibility, and confirming that SEO performance exhibits meaningful variation across industry contexts.

The relationship between SEO performance and brand equity has also received scholarly attention. Keller (1993), in his foundational model of customer-based brand equity, identified brand awareness and brand associations as the primary determinants of brand strength in consumer memory. Digital marketing scholars have subsequently argued that search visibility functions as a structural driver of brand awareness in online environments, with organic ranking position influencing both the likelihood of initial brand encounter and the credibility attributions consumers form upon that encounter. High SERP rankings have been shown to convey implicit quality signals that influence consumer trust and purchase intention independently of the content of the ranked page (Mladenovic et al., 2023; Özkan et al., 2020; Salvi et al., 2021).

2.2 The emergence of large language models as search intermediaries

Large language models represent a qualitatively distinct class of information retrieval system that diverges from conventional search engines in its fundamental operational architecture. Aggarwal et al. (2024) formalised this distinction through the concept of the generative engine, defined as a system that retrieves relevant documents from an indexed knowledge base and employs a large neural language model to synthesize a response grounded in those retrieved sources. This two-stage retrieve-then-generate architecture produces outputs that are qualitatively different from the ranked link lists returned by traditional search engines: they are direct, synthesized, natural-language answers that reduce the need for users to evaluate multiple sources independently.

The competitive implications of this architecture for brand visibility were examined empirically by Aggarwal et al. (2024) through the introduction of GEO-bench, a large-scale benchmark comprising 10,000 diverse user queries across multiple domains. Their experimental analysis demonstrated that classical SEO signals, particularly keyword density, showed minimal influence on citation probability within generated responses, while content-level characteristics associated with epistemic authority, including the presence of statistics, source citations, and quotations from credible authorities, improved visibility by up to 40%. These findings established that the optimization logic governing LLM-based visibility is substantively distinct from that governing traditional SERP ranking.

A complementary strand of research has examined the vulnerability of LLM-based recommendation systems to strategic manipulation. Kumar and Lakkaraju (2024) demonstrated that the insertion of carefully optimized strategic text sequences into product information pages could systematically increase a product's

likelihood of being listed as the top LLM recommendation, even when the product's underlying quality did not merit such prominence. Their findings raise significant concerns regarding the integrity of AI-mediated brand recommendations as a consumer guidance mechanism, and introduce a parallel to the black-hat SEO tactics that have historically required regulatory and algorithmic countermeasures. Wen et al. (2025) further documented the emergence of commercial GEO service providers that offer systematic content manipulation for LLM visibility, positioning this as an emerging advertising and competitive intelligence surface with material implications for market fairness.

2.3 Consumer behaviour in AI-mediated discovery environments

Consumer adoption of LLM-based search interfaces has proceeded at a pace exceeding that of prior transformative consumer technologies. The generative AI user base grew substantially in 2024-2025, with ChatGPT accumulating an estimated 800 to 900 million weekly active users across platforms by 2025 (Moore and Moore, 2025). Consumer survey data indicates that 36% of American consumers had discovered a new product or brand through ChatGPT, a proportion rising to 47% among Gen Z respondents. These figures are consistent with a pattern of AI-mediated discovery that operates in parallel with, rather than as a complete substitute for, traditional search behaviour.

The quality of LLM-mediated traffic has been identified as a distinguishing characteristic with strategic relevance for brand managers. Session analytics from large-scale behavioural studies have indicated that LLM-referred visitors exhibit substantially higher conversion rates compared to organic search-referred visitors, with ChatGPT-originating sessions converting at rates approximately nine times higher than Google organic traffic (Bell, 2025). This quality differential positions LLM-mediated discovery as a high-value segment of the consumer journey, despite its currently smaller volume relative to traditional search.

The structural conditions that have produced this behavioural shift include the rise of zero-click search, where AI-generated summaries answer the user's query on the spot, removing any reason to click through to a website (Mukherjee, 2026). In 2025, more than six in ten Google searches ended before the user ever clicked a link (Jelenik, 2025), and Google's own AI Overviews were appearing in over 57% of search results pages. These conditions collectively reduce the traffic-generating value of traditional SERP ranking while elevating the strategic importance of citation within AI-generated responses. Gartner (2024) has argued that these trends represent not a transient technological shift but a structural reorientation of the information retrieval market, with projections indicating a 25% decline in traditional search engine volume by 2026.

3. Research Methodology

The study employed a conceptual research design, combining purposive literature synthesis with secondary data analysis of industry behavioural analytics. This design is appropriate given the nascent character of the research domain: GEO as a formalised concept was introduced only in 2024 (Aggarwal et al., 2024), and the empirical evidence base remains concentrated in technical computer science publications and industry-generated observational datasets rather than in established management journal literature. Conceptual papers of this type contribute to knowledge by synthesising emerging evidence streams into practitioner-oriented frameworks, an approach well established in the digital marketing and brand

management literature (Keller, 1993; Erdmann et al, 2022; Cenamor et al., 2019; Dotson et al., 2017).

Academic sources were selected purposively, focusing on peer-reviewed work covering SEO strategy, digital brand visibility, and LLM-based information retrieval. Two foundational contributions predating the LLM era were retained given their lasting theoretical relevance to the field: Keller's (1993) framework of customer-based brand equity and Berman and Katona's (2013) treatment of competitive dynamics in search engine markets. These were complemented by a set of recent empirical works, most notably the GEO framework developed by Aggarwal et al. (2024), the product visibility manipulation analysis by Kumar and Lakkaraju (2024), and the SEO strategy research of Erdmann et al (2022) and Mladenovic et al. (2023), were selected based on methodological rigour and direct relevance to the three research questions. Sources were excluded where empirical claims were unsupported by transparent methodology or where findings had been superseded by more recent primary work.

Secondary data analysis was conducted on industry behavioural analytics derived from three principal sources: the 2025 State of AI Discovery Report, based on 1.96 million LLM-mediated sessions observed between November 2024 and November 2025 across e-commerce, SaaS, legal, health, and publishing sectors (Bell, 2025); the TheDigitalBloom (2025) AI Visibility Report, synthesising citation pattern data from over 680 million LLM citations; and the Omnius (2025) comparative platform analysis covering citation behaviour across Google AI Overviews, ChatGPT, and Perplexity during the first quarter of 2025. These secondary datasets were selected on the basis of sample size, methodological transparency, and multi-platform coverage.

The analytical approach consisted of thematic synthesis across the selected literature and secondary data, organised around the three research questions guiding the study. Findings were triangulated across academic and industry sources where parallel evidence was available. The conceptual framework developed in the results section was derived inductively from the synthesised signal taxonomy and grounded in the experimental evidence of Aggarwal et al. (2024). As with all conceptual contributions, the framework's propositions remain subject to empirical validation; future research employing primary experimental, or survey designs would strengthen the causal claims inherent in the audit dimensions proposed.

4. Results and Discussion

4.1 Mechanisms governing brand visibility in generative engine responses (RQ1)

The analysis revealed that generative engine architecture imposes a fundamentally different competitive structure on brand visibility compared to traditional search engines. The retrieve-then-generate pipeline identified by Aggarwal et al. (2024) operates in two sequential stages: a retrieval stage, functionally analogous to conventional search, in which candidate documents are fetched from the indexed web, and a generation stage, in which the LLM synthesizes a response from the top-ranked retrieved documents. Due to context window constraints, generation typically draws on three to five retrieved documents, producing a qualitative compression of the visibility surface from the ten or more positions available on a traditional SERP to a dramatically smaller citation set.

The findings indicate that the retrieval stage exhibits meaningful overlap with traditional SEO signals: domain authority, page indexability, and keyword relevance continue to influence which pages enter the retrieved candidate set. However, the generation stage introduces a distinct additional filtering layer governed by content-level evidential characteristics. The analysis of the Aggarwal et al. (2024) experimental results revealed that content modifications emphasizing statistics, authoritative source citations, and quotations from credible references produced visibility gains of up to 40% in generated responses, while keyword-based modifications characteristic of conventional SEO showed negligible effect. This finding was corroborated by the TheDigitalBloom (2025) citation analysis, in which brand search volume exhibited the highest correlation with citation frequency at 0.334, surpassing the influence of backlink profiles.

The results further revealed that the generation stage introduces a winner-takes-most dynamic absent from traditional SERPs. Because generative responses typically present a single synthesized recommendation or a small named set of brands, the commercial value of a citation is concentrated in the first-mentioned or most prominently featured brand within the response. The binary nature of AI visibility, where a brand is either cited or absent, with no equivalent of the second-page SERP position, represents a structural intensification of visibility inequality that brand managers have not previously encountered in the digital environment.

4.2 Divergence between LLM citation determinants and traditional SEO signals (RQ2)

The analysis identified five principal dimensions along which LLM citation determinants diverge from traditional SEO ranking signals. The first dimension concerns content evidentiality. Traditional SEO rewards topical relevance, keyword density, and content comprehensiveness. LLM citation, by contrast, rewards epistemic authority: the presence of verifiable statistics, original data, explicit source attributions, and semantically self-sufficient paragraphs capable of functioning as standalone citable answers. The TheDigitalBloom (2025) data indicated that adding statistics increased AI visibility by 22% and adding authoritative quotations increased it by 37%, while keyword-centric modifications produced no measurable improvement.

The second dimension concerns earned media primacy. Traditional SEO treats backlinks as the primary off-page authority signal, with direct brand ownership of the linking domain irrelevant. The findings indicate that LLM citation is substantially driven by earned media breadth: coverage in high-authority third-party sources including news publications, Wikipedia, and industry review platforms. Taylor (2025) reported that up to 90% of citations driving LLM brand visibility can originate from earned rather than owned content. This represents a significant inversion of the owned-versus-earned resource allocation logic that has characterised digital marketing investment in the SEO era.

The third dimension concerns platform-specific source preferences. Traditional SEO optimization is largely platform-agnostic, as Google's dominant market share has meant that a single optimization strategy addresses the majority of search traffic. The analysis revealed that ChatGPT, Google AI Overviews, and Perplexity exhibit substantially different source citation patterns: ChatGPT draws approximately 48% of citations from Wikipedia and community platforms including Reddit, while Google AI Overviews favour structured professional content from LinkedIn and YouTube (Stojadinovic, 2025). The finding that only 11% of domains are cited by both

ChatGPT and Perplexity underscores the platform-specific character of LLM citation logic and the inadequacy of a single-platform optimization strategy.

The fourth dimension concerns technical extractability. Both traditional SEO and LLM optimization reward technically accessible, well-structured pages. However, the analysis identified page load time as a particularly salient variable in the LLM context: the TheDigitalBloom (2025) data showed that pages with first contentful paint times below 0.4 seconds received an average of 6.7 AI citations, compared to 2.1 for pages exceeding 1.13 seconds. This finding suggests that LLM retrieval systems weight page performance signals more heavily than traditional search crawlers, likely reflecting the real-time nature of live-web retrieval in conversational AI contexts.

The fifth dimension concerns entity recognition consistency. Traditional SEO does not require a brand to be described consistently across external sources; ranking is determined by the content of the optimized page itself. The analysis indicated that LLM citation is substantially influenced by the coherence of a brand's entity representation across the web: brands described inconsistently across platforms, as commonly occurs following rebranding, mergers, or international market entry, risk being omitted or inaccurately represented in generated responses. This finding introduces a new dimension of brand governance not previously recognized in the digital marketing literature.

4.3 A GEO Readiness Audit framework (RQ3)

Based on the synthesised evidence, this study proposes a five-dimension GEO Readiness Audit framework designed to enable brands to systematically assess their current AI visibility posture and identify prioritised remediation actions. The framework was developed inductively from the signal taxonomy identified in Sections 4.1 and 4.2, and is organized around the five divergent dimensions identified in the comparative analysis.

Dimension 1, AI Citation Baseline, constitutes the diagnostic entry point. The audit requires systematic prompt testing across the principal generative engines, using prompts representative of the brand's core consumer intent categories: category entry queries, brand-specific evaluation queries, and competitive comparison queries. For each prompt, the auditor records citation presence, response position, descriptive framing, and referenced sources. Commercial monitoring platforms including Semrush AI Visibility Toolkit, Profound, and Peec AI provide longitudinal tracking capabilities at scale. Brands absent from responses to high-intent queries in their categories have identified the primary expression of a GEO visibility gap.

Dimension 2, Content Evidentiality, assesses whether owned content satisfies the epistemic authority criteria that favour LLM citation. The audit evaluates: the proportion of content pages incorporating verifiable statistics or original data; the presence of direct quotations from authoritative external sources; the degree to which individual paragraphs are semantically self-sufficient; and whether explicit answers to predictable consumer queries appear near page headings rather than embedded in promotional prose. Pages structured as marketing copy rather than informational material were found to consistently underperform in LLM retrieval contexts (Aggarwal et al., 2024).

Dimension 3, Earned Media Coverage, examines third-party reference breadth and quality. The audit assesses Wikipedia presence, coverage density in sector-relevant authoritative publications, review platform activity across Trustpilot, G2,

Capterra, and Yelp, and brand mention frequency on Reddit and Quora. The findings establish that a brand with strong owned content but limited earned media breadth is likely to underperform in LLM citation despite high SERP rankings, as generative engine credibility assessment weights external corroboration substantially (Taylor, 2025).

Dimension 4, Technical Extractability, examines whether page performance and structural characteristics enable LLM retrieval and parsing. Key indicators include first contentful paint time, structured data markup coverage across product, organisation, and FAQ schema types, AI crawler accessibility, and heading hierarchy granularity. Technical extractability functions as a prerequisite layer: content that cannot be retrieved cannot be cited regardless of its evidential quality.

Dimension 5, Multi-Platform Presence Architecture, maps brand coverage against the documented source preferences of each major generative engine. It further assesses entity recognition consistency: whether the brand is described with coherent naming, category associations, and attribute framing across the full range of external sources likely to be retrieved by LLMs. Inconsistencies at the entity level were found to create systematic citation risk that content-level optimizations cannot remedy.

The audit produces a gap matrix correlating dimension-level deficiencies with citation performance data from Dimension 1, enabling prioritised remediation roadmap development. The cross-functional nature of the required interventions, spanning content strategy, public relations, technical marketing, and brand governance, represents a distinguishing characteristic of GEO management relative to conventional SEO, which has historically been concentrated within digital marketing or IT functions.

4.4 Strategic implications

The findings carry several implications for brand strategy and marketing resource allocation. The study revealed that the GEO competitive landscape structurally favours challengers: Aggarwal et al. (2024) reported that pages ranked around position five on traditional SERPs experienced proportionally larger visibility gains from GEO optimization than pages already ranked first. This finding implies that brands with modest SEO standings can achieve disproportionate AI visibility through targeted evidential content investment and earned media development, disrupting incumbent brand hierarchies in ways that were less accessible under the previous paradigm.

The measurement implications are equally significant. The analysis established that traditional digital marketing KPIs, including sessions, pageviews, and SERP ranking position, do not capture AI visibility. New measurement constructs are required: prompt-level citation frequency, share of voice in AI-generated responses across competitor sets, entity description accuracy and sentiment, and source attribution patterns. The finding that only 16% of brands were systematically monitoring their AI search presence as of early 2026 (Natarajan, 2026) indicates that most organizations have not yet adapted their performance management frameworks to the generative search environment, representing a material first-mover opportunity for early adopters.

The ethical dimension of GEO practice also warrants attention at the strategic level. The manipulation evidence presented by Kumar and Lakkaraju (2024), together with the documented emergence of commercial GEO service providers offering systematic content engineering for LLM visibility (Wen, Zhang et al, 2025),

indicates that adversarial optimization practices are already operating in this space. Brands that invest in transparent, merit-based GEO practices grounded in genuine content quality and earned media breadth are positioned to benefit from anticipated regulatory responses to manipulative practices, which parallel the algorithmic and regulatory countermeasures deployed against black-hat SEO in the traditional search environment.

5. Conclusions

This study examined the mechanisms governing brand visibility in AI-mediated consumer discovery environments and developed a structured GEO Readiness Audit framework in response to the strategic management gap identified in the literature. The findings confirm that the transition from traditional search to generative engine-mediated discovery constitutes a structural shift in the competitive conditions governing digital brand visibility, driven by architectural differences in how LLMs retrieve and synthesize information compared to conventional search engines.

The study established that LLM citation is governed by a distinct signal taxonomy, including content evidentiality, earned media breadth, entity consistency, platform-specific source preferences, and technical extractability, that diverges substantially from the keyword relevance, backlink authority, and on-page optimization signals that have underpinned SEO practice for over two decades. The five-dimension GEO Readiness Audit framework proposed by this study provides brand managers with a structured diagnostic instrument to assess their current AI visibility posture and prioritize remediation actions across these dimensions.

The study contributes to the digital marketing literature in three respects. First, it synthesises the emerging computer science literature on GEO with established brand management and digital marketing research, establishing a cross-disciplinary foundation for the management treatment of AI visibility. Second, it proposes a practitioner-oriented audit framework grounded in synthesised empirical evidence, addressing the absence of structured guidance for GEO readiness assessment in the management literature. Third, it identifies the earned media primacy of LLM citation as a finding with significant implications for the relative strategic weight assigned to owned, earned, and paid media in digital marketing channel allocation.

Several limitations of this study should be acknowledged. The reliance on secondary data and observational industry analytics means that causal claims regarding the effect of specific GEO interventions on citation rates remain inferential. The rapidly evolving nature of LLM architectures and training data means that citation signal weights may shift as platform operators update their systems. The proposed audit framework has not yet been validated through longitudinal empirical testing across diverse industry contexts.

Future research should address these limitations through experimental studies tracking the citation impact of targeted GEO interventions across controlled brand pairs over time. Sector-level analysis of GEO effectiveness variation would refine the framework's applicability across industry contexts. The governance and regulatory dimensions of AI-mediated brand recommendation represent a pressing research agenda given the documented emergence of adversarial optimization practices. As Aggarwal et al. (2024) noted, the transition to generative engines opens a new frontier in information discovery systems with profound implications that the management discipline is only beginning to examine.

REFERENCES

- Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K. and Deshpande, A., 2024. GEO: Generative Engine Optimization. *KDD 2024 – Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. New York: ACM, pp.5-16. <https://doi.org/10.1145/3637528.3671900>
- Bell, D., 2025. 2025 State of AI Discovery Report: What 1.96 Million LLM Sessions Tell Us About the Future of Search. *Previsible.io*. Available at: <<https://previsible.io/seo-strategy/ai-seo-study-2025/>> [Accessed April 2026].
- Berman, R. and Katona, Z., 2013. The Role of Search Engine Optimization in Search Marketing. *Marketing Science*, 32(4), pp.644-651. <https://doi.org/10.1287/mksc.2013.0783>
- Cenamora, J., Parida, V. and Wincent, J., 2019. How entrepreneurial SMEs compete through digital platforms: The roles of digital platform capability, network capability and ambidexterity. *Journal of Business Research*, 100, pp.196-206. <https://doi.org/10.1016/j.jbusres.2019.03.035>
- Dotson, J.P., Fan, R.R., Feit, E.M., Oldham, J.D. and Yeh, Y-H., 2017. Brand Attitudes and Search Engine Queries. *Journal of Interactive Marketing*, 37, pp.105-116. <https://doi.org/10.1016/j.intmar.2016.10.002>
- Erdmann, A., Arilla, R. and Ponzoa, J.M., 2022. Search engine optimization: The long-term strategy of keyword choice. *Journal of Business Research*, 144, pp.650-662. <https://doi.org/10.1016/j.jbusres.2022.01.065>
- Gartner, 2024. Gartner Predicts Search Engine Volume Will Drop 25% by 2026, Due to AI Chatbots and Other Virtual Agents. Available at: <<https://www.gartner.com/en/newsroom/press-releases/2024-02-19-gartner-predicts-search-engine-volume-will-drop-25-percent-by-2026-due-to-ai-chatbots-and-other-virtual-agents>> [Accessed April 2026].
- Jelenik, J., 2025. AI Search: Traffic Killer or the Biggest Opportunity Yet? *Insightland*. Available at: <<https://insightland.org/blog/ai-search-traffic-killer-or-the-biggest-opportunity-yet/>> [Accessed April 2026].
- Keller, K.L., 1993. Conceptualizing, Measuring, and Managing Customer-Based Brand Equity. *Journal of Marketing*, 57(1), pp.1-22. <https://doi.org/10.1177/00222429930570010>
- Kuriatnyk, V., 2025. 2025 AI Visibility Report: How LLMs Choose What Sources to Mention. *TheDigitalBloom*. Available at: <<https://thedigitalbloom.com/learn/2025-ai-citation-llm-visibility-report/>> [Accessed April 2026].
- Kumar, A. and Lakkaraju, H., 2024. Manipulating Large Language Models to Increase Product Visibility. *arXiv* [Preprint], arXiv:2404.07981. <https://doi.org/10.48550/arXiv.2404.07981>
- Mladenovic, D., Rajapakse, A., Kozuljevic, N. and Shukla, Y., 2023. Search engine optimization (SEO) for digital marketers: exploring determinants of online

- search visibility for blood bank service. *Online Information Review*, 47(4), pp.661-679. <https://doi.org/10.1108/OIR-05-2022-0276>
- Moore, O. and Moore, J. 2025. State of Consumer AI 2025: Product Hits, Misses, and What's Next. *Andreessen Horowitz* Available at: <<https://a16z.com/state-of-consumer-ai-2025-product-hits-misses-and-whats-next/>> [Accessed April 2026].
- Mukherjee, S., 2026. The Science of AI Citations: How LLMs choose what sources to reference. *Ekamoira Digital GmbH*. Available at: <https://www.ekamoira.com/blog/ai-citations-llm-sources> [Accessed April 2026].
- Natarajan, P., 2026. LLM Brand Visibility: How to Track and Improve What AI Says About You. *Erlin.ai*. Available at: <<https://www.erlin.ai/blog/llm-brand-visibility/>> [Accessed April 2026].
- Salvi, A., Vitolla, F., Rubino, M., Giakoumelou, A. and Raimo, N., 2021. Online information on digitalisation processes and its impact on firm value. *Journal of Business Research*, 124, pp.437-444. <https://doi.org/10.1016/j.jbusres.2020.10.025>
- Stojadinovic, S., 2025. Google vs ChatGPT & AI Platforms Q1 2025: Trends and Insights. *Omnium*. Available at: <<https://www.omnium.so/blog/google-vs-chatgpt-and-ai-platforms>> [Accessed April 2026].
- Swanson, S., 2026 How to Improve Visibility in AI Results with GEO Strategies. *WebFX*. Available at: <https://www.webfx.com/blog/seo/improve-visibility-ai-results/> [Accessed April 2026].
- Taylor, N., 2025. How Brands Can Stay Visible in an AI-Driven Search World. *Edelman*. Available at: <<https://www.edelman.com/insights/how-brands-stay-visible-ai-search>> [Accessed April 2026].
- Wen, Y., Zhang, N., Yuan, H., Chen, X., Zhang, H. and Guo, H., 2025. Position: On the Risks of Generative Engine Optimization in the Era of LLMs. *TechRxiv*, 20 December. <https://doi.org/10.36227/techrxiv.176620816.64043115/v1>