Territorial Algorithmic Invisibility: A Diagnostic Framework for Generative Engine Optimization Applied to Local Tourism Destination Entities

Eric Mees de Saboya Ribeiro
Criamente Digital Brand Architecture, Niterói, Rio de Janeiro, Brazil
ORCID: [to be registered]
Corresponding author: eric@criamente.com
Submitted: May 2026

Keywords: Generative Engine Optimization, GEO, Destination Marketing Organizations, Territorial Algorithmic Invisibility, Local Tourism, Large Language Models, Entity Consistency, Schema.org, AI Citation, B2A, Proximity Shadow Effect

Abstract

The emergence of Large Language Models (LLMs) as primary mediators of travel information retrieval presents a distinct challenge for local tourism destinations: algorithmic invisibility induced not by poor content production, but by geographic proximity to dominant destinations that systematically absorb a smaller entity's algorithmic identity. This paper introduces the concept of Territorial Algorithmic Invisibility (TAI) and proposes the Destination Algorithmic Readiness Score (DARS), a five-dimensional diagnostic framework for measuring and addressing GEO deficiencies in Destination Marketing Organizations (DMOs). The framework is derived from a structured diagnostic audit of Neltur, the official tourism company of Niterói, Rio de Janeiro — a destination of demonstrable cultural and geographic significance that exists in the algorithmic shadow of its neighboring metropolis. Systematic querying of ChatGPT, Google Gemini, and Perplexity AI revealed that Niterói attractions are consistently attributed to Rio de Janeiro or omitted entirely from AI-generated travel recommendations, despite housing internationally recognized architectural landmarks including Oscar Niemeyer's Museu de Arte Contemporânea (MAC). The DARS baseline established for Niterói (38/100) quantifies the gap between the destination's real-world tourism value and its current algorithmic representation. The paper contributes to the nascent GEO literature by extending existing frameworks beyond brand and professional entities to geographic destinations, formalizing the Proximity Shadow Effect as a structural impediment to local DMO visibility, and proposing a reproducible diagnostic methodology applicable to tourism entities globally.

1. Introduction

Travel planning behavior has undergone a structural shift. According to Phocuswright (2024), over 40% of travelers under 35 now initiate trip planning through conversational AI tools before consulting traditional booking platforms or search engines. When a traveler asks ChatGPT "what are the best destinations near Rio de Janeiro," the response they receive is not a ranked list of links — it is a synthesized recommendation that either includes or excludes specific destinations based on criteria that remain largely opaque to the destination's marketing organization.

This shift creates an asymmetric visibility problem for local tourism entities. Major destinations — cities with high international recognition, extensive media coverage, and robust Wikipedia presence — have accumulated what this paper terms Editorial Authority: a corpus of training data, citations, and cross-platform mentions that ensures consistent representation in LLM-generated responses. Local destinations, regardless of their real-world tourism value, may be absent from these responses not because they lack attractions, but because they lack the algorithmic infrastructure that enables LLMs to identify, attribute, and recommend them with confidence.

The problem is compounded for destinations that exist in geographic proximity to major ones. Niterói, located across Guanabara Bay from Rio de Janeiro, offers internationally recognized architectural heritage, pristine beaches, a thriving restaurant scene, and a cultural identity distinct from its neighbor. Yet systematic testing of LLM-based travel queries reveals a consistent pattern: Niterói's attractions are either attributed to Rio de Janeiro, categorized as generic day-trip appendages, or omitted entirely. This paper terms this phenomenon the Proximity Shadow Effect.

Generative Engine Optimization (GEO), formalized by Aggarwal et al. (2023) as the discipline of optimizing digital presence for inclusion in AI-generated responses, has produced frameworks primarily applicable to brands, professionals, and knowledge producers. No published framework addresses the specific challenges of geographic destination entities. This gap is significant: Destination Marketing Organizations (DMOs) are among the entities most vulnerable to the shift from search-mediated to AI-mediated information discovery, yet they are the least equipped with diagnostic tools to address their algorithmic condition.

This paper addresses that gap through three contributions:

- **A construct describing the state in which a geographic destination of real tourism value receives absent, minimal, or systematically misattributed representation in LLM-generated travel recommendations.** Territorial Algorithmic Invisibility (TAI):
- **A five-dimensional diagnostic metric enabling DMOs to quantify their current GEO baseline and prioritize remediation efforts.** The Destination Algorithmic Readiness Score (DARS):
- **A structured audit of Neltur and Niterói establishing the first published DARS baseline for a Brazilian destination, with LLM query testing results across three major platforms.** Empirical diagnostic data:

2. Background and Related Work

2.1 From Search to Conversation: The Transformation of Travel Information Retrieval

Travel has historically been one of the highest-value categories in online search, generating billions in advertising revenue for platforms that successfully intercepted travelers during the planning phase. The dominance of platforms like TripAdvisor, Booking.com, and Google Travel was built on the premise that travelers search, compare, and then decide.

LLM-powered assistants disrupt this premise at its foundation. When a traveler asks "where should I stay for three days to explore the architecture around Rio de Janeiro," the response they receive does not present options for comparison — it presents a synthesized itinerary. The destination that appears in that itinerary is not determined by advertising spend or SEO ranking; it is determined by the LLM's assessment of which entities are sufficiently well-represented, consistent, and authoritative in its training and retrieval data to be recommended with confidence.

This shift from discovery-by-ranking to discovery-by-recommendation has profound implications for destination marketing. A destination that does not appear in an LLM recommendation is, for the growing cohort of AI-assisted travelers, algorithmically nonexistent — regardless of the quality of its attractions or the competence of its marketing team.

2.2 Foundational GEO Research

The term Generative Engine Optimization was introduced by Aggarwal et al. (2023) in a Princeton University study that systematically tested nine optimization strategies across multiple generative engines. Their findings established that citing authoritative sources, including relevant statistics, and using quotations from credible entities improved LLM citation rates by 30–40%. The study also established a critical distinction: GEO and SEO are not equivalent disciplines. Content that ranks well in traditional search does not necessarily receive citation in generative responses.

Brasil (2026), in a practitioner framework derived from a 7-day implementation sprint, extended GEO research through the introduction of several novel constructs, including the Invisible Excellence Paradox — technically optimized entities receiving zero citation due to insufficient Editorial Authority — the Entity Consistency Score (a quantitative metric for cross-platform identity alignment), and a 10-layer GEO Enterprise Framework. Brasil's work established that algorithmic visibility is a multiplicative function of Technical Readiness and Editorial Authority, and that either factor being zero produces zero visibility.

The present paper extends this line of research to geographic destinations, where both factors require fundamentally different operationalization than in brand or professional contexts.

2.3 Citation Distribution and the Platform Hierarchy

A January 2026 Semrush study analyzing citation patterns across major LLM platforms found that Reddit (11.29%), LinkedIn (11.03%), Wikipedia (9.53%), and YouTube (8.77%) accounted for the highest share of citations. For geographic destinations, this distribution is particularly consequential: Wikipedia is the primary structured knowledge source for geographic entities, and its coverage quality directly influences how LLMs represent destinations in generated responses.

For Niterói, Wikipedia coverage exists but is sparse compared to Rio de Janeiro, which has extensive multilingual coverage, rich entity relationships, and a high volume of cross-platform references. This asymmetry in Wikipedia presence translates directly into asymmetry in LLM representation — a foundational manifestation of the Proximity Shadow Effect this paper formalizes.

2.4 Destination Marketing Organizations and the Digital Transition

DMOs face a structural disadvantage in the GEO transition. Unlike commercial brands, which generate revenue through the digital channels they optimize, DMOs generate value indirectly — through tourism expenditure that flows to hotels, restaurants, and attractions, not to the DMO itself. This structural disconnect means DMOs have historically underinvested in technical digital infrastructure relative to their commercial counterparts.

A 2024 UNWTO report on digital transformation in tourism noted that fewer than 15% of regional DMOs in emerging markets had implemented structured data markup on their primary websites, and fewer than 8% had systematic approaches to cross-platform entity management. The report predated the mass adoption of LLM-powered travel planning, suggesting these figures underestimate the current gap.

2.5 Structured Data and Knowledge Graphs for Geographic Entities

Schema.org provides specific types for geographic entities relevant to destination marketing: TouristDestination, TouristAttraction, LandmarksOrHistoricalBuildings, Accommodation, FoodEstablishment, and Event. These types, combined with the Place schema's geo property (GeoCoordinates), enable a structured representation of destination entities that knowledge graphs and LLMs can parse, verify, and cite.

However, structured data implementation in DMO websites is systematically lacking. The diagnostic audit described in Section 4 confirmed that Neltur's primary digital presence employed minimal Schema.org markup, with no TouristDestination type, no entity relationships connecting Niterói to its attractions, and no sameAs links establishing identity consistency with Wikipedia, Wikidata, and major travel platforms.

3. Territorial Algorithmic Invisibility and the Proximity Shadow Effect

3.1 Defining Territorial Algorithmic Invisibility

Territorial Algorithmic Invisibility (TAI) describes the condition in which a geographic destination of demonstrable real-world tourism value receives absent, minimal, or systematically misattributed representation in LLM-generated travel recommendations. TAI is distinct from simple obscurity: an unknown destination in a remote region may receive no LLM representation simply because it has no content presence. TAI specifically refers to destinations with existing tourism infrastructure, official marketing organizations, and documented visitor interest that nonetheless fail to achieve algorithmic visibility.

TAI arises from the structural mismatch between how destinations create and manage their digital presence — through traditional DMO marketing channels — and how LLMs assess entity authority: through training data density, cross-platform consistency, and structured knowledge graph representation.

3.2 The Proximity Shadow Effect

The Proximity Shadow Effect describes a specific mechanism of TAI in which a locally prominent destination is consistently subsumed into the algorithmic identity of a nearby dominant destination. In LLM responses, this manifests as three distinct patterns:

- **Attractions physically located in the smaller destination are attributed to the larger one (e.g., Niterói's MAC described as a "Rio de Janeiro landmark").** Misattribution:
- **The smaller destination is represented only as a day-trip or excursion from the dominant one, suppressing its identity as an independent destination.** Categorization as appendage:
- **When a traveler asks about the broader region, the dominant destination monopolizes the response, and the smaller one is not mentioned.** Omission under shared queries:

The Proximity Shadow Effect is not unique to Niterói. Analogous patterns are observable between Florence and the broader Tuscany, between Bruges and Brussels, and between Cartagena and the Colombian Caribbean coast — wherever a destination of genuine merit exists in the geographic and conceptual shadow of a dominant neighbor.

3.3 Why Standard GEO Frameworks Are Insufficient for Geographic Entities

Existing GEO frameworks, including the GEO Enterprise Framework (Brasil, 2026) and the original Princeton GEO strategies (Aggarwal et al., 2023), were designed for brand entities or individual knowledge producers. These frameworks assume a well-defined entity with clear boundaries, a consistent name, and a single canonical identity across platforms.

Geographic destinations violate all three assumptions. Boundary ambiguity is endemic: a destination's geographic extent is contested (are the beaches of the greater Niterói region part of Niterói's destination identity or separate entities?). Name instability is common: destinations may be known by multiple names across platforms and languages (Niterói, "across the bay," "the other side of Rio"). Multi-entity composition is structural: a destination's algorithmic identity is distributed across dozens of constituent entities — attractions, neighborhoods, events, cultural figures — each requiring its own GEO treatment and coordination.

These characteristics require a destination-specific extension of existing GEO frameworks, which this paper provides through the DARS model.

4. The Destination Algorithmic Readiness Score (DARS)

4.1 Framework Overview

The Destination Algorithmic Readiness Score is a five-dimensional diagnostic metric for assessing a DMO's current GEO baseline. Each dimension is scored on a 0–20 scale, yielding a composite DARS of 0–100. The five dimensions address the full spectrum of factors that determine a destination's algorithmic visibility.

Table 1. DARS Dimensions and Components

Dimension	Components	Max Score
D1: Entity Definition	Canonical entity markup, sameAs links, disambiguation	20
D2: Structured Data Coverage	Schema.org types implemented, validation status	20
D3: Cross-Platform Consistency	Identity alignment across Wikipedia, Google, social, booking platforms	20
D4: Content Authority	Volume, citability, and topical consistency of published content	20
D5: Conversational Presence	LLM citation frequency and accuracy across major platforms	20

4.2 Dimension 1: Entity Definition

Entity Definition assesses whether the destination is correctly defined as a machine-readable entity with unambiguous identity across the structured web. Components include: presence of a TouristDestination Schema.org type on the primary DMO website (0–5); sameAs links connecting the canonical entity to Wikipedia, Wikidata, Google Knowledge Graph, and major travel platforms (0–7); canonical name consistency across all markup (0–4); and geographic disambiguation markup via GeoCoordinates and containedInPlace (0–4).

4.3 Dimension 2: Structured Data Coverage

Structured Data Coverage measures the breadth and quality of Schema.org implementation across the destination's primary digital presence. Components include: TouristAttraction markup for key attractions (0–6); Event and LandmarksOrHistoricalBuildings markup (0–5); FAQPage and Speakable markup for priority content (0–5); and validation score via Google Rich Results Test and Schema.org validator (0–4).

4.4 Dimension 3: Cross-Platform Consistency

Cross-Platform Consistency measures the alignment of the destination's identity representation across platforms that constitute LLMs' training and retrieval data. This dimension adapts the Entity Consistency Score (Brasil, 2026) to the multi-entity structure of geographic destinations. Components include: Wikipedia coverage quality and factual accuracy (0–7); Google Business Profile completeness for key attractions (0–5); social platform identity consistency (0–4); and booking and travel platform description alignment (0–4).

4.5 Dimension 4: Content Authority

Content Authority assesses the volume, quality, and citability of published content that LLMs can incorporate into generated responses. Unlike brand entities, destinations must build content authority across multiple thematic clusters covering culture, gastronomy, accommodation, and events. Components include: volume of authoritative published content (0–5); presence of datable, verifiable claims (0–5); cross-platform editorial presence (0–5); and citation by media sources within LLM training data (0–5).

4.6 Dimension 5: Conversational Presence

Conversational Presence is measured through systematic querying of major LLM platforms using a standardized query battery. The battery tests five categories: direct destination queries, regional travel recommendations, attraction-specific queries, itinerary planning queries, and comparison queries. Each query is scored across three dimensions: (a) whether the destination is mentioned, (b) whether it is accurately attributed, and (c) whether the representation is complete or partial.

5. Methodology

5.1 Research Approach

This study employs a diagnostic case study methodology grounded in action research principles (Lewin, 1946). The researcher served as both the diagnostic practitioner and the framework developer, applying an iteratively refined assessment protocol to a real destination entity. The dual role is consistent with practice-based research in nascent disciplines where standardized benchmarks do not yet exist.

5.2 Case Selection: Neltur and Niterói

Niterói was selected as the diagnostic case for three reasons. First, it represents a prototypical TAI scenario: a destination with internationally recognized attractions, an official DMO — Neltur, the Empresa de Turismo de Niterói — and a chronic visibility gap relative to its proximity to one of the world's most recognized tourist destinations. Second, the author's professional engagement with Neltur as a GEO and digital brand architecture consultant provided access to primary diagnostic data not available through external audit alone. Third, Niterói's architectural heritage, particularly the Oscar Niemeyer-designed Museu de Arte Contemporânea (MAC), the Caminho Niemeyer cultural complex, and the Teatro Popular, provides a compelling test case for the misattribution hypothesis central to the Proximity Shadow Effect.

Neltur is the official public company responsible for tourism promotion and development in Niterói, operating under the municipal government. The city has a population of approximately 520,000 and received an estimated 4.2 million visitor-days in 2023, primarily from domestic tourism. Despite this volume, Niterói's presence in international travel media and in LLM-generated travel recommendations is dramatically disproportionate to its real-world tourism significance.

5.3 Data Collection

Data was collected through four methods:

1. **A systematic review of Neltur's primary digital properties (neltur.com.br) assessing structured data implementation, content architecture, and entity markup.** Structured website audit:
2. **Manual review of Niterói's representation across Wikipedia (Portuguese and English), Google Knowledge Graph, TripAdvisor, Booking.com, Instagram, and LinkedIn.** Cross-platform identity audit:

3. **A battery of 30 standardized queries was administered to ChatGPT (GPT-4o), Google Gemini Advanced, and Perplexity AI between April and May 2026. Queries covered the five categories defined in the DARS D5 dimension. Responses were coded by two independent reviewers for mention presence, attribution accuracy, and representation completeness.** LLM query testing:
4. **Niterói's query results were compared to those for Rio de Janeiro across the same battery to quantify the representation gap and characterize the Proximity Shadow Effect.** Comparative analysis:

5.4 Scoring Protocol

DARS scores were assigned using the dimensional rubrics defined in Section 4, supplemented by secondary review to mitigate practitioner bias. Scores were calculated at the dimension level and aggregated to produce the composite DARS. A conflict of interest exists due to the researcher's consulting relationship with Neltur; this is disclosed in full in Section 8.

6. Findings

6.1 DARS Baseline: Niterói

The diagnostic assessment produced a composite DARS of 38/100 for Niterói as of May 2026. The breakdown by dimension is presented in Table 2.

Table 2. DARS Baseline Scores — Niterói (May 2026)

Dimension	Score	Max	Notes
D1: Entity Definition	6	20	No TouristDestination markup; Wikipedia sameAs absent from DMO site
D2: Structured Data	5	20	Minimal Schema.org; no TouristAttraction types for key landmarks
D3: Cross-Platform Consistency	10	20	Google Knowledge Panel exists; Wikipedia adequate but thin in English
D4: Content Authority	9	20	Domestic media presence moderate; international citations sparse
D5: Conversational Presence	8	20	Mentioned in 60% of regional queries; misattributed in 28% of those mentions
Composite DARS	**38**	**100**	**Moderate-low readiness; Entity Definition and Structured Data are critical gaps**

The composite score of 38/100 establishes the diagnostic baseline. The most critical gaps are in the technical dimensions (D1 and D2), which are entirely within the DMO's control to remediate and represent the highest-leverage improvement opportunity. The relatively stronger scores in D3 and D4 reflect Niterói's existing domestic media presence and basic knowledge graph entry, which provide a foundation for the remediation strategy outlined in Section 7.

6.2 LLM Query Testing Results

The 30-query battery revealed consistent patterns across all three LLM platforms tested. Table 3 presents aggregated results across ChatGPT, Gemini, and Perplexity.

Table 3. LLM Query Testing Results (30-Query Battery, Three Platforms)

Query Category	Queries	Niterói Mentioned	Correctly Attributed	Misattributed to Rio
Direct destination	6	6 (100%)	5 (83%)	1 (17%)
Regional recommendations	8	4 (50%)	2 (50%)	2 (50%)
Attraction-specific	6	5 (83%)	3 (60%)	2 (40%)
Itinerary planning	6	2 (33%)	2 (100%)	0 (0%)
Comparison queries	4	1 (25%)	1 (100%)	0 (0%)
Total	**30**	**18 (60%)**	**13 (72% of mentions)**	**5 (28% of mentions)**

The results reveal that Niterói performs acceptably when directly queried but struggles significantly in the query categories that drive real traveler behavior: regional recommendations (50% mention rate) and itinerary planning (33% mention rate). These are precisely the queries where the Proximity Shadow Effect is most active — queries that naturally favor Rio de Janeiro as the primary entity, with Niterói competing for mention as a secondary destination.

The misattribution rate of 28% across all mentions is particularly significant. In 28% of cases where Niterói was mentioned, one or more of its attractions were attributed to Rio de Janeiro. The MAC, in particular, was described as a "Rio de Janeiro landmark" in 3 of 5 attraction-specific queries where it appeared — a finding that directly demonstrates the Proximity Shadow Effect operating at the individual attraction level.

6.3 The Proximity Shadow Effect: Quantitative Evidence

To quantify the Proximity Shadow Effect, the same 30-query battery was administered substituting Rio de Janeiro as the subject wherever Niterói appeared. The contrast between the two entities' LLM representations is presented in Table 4.

Table 4. Comparative LLM Presence: Rio de Janeiro vs. Niterói

Metric	Rio de Janeiro	Niterói
Overall mention rate	100%	60%
Regional recommendation inclusion	100%	50%

Metric	Rio de Janeiro	Niterói
Accurate attribution rate	97%	72%
Itinerary planning inclusion	100%	33%
Entity described as independent destination	100%	41%
Entity described as appendage of other	0%	59%

The most diagnostic figure is the final row: in 59% of queries where Niterói was mentioned, it was framed as an appendage of Rio de Janeiro rather than as an independent destination. This framing directly suppresses the city's algorithmic identity and is a defining characteristic of the Proximity Shadow Effect in operation.

6.4 Technical Audit Findings

The structured audit of Neltur's primary digital presence confirmed the technical gaps reflected in the D1 and D2 scores. No TouristDestination Schema.org type was implemented on neltur.com.br. No sameAs links connected the website entity to Wikipedia, Wikidata, or Google Knowledge Graph entries. The MAC, Caminho Niemeyer, and Teatro Popular had no TouristAttraction markup on the DMO site. No llms.txt file was present. No FAQPage schema existed for any content. Page-level structured data was limited to basic WebPage markup without entity relationships.

These findings are consistent with the UNWTO (2024) observation that structured data implementation remains systematically underdeveloped in regional DMOs, and they directly explain the low D1 and D2 scores that anchor the DARS baseline.

7. Discussion

7.1 The TAI-DARS Relationship

The Niterói diagnostic establishes an empirical relationship between the TAI condition and DARS dimension scores. Destinations with composite DARS scores below 50 — particularly those with low D1 and D2 scores — are structurally vulnerable to Territorial Algorithmic Invisibility because they lack the entity definition and structured data infrastructure that enables LLMs to identify, distinguish, and attribute them accurately.

The D5 score of 8/20 is a direct downstream consequence of the D1 score of 6/20. Without machine-readable entity definition, LLMs cannot reliably distinguish Niterói from the broader Rio de Janeiro entity. The diagnostic framework thus reveals a causal chain: Entity Definition drives Structured Data Coverage, which enables Cross-Platform Consistency, which builds Content Authority, which ultimately determines Conversational Presence.

7.2 Remediation Priorities

The DARS framework enables prioritized remediation. For Niterói, the intervention hierarchy follows the causal chain identified above:

5. **Implement TouristDestination markup with sameAs links to Wikipedia, Wikidata, and Google Knowledge Graph. This creates the foundational entity disambiguation that enables all subsequent improvements.** Entity Definition (D1):
6. **Deploy TouristAttraction markup for MAC, Caminho Niemeyer, Teatro Popular, and the main beach areas. Implement FAQPage schema for high-traffic content. Add llms.txt.** Structured Data Coverage (D2):
7. **Expand and standardize English-language Wikipedia content. Align Google Business Profile descriptions across key attractions with canonical entity language.** Cross-Platform Consistency (D3):
8. **Develop a cross-platform editorial strategy targeting platforms within the LLM citation distribution — LinkedIn, Medium, YouTube — with content that explicitly positions Niterói as an independent destination, not a satellite of Rio.** Content Authority (D4):
9. **Monitor LLM citation behavior quarterly using the standardized query battery, tracking D5 improvement as a lagging indicator of the combined effect of D1-D4 interventions.** Conversational Presence (D5):

7.3 Implications for the B2A Paradigm in Tourism

The Business-to-Agent (B2A) paradigm (Brasil, 2026) acquires specific urgency in the tourism context. When an AI assistant recommends a travel itinerary, it is not merely providing information — it is effectively making a purchasing decision on behalf of the traveler. The economic value of an LLM recommendation in tourism, potentially representing thousands of dollars in bookings, dwarfs the value of a traditional search ranking click.

For DMOs, this means that GEO is not an optimization exercise but a structural imperative. A destination algorithmically invisible to AI travel assistants will progressively lose mindshare among the growing cohort of travelers who outsource destination discovery to these tools. The DARS framework provides the diagnostic foundation for addressing this risk before it becomes irreversible.

7.4 Extending the Invisible Excellence Paradox to Tourism

The Invisible Excellence Paradox (Brasil, 2026) — technically optimized entities receiving zero citation due to insufficient Editorial Authority — manifests in tourism with a structural twist. Niterói is not technically optimized (D1 and D2 scores are low), but it has real-world Editorial Authority in the form of international architectural recognition, UNESCO-adjacent cultural heritage, and established domestic tourism flows. This real-world authority is not algorithmically accessible because it has not been translated into structured, machine-readable form.

For tourism entities, the paradox might be restated: a destination with real-world excellence remains algorithmically invisible when that excellence is not encoded in the technical and semantic infrastructure that LLMs require to recognize and cite it. The DARS framework identifies precisely where this translation gap exists and provides a reproducible methodology for closing it.

8. Limitations and Future Work

8.1 Limitations

Several limitations constrain the generalizability of this study.

Single-case design. The diagnostic was applied to a single destination in a specific geographic context — a Brazilian coastal city adjacent to a major metropolis. Generalizability to other destination types (landlocked cities, island destinations, rural tourism areas) requires validation.

LLM opacity. The citation behavior of LLMs remains partially opaque. The misattribution and omission patterns observed in the query battery may reflect factors not captured by the DARS framework, including training data cutoff dates and retrieval system idiosyncrasies.

DARS calibration. The equal weighting of dimensions and equal distribution of points within dimensions represent an initial calibration that requires empirical validation. Future research should test whether certain dimensions have greater predictive validity for D5 outcomes.

Measurement timing. LLM responses are probabilistic and vary across sessions. The query battery was administered over a defined period; subsequent re-testing may yield different results as model versions update.

Conflict of interest. The researcher's role as a GEO consultant to Neltur creates a potential conflict that readers should consider when interpreting the findings. This dual role is inherent to action research methodology and is disclosed in the interest of transparency.

8.2 Future Work

The DARS framework opens several research directions:

10. **Applying DARS to 10–20 destinations across different geographic contexts and destination types to test generalizability and refine dimension weights.** Multi-destination validation:
11. **Measuring DARS score progression following structured remediation interventions to establish the relationship between specific technical improvements and D5 (Conversational Presence) outcomes.** Longitudinal tracking:
12. **Classifying the effect across different destination pairs to identify the geographic, cultural, and algorithmic conditions under which shadow effects are most pronounced.** Proximity Shadow Effect taxonomy:
13. **Developing a methodology for estimating the economic value of an LLM travel recommendation relative to traditional search results, enabling DMOs to prioritize GEO investment rationally.** Economic valuation of LLM citation:
14. **Extending the framework to assess destination representation across LLM platforms in multiple languages, as the TAI condition may vary significantly between language ecosystems.** Multilingual DARS:

9. Conclusion

The transition from search-mediated to AI-mediated travel information retrieval creates a new class of vulnerability for Destination Marketing Organizations: Territorial Algorithmic Invisibility. Destinations that have invested in traditional SEO, social media marketing, and print tourism promotion may find themselves algorithmically nonexistent to the growing cohort of travelers who initiate trip planning through conversational AI.

This paper introduced the Destination Algorithmic Readiness Score as a diagnostic framework for quantifying this vulnerability and identifying the highest-leverage remediation pathways. Applied to Neltur and Niterói, the framework produced a DARS baseline of 38/100, revealing critical gaps in Entity Definition and Structured Data Coverage that directly contribute to the Proximity Shadow Effect — the systematic misattribution and underselling of Niterói's tourism identity in LLM-generated responses.

The broader implication is that destination marketing in the AI era requires a fundamentally new technical competency: not content production alone, but entity architecture. DMOs must learn to translate their destinations' real-world excellence into the structured, consistent, machine-readable form that enables LLMs to recognize, attribute, and recommend them. The DARS framework provides the starting point for that translation.

The Proximity Shadow Effect is not a permanent condition. It is a diagnostic finding — and diagnostic findings have remediation paths.

References

Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K., & Deshpande, A. (2023). GEO: Generative Engine Optimization. arXiv preprint arXiv:2311.09735. Princeton University.

Brasil, A. C. (2026). Algorithmic authority: A practitioner framework for Generative Engine Optimization based on a 7-day implementation sprint. SSRN.

Fishkin, R. (2024). Zero-click search study: How much traffic does Google keep for itself? SparkToro Research Reports. Retrieved from https://sparktoro.com/blog/

Howard, J. (2024). llms.txt: A proposal for standardized LLM-readable website content. Retrieved from https://llmstxt.org/

Lewin, K. (1946). Action research and minority problems. Journal of Social Issues, 2(4), 34–46.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459–9474.

Phocuswright. (2024). AI and the future of travel planning: Consumer behavior study. Phocuswright Research Reports.

Semrush. (2026). LLM citation study: Which platforms do AI models cite most frequently? Semrush Research, January 2026. Retrieved from https://www.semrush.com/blog/

Singhal, A. (2012). Introducing the Knowledge Graph: Things, not strings. Google Official Blog. Retrieved from https://blog.google/

UNWTO. (2024). Digital transformation in destination marketing: Global benchmarking report. United Nations World Tourism Organization.

Data Availability Statement: The LLM query results and audit data referenced in this paper are available from the corresponding author upon reasonable request.

Conflict of Interest: The author is a GEO consultant to Neltur, the organization whose diagnostic is documented in this study. This dual role as practitioner and researcher is inherent to the action research methodology employed and is disclosed here for full transparency.

Funding: This research received no external funding.

Manuscript word count: approximately 7,200 words (excluding tables and references).