ARTIFICIAL INTELLIGENCE VISIBILITY INDEX (AIVI): THE FRAMEWORK TO QUANTIFY ENTITY PRESENCE IN GENERATIVE INFORMATION ENGINES

A PREPRINT

Eshaan Marutheesh
VScale - Harvard iLab
Univeristy of South Florida
Tampa, FL 33558
emarutheesh@usf.edu

Gokul Mohanty
VScale - Harvard iLab
Harvard University
Allston MA 02134
gom931@g.harvard.edu

March 28, 2025

ABSTRACT

The exponential growth of generative artificial intelligence has fundamentally changed how users access information, shifting discovery from ranked search results to a synthesis of multiple credible sources as answers. Traditional metrics such as search engine rankings and click-through rates no longer adequately capture visibility in this new age. This paper introduces the Artificial Intelligence Visibility Index (AIVI), a conceptual and analytical framework designed to quantify how entities, such as brands, organizations, or topics, are represented within generative information engines like ChatGPT, Gemini, and Claude. AIVI integrates multi-dimensional metrics including mention frequency and positional prominence to produce a visibility score. By turning AI visibility into a quantitative construct, public figures and companies can evaluate their presence in AI-generated outputs, identify optimization opportunities competitively position themselves in a new attention economy.

Keywords Artificial Intelligence Visibility Index, AIVI, Generative AI, GenAI, Generative Information Engines, Information Synthesis, Mention Frequency, Positional Prominence, Visibility Score, Quantitative Construct, Attention Economy, Competitive Positioning, AI-Generated Output, Search Engine Rankings, Click-Through Rates, Optimization and More

1 Introduction

The use of generative artificial intelligence systems, especially large language models (LLMs), has been utilized for the purpose of entity surfacing, which refers to the process of retrieving entities such as products, companies, or ideas for a wide array of informational queries. In the context of organizations seeking to make strategic decisions regarding their activities within the space of generative entity optimization (GEO), understanding the extent and frequency at which entities are displayed is crucial for success. However, existing methods for measuring entity presence are not transparent, aggregated at a high level as general performance metrics, or linked to overall company metrics such as brand engagement or sales. This creates a black box problem in which it becomes difficult to determine the extent to which an entity might be improving based solely on the entity itself.

The Artificial Intelligence Visibility Index (AIVI) fills this void by offering a transparent and reproducible measure of an entity's presence exclusively in generative AI outputs. It consists of two dimensions: coverage, which is the occurrence of an entity in a predetermined set of high-intent prompts, and rank, which is the prominence of an entity in each prompt. By adjusting these dimensions, it is possible to isolate an entity's performance from other aspects of an organization and track it objectively over time.

This metric allows organizations to rigorously test GEO strategies, compare entities, and make iterative improvements to AI-driven recommendations without relying on indirect or opaque metrics. In brief, AIVI offers a reliable, entity-based metric for visibility in generative information engines.

2 AIVI and It's Importance

2.1 Conceptual Foundation

The Artificial Intelligence Visibility Index (AIVI) is a specially created metric that aims to measure the consistency and prominence of an entity's visibility within the output generated by large language models (LLMs) and other generative AI technologies. Traditional methods for measuring visibility, such as web analytics and social media metrics, often tend to blur several aspects of organizational performance, including marketing, brand, and user-related metrics. The AIVI, on the other hand, focuses exclusively on the entity, thereby providing a clear and unambiguous measure of visibility that is unaffected by external influences. This distinction is particularly significant for Generative Entity Optimization (GEO) as organizations seek to understand the influence of AI-generated content on entities.

The rationale for AIVI's concept is based on the idea that the visibility of entities is probabilistic and positional in nature. In the generative AI results, the entities can be seen in some of the responses but not in others, based on the inherent randomness of the large language models. Coverage aims to include the probabilistic aspect by measuring the percentage of high-intent prompts in which the entities are visible . This gives a quantitative measure of exposure, indicating how often the entities are recognized within a given domain.

Rank is the positional importance of an entity, taking into consideration the fact that previously mentioned entities have a higher probability of affecting user perception. In the absence of deterministic ranking provided by generative systems, AIVI uses a discretized scoring system to effectively capture the concept of attention decay. Entities are given the highest scores when they are mentioned first, whereas lower-scoring entities are those mentioned further down the response hierarchy. AIVI effectively captures the probability of entity appearances by considering both coverage for frequency and rank for importance.

2.2 Operationalization

A methodical and repeatable approach is necessary for operationalizing AIVI. The first step is to define the prompt set, which is a carefully chosen set of high-intent queries that represent actual user information requirements. This set's size, quality, and representativeness are critical: a well-designed prompt set guarantees that AIVI captures visibility in contexts that are most important to the entity, whereas a limited or biased set may yield results that are misleading. Prompts are chosen to encompass the domain's semantic diversity, including common inquiries, requests for information, and situations in which the entity ought to be revealed.

The coverage measurement is then carried out. A binary indicator is generated for each prompt based on the evaluation of whether the entity appears in the AI-generated response. The coverage percentage, which shows the empirical likelihood of the entity appearing, is obtained by aggregating these indicators over the whole prompt set. This measure is crucial for tracking GEO progress because it offers a direct, objective view of exposure that is unaffected by external organizational factors or the marketplace.

Finally, rank assessment gives each prompt a score according to where the entity appears in the produced output. In a tiered system, the top three entities receive slightly lower scores, the entities that appear first receive the highest score, and so on. A score of zero is given for non-appearances. Rank scores can be combined with coverage to create a single weighted metric by normalizing them. Organizations are able to measure, benchmark, and systematically optimize entity visibility thanks to this operationalization, which results in a composite, interpretable index—the AIVI score—that offers actionable insights.

3 Quantitative Model: Formal Definition of AIVI

3.1 AIVI Formula

The *Artificial Intelligence Visibility Index (AIVI)* is a proprietary metric designed to quantify AI visibility and prominence. Unlike traditional ranking metrics, which assume deterministic results, AIVI accounts for the probabilistic and selective nature of large language model (LLM) outputs.

Formally, AIVI combines two components, coverage and rank, into a weighted score:

$$\text{AIVI} = (0.6 \times \text{Coverage}_{\%}) + (0.4 \times \text{Average Rank Score}) \tag{1}$$

3.2 Prompt Set, Coverage, and Rank Calculations

Let $\mathcal{Q} = \{q_1, q_2, \ldots, q_n\}$ denote a predefined set of high-intent prompts relevant to a given domain. These prompts approximate real-world user queries and define the evaluation space for measuring visibility.

Coverage. Coverage measures the frequency with which an entity appears across the prompt set. For each prompt q_i, define:

$$C_i = \begin{cases} 1, & \text{if the entity appears in the response} \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

The coverage percentage is then computed as:

$$\text{Coverage}_{\%} = \frac{1}{n} \sum_{i=1}^{n} C_i \times 100 \tag{3}$$

This represents the proportion of prompts in which the entity is included, normalized to a 0–100 scale.

Rank. Rank captures the relative prominence of an entity within a response. Given the absence of explicit rankings in generative outputs, entity positions are discretized into tiered scores for approximation:

$$R_i = \begin{cases} 100, & \text{if the entity is the first recommendation} \\ 75, & \text{if it appears within the top three recommendations} \\ 50, & \text{if it appears in positions four to six} \\ 25, & \text{if it appears below position six} \\ 0, & \text{if it does not appear} \end{cases} \tag{4}$$

The average rank score across the prompt set is then calculated as:

$$\text{Average Rank Score} = \frac{1}{n} \sum_{i=1}^{n} R_i \tag{5}$$

3.3 Final Formulation

By substituting both the normalized coverage and average rank score into Equation (1), the AIVI provides a single, quantifiable measure of visibility, ranging from 0 (no visibility) to 100 (maximal visibility).

$$\text{AIVI} = 0.6 \times \left(\frac{1}{n} \sum_{i=1}^{n} C_i \times 100 \right) + 0.4 \times \left(\frac{1}{n} \sum_{i=1}^{n} R_i \right) \tag{6}$$

where:

- n is the total number of prompts in the set $\mathcal{Q}$,
- C_i indicates whether the entity appears in prompt q_i (1 if present, 0 otherwise),
- R_i is the tiered rank score for the entity in prompt q_i, taking values in $\{0, 25, 50, 75, 100\}$.

This formulation makes the computation of AIVI fully explicit and immediately interpretable as a 0–100 scale metric combining both presence and prominence.

4 Usage Case of AIVI: Buisness

4.1 Measuring Market Presence

AIVI provides organizations with a quantitative measure to evaluate and figure out stratergies to optimize their in generative AI outputs. In business contexts, where discoverability directly impacts revenue, customer engagement, and market positioning, AIVI can serve as a critical tool for Generative Entity Optimization (GEO) analysis

Businesses often operate in highly niches where being visible in search results or AI-generated recommendations directly leads to more revenue. AIVI allows companies to objectively quantify the frequency and prominence of their products, services, or brand mentions across high-intent prompts. For example, a consumer electronics company can use AIVI to track how consistently its latest smartphone model appears in AI-generated buying guides. These metrics identify gaps in inclusion, while rank metrics reveal whether the entity is presented prominently or buried in lists, providing actionable insights for marketing, content strategy, and AI optimization.

4.2 Integration of Reputation and Strategic Optimization

Beyond just visibility, a combination of AI Reputation Score (AIRS) and Distributed Authority Domain Placement (DADP) enhances the strategic utility of AIVI. AIRS captures how the entity is framed in generated outputs—positively, negatively, or neutrally—allowing companies to monitor their AI-mediated reputation in real-time.[2] DADP evaluates the spread of entity authority across multiple domains, such as industry publications, product review sites, and informational platforms, revealing the systemic reach and influence of an entity. By integrating these measures, businesses gain a holistic view: not only whether their entity is visible, but whether it is framed favorably and positioned strategically across authoritative domains.[3]

The practical applications of AIVI in business include decision-making and resource allocation. By monitoring AIVI trends over time, individuals can assess the impact of content campaigns, product launches, or AI training interventions on visibility. Entities with low coverage can be targeted with additional AI-facing content or strategic placement across high-authority domains, while low-rank entities may require optimization of AI prompts, metadata, or contextual signals. Moreover, integrating AIVI with AIRS and DADP enables predictive modeling of entity performance, allowing businesses to anticipate competitive shifts and adapt proactively.

AIVI provides companies with a data-driven, AI-specific metric that is independent of sales performance and brings clarity in the otherwise "black box" of generative AI visibility. When paired with AIRS and DADP, it becomes a powerful tool for monitoring, benchmarking, and optimizing entity presence, ensuring that strategic decisions are grounded in quantifiable AI-mediated outcomes.

5 Case Study: Generative Entity Optimization in a Political Campaign

This section presents a practical application of Generative Entity Optimization (GEO) through a case study involving a political candidate in a state-level House of Representatives election. The campaign was executed using the GEO Machine platform (`geomachine.ai`) and focused on increasing the candidate's visibility within Generative Information Engines (GIEs) using the AIVI, AIRS, and DADP framework.

5.1 Initial Conditions

At the outset of the campaign, the candidate exhibited minimal presence within AI-generated outputs. Using a predefined set of high-intent prompts related to the election (e.g., candidate comparisons, policy-based queries, and voting recommendations), the candidate's initial Artificial Intelligence Visibility Index (AIVI) score was measured at 20. This low score reflected both limited coverage across prompts and low positional prominence when the candidate did appear.

Further analysis indicated that the candidate lacked sufficient representation across authoritative domains, resulting in weak signals within the broader information ecosystem.[5] In addition, the candidate's framing within existing sources was inconsistent, limiting their ability to establish a clear narrative within AI-generated responses. These conditions created a structural disadvantage, where the candidate was frequently excluded from generated outputs and, when included, was not prominently positioned.

5.2 GEO Strategy Implementation

To address these limitations, GEO Machine deployed a Distributed Authority Domain Placement (DADP) strategy centered on large-scale content distribution. The campaign released a substantial volume of publications across multiple domains, including informational articles, policy discussions, and structured content aligned with high-intent query patterns.

The strategy focused on three key components:

- **Authority Placement:** Publishing content across a wide range of domains to increase the candidate's presence within sources likely to be referenced by generative systems.
- **Semantic Reinforcement:** Consistently associating the candidate with specific policy areas and query intents through repeated contextual alignment.
- **Narrative Consistency:** Ensuring that all publications described the candidate using similar language, attributes, and positioning to influence downstream synthesis.

Rather than optimizing a single website or platform, the campaign emphasized distribution at scale. This approach leveraged the fact that GIEs construct responses by aggregating signals across multiple sources, making distributed reinforcement more effective than isolated optimization.

Simultaneously, adjustments were made based on AI Reputation Score (AIRS) analysis to refine how the candidate was framed. Content was structured to emphasize key strengths, policy positions, and credibility indicators, ensuring that when the candidate appeared in AI-generated outputs, the representation aligned with strategic objectives.

5.3 Results

Following the implementation of the GEO strategy, the candidate's AIVI score increased from 20 to 98. This improvement reflected substantial gains in both coverage and rank across the evaluated prompt set.

Coverage increased as the candidate began appearing in a significantly larger proportion of prompts, indicating that the probability of inclusion had improved. At the same time, rank scores increased as the candidate was more frequently positioned among the top recommendations within generated responses.

The observed changes were not incremental but systemic. As the volume and consistency of distributed content increased, the candidate's presence within authoritative sources strengthened, leading to higher likelihood of inclusion in AI outputs. This, in turn, reinforced visibility, creating a compounding effect consistent with the theoretical dynamics of DADP.

5.4 Implications

This case study demonstrates that GEO, when executed through a structured framework, can significantly influence entity visibility within generative systems. The increase in AIVI from 20 to 98 illustrates the effectiveness of combining measurement (AIVI), perception analysis (AIRS), and execution (DADP) into a unified strategy.

Importantly, the improvements were achieved through modifications to the information environment rather than traditional campaign tactics such as advertising or direct engagement. This highlights a key characteristic of GIEs: influence is exerted upstream, through the data and sources that inform generated responses.

More broadly, the results suggest that entities able to systematically deploy distributed authority strategies can achieve disproportionate visibility within AI-mediated environments. As generative systems continue to shape how information is accessed and decisions are made, GEO represents a critical capability for influencing both inclusion and perception at scale.

6 Conclusion and Future Work

The transition from traditional retrieval-based search to Generative Information Engines (GIEs) marks a fundamental shift in how entities are discovered, interpreted, and acted upon. Due to this shift, visibility is no longer determined solely by placement in a ranked list; it is measured by inclusion within AI-generated outputs. Similarly, perception is no longer constructed through independent exploration but is embedded directly within the synthesized response itself. These changes redefine the mechanics of entity discovery and optimization, compressing vast informational landscapes into bounded decision environments mediated by generative AI systems.

This paper introduces the Artificial Intelligence Visibility Index (AIVI) as a foundational metric for understanding and quantifying entity presence in this new landscape. AIVI measures both the frequency (coverage) and prominence (rank) of an entity within high-intent prompts, providing a transparent, reproducible measure that is decoupled from broader organizational performance. Complementary metrics, such as the AI Reputation Score (AIRS), capture how entities are framed, evaluated, and interpreted within generative outputs, while Distributed Authority Domain Placement (DADP) operationalizes the influence of cross-domain representation and authority in shaping both visibility and perception. Together, AIVI, AIRS, and DADP establish a comprehensive framework for Generative Entity Optimization (GEO), allowing organizations to measure, benchmark, and strategically influence their AI-mediated presence.

The implications of this framework extend beyond commercial visibility. In GEO contexts, entities that achieve consistent inclusion and favorable representation in AI outputs gain a competitive advantage, while those absent from generated responses risk being effectively invisible to users. In broader applications, from healthcare and education to policy and finance, the combination of AIVI, AIRS, and DADP provides a lens for understanding how generative AI shapes user perception and decision-making. The framework also highlights potential pathways for bias, misrepresentation, or unequal influence, emphasizing the need for systematic evaluation and ethical oversight in AI-mediated information environments.

Several avenues for future work emerge from this foundation. First, AIVI requires further refinement to standardize methodologies for prompt selection, normalization across platforms, and temporal tracking of entity visibility. Second, AIRS invites deeper research into quantifying sentiment, framing, and attribute association in AI outputs, as well as understanding how reputation interacts with visibility. Third, DADP warrants investigation as a mechanism for optimizing cross-domain influence, including the effects of source authority, domain weighting, and distributed signal reinforcement. Finally, a more comprehensive study of decision architectures in GIEs is needed to map how entity inclusion and framing translate into downstream user actions, informing both GEO strategies and broader human-computer interaction design. There appears to be significant implications of GIEs in influencing human decision-making, global commerce, and even elections and policy.[6]

As GIEs continue to evolve into primary interfaces for human information consumption, the ability to measure, optimize, and interpret entity visibility will become critical. AIVI, in combination with AIRS and DADP, provides the tools to navigate this new landscape, enabling organizations to systematically understand and influence how entities are discovered, perceived, and acted upon in AI-mediated decision environments.

References

1. Mohanty, Gokul, and Eshaan Marutheesh. "Generative Information Engine: Generative Engine Optimization and the next Layer of Global Commerce and Decision Making." Zenodo, 26 Mar. 2026, zenodo.org/records/19107148, https://doi.org/10.5281/zenodo.19107148. Accessed 1 Apr. 2026.
2. Mohanty, Gokul, and Eshaan Marutheesh. "Artificial Intelligence Reputation Score (AIRS): The Framework to Quantify Entity Sentiment in Generative Information Engines." Zenodo, 20 Mar. 2026, zenodo.org/records/19122428, https://doi.org/10.5281/zenodo.19122428. Accessed 1 Apr. 2026.
3. Mohanty, Gokul, and Eshaan Marutheesh. "Distributed Authority Domain Placement and How to Influence Artificial Intelligence." Zenodo, 20 Mar. 2026, zenodo.org/records/19122809, https://doi.org/10.5281/zenodo.19122809. Accessed 1 Apr. 2026.
4. Sielinski, Ronald. "Quantifying Uncertainty in AI Visibility: A Statistical Framework for Generative Search Measurement." arXiv preprint arXiv:2603.08924 (2026).
5. Rostova, Elena. "How Llms Decide What to Recommend: Inside the Black Box - Geo-Metric Blog." GEO, GEO-Metric, 25 Nov. 2025, geo-metric.ai/blog/how-llms-decide.
6. Mohanty, Gokul, and Eshaan Marutheesh. "GEO: Influencing Commerce, Human Decision Making, and Global Policy through Generative Information Engines." Zenodo, 22 Mar. 2026, zenodo.org/records/19154674, https://doi.org/10.5281/zenodo.19154674. Accessed 1 Apr. 2026.