RESEARCH SYNTHESIS · 2023–2026

The Disappearing Hotel

Generative Discovery, AI Visibility, and the Case for Generative Engine Optimization in Hospitality

Prepared by Victor Frenzel — AI Visibility & GEO for design-led hospitality

VICTORFRENZEL.COM · KÖLN · JUNE 2026 · INDEPENDENT SYNTHESIS OF INDUSTRY & PEER-REVIEWED RESEARCH

Abstract

Travel discovery is migrating from keyword search and online travel agency (OTA) browsing toward conversational, AI-mediated recommendation. As travelers increasingly open a large language model (LLM) and ask where to stay, generative engines answer by naming a handful of properties rather than returning a ranked list of links. This article synthesizes recent industry research and peer-reviewed and preprint scholarship (2023–2026) to characterize the resulting visibility problem. Survey evidence indicates that roughly four in ten travelers now use AI to research trips, while traditional search use has declined (Phocuswright, cited in PhocusWire, 2026). A large prompt-based audit found that across 4,545 ChatGPT prompts, only 2,721 unique hotels were ever named, with recommendations skewing heavily toward four- and five-star and chain properties (Lighthouse, 2026). These patterns expose independent, boutique, and design-led hotels disproportionately, because the third-party authority signals generative engines rely upon are concentrated among large brands. Generative Engine Optimization (GEO), formalized by Aggarwal et al. (2024), offers a partial response, with controlled experiments showing visibility gains of up to 40%. The discipline remains young, its measurement contested, and its long-run efficacy unproven. We outline mechanisms, exposures, tactics, and open questions for hospitality decision-makers.

1. Introduction

For two decades, a hotel's commercial visibility was governed by two systems it could partly influence: search engine results and OTA marketplaces. A new intermediary is now inserting itself ahead of both. Travelers are beginning to plan trips by asking an AI assistant a question and accepting a short, synthesized answer rather than scanning a page of links. This shift matters because of a structural difference in the output: a search engine returns many results and lets the user choose, whereas a generative engine returns a curated few and makes the choice on the user's behalf. The consideration set shrinks from a page to a sentence.

The consequence for hotels is stark. If a property is not named in that synthesized answer, it is not merely ranked lower; it is functionally absent from the traveler's awareness. This article examines the evidence for that claim, the mechanisms that drive which hotels appear, the specific exposure of independent and design-led properties, and GEO as the emerging strategic response. It treats GEO as a nascent and unsettled discipline rather than a proven solution.

2. Background: from SEO and OTAs to conversational discovery

The dominant pre-AI architecture of hotel discovery rested on search engine optimization (SEO) and OTA intermediation. Roughly 80% of global travelers used OTAs such as Booking.com and Expedia in the lead-up to a trip (APH, 2025), and hotels competed for organic and paid placement on Google. That system rewarded link-based authority, keyword relevance, and OTA listing quality.

Generative engines operate on a different principle. Aggarwal et al. (2024) formalize a *generative engine* (GE) as a system that uses generative models to “gather and summarize information to answer user queries,” synthesizing across multiple retrieved sources rather than returning them as a list. ChatGPT, Google’s AI Overviews and AI Mode, Gemini, and Perplexity all fit this description. The commercial migration into travel is now explicit: in late 2025, OpenAI began integrating third-party apps into ChatGPT conversations, with Booking.com and Expedia among the early pilot partners, allowing users to obtain travel results without leaving the chat; Tripadvisor later joined and also partnered with Perplexity (Statista, 2026).

The user-behavior effect of synthesis is measurable. A Pew Research Center study tracking the browsing of 900 U.S. adults across 68,879 Google searches in March 2025 found that when an AI summary appeared, users clicked a result link only 8% of the time, versus 15% without a summary, and clicked links *within* the AI summary just 1% of the time (Pew Research Center, 2025). This “zero-click” dynamic is the mechanism by which visibility, not ranking, becomes the decisive variable.

3. The scale of the shift

Adoption figures should be read with care: they come largely from commissioned industry surveys with differing samples and definitions, not from a common instrument. With that caveat, the direction is consistent across providers.

Indicator	Figure	Source (year)
Travelers using AI to research/plan travel (U.S.)	39%	Phocuswright, cited in PhocusWire (2026)
Traditional search use for travel (U.S.), late 2024 → H2 2025	51% → 36%	Phocuswright, cited in PhocusWire (2026)
Travelers using AI in some aspect of travel	67%	Booking.com (2025)
Consumers wanting to use AI in future travel planning	89%	Booking.com (2025)
Comfortable letting AI suggest travel options (UK/US/India)	53%	Expedia Group (2026)
Would not trust an AI assistant to book for them	66%	Expedia Group (2026)
Comfortable booking through an AI platform	8%	Expedia Group (2026)

Two patterns emerge. First, AI has become a mainstream *discovery and planning* tool: Booking.com’s Global AI Sentiment Report, based on 37,325 respondents across 33 markets surveyed in April–May 2025, found that 67% of travelers already use AI in some aspect of their travels and that AI assistants (24%) are now trusted more than travel bloggers (19%) or social media influencers (14%) (Booking.com, 2025). Second, a pronounced *trust gap* persists at the point of transaction: Expedia Group's survey of more than 5,700 adults in the U.S., U.K., and India found that while 53% are comfortable letting AI suggest options, 66% would not trust an AI assistant to book on their behalf, and only 8% are comfortable booking through an AI platform (Expedia Group, 2026).

The practical implication is that AI is reshaping the *top* of the funnel — inspiration and shortlisting — more than the booking click itself. That is precisely the stage at which a property either enters or never enters the consideration set. The macro stakes are large: the World Travel & Tourism Council estimated travel and tourism's 2025 global GDP contribution at a record US$11.6 trillion, about 9.8% of the global economy (WTTC, cited in Hospitality Net, 2026). Forward-looking claims — for example, Gartner's projection of a 25% decline in traditional search volume by 2026 (cited in ZipTie, 2026) — are forecasts, not measured outcomes, and should be treated as such.

4. The visibility problem: most properties are invisible

The clearest evidence that generative discovery concentrates visibility comes from prompt-based audits. These are vendor studies rather than peer-reviewed work, and their methodologies and incentives warrant scrutiny; nonetheless their core findings are mutually corroborating.

Lighthouse, a hospitality intelligence platform, ran 4,545 distinct ChatGPT prompts across nine global destinations and five traveler personas. Across those prompts ChatGPT named hotels 49,707 times, but those mentions resolved to only 2,721 unique properties — the remainder were repeats of a small recommended core (Lighthouse, 2026). The study reported "market coverage rates" — the share of a market's hotels ever surfaced — of about 10% in Tokyo and 13% in Paris, with roughly two-thirds of hotels never mentioned even in a small market such as Park City. The 100 most-mentioned hotels worldwide accounted for over 13% of all mentions (Lighthouse, 2026).

A separate multi-platform study, conducted by LuxDirect across 25 London luxury hotels on six AI platforms (ChatGPT, Claude, Gemini, Grok, Perplexity, and Google AI Mode) using 2,700 standardized queries, found that five hotels captured 57% of all AI recommendations (LuxDirect, cited in Hospitality Net, 2026). The concentration is therefore observable across engines, not unique to one model.

Crucially, the bias is toward scale and luxury rather than guest satisfaction. Lighthouse reported only a weak correlation between a hotel's Booking.com guest-review score and its ChatGPT visibility; star rating was a far stronger predictor (Lighthouse, 2026). Even on generic prompts — "I need a hotel in X," with no budget or preference stated — four- and five-star properties dominated, and three-star hotels were nearly absent. Business-travel prompts returned 83% four- and five-star recommendations and family prompts 73%, despite mid-tier hotels being appropriate for many such stays (Lighthouse, 2026). On the chain side, more than a quarter of branded U.S. hotel mentions went to Marriott, and the top three brand families accounted for 54% of branded mentions (Lighthouse, 2026).

This concentration is not merely an industry observation; it is reproduced under experimental control. An algorithm audit using a randomized conjoint design across a panel of LLMs found that chain affiliation increased the probability of recommendation, consistent with a “brand-as-risk-reduction” heuristic, and documented systematic position and reputation-signal effects in the selection stage (Ulrich et al., 2026, preprint). Together the audit and vendor studies support a single conclusion: invisibility is the default state for most hotels, and a property must actively earn its way out of it.

5. Mechanisms: how generative engines select, trust, and cite

Understanding the bias requires understanding how generative engines assemble an answer. Three mechanisms are central.

Source provenance. Generative engines ground recommendations in retrieved third-party content. Lighthouse found that 82% of the sources ChatGPT cross-referenced when forming hotel recommendations fell into two categories: OTA and metasearch platforms (Booking.com, Expedia, and similar) and editorial or media sites (Forbes, Lonely Planet, Condé Nast Traveler, and equivalents) (Lighthouse, 2026). A hotel’s presence in “best of” lists and editorial features therefore functions as a distribution input, not merely brand-building. The same dynamic applies to Google AI Overviews, where Pew found Wikipedia, YouTube, and Reddit among the most-cited sources overall (Pew Research Center, 2025).

Content characteristics rewarded by GEO experiments. Aggarwal et al. (2024) tested nine content modifications against a generative-engine benchmark (GEO-bench) and found that adding verifiable statistics, incorporating credible quotations, and citing reliable sources produced the largest visibility gains — on the order of 30–40% on a position-adjusted word-count metric — while improving fluency and readability yielded smaller but significant gains of roughly 15–30%. Notably, keyword stuffing, the legacy SEO tactic, produced negligible or negative effects, and a merely “authoritative” tone did not help (Aggarwal et al., 2024). This implies generative engines reward content that is credible, evidenced, and clearly written — not content that is merely keyword-dense.

Entity and structured signals. Because engines retrieve and reconcile structured data about entities, consistent and machine-readable property information (across one’s own site, OTA listings, and authoritative directories) reduces ambiguity about what a property is. Lighthouse explicitly advises auditing OTA and metasearch listings because these are “what AI uses to verify what it already knows,” and stale or inconsistent listings degrade the input (Lighthouse, 2026). The convergence of these mechanisms explains the chain-and-luxury skew: large brands already possess dense editorial coverage, voluminous structured data, and consistent cross-platform listings.

6. Why independent and design-led properties are disproportionately exposed

The exposure of independent and boutique hotels follows directly from the mechanisms above. Each of the signals that generative engines weight most heavily is one in which large chains hold a structural advantage.

First, editorial and “best of” authority is unevenly distributed. Established luxury brands accumulate press coverage as a function of marketing budget and longevity; a newly opened design-led property may be architecturally distinctive yet thinly covered. Second, structured-data and listing consistency are easier to maintain at scale, where brands deploy central distribution teams; an independent owner-operator may have inconsistent representations across channels. Third, the engines’ demonstrated star-rating and chain heuristics directly penalize the three- and four-star independent segment even when guest satisfaction is high — recall that review scores correlated only weakly with visibility (Lighthouse, 2026).

The evidence does not, however, support fatalism. The London study found that a 26-room independent property appeared more frequently than a 174-room chain hotel of the same quality tier in the same city, and Paris was the one audited market where independents collectively exceeded chains in share of mentions (LuxDirect, cited in Hospitality Net, 2026; Lighthouse, 2026). Visibility is patterned and concentrated, but it is not strictly a function of size — which is the opening that GEO seeks to exploit.

This segment is also commercially material. Estimates of the independent and boutique market vary widely by vendor and definition — the global boutique hotel market alone was valued at roughly US$26.7 billion in 2024 (Grand View Research, 2025) — and such figures should be read as indicative rather than precise. The relevant point is qualitative: the properties most exposed to AI invisibility are concentrated in exactly the experiential, design-led category that travelers increasingly say they want.

7. Generative Engine Optimization (GEO)

Definition and origin. GEO is the practice of optimizing content to increase its visibility in generative-engine responses, as distinct from SEO’s goal of ranking in link-based results. The term and the first systematic framework originate with Aggarwal, Murahari, Rajpurohit, Kalyan, Narasimhan, and Deshpande, whose paper “GEO: Generative Engine Optimization” was first posted to arXiv in November 2023 and published at the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD ‘24) in 2024 (Aggarwal et al., 2024). The authors introduced GEO-bench, a benchmark of diverse queries with associated sources, and a black-box optimization framework for defining and improving visibility metrics.

Core tactics and measured effect. The paper’s headline finding is that GEO methods can boost visibility “by up to 40%” in generative-engine responses, with efficacy varying by domain (Aggarwal et al., 2024). The most effective tested tactics were Cite Sources, Statistics Addition, and Quotation Addition; fluency and readability improvements helped moderately; keyword stuffing did not. A finding of particular relevance to independents is that lower-ranked sources benefited more from GEO than already-dominant ones — suggesting the technique can partially offset incumbency. Independent replications have reported directionally similar results: one audit of generative search visibility reported relative improvements of roughly +40% (Cite Sources), +37% (Statistics Addition), and +22% (Quotation Addition), with keyword stuffing near zero (Roumeliotis et al., 2026, preprint).

Translation to hospitality. Applied to hotels, GEO implies several concrete moves, each grounded in the mechanisms of Section 5. Lighthouse's practitioner recommendations align closely with the academic findings: treat PR and editorial coverage as a distribution channel; audit OTA and metasearch listings for accuracy and consistency; align on-site and listing language with the target guest segment, because that language teaches the engine which travelers to route to a property; and measure current AI visibility directly rather than assuming it (Lighthouse, 2026). The throughline is that descriptive content has become "algorithm input," not only marketing copy.

8. Limitations, measurement challenges, and open questions

GEO should be presented as an emerging, not settled, discipline, for several reasons.

Evidence quality is uneven. The foundational experimental result (Aggarwal et al., 2024) is peer-reviewed, but much of the hospitality-specific evidence comes from vendor studies whose methods are not fully transparent and whose authors sell related products (Lighthouse, 2026; LuxDirect, cited in Hospitality Net, 2026). Their findings are mutually corroborating and broadly consistent with controlled audits, which lends them credibility, but they are not independent peer-reviewed research and may carry commercial incentives.

The target is non-stationary. Generative engines change frequently and opaquely. Recommendation patterns differ across models — one audit noted Perplexity skews toward luxury brands while other engines diverge (Hotelrank, 2026, vendor analysis) — and a tactic that works on one engine or version may not generalize or persist. Unlike Google's relatively legible ranking signals, the criteria driving AI recommendations are difficult to audit, so hotels often cannot determine why they do or do not appear (LuxDirect, cited in Hospitality Net, 2026).

Measurement is unsettled. There is no standard, the way "search rank" was standardized. Visibility is being operationalized variously as share of voice, market coverage, mention counts, and citation share, and these metrics are not directly comparable across studies. Outcome attribution is harder still, because the booking may occur on a later channel.

Manipulation and integrity risks. The same research line that established GEO has documented adversarial methods — prompt injection, content "hijacking," and poisoning attacks — designed to manipulate generative-engine outputs (e.g., methods discussed in Aggarwal et al., 2024, and subsequent work). This raises the prospect of an arms race and of platform countermeasures that could neutralize current tactics, as happened with manipulative SEO.

Open questions therefore include: How durable are GEO gains as models update? Do visibility gains convert to bookings and revenue, given the documented trust gap at the point of purchase (Expedia Group, 2026)? Will engines move toward transactional integration (the Booking.com/Expedia ChatGPT pilots) in ways that re-advantage large OTA-partnered brands? And how will regulators treat algorithmic gatekeeping in discovery, an issue already surfacing in publisher antitrust actions over AI Overviews (ALM Corp, 2026)?

9. Conclusion and practical implications

The shift from search to generative discovery changes the unit of competition from rank to mention. The evidence assembled here — survey data showing AI is now a mainstream planning tool, behavioral data showing AI summaries suppress click-through, and audit data showing recommendations concentrate on a small core of large and luxury properties — points to a consistent conclusion: most hotels are currently invisible to AI recommendation, and the properties most exposed are independent, boutique, and design-led ones whose third-party authority signals are thin relative to chains.

For hospitality decision-makers, several implications follow, each appropriately hedged. The window for early action appears genuinely open but is unlikely to remain so, because machine-learning systems tend to reinforce existing recommendation patterns (Lighthouse, 2026). Practitioners can reasonably act now on the lowest-regret measures, which align across academic and industry evidence: ensuring listing and structured-data accuracy across channels; pursuing credible editorial and “best of” coverage as a distribution input; writing property descriptions in clear, evidenced, segment-aligned language; and establishing direct measurement of AI visibility against a competitive set. These steps are defensible even if specific GEO tactics prove transient, because they improve the underlying signals on which any future engine is likely to rely. What practitioners should *not* do is treat a 40% visibility figure, drawn from a 2024 benchmark in non-travel domains, as a guaranteed hospitality outcome. GEO is a promising and partially validated response to a real and measurable problem — but it is an early-stage discipline operating against a moving, opaque target, and it should be resourced as an experiment with measurement attached, not as a solved channel.

References

ALM Corp. (2026, February 18). *Google’s AI Overviews and publisher traffic: How antitrust filing reveals 58% click decline and the breakdown of the web ecosystem.*

Retrieved June 17, 2026, from https://almcorp.com/blog/google-ai-overviews-publisher-traffic-decline-antitrust-lawsuit-analysis/

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). Association for Computing Machinery. https://doi.org/10.1145/3637528.3671900 (Preprint: arXiv: 2311.09735, posted November 16, 2023.)

APH. (2025, December 18). *Travel booking statistics — 2025 report and trends.* Airport Parking and Hotels. Retrieved June 17, 2026, from https://www.aph.com/community/holidays/travel-booking-statistics/

Booking.com. (2025, July 23). *Booking.com releases the Global AI Sentiment Report.* Retrieved June 17, 2026, from https://news.booking.com/bookingcom-releases-the-global-ai-sentiment-report/

Expedia Group. (2026, April). *AI Trust Gap report* [Survey of 5,700+ adults in the U.K., U.S., and India]. As reported in PhocusWire (2026, April 14), *Losing control, data privacy top concerns around AI travel booking.* Retrieved June 17, 2026, from https://www.phocuswire.com/news/online/ai-trust-discovery-planning-expedia-dune7

Grand View Research. (2025). *Boutique hotel market size & share | Industry report, 2030.* Retrieved June 17, 2026, from https://www.grandviewresearch.com/industry-analysis/boutique-hotel-market-report

Hospitality Net. (2026, April 16). *Travelers use AI to dream, trusted brands to book; travel hit $11.6 trillion in 2025; SiteMinder connects 53,000 hotels to AI platforms.* Retrieved June 17, 2026, from https://www.hospitalitynet.org/editorial/4131925/

Hotelrank. (2026, January 27). *AI hotel landscape 2026.* Retrieved June 17, 2026, from https://hotelrank.ai/research/ai-hotel-landscape-2026

Lighthouse. (2026, June 12). *The AI recommendation is the new battleground for hotel distribution* [Research presented at Luminate 2026 by B. Reiter]. Hospitality Net. Retrieved June 17, 2026, from https://www.hospitalitynet.org/news/4132964/

LuxDirect. (2026). *CS7 London study: Five London hotels capture 57% of AI recommendations.* As reported in Hospitality Net (2026, March 16). Retrieved June 17, 2026, from https://www.hospitalitynet.org/opinion/4131422/

Pew Research Center. (2025, July 22). *Google users are less likely to click on links when an AI summary appears in the results.* Retrieved June 17, 2026, from https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/

PhocusWire. (2026, April 14). *Losing control, data privacy top concerns around AI travel booking* [Reporting Phocuswright U.S. traveler research]. Retrieved June 17, 2026, from https://www.phocuswire.com/news/online/ai-trust-discovery-planning-expedia-dune7

Roumeliotis, K. (2026). *Algorithmic trust and compliance: Benchmarking brand notability for UK iGaming entities in generative search engines* [Preprint]. arXiv:2603.12282. Retrieved June 17, 2026, from https://arxiv.org/pdf/2603.12282

Statista. (2026). *Artificial intelligence (AI) use in travel and tourism — Statistics & facts.* Retrieved June 17, 2026, from https://www.statista.com/topics/10887/artificial-intelligence-ai-use-in-travel-and-tourism/

Ulrich, F. (2026). *Whose hotel does the AI recommend? An algorithm audit of reputation signals in LLM-assisted hotel selection* [Preprint]. arXiv:2606.16344. Retrieved June 17, 2026, from https://arxiv.org/html/2606.16344

ZipTie. (2026, May 16). *Click behavior in zero-click search: Why rankings no longer predict traffic (and what does).* Retrieved June 17, 2026, from https://ziptie.dev/blog/click-behavior-in-zero-click-search/

Author's note on sourcing. Statistics are attributed to the organization that produced the underlying data wherever possible, with the reporting outlet named separately when the primary release was not directly retrieved. Several hospitality-specific figures derive from vendor studies (Lighthouse, LuxDirect, Hotelrank) rather than peer-reviewed research; these are corroborated where possible against controlled audits (Ulrich, 2026; Roumeliotis, 2026) and the foundational GEO experiment (Aggarwal et al., 2024). Two author names for arXiv preprints (Ulrich; Roumeliotis) reflect the first-listed author as displayed at the time of retrieval and should be verified against the published versions before formal citation. Forecasts are labeled as such and distinguished from measured results.