

A GEO-First Framework: Integrating Search Visibility, Sentiment, and Digital Authority for Organic Growth in the AI Era

Uri Samet *

CEO, Buzz Dealer, Limassol Cyprus.

World Journal of Advanced Research and Reviews, 2026, 29(01), 2031-2040

Publication history: Received on 11 December 2025; revised on 24 January 2026; accepted on 29 January 2026

Article DOI: <https://doi.org/10.30574/wjarr.2026.29.1.0152>

Abstract

The transition from traditional "Retrieve-and-Rank" search architectures to what can be termed as "Retrieve-and-Synthesize" in generative environments has structurally changed the paradigm of digital discovery. This paper proposes a "GEO-First" framework, Generative Engine Optimization as the primary strategic layer, to unify search visibility, online reputation management (ORM), and digital authority. The paper aims to analyze how a GEO-First approach facilitates a continuous discovery-to-branded-search loop by leveraging Large Language Models (LLM) retrieval patterns and Retrieval-Augmented Generation (RAG).

Through a systematic review of empirical literature between the years 2024 and 2026, including metrics such as "Share of Model" (SoM) and citation density, this study identifies two operational pillars, Discovery and Sentiment, grounded in a foundational layer of Digital Authority. The analysis synthesizes retrieval and citation patterns across dominant platforms, including OpenAI's GPT-5.1, Google's Gemini 3, Anthropic's Claude 4.5, and specialized systems like Perplexity AI and DeepSeek. The methodology follows a structured literature review approach, based on cross-platform visibility audits, citation correlation studies, and conceptual work on proper context attribution (e.g., the "MaxShapley" algorithm) to propose an integrated conceptual framework. Findings from the literature review suggest that narrative inclusion within the AI's "retrieval set" acts as a catalyst for subsequent high-intent branded searches, provided that the sentiment integrity of the mention remains positive. The framework concludes that organic growth in the AI era requires a transition from traditional click-oriented metrics to influence-oriented KPIs, to ensure brand resilience in a zero-click, agentic landscape.

Keywords: Generative Engine Optimization (GEO); Online Reputation Management (ORM); Search Engine Optimization (SEO); Digital PR; Digital Authority; Search Visibility

1. Introduction: The Evolution of the Discovery Journey

The digital ecosystem is currently undergoing its most significant structural shift since the rise of the World Wide Web. For over two decades, the primary mechanism of information consumption was defined by the "Retrieve-and-Rank" model, where search engines acted as guides providing a list of potentially relevant URLs.⁵ In this legacy environment, the success of digital presence was measured through rankings, impressions, and click-through rates (CTR). However, the evolution of Generative Search Engines (GSEs) has introduced what can be termed as a "Retrieve-and-Synthesize" paradigm.² In this new era, the engine does not purely point to information. Instead, it consumes, filters, and reconstructs this information into a cohesive, conversational response, often resolving the user's intent without a single click to an external source.⁶

* Corresponding author: Uri Samet

1.1. The Transition: From Retrieve-and-Rank to Retrieve-and-Synthesize

The transition from a linear search process to a generative synthesis process represents a fundamental change in the incentive structure of the internet. Content providers, ranging from news organizations to educational blogs, traditionally relied on search engines to direct traffic to their websites, which was then monetized mainly through advertising.² Generative search engines, however, allow users to obtain comprehensive answers directly through an AI-generated summary. Recent data suggests that the fragment of global web traffic produced by traditional search fell by approximately 5% between June 2024 and June 2025, with some estimates suggesting a drop of up to 25% in specific high-intent categories.²

This shift is characterized by the emergence of Retrieval-Augmented Generation (RAG). RAG systems do not rely solely on the static training data of an LLM, and instead they actively "Google" the web in real-time to ground their answers with current facts.⁶ This technical evolution means that a brand's website is no longer the final destination for the majority of users, but rather a source of raw data for the AI's synthesis engine. As a result, the "zero-click" reality has required a pivot from optimizing for traffic to optimizing for "Share of Model" (SoM), defined as the proportion of times a brand is featured as a trusted recommendation within these synthesized responses.⁹

1.2. The Feedback Loop: Discovery-to-Search

In the generative search environment, the traditional user journey is being replaced by a recursive "Discovery-to-Search" loop, where the AI-generated mention serves as the primary awareness layer. When an LLM recommends a specific product or service in response to a categorical query (e.g., "What is the most reliable CRM for a scaling fintech company?"), it establishes an immediate brand salience.¹¹ This mention triggers a subsequent, high-intent branded search on traditional search engines or social platforms, as users seek to get additional validation to the AI's recommendation.⁴

This loop is powered by the psychological phenomenon of brand salience. Even if a user does not click on the provided citations within a chat response, the inclusion of the brand name in the generated text acts as a mental anchor. Research indicates that branded search volume acts as a proxy for wider brand recognition, in which if the AI mentions a brand, users are significantly more likely to search for that brand directly later in their decision cycle. It should be noted, however, that while this causal mechanism is strongly supported by established brand salience theory, direct empirical studies isolating the AI-mention-to-branded-search pathway remain limited. Much of the current evidence is correlational, and further controlled research is needed to quantify the effect with precision.^{28,30,11}

Table 1: User Journey in Traditional Search Vs. AI-Era Discovery

Traditional Search Journey	AI-Era Discovery Journey
Trigger: Keyword search for information.	Trigger: Conversational prompt for a solution.
Action: Scanning "ten blue links" for relevance.	Action: Consuming a synthesized recommendation.
Destination: Clicking a URL to visit a website.	Destination: Narrative inclusion; zero-click resolution.
Outcome: Conversion on the owned domain.	Outcome: Branded search validation and agentic purchase.

Source: ⁵

2. Literature Review: Digital Authority as the Foundation

This study adopts a structured review methodology to synthesize the recent common knowledge on generative engine optimization and its implications for brand visibility. The review draws on academic publications indexed in ACM Digital Library, arXiv, and Google Scholar, supplemented by high-authority industry reports from relevant organizations including Ahrefs, Forrester, and Jellyfish/INSEAD, as well as practitioner-level publications that capture phenomena which are too recent for peer-reviewed coverage. The literature search spanned the period from 2024 to early 2026, reflecting the rapid emergence of RAG-based search systems and their commercial implications. Sources were selected based on relevance to the core research question, which is how brands can systematically optimize for narrative inclusion in LLM-generated responses, and they were evaluated for methodological precision, recency, and contribution to the proposed conceptual framework. Given the developing and fast-evolving nature of the field, practitioner sources were included where they could contribute empirical data (e.g., citation correlation analyses, platform-specific retrieval audits) not yet available in peer-reviewed form, and are distinguished from academic sources throughout the discussion.

2.1. From Keywords to Topical Entities

The evolution of search authority has moved beyond keyword density toward the recognition of topical entities within a semantic network. Traditional SEO was based on matching exact strings of text.⁶ In contrast, generative engines utilize vector embeddings, which are numerical representations of meaning within a multi-dimensional space, to understand conceptual relationships.⁶ For example, a system recognizing "customer loyalty" will mathematically associate it with related terms such as "CLTV" and "retention programs," even if those terms are not explicitly paired on the page.⁶

This transition is demonstrated by systems like Amazon's COSMO, which mines user-centric knowledge from behaviors to build wide range knowledge graphs.¹³ These graphs allow search systems to move from "strings to things," prioritizing entities that demonstrate deep topical authority over entities that simply rank for a specific term.¹¹ Authority in the GEO-First framework is therefore defined by "Entity Solidification", which stands for the degree to which a brand is consistently associated with specific capabilities, experts, and industries across the digital landscape.⁵

2.2. Brand Mentions as Trust Vectors

In the RAG-driven retrieval environment, brand mentions have become the "new link building." While traditional backlinks still function as signals of relevance, the AI training and retrieval process increasingly weights unlinked mentions across high-trust third-party publications.^{6,29,30} When an AI observes a brand being cited repeatedly as a source of truth for a particular topic (e.g., a "cybersecurity firm" + "threat intelligence reports") across diverse domains, then the confidence score for that entity increases.⁶

Research into the "MaxShapley" algorithm further explains this shift. MaxShapley provides an efficient mechanism for fair attribution in RAG pipelines, compensating content providers based on their specific contribution to a generated answer.² This suggests that "Citation Share" is a critical measure of authority. Being the cited source for a specific "ground truth" fact might be the strongest form of authority in an environment where AI engines synthesize multiple perspectives into a single voice.²

2.3. The Mechanics of RAG-Based Authority

To understand how authority is calculated in 2026, we must analyze the interaction between the "Knowledge Layer" (the LLM's training data) and the "Retrieval Layer" (the live web indexing mechanism).

- **Vector Embeddings and Semantic Density:** Search engines convert content into vector embeddings. The more a certain content piece is "conceptually complete", the stronger its vector signal becomes, making it more likely to be retrieved.⁶
- **Information Gain:** AI engines prioritize content that offers unique data or perspectives that the LLM does not already possess in its training set. This includes proprietary research, survey results, or first-person experience-based narratives.⁸
- **The Inverted Pyramid:** For machine readability, authority is signaled through a structural hierarchy. Direct answers at the top of a page offer "Grounding," while additional data (tables, lists) acts as a "Trust Anchor" for the model.⁶

3. Pillar I: Discoverability

3.1. The Mechanics of Citations: Securing a Place in the Retrieval Set

Discoverability in the generative era is a technical battle. For a brand to be included in the final synthesis, its content must first be "retrieved" from the vector database. This retrieval process favors content that is technically accessible, modular, and optimized for "chunking", which can be described as the process of breaking text into logical pieces for AI 'ingestion'.¹⁶

Technical SEO has evolved into serving as a precondition for "Extreme Reliability." AI engines tend to be "impatient" and prioritize pages with low latency and server-side rendering (SSR), saving the wait for client-side JavaScript to load the primary content.⁸

Table 2 Technical SEO Requirement's impact on AI retrieval

Technical Requirement	Impact on AI Retrieval
Server-Side Rendering (SSR)	Ensures full text is available in the initial HTML data for immediate parsing.
Low Latency (<200ms)	Prevents crawlers from moving to the next source in the index.
Interaction to Next Paint (INP)	Measures the responsiveness of a page for agentic scraping.
Modular Content Design	Enables the "chunking" of facts into independent units of information.

Source: ⁶

A brand that excels at technical structure but lacks trust will be taken by the AI but ultimately ignored in the generation phase. On the other hand, a highly trusted brand with a messy technical structure remains invisible to AI's real-time agents.¹⁶

3.2. Digital PR and the Seeding of Ground Truth

In the GEO-First framework, Digital PR shifts from obtaining brand mentions and backlinks to seeding the "ground truth" across the "Reference Layer", the third-party media and review sites that LLMs use to validate claims.¹⁰ AI engines do not just "remember" facts, they actively verify them against a variety of authoritative sources.

To win this "Grounding Game," brands must ensure their information (pricing, features, leadership etc.) is accurate across their entire digital footprint. If an AI encounters conflicting information about a brand's CEO or product specifications, it may experience a "hallucination" or move to an exclusionary logic, omitting the brand to avoid inaccuracies.⁸ Strategic Digital PR thus involves identifying the "Kingmaker" URLs that AI engines currently cite for priority keywords and ensuring that the brand is accurately represented on those pages.¹⁰

3.3. The Role of Schema markup and Entity Solidification

Schema markup (Schema.org) has become the primary language for "constraining the AI's creativity." By publishing "Source of Truth" data through 'Organization', 'Product' or 'FAQPage' schemas, brands provide a structured Knowledge Graph that the AI can use to ground its generated responses.⁵

- **FAQPage Schema:** Provides direct, natural language answers that match the conversational prompts of LLMs.⁹
- **Organization and LocalBusiness Schema:** Establishes a clear entity profile, linking social profiles and official domains ('SameAs') to ensure the AI recognizes the brand as a distinct and authoritative entity.⁵
- **Article and Author Schema:** Signals E-E-A-T by verifying the human expertise behind the content, which serves as a critical filter for AI engines seeking credible "ground truth".⁵

4. Pillar II: Sentiment

4.1. The Warning Effect: Exclusionary Logic and Negative Social Proof

Sentiment acts as the "Gatekeeper" for AI recommendations. While Discoverability ensures narrative inclusion, Sentiment determines the *tone* of that inclusion. AI engines are built to avoid recommending brands which are associated with high risk or with a negative user sentiment. This "Warning Effect" is triggered when the AI retrieves negative social proof, such as recurring complaints about customer service or product failures, from its vector database.⁴

Research of "asymmetric propagation" has shown that LLMs often have strong guardrails against negative misinformation regarding certain entities, yet they are highly influenced by the "foundational favorability" of the data they ingest.¹⁸ In some cases, LLMs demonstrate "exclusionary logic," where a brand is filtered out of a "Top 10" list not because it lacks relevance, but because the model's guardrails identify a high rate of "Hallucinations" or a poor sentiment velocity, which is defined as the rate and direction in which public sentiment about a brand is changing across indexed sources over a given period.⁴

4.2. Sentiment Integrity for Agentic Systems

As we move toward agentic commerce, where AI agents act on behalf of users to research and purchase products, "Sentiment Integrity" becomes a critical requirement. Agentic systems require "Clean Recommendations", made out of

data that is free from contradictory reviews or reputational red flags.¹⁷ An AI agent authorized to purchase an enterprise software will bypass any vendor that the model perceives as having a "chaotic" or "incomplete" digital footprint.¹⁵

Managing sentiment in 2026 involves:

- **Inference Optimization:** Monitoring what AI tools are telling customers about the brand and correcting inaccuracies at their source.¹⁵
- **The "Fortress of Facts":** Curation of a deep, factual knowledge base on the brand's owned domain to serve as the corrective data against the "fog" of online information.¹⁵
- **Sentiment Score Monitoring:** Tracking whether the AI positions the brand as "best value" or "expensive and buggy," and adjusting content strategies to shift these semantic assumptions.¹⁰

4.3. The Conversion Paradox: Sentiment vs. Visibility

A high "Share of Model" (SoM) is only valuable if the sentiment is positive, thanks to a structured "Conversion Paradox" in generative search: while visibility can be achieved through technical optimization, value is only captured through sentiment clarity.¹⁰ Visitors from LLMs convert at significantly higher rates because the AI has already performed a qualitative filtering. However, if that filter includes a warning, the conversion rate naturally collapses.

Table 3 The impact of Sentiment on GEO metrics

Metric	Impact of Positive Sentiment	Impact of Negative Sentiment
Share of Model	High recommendation frequency.	High warning frequency or exclusion.
Citation Share	Trusted as a source of truth.	Cited as an example of a common issue.
Branded Search	"Validation" searches.	"Alternative to" or "problem with" searches.
Conversion Rate	Up to 4.4x traditional search.	Near zero for agent-led transactions.

Source: ^{9,31}

5. Platform-Specific Benchmarking (Multi-Model Analysis)

As of early 2026, the AI landscape has split into specialized models with distinct retrieval and synthesis styles. A GEO-First framework must account for these variations. The following comparison represents a point-in-time snapshot. Given the speed of model development, specific benchmark scores and capability characterizations are subject to a rapid change, and practitioners should treat this as directional rather than definitive.

5.1. Comparative Evaluation of Leading Models

Table 4 Comparative Evaluation of Leading Models

AI Platform	Retrieval/Synthesis Logic	Key Benchmark Strength
Gemini 3 (Google)	Generative UI & Reasoning: Uses "Deep Think" mode to plan steps; renders widgets instead of just text.	93.8% on GPQA Diamond; leading reasoning for science. ²⁰
GPT-5.1 (OpenAI)	Balanced Fluidity: Split into "Instant" (speed) and "Thinking" (deep logic) modes.	94.6% on AIME 2025; largest context window at 400k tokens. ²⁰
Claude 4.5 (Anthropic)	Extended Reasoning & Safety: Combines strong analytical reasoning and code generation with "Constitutional AI" safety alignment, emphasizing factual completeness and nuanced synthesis.	77.2% on SWE-bench Verified; strong performance across reasoning, coding, and analytical tasks. ²⁰
Perplexity AI	Search-First UX: Optimized for real-time web retrieval with automatic source citation.	Runaway favorite for academic research and market fact-gathering. ²¹

Grok 4.1 (xAI)	Real-Time Social: Wired into the X data stream for live news and social sentiment analysis.	#1 on EQ-Bench3; highest emotional intelligence and real-time news access. ²⁰
DeepSeek V3.2	Efficient Classification: High-value, cost-effective model with strong sentiment and predictive analytics.	Outperforms competitors in classification accuracy for finance. ²³

Source: ²⁰

5.2. Retrieval Patterns and Citation Density

LLMs vary significantly in how they attribute sources. Perplexity AI and Gemini are known for "link-rich" answers, while GPT-5.1 and Claude often provide more synthesized responses with fewer external citations unless specifically prompted.²⁰ Furthermore, "Share of Model" is extremely varied between models. For example, research into the laundry detergent market showed a brand commanding 24% of mentions on Meta's Llama but dropping to less than 1% on Google's Gemini, demonstrating that a brand's authority is platform-dependent.¹

5.3. Synthesis Strategies: Generative UI Vs. Static Text

Google's Gemini 3 has pioneered "Generative UI," where information is synthesized into interactive widgets, graphs, and buttons rather than just static paragraphs.²⁰ This shift demands that brands provide structured data that can be dynamically re-rendered by the AI. Meanwhile, OpenAI's GPT-5.1 implements "surgical diffs" for technical synthesis (e.g., the `apply_patch` tool), which favors modular, documentation-style content that can be easily integrated into larger codebases or reports.²⁰

6. Empirical Evidence & Aggregated Analysis

6.1. Share of Model (SoM) Analysis: Variability and Correlation

The "Share of Model" concept, pioneered by researchers at Jellyfish, has provided a large-scale evidence for the variability of brand visibility across different LLMs.¹ An analysis of the automotive and retail sectors reveals a significant disconnection between traditional brand equity and AI-perceived relevance. By mapping brands along two axes, human awareness (measured through conventional brand recall surveys) and AI awareness (measured through Share of Model across leading LLMs), four distinct categories of brand visibility can be identified, each with different strategic imperatives.

Table 5 Brand Visibility Categories by Human Vs. AI Awareness

Brand Category	Human Awareness	AI Awareness	Strategic Focus
Cyborgs	High	High	Maintain "feature-dense" marketing to support AI reasoning.
AI Pioneers	Low	High	Leverage "solution-oriented" content to bypass traditional competitors.
High-Street Heroes	High	Low	Shift from "heritage" narratives to "functional data" for LLM ingestion.
Emergent Brands	Low	Low	Primary risk of "digital irrelevance" in the zero-click era.

Source: ¹

This categorization exposes a "Human-AI Awareness Gap" that challenges long-standing assumptions about brand strength. Most notably, High-Street Heroes, which stands for brands with strong consumer recognition, often fail to translate that equity into AI visibility because their marketing emphasizes intangible qualities (such as elegance or heritage) that LLMs do not weight as heavily as they weight functional, structured data such as battery life or specific software features. Conversely, AI Pioneers demonstrate that brands with limited traditional awareness can achieve outsized influence in the generative discovery layer by producing solution-oriented, data-rich content that aligns with

the retrieval logic of LLMs. The implication is clear: legacy brand equity is no longer a reliable criterion for discoverability in the AI era, and each quadrant demands a fundamentally different optimization strategy.¹

6.2. The Influence of Digital PR on AI Citations

The correlation between a brand's presence in high-authority third-party publications and its citation frequency in LLMs is well-documented. Princeton research on evaluating GEO methods found that the inclusion of statistics, authoritative citations, and quotations boosted a source's visibility in generative responses by up to 40%.^{9,28}

In a study of the Italian laundry detergent market, the brand **Chanteclair** held a 19% SoM on Perplexity but was entirely absent from Llama.¹ This suggests that "model-specific optimization" is a requirement for modern Digital PR, which means that brands must identify which models their target audience uses and make efforts to seed content in the specific URLs those models' retrieval systems prioritize.¹

6.3. Aggregate Impact on Branded Search Volume

Case studies across B2B and B2C sectors demonstrate that a successful "GEO-First" campaign leads to a significant increase in branded search volume. As users encounter a brand as a recommended solution in ChatGPT, Gemini or other AI engines, they perform follow-up searches to verify the recommendation.⁷ One B2B Webflow agency reported that restructuring their site for LLM discoverability led to 10% of their organic traffic coming directly from generative engines, with 27% of this traffic converting into Sales Qualified Leads (SQLs).⁹

7. Strategic Implementation across Online Industries

7.1. Universal Applications: B2B SaaS, E-commerce, and YMYL

The GEO-First approach requires industry-specific tailoring.

- **B2B SaaS:** Success is driven by "Fact Density" and technical documentation. Therefore, SaaS brands must focus on being the "Ground Truth" for technical specifications and comparison queries.⁶
- **E-commerce:** Optimization focuses on the "Universal Commerce Protocol" (UCP) and agent-friendly product feeds. AI agents making purchasing decisions have "zero tolerance" for inaccurate inventory or pricing.¹⁷
- **YMYL (Your Money Your Life):** In regulated industries, the priority is "Radical Transparency" as well as the use of cryptographic signatures (C2PA) to prove that content is human-verified and authoritative.⁸

7.2. The Transition from Clicks to Influence

The fundamental shift in KPIs involves moving from measuring *traffic* to measuring *salience* and *sentiment clarity*.

Table 6 Legacy KPIs Vs. GEO-First KPIs

Legacy KPI	GEO-First KPI	Strategic Purpose
Keywords in Top 10	Share of Model (SoM)	Measures narrative inclusion across LLM responses.
Organic Traffic (CTR)	Citation Share/Frequency	Measures the degree to which a brand is the "source of truth."
Sentiment Analysis	Sentiment Clarity Index	Measures the consistency of positive social proof for agents (i.e., the degree to which AI-retrieved sentiment about a brand is uniformly favorable across sources, without contradictory signals).
Backlink Count	Entity Mention Count	Measures "Trust Vectors" across the reference layer.

Source: ⁴

7.3. Implementation Measures

Organizations should implement GEO-First monitoring on a structured basis to ensure long-term salience:

- **Weekly:** Visibility audits on the top 20 "Golden Prompts", which accounts for the highest-value conversational queries that a brand's target audience is most likely to pose to an LLM, in order to track appearances and competitive movements.⁴
- **Monthly:** Aggregation of citation frequency and documentation of hallucination fixes.⁴
- **Quarterly:** Deep-dive analysis of sentiment velocity and cross-platform presence, compared to traditional brand awareness surveys.⁴
- **Bi-annually:** Connection of GEO metrics to business indicators, specifically branded search trends and sales conversion themes.⁴

While the above implementation measures provide the operational foundation for today's GEO-First strategy, the landscape itself is evolving at a speed that demands forward-looking preparation. The monitoring and optimization practices outlined above will need to adapt as the baseline technology shifts, particularly as AI agents move beyond information retrieval into autonomous transactional roles. The following section examines the near-term developments that are expected to reshape these strategic requirements.

8. Future Outlook (Near-Term 2026-2027)

8.1. Agentic Commerce and the Agent-to-Agent Economy

By late 2026, the market is expected to see the rise of "Agentic Commerce," where buyer bots negotiate with seller bots.²⁶ Google's Universal Commerce Protocol (UCP) and Amazon's Alexa+ are already setting the stage for a world where AI agents handle the entire shopping journey from discovery to checkout.¹⁷ This shift will force vendors to build their own "Business Agents" in order to provide real-time quote negotiations and inventory verification.¹⁷

8.2. The "Human-Verified" Premium

In an environment saturated with AI-generated content (AIGC), the value of human expertise might reach a premium. AI engines are increasingly seeking for "experience-based content" that models cannot easily replicate.¹⁴ Techniques such as "Radical Transparency", using C2PA to sign content as human-verified, might become essential for maintaining "Entity Trust".⁸

8.3. Security and Reputation in Agentic Discovery

The "SENTINEL Framework" highlights the emerging security challenges for AI agents in cyber-physical systems, emphasizing the need for source-grounded system design.²⁷ For brands, this means that security and protection against "Identity Spoofing" will become a part of the reputation management stack. Ensuring that an AI agent interacts with a legitimate and authenticated "Source of Truth" will be the next frontier for digital authority.¹⁷

Limitations

Several limitations should be acknowledged when interpreting the findings of this review. First, the GEO-First framework is a conceptual model that was synthesized from existing literature, and has not been validated through controlled experimentation. While the reviewed studies provide correlational evidence supporting the proposed mechanisms, and particularly the relationship between AI brand mentions and branded search volume, causal claims remain provisional. Second, this review is based on a mix of peer-reviewed academic research, industry reports, and practitioner-level publications. The third category, while valuable for capturing emerging phenomena in a rapidly evolving field, does not undergo the same level of methodological inspection as peer-reviewed work, and readers should weigh the evidence accordingly. Third, the platform-specific benchmarks and the different model characterizations presented in Chapter 6 reflect a point-in-time snapshot of a landscape in continuous flux, as model capabilities, retrieval architectures and citation behaviors are subject to change with each major model release. Fourth, the framework's applicability may vary across industries, geographies, and brand sizes in ways which are not fully captured by the reviewed literature, which might include a structured bias toward English-language, Western-market contexts. Finally, key metrics proposed in this framework, including Sentiment Clarity Index and Sentiment Velocity, are emerging indicators that still lack standardized measurement methodologies. Future empirical work is needed to validate these measures before they can be reliably referred to as KPIs.

9. Conclusion

The "GEO-First" framework provides a unified strategic roadmap for organic growth in an era dominated by generative engines and autonomous agents. By positioning Generative Engine Optimization as the primary layer for narrative

inclusion, brands can facilitate a continuous discovery-to-branded-search loop that prioritizes influence over clicks. The two operational pillars of Discovery and Sentiment, grounded in a foundational layer of Digital Authority, ensure that a brand is not only retrieved by the AI's RAG systems but also gets recommended as a trusted entity with positive sentiment integrity. Authority, in this model, is not a parallel objective but the enabling condition: it feeds both the likelihood of being discovered and the favorability with which the brand is discussed online.

Ultimately, the transition from traditional SEO to a GEO-First approach requires a radical commitment to technical reliability, structured data (Schema), and a growth of unlinked brand mentions as the new "trust vectors" of the AI age. Brands that successfully achieve the status of being perceived as a "ground truth" by AI models will achieve a sustainable competitive advantage in the zero-click economy of 2026 and beyond.

References

- [1] Dubois, D., Dawson, J., & Jaiswal, A. (2025, June). Forget what you know about search. Optimize your brand for LLMs. Harvard Business Review. <https://hbr.org/2025/06/forget-what-you-know-about-seo-heres-how-to-optimize-your-brand-for-llms>
- [2] MaxShapley: Towards incentive-compatible generative search with fair context attribution. (2025). ResearchGate. https://www.researchgate.net/publication/398430304_MaxShapley_Towards_Incentive-compatible_Generative_Search_with_Fair_Context_Attribution
- [3] Caption injection for optimization in generative search engine. (2025). arXiv. <https://arxiv.org/pdf/2511.04080.pdf>
- [4] GEO metrics: How to measure visibility, trust, and brand presence in AI search. (2026). Foundation Marketing. <https://foundationinc.co/lab/geo-metrics>
- [5] 10 proven AEO strategies for higher rankings in 2026. (2026). Woonyb. <https://woonyb.com/blog/seo-marketing/10-proven-aeo-strategies-for-higher-rankings/>
- [6] Best 18 tips for ranking in AI search in 2026. (2026). Yotpo. <https://www.yotpo.com/blog/ranking-ai-search-geo-tips/>
- [7] Roach, T. (2026). Digital marketing (Tag archive). <https://thetomroach.com/tag/digital-marketing/>
- [8] Search engine optimisation (SEO) & GEO. (2026). Coaley Peak®. <https://www.coaleypeak.co.uk/expertise/digital-marketing/search-engine-optimisation/>
- [9] Engineering brand entity status for the age of generative answer engines. (2026). 1827 Marketing. <https://1827marketing.com/smart-thinking/engineering-brand-entity-status-for-the-age-of-generative-answer-engines/>
- [10] SEO vs. GEO: How to optimize for search engines and AI. (2026). Foundation Marketing. <https://foundationinc.co/lab/seo-vs-geo>
- [11] What are AEO and GEO? How AI is redefining search. (2026). Jellyfish. <https://www.jellyfish.com/en-gb/training/blog/what-are-aeo-and-geo>
- [12] Share of model: The essential marketing metric for the AI era. (2026). Symphonic Digital. <https://www.symphonicdigital.com/blog/understanding-share-of-model>
- [13] Cosmo: A large-scale e-commerce common sense knowledge generation and serving system at Amazon. (n.d.). Scribd. <https://www.scribd.com/document/830397246/Cosmo-a-Large-Scale-e-Commerce-Common-Sense-Knowledge-Generation-and-Serving-System-at-Amazon>
- [14] Hassan, S. (2026). How to choose the best keyword research tools in 2026. <https://blog.shayakehassan.com/how-to-choose-the-best-keyword-research-tools-in-2026>
- [15] Inference optimisation: Build your own perpetual AI-powered content machine. (2026). Hobo. <https://www.hobo-web.co.uk/inference-optimisation-build-your-own-legitimate-perpetual-ai-powered-content-machine/>
- [16] How to optimize content for AI search engines in 2026. (2026). Elementor. <https://elementor.com/blog/how-to-optimize-content-for-ai-search-engines/>
- [17] What is agentic commerce? 45% of shoppers use AI (2026). (2026). Ekamoira Blog. <https://www.ekamoira.com/blog/what-is-agentic-commerce-the-complete-2026-guide-to-ai-shopping-agents>

- [18] Do language models favor their home countries? Asymmetric propagation of positive misinformation and foreign influence audits. (2026). Harvard Kennedy School Misinformation Review. <https://misinforeview.hks.harvard.edu/article/do-language-models-favor-their-home-countries-asymmetric-propagation-of-positive-misinformation-and-foreign-influence-audits/>
- [19] Hobo strategic AiSEO 2025. (2025). Scribd. <https://www.scribd.com/document/910266016/Hobo-Strategic-AiSEO-2025>
- [20] The best AI of November 2025: Gemini 3 vs GPT-5.1 vs Grok 4.1 vs others. (2025). FelloAI. <https://felloai.com/the-best-ai-of-november-2025/>
- [21] The complete guide to choosing AI platforms in 2026: ChatGPT, Claude, Gemini, Perplexity, Copilot, DeepSeek, Grok & Mistral compared. (2026). First AI Movers. <https://www.firstaimovers.com/p/complete-eight-ai-platform-comparison-guide-2025>
- [22] Perplexity: Best chatbot for academic papers. (2025, August 14). ETC Journal. <https://etcjournal.com/2025/08/14/perplexity-best-chatbot-for-academic-papers/>
- [23] From trading bot to trading agent: How to build an AI-based investment system. (2026). Medium. <https://medium.com/@gwrx2005/from-trading-bot-to-trading-agent-how-to-build-an-ai-based-investment-system-313d4c370c60>
- [24] A comparison of DeepSeek and other LLMs. (2025). ResearchGate. https://www.researchgate.net/publication/388791877_A_Comparison_of_DeepSeek_and_Other_LLMs
- [25] Retail's risky AI commerce bet. (2026). Retail Dive. <https://www.retaildive.com/news/retails-ai-bet-risks-data-market-share/810202/>
- [26] Predictions 2026: The agentic commerce race in digital commerce. (2026). Forrester Research. <https://www.forrester.com/blogs/predictions-2026-the-agenetic-commerce-race-and-some-potential-regrets-in-digital-commerce/>
- [27] A survey of environmental interactions, deepfake threats, and defenses. (2026). arXiv. <https://arxiv.org/html/2601.20184v1>
- [28] Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2024). GEO: Generative engine optimization. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24) (pp. 5–16). ACM. <https://dl.acm.org/doi/10.1145/3637528.3671900>
- [29] Chen, M., Wang, X., Chen, K., & Koudas, N. (2025). Generative engine optimization: How to dominate AI search. arXiv preprint arXiv:2509.08919. <https://arxiv.org/abs/2509.08919>
- [30] Guan, X., et al. (2025, August). An analysis of AI overview brand visibility factors (75K brands studied). Ahrefs Blog. <https://ahrefs.com/blog/ai-overview-brand-correlation/>
- [31] GEO industry report 2025: Trends in AI & LLM optimization. (2025). Omnis. <https://www.omnis.so/blog/geo-industry-report>
- [32] Buzz Dealer. (2025, October). GEO: The key to reputation in the age of AI. BuzzDealer.com. <https://buzzdealer.com/geo-the-key-to-reputation-in-the-age-of-ai/>