

# Concept–Reference Emergence under Generative Search: An Observational Single–Case Study of GhostDrift Mathematical Institute

GhostDrift Mathematical Institute

Correspondence: <https://www.ghostdriftresearch.com/>

March 2026

## Abstract

Generative search engines synthesize information by retrieving, re-ranking and producing new content rather than simply listing hyperlinks. This shift challenges conventional visibility metrics: retrieved sources may not survive into the final answer, and user clicks decline when AI summaries appear. We document an observational single–case study of concept–reference emergence under generative search, focusing on GhostDrift Mathematical Institute (GMI). Despite disadvantageous initial conditions—recent incorporation, low public recognition, no peer-reviewed authority, limited human traffic and no mainstream coverage—GMI reported that its definitional pages were observed among referenced sources in generative search outputs and external media. We reconstruct a timeline of concept definitions, external mentions and AI-related observations; classify evidence into official documentation, third-party sources, self-published logs and author interpretation; and outline inclusion and exclusion criteria. Results show that AI-related queries far exceeded human sessions, that definitional pages were available prior to external citations and AI references, and that GMI’s concepts were observed among referenced sources in AI summaries. While causal inference is not possible, the case illustrates that clear concept definitions and stable hosting can coincide with observable reference signals in generative search, even without high traffic. This finding contributes an empirically grounded instance to ongoing discussions of generative engine optimization.

**Keywords:** generative search; generative engine optimization; concept referencing; case study; GhostDrift

## 1 Introduction

Generative search integrates information retrieval, re-ranking and answer synthesis to deliver consolidated responses to user queries. In this paradigm, a large language model consults multiple sources, organizes them internally and produces a synthesized answer with citations. Research on Search–Augmented Generative Engines (SAGE) notes that these systems bridge web–scale retrieval with generative capabilities and thus reshape how web content gains exposure[1]. Generative Engine Optimization (GEO) extends beyond

conventional search engine optimization (SEO) by addressing how pages are selected for inclusion in AI-generated answers[1].

Two empirical observations motivate this study. First, retrieval does not guarantee citation: a publicly reported analysis of generative search behaviour found that only about 15% of pages retrieved by ChatGPT were cited in final answers[6]. Second, user engagement shifts when AI summaries appear; a Pew Research study reported that users click on traditional search results in 15% of visits without an AI summary but only 8% when a summary is present, and clicks on links within the summary occur in just 1% of such visits[5]. These findings imply that being selected as a source in generative answers is valuable even if subsequent clicks are scarce.

This paper examines how reference visibility under generative search can emerge from an unusually weak starting position. In the GhostDrift case, public-facing activity began from a personal website and individual publishing channels in March 2025, before formal incorporation in February 2026. At the observation stage, the case still lacked peer-reviewed authority, major institutional backing, mainstream media exposure, and large-scale human traffic. This makes the case analytically useful because generative-search reference signals are harder to attribute to pre-existing authority.

We adopt an observational single-case study design. Rather than testing a hypothesis, we document patterns and assess whether they align with mechanisms discussed in GEO and generative search research. Following case study guidelines, we explicate our unit of analysis, case selection logic, evidence classification, inclusion criteria and limitations[9]. The goal is to provide a transparent empirical record that others can interpret in light of broader theory.

## 2 Background

### 2.1 Generative search and GEO

Search-augmented generative engines combine retrieval, re-ranking and generation; their integrated pipeline means that retrieved documents compete not only for ranking but also for selection in the synthesized answer[1]. Traditional SEO aims to improve ranking visibility, whereas GEO emphasizes the additional steps of being re-ranked and cited. The SAGEO Arena benchmark demonstrates that existing optimization methods often degrade performance in retrieval and re-ranking under realistic conditions, and that structural information (such as schema markup) can improve selection[1]. Google’s AI features documentation explains that AI Overviews and AI Mode may issue multiple related searches across subtopics—a process called query fan-out—to identify additional supporting pages[3]. Importantly, Google states that appearing in AI features requires only that a page meet the standard technical requirements for search; no special markup or schema is needed[3].

### 2.2 Reference signals versus clicks

Generative search devalues click metrics. Public analyses of generative search behaviour report that about 85% of pages retrieved by ChatGPT are never cited in final answers[6]. Pew Research demonstrated that when AI summaries appear, users are less likely to click on search results and more likely to end their browsing sessions; click-through rates drop from 15% to 8%, and clicks on links within the AI summary occur in only 1% of visits[5].

Therefore, in generative search environments, being referenced may provide exposure even without high click-through rates. This motivates the focus on concept-reference emergence rather than traffic.

## 2.3 Platform-level distinctions

Interpreting analytics requires distinguishing between different kinds of AI interactions. OpenAI’s documentation describes distinct crawlers: *OAI-SearchBot* is used to surface websites in ChatGPT’s search features, whereas *GPTBot* crawls content for training generative models[2]. Sites can allow one and disallow the other; allowing OAI-SearchBot affects search visibility but not training[2]. Wix’s analytics explain that *sessions from AI platforms* count when an AI bot includes a site in its response and a user clicks the link, whereas *user queries on AI* count whenever an AI bot scans a site in response to a question; the latter does not indicate that a link was shown or clicked[4]. Google’s AI features documentation clarifies that query fan-out may yield a broader set of supporting pages and that no special schema is required[3].

## 2.4 Case study methodology

Case study research is a systematic strategy, not merely descriptive storytelling. Yin emphasises that case studies require structured data collection and multiple sources of evidence to ensure validity and reliability[9]. Triangulation—cross-checking findings across independent sources—enhances credibility[9]. Researchers must state their assumptions and propositions explicitly and maintain a chain of evidence linking raw data to conclusions[9]. Our study follows these guidelines by classifying evidence tiers, defining inclusion and exclusion criteria, and acknowledging limitations.

# 3 Case Selection and Research Question

## 3.1 Unit of analysis

The unit of analysis is not the organisation as such but the emergence of concept-reference signals in generative search within a single low-recognition case. We focus on how GMI’s definitional pages and external references became associated with AI-generated summaries.

## 3.2 Case selection

GMI was selected because its initial conditions were unusually weak yet publicly traceable. The case began from personal-level publishing in March 2025 through a personal website and individual media channels, and only later became a formally incorporated institute in February 2026. At the time of observation, it still lacked peer-reviewed authority, major institutional affiliation, mainstream media backing, and high human traffic. For that reason, the case is analytically useful: observed reference signals are less easily explained by pre-existing brand or academic authority and therefore provide a clearer GEO-relevant case of concept-reference emergence.

### 3.3 Research question

Guided by these considerations, the study poses a single question: *How did concept-definition pages and external references become associated with observable generative-search reference signals in a case that began from personal-level publishing in March 2025 and still lacked peer-reviewed authority, major institutional backing, mainstream media exposure, and large-scale human traffic at the time of observation?*

## 4 Materials and Methods

### 4.1 Evidence classification and data sources

Following case study guidance, we classified evidence into four tiers before analysis. **Tier 1** consists of official platform documentation, including Google’s AI features guide explaining query fan-out and stating that no special markup is required[3], OpenAI’s crawler documentation distinguishing *OAI-SearchBot* from *GPTBot*[2], and Wix’s definitions of AI metrics[4]. **Tier 2** encompasses third-party external references, such as the Reinforz.ai and Artificial ALICE articles that define Ghost Drift and related concepts[8, 7]. **Tier 3** includes self-published observation logs, notably GMI’s Note and Zenn posts documenting publication dates, AI observations and analytics (e.g., the AEO log reporting 1,077 AI queries vs 306 human sessions)[14]. **Tier 4** denotes author interpretation—analytical inferences drawn from the preceding tiers. Strong claims in this paper rely on Tier 1 and Tier 2 evidence; Tier 3 evidence is treated as descriptive, and Tier 4 statements are clearly separated from factual reporting.

### 4.2 Inclusion and exclusion criteria

Materials were included if they were publicly accessible, timestamped or datable, and directly relevant to GMI’s concept definitions, external mentions or AI observations. Private server logs, unverifiable screenshots and records lacking stable timestamps were excluded from strong inference. Claims without at least one public source were marked as unverified and used only descriptively. Page-level traffic data were not included because raw logs were unavailable.

### 4.3 Public verifiability and reproducibility

The timeline of publications and external references can be independently verified via the cited sources. However, AI-generated summaries and platform metrics are time-sensitive; reproduction may yield different outputs as algorithms evolve. Our study is therefore partially reproducible: document-level events are verifiable, but platform outputs may not be.

### 4.4 Limits of causal interpretation

This study is observational; we do not infer causal effects. We cannot determine whether GMI’s actions caused AI referencing. Alternative explanations, such as algorithmic experimentation or random crawler behaviour, remain possible. Our aim is to document a pattern and assess its consistency with GEO theory.

## 4.5 Metric distinctions

To avoid conflating different signals, we distinguish metrics (Table 1). *Human sessions* reflect visits by people. *User queries on AI platforms* count when an AI bot scans the site in response to a question and may not correspond to any citation or click[4]. *Sessions from AI platforms* record when an AI bot includes a link and a user clicks it[4]. *Crawler activity* refers to bot hits recorded in server logs but does not imply inclusion or citation. *Observed AI summary reference* indicates that a page appeared among cited sources in an AI-generated answer; this requires manual observation and cannot be inferred from analytics alone.

Table 1: Distinction among metrics used in this study

Metric	Interpretation & limitations
Human sessions	Visits by human users; indicates traffic but not necessarily AI engagement.
User queries on AI platforms	Counts when an AI bot scans the site in response to a question; does <i>not</i> indicate that a link was shown or clicked[4].
Sessions from AI platforms	Counts when an AI bot includes the site in its response and a user clicks the link[4].
Crawler activity	Bot hits recorded in logs; may reflect indexing or training; does not imply appearance in search results.
Observed AI summary reference	Manual observation that a page appeared among cited sources in an AI-generated answer; time-sensitive and not captured by analytics.

## 5 Results

### 5.1 Initial conditions

Table 2 summarizes the initial conditions that make this case analytically important. Public-facing activity related to the case began from personal-level publishing in March 2025 through a personal website and individual channels, whereas formal incorporation of GMI occurred later on 10 February 2026[16]. At the observation stage, no peer-reviewed publications attributable to GMI were found in the examined public record, no major institutional backing or mainstream media exposure was identified, and a 30-day observation log reported only 306 human sessions against 1,077 AI-related queries[14]. These conditions matter methodologically because the observed reference signals are difficult to explain by pre-existing academic, institutional, or traffic authority.

### 5.2 Timeline of publications and references

We reconstructed the timeline of concept definitions, external mentions and AI observations (Appendix A). The timeline begins not with formal incorporation but with personal-level public publishing in March 2025, followed by later concept-definition pages, external

Table 2: Initial disadvantageous conditions for GMI

Condition	Publicly documented evidence
Personal-level start	Public-facing activity began from a personal website and individual publishing channels in March 2025 (to be cited with earliest public records).
Pre-corporate phase	The case existed publicly before formal incorporation; GMI was incorporated only on 2026-02-10, with registration number assigned on 2026-02-17[16].
Absence of peer-reviewed authority	No peer-reviewed publication attributable to GMI was found in the examined public record at the time of observation.
Absence of institutional/media backing	No major university affiliation, large research-institute backing, or mainstream media coverage was identified in the examined record.
Limited human traffic	Observation log reports 306 human sessions vs 1,077 AI queries over 30 days[14].

mentions, AI observations, and finally formal incorporation in February 2026. Key milestones include the early personal publishing phase in March 2025 (to be cited with the earliest public records), publication of the “Finite Closure” concept on 2025-12-06[10], definition of Ghost Drift and its recording as a theoretical framework on 2025-12-25[12], external articles by Artificial ALICE (2026-02-15)[7] and Reinforz.ai (2026-02-28)[8], and AI Overview observations recorded by GMI for Beacon Architecture on 2026-03-10[17]. We note that the definitional publications were available before both external citations and AI references, but this temporal ordering alone does not establish causality.

### 5.3 AI-side observations

GMI’s AEO log reports that user queries on AI platforms (1,077) exceeded human sessions (306) by a factor of approximately three[14]. According to Wix definitions, user queries count when bots scan pages and do not necessarily correspond to citations[4]. GMI noted that AI traffic concentrated on definitional pages, whereas human traffic was more evenly distributed; however, raw page-level data were not publicly released. A spike on 2026-01-14 remains unexplained, illustrating the difficulty of attributing causation. We interpret these observations as descriptive signals of AI interest rather than proof of prominence.

Figure 1 provides an illustrative screenshot of the Wix analytics panel reporting 1,077 AI-related queries and 306 human sessions. Consistent with our evidence policy, this figure is presented as descriptive self-published evidence rather than independently auditable proof.



Figure 1: AI Queries vs. Human Sessions in Wix Analytics. Illustrative screenshot of a self-published analytics panel reporting 1,077 AI-related queries and 306 human sessions over the observed period. This figure is used descriptively and not as independently auditable evidence.

## 5.4 External references

Table 3 summarises third-party articles that referenced GMI concepts. Artificial ALICE described Ghost Drift as “responsibility evaporation” and introduced Pre-decision Constraint and Post-hoc Impossibility[7]. Reinforz.ai’s article defined Ghost Drift as responsibility ambiguity in self-updating AI and proposed layered responsibility allocation[8]. These references indicate that GMI’s terminology entered broader discourse.

Table 3: Third-party articles referencing GMI concepts

Source	Date	Summary of reference
Artificial ALICE	2026-02-15	Defines Ghost Drift as responsibility evaporation and discusses Pre-decision Constraint and Post-hoc Impossibility[7].
Reinforz.ai	2026-02-28	Section “Ghost Drift problem and AI governance” describes responsibility ambiguity in self-updating AI and suggests three-layer responsibility design[8].

Figure 2 presents one illustrative example of third-party definition, showing that Ghost Drift was described externally in governance-related terms rather than only within GMI’s own publications.



Figure 2: Independent External Definition of Ghost Drift. Illustrative screenshot of an external article describing Ghost Drift in AI governance terms. This figure is included as an example of third-party conceptual uptake.

## 5.5 Evidence ledger

To transparently link claims to evidence, we compiled an evidence ledger (Appendix B). Each claim is associated with its source tier, public verifiability and permissible inference. For example, the claim that AI queries exceeded human sessions is supported by a Tier 3 self-published log; it is publicly viewable but not independently auditable, so we treat it as descriptive. Claims supported by Tier 1 or Tier 2 sources (e.g., definitions of metrics) are treated as more robust.

## 5.6 Page-level visibility

GMI reports that AI traffic focused on definitional pages, whereas narrative posts received less AI attention. Because raw data were not published, we cannot quantify this pattern. We therefore state only that page-level concentration is suggestive and may reflect the tendency of generative models to prioritise clear, authoritative sources.

# 6 Discussion

## 6.1 Consistency with GEO structure

The chronological ordering of events—definitional publication, external adoption, AI observation—mirrors the stages of retrieval, re-ranking and generation. GMI’s explicit concept definitions provided structured content that could be retrieved and ranked. The subsequent external articles acted as earned media, which generative models are known to favour. AI-related queries and AI summary references suggest that GMI’s pages entered the retrieval pool and occasionally survived into the final answer. This pattern is consistent with GEO theory but does not confirm causality.



## 6.2 Distinction from SEO outcomes

Unlike traditional SEO success, GMI’s case did not involve high human traffic or prominent ranking in standard search results. The absence of mainstream coverage and the limited number of human sessions indicate that conventional popularity was low. Instead, being cited in generative answers provided visibility. As Pew’s data show, user clicks drop sharply when AI summaries appear[5], but citation may still shape knowledge representation. The case illustrates that definitional clarity and external placements can create reference signals independent of click-based metrics.

## 6.3 Implications for small actors

For low-recognition entities, this case suggests that publishing well-structured concept definitions on stable platforms and ensuring that content is indexable may increase the likelihood of being referenced by AI systems. External validation through third-party articles can further strengthen authority signals. However, we caution against overgeneralization: generative engines are evolving, and what works in one instance may not in another. Citation does not guarantee influence, and concept novelty could play a role. Comparative studies across multiple cases are needed to draw broader conclusions.

# 7 Limitations

This study has several limitations. It examines only a single case and lacks a counterfactual comparison with other low-recognition actors. Many observations rely on self-published logs, which cannot be independently audited. AI-generated summaries and analytics are time-sensitive; results may differ if repeated at a later date. We do not have access to engine-internal ranking or citation logic, nor do we consider other generative search engines beyond Google. The novelty of GMI’s concepts may itself attract AI attention, confounding the role of structural practices. These limitations mean that our findings are suggestive, not definitive.

# 8 Conclusion

This paper presented an observational single-case study of GhostDrift Mathematical Institute to explore how concept definitions can enter generative search reference structures under disadvantageous conditions. We distinguished metrics, classified evidence tiers and reconstructed a timeline indicating that clear definitions were published earlier than external citations and AI references. We emphasise that this temporal order does not imply causation; it merely reflects the chronological sequence of events. While the data do not allow causal claims, the case illustrates that small actors can achieve reference visibility in generative search without high traffic. Future research should compare multiple cases, examine other platforms and investigate how generative engines decide which sources to cite.

# A Timeline of publications and observations

Table 4: Timeline of GMI publications, external mentions and AI observations (selected events)

Date	Event
2025-03-15	Early individual publishing on note or equivalent personal channels begins; this establishes that the case originated from personal-level dissemination rather than institutional authority (cite earliest public record).
2025-10-30	Public-facing activity begins from a personal website under an individual-level publishing configuration; this marks the start of the case before institutional incorporation (cite earliest public record).
2025-12-06	Publication of “Finite Closure” concept; article defines closing infinite domains in finite windows and notes that Google AI Overview began citing it[10].
2025-12-13	Medium article introduces the Three Core Principles (Finite Respect, Beacon, Consistency) and describes Beacon Architecture[11].
2025-12-25	GMI publishes “How is GhostDrift understood?” reporting that Google AI Overview characterises Ghost Drift as a theoretical framework for AI safety, reliability and accountability[12].
2026-01-08	Zenn article announces the Crisis Management Investment Task Force and outlines missions to fix responsibility via standard identifiers and evidence chains[13].
2026-01-19	Zenn post releases AEO observation log: 306 human sessions vs 1,077 AI queries, with AI traffic concentrated on definitions[14].
2026-01-21	Publication of Algorithmic Legitimacy Shift, discussing how risk and legitimacy shift from humans to algorithms[15].
2026-02-10	Formal incorporation of GMI; corporate number assigned on 2026-02-17[16].
2026-02-15	Artificial ALICE article introduces Ghost Drift, describing responsibility evaporation and related constraints[7].
2026-02-28	Reinforz.ai article includes section “Ghost Drift problem and AI governance,” defining Ghost Drift and proposing layered responsibility allocation[8].
2026-03-10	GMI records that Google AI Overview describes Beacon Architecture as a next-generation AI design concept based on the Protect-then-Select principle[17].

## B Evidence ledger

Table 5: Evidence ledger: claims, sources and inference levels

Claim	Evidence	Source tier	Inference level
AI queries exceeded human sessions	Zenn AEO log reports 1,077 AI-related queries vs 306 human sessions[14]	Tier 3	Descriptive only (publicly viewable but not auditable)
Definition of metric categories	Wix help page explains sessions from AI platforms count when a user clicks after an AI includes a link; user queries on AI count when a bot scans the site and do not indicate a click[4]	Tier 1	Direct descriptive
Search vs training bots	OpenAI bot documentation states that OAI-SearchBot surfaces websites in ChatGPT search features, while GPTBot crawls for model training[2]	Tier 1	Direct descriptive
Query fan-out	Google documentation notes that AI features may use query fan-out to issue multiple related searches and identify more supporting pages[3]	Tier 1	Direct descriptive
No special markup required	Google documentation states that AI features follow standard SEO requirements and require no special schema or new files[3]	Tier 1	Direct descriptive
Clicks drop when AI summaries appear	Pew research reports that click-through rates drop from 15% to 8% when AI summaries appear and that only 1% of visits click links in the summary[5]	Tier 2	Descriptive

(continued)				
Claim	Evidence	Source tier	Inference level	
Retrieval $\neq$ citation	Public analysis found that only 15% of pages retrieved by ChatGPT were cited in final answers[6]	Tier 2	Descriptive	
External adoption of Ghost Drift	Articles from Artificial ALICE and Reinforz.ai define Ghost Drift and related concepts[7, 8]	Tier 2	Descriptive	

## C Supplementary screenshot of AI overview observation

The following figure is included as a supplementary, time-sensitive manual observation of AI Overview behaviour. Because the present paper’s main Results section does not centre on Prime Gravity, we place this screenshot in the appendix as contextual supplementary material rather than as core evidence.



Figure 3: Google AI Overview Observation for “Prime Gravity.” Supplementary screenshot showing a time-sensitive manual observation in which GMI-related sources appeared in Google AI Overview results for the query “Prime Gravity.” This figure is illustrative and is not used for strong causal inference.

## References

- [1] S. Kim, W. Jeong, S. Kim, S. Lee, and D. Lee. “SAGEO Arena: A Realistic Environment for Evaluating Search–Augmented Generative Engine Optimization.” *arXiv preprint arXiv:2602.12187*, 2026. Available at: <https://arxiv.org/abs/2602.12187>.
- [2] OpenAI. “Overview of OpenAI Crawlers.” OpenAI Developer Documentation. Available at: <https://developers.openai.com/api/docs/bots/> (accessed 15 March 2026).
- [3] Google. “AI Overviews and AI Mode in Search.” Google Search Central Documentation. Available at: <https://developers.google.com/search/docs/appearance/ai-overviews> (accessed 15 March 2026).
- [4] Wix. “Wix Analytics: Reports Glossary.” Wix Help Center. Available at: <https://support.wix.com/en/article/wix-analytics-reports-glossary> (accessed 15 March 2026).
- [5] A. Chapekis and A. Lieb. “Google users are less likely to click on links when an AI summary appears in the results.” *Pew Research Center*, 22 July 2025. Available at: <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/>.
- [6] D. Goodwin. “Only 15% of pages retrieved by ChatGPT appear in final answers: report.” *Search Engine Land*, 13 March 2026. Available at: <https://searchengineland.com/chatgpt-retrieved-vs-citations-study-471606>.
- [7] Artificial ALICE. “AI Ethics 2026: State of the Art, Limits, and Breakthroughs (GhostDrift Perspective).” Independent media article [Japanese original], 15 February 2026. Available at: <https://www.artificialalice.com/133/> (accessed 15 March 2026).
- [8] Reinforz.ai. “What Will Change in AI by 2026? A Thorough Explanation of Natural Language Understanding (NLU) in the Social Implementation Phase and Japan’s Strategic Opportunities.” *Reinforz Insight* [Japanese original], 7 June 2024. Available at: <https://ai.reinforz.co.jp/87> (accessed 15 March 2026).
- [9] R. K. Yin. *Case Study Research and Applications: Design and Methods*. Sixth Edition. Sage Publications, 2018.
- [10] GhostDrift Mathematical Institute. “Finite Closure.” GhostDrift website. Available at: <https://www.ghostdriftresearch.com/%E6%9C%89%E9%99%90%E9%96%89%E5%8C%85> (accessed 15 March 2026).
- [11] GhostDrift Mathematical Institute. “[Official] The Three Core Principles of GhostDrift — Finite Respect, Beacon, Consistency.” *Medium*. Available at: <https://medium.com/@hide.pollini.kenkou/official-the-three-core-principles-of-ghostdrift-finite-respect-beacon-consistency-github-49faa5cbf1a4> (accessed 15 March 2026).

- [12] GhostDrift Mathematical Institute. “How is GhostDrift Understood? Record of its Positioning as a ‘Theoretical Framework’ in Google AI Overview.” *note*, 25 December 2025. Available at: [https://note.com/legit\\_cobra5046/n/nf6e7375e53cc](https://note.com/legit_cobra5046/n/nf6e7375e53cc).
- [13] GhostDrift Mathematical Institute. “Crisis Management Investment Mathematical Task Force.” GhostDrift website. Available at: <https://www.ghostdriftresearch.com/%E5%8D%B1%E6%A9%9F%E7%AE%A1%E7%90%86%E6%8A%95%E8%B3%87%E6%95%B0%E7%90%86%E5%AF%BE%E7%AD%96%E6%9C%AC%E9%83%A8> (accessed 15 March 2026).
- [14] GhostDrift Mathematical Institute. “AEO: Current Status, Limits, and Observation Log as a Breakthrough.” *Zenn*, 19 January 2026. Available at: <https://zenn.dev/ghostdrift/articles/06db95ccfeedde>.
- [15] GhostDrift Mathematical Institute. “Algorithmic Legitimacy Shift (ALS) and the Rewriting of Responsibility and Legitimacy.” *Zenn*, 21 January 2026. Available at: <https://zenn.dev/ghostdrift/articles/65384263288acd>.
- [16] GhostDrift Mathematical Institute. “Announcement of the Establishment of Ghost-Drift Mathematical Institute.” *note*, 18 February 2026. Available at: [https://note.com/legit\\_cobra5046/n/ndfc671444b3d](https://note.com/legit_cobra5046/n/ndfc671444b3d).
- [17] GhostDrift Mathematical Institute. “Observation Record of Beacon Architecture in Google AI Overview.” *note*, 10 March 2026. Available at: [https://note.com/legit\\_cobra5046/n/nac7ea1639005](https://note.com/legit_cobra5046/n/nac7ea1639005).