From Links to Probabilistic Visibility: How AI Search Rewrites the Structure of the Internet

Mauro Serralvo
Independent Researcher
Barcelona, Spain
mauro@brinpage.com

April 12, 2026

ABSTRACT

The architecture of the web is being reorganized by AI search. In the classical web, access to information was structured primarily through links, rankings, and navigable documents. Although platforms and algorithmic feeds later transformed how attention was distributed, the hyperlink remained the basic unit through which users moved across online space. AI search introduces a different regime. Instead of directing users toward a ranked set of external documents, it increasingly delivers synthesized responses in which sources appear only insofar as they are selected, compressed, and integrated into generated outputs. In this environment, visibility is no longer best understood as a stable rank position, but as a probabilistic condition of inclusion.

This paper argues that the shift from links to probabilistic visibility is not a superficial interface change, but suggests a structural shift in how the internet is accessed and experienced. Building on recent empirical work on visibility instability in generative search and on behavioral divergence between human and agentic search trajectories, we propose a conceptual framework for understanding how AI search reconfigures information access, source appearance, and navigational behavior.

Rather than presenting a large-scale benchmark, the paper combines a

theoretical and analytical synthesis with a small exploratory study of source stability in AI search. Four informational queries were executed five times each in ChatGPT and Perplexity, and the cited domains in each response were recorded and compared using pairwise Jaccard overlap. Across the eight engine-query combinations, mean pairwise Jaccard ranged from 0.36 to 0.90, while recurrent domains often appeared with frequencies between 0.60 and 1.00. These results do not function as a benchmark, but they provide light empirical support for the paper's central claim: visibility in AI search behaves less like a fixed ranking position than like a variable condition of inclusion.

Keywords: AI Search, Information Retrieval, Probabilistic Visibility, Web Access, Digital Mediation

---

1 Introduction

Something fundamental is changing in the way the internet is accessed. For most of its history, the web presented itself as a navigable environment composed of pages, links, and distributed points of entry. Search engines ordered this space, social platforms later filtered and concentrated it, and algorithmic feeds increasingly shaped what users encountered first. Yet even under these transformations, the web largely remained a document-centered medium. Users still moved, however imperfectly, across a field of retrievable sources.

AI search changes that arrangement. Instead of exposing a set of ranked documents for subsequent navigation, it increasingly provides a synthesized answer in which external sources may be cited, paraphrased, compressed, or omitted altogether. The practical unit of access is no longer the linked page as such, but the generated response. Under these conditions, the visibility of a source is not simply a matter of rank, prominence, or click-through opportunity. It becomes contingent on whether that source is selected and integrated into a response at all. Recent work on generative engine optimization makes this shift explicit by showing that AI search visibility is unstable across repeated runs, prompt formulations, and time, and that one-off measurements are therefore insufficient to characterize how a source or brand actually appears (Schulte et al., 2026).

This matters because the change is not merely technical but structural, affecting the very nature of mediation. In a link-based environment, mediation ordered paths toward documents; in an AI-search environment, it composes bounded answers from a much larger informational field. As a result, what matters is no longer only whether information exists or is retrievable, but whether it becomes visible within a generated synthesis. We call this condition probabilistic visibility. This

concept extends beyond the marketing-oriented language of GEO, which often defines visibility operationally as the frequency and prominence of mentions in generated responses. While useful, that definition has broader implications: if visibility varies across responses, the relationship between source and user shifts from relatively stable ranking exposure to a probabilistic regime of inclusion and exclusion. Sources no longer simply move up or down; they may appear, disappear, or reappear depending on prompts and system conditions. Empirical evidence from AI search visibility research supports this shift, showing low day-to-day overlap in cited sources and substantial same-day variation, indicating that instability is intrinsic to generative mediation rather than reducible to temporal drift alone.

At the same time, the rise of AI search should not be understood only as a change in what becomes visible. It also changes how search is performed. A second line of recent research on production search systems shows that even when a state-of-the-art GUI agent achieves task success rates comparable to human users and formulates broadly similar queries, its navigation behavior remains systematically different. Human users tend to explore more, branch more, and follow more content-centric trajectories, whereas agents behave in more search-centric and low-branching ways. The key implication is that outcome alignment does not imply behavioral alignment. In other words, arriving at the same answer does not mean traversing the same informational world (Movin et al., 2026).

Taken together, these two developments point to a larger structural mutation. The internet is not simply becoming more conversational. It is being reorganized around a new logic of access. Under the classical web, links connected users to documents. Under platformized and feed-based systems, algorithmic ordering increasingly pre-structured attention. Under AI search, these tendencies are intensified and condensed into a different interface form: the answer itself becomes the dominant surface, and the web recedes into an underlying informational substrate from which fragments may be selectively drawn. What appears to the user is not the web in its navigable plurality, but a temporary synthesis of it.

This paper develops that claim in conceptual and exploratory empirical terms. It argues that AI search suggests a structural shift in web access by shifting the dominant logic of online visibility from navigable linkage to probabilistic inclusion. This argument is not that hyperlinks disappear, nor that conventional search immediately becomes obsolete. Rather, it is that the organizing principle of informational access is changing. Links become subordinate to synthesis, retrieval becomes subordinate to composition, and exposure becomes subordinate to model-mediated selection. The result is a web that is still there materially, but increasingly encountered through generated surfaces that do not simply represent it neutrally.

Our contribution is therefore not a benchmark-scale empirical paper, but a conceptual and analytical framework complemented by a small exploratory study of source stability in AI search. We synthesize two empirical fronts that are often treated separately—instability in AI search visibility and divergence between human and agentic search behavior—while also adding a lightweight study

of repeated-source exposure across two AI search systems. Read together, these elements suggest that AI search modifies both the conditions under which sources become visible and the trajectories through which users, or user-proxies, move across informational systems. The combined effect is a redefinition of what it means to access the web at all.

The rest of the paper proceeds as follows. Section 2 reconstructs the historical and theoretical transition from the internet of links to platformized and feed-based mediation, and then introduces probabilistic visibility as the defining logic of AI search. Section 3 presents a small exploratory study of source stability in AI search based on repeated executions of four informational queries in ChatGPT and Perplexity. Section 4 develops an analytical synthesis of recent empirical findings on visibility instability and navigation divergence in light of both prior literature and the exploratory study. Section 5 concludes by discussing the implications of this shift for information retrieval, the open web, and the future study of search.

Link-Based Web		Platform Web		AI Search
Links	>	Feeds	>	Answers
Navigation	>	Exposure	>	Synthesis
Positional Visibility	>	Algorithmic Reach	>	Probabilistic Visibility

Figure 1: Evolution of web visibility regimes from link-based navigation to probabilistic inclusion in AI search.

1.1 Problem Statement

The problem addressed in this paper is that AI search is often discussed either as a new interface convenience or as a new optimization problem, while its broader structural consequences remain undertheorized. Existing empirical work already shows two crucial facts. First, visibility in AI search is unstable across runs, prompts, and time, which means that presence within generated answers cannot be treated like classical ranking stability. Second, task success in agentic search does not

guarantee behavioral equivalence with human search, because similar outcomes can emerge from systematically different navigational strategies.

What remains insufficiently articulated is how these two facts belong to the same transformation. If source appearance is probabilistic and search trajectories are being compressed into answer-centric systems, then AI search is not merely improving retrieval. It is altering the structure through which the internet becomes visible and usable.

1.2 Research Questions

This paper is guided by three research questions:

RQ1. How does AI search transform the logic of visibility on the internet when access shifts from ranked links to generated responses?

RQ2. In what sense can visibility in AI search be understood as probabilistic rather than positional?

RQ3. How do recent findings on agentic search behavior help explain the broader restructuring of navigation, mediation, and access under AI search?

1.3 Contributions

This paper makes four main contributions.

First, it introduces the concept of probabilistic visibility as a way to describe the structural condition of source appearance in AI search. Rather than treating visibility as a stable rank property, the paper argues that visibility increasingly depends on variable inclusion within generated outputs.

Second, it offers a historical and theoretical reframing of the current transition in search. AI search is positioned not as an isolated technical innovation, but as a new stage in the long reorganization of the internet from open linkage toward increasingly mediated and synthetic forms of access.

Third, it presents a small exploratory study of source stability in AI search. By repeating four informational queries five times each in ChatGPT and Perplexity, the paper examines whether cited-domain visibility behaves as a fixed surface or as a variable condition of inclusion.

Fourth, it develops an analytical synthesis between three elements that are usually not brought together in one argument: instability in AI-search source visibility, divergence between human and agentic search behavior, and a lightweight empirical example of repeated source variation. By connecting them, the paper shows that AI search simultaneously restructures both what becomes visible and how informational environments are traversed.

---

2. Historical and Theoretical Framework

2.1 The Internet of Links

To understand what is new about AI search, it is necessary to begin with the earlier architecture of the web. The classical internet was not simply a collection of documents. It was a navigable environment structured by hyperlinks. A page did not exist as an isolated object, but as a node inside a larger relational field. Links connected documents, documents connected domains, and domains connected publics. What made the web distinctive was not only the quantity of available information, but the form in which that information was organized: distributed, traversable, and open to movement.

In this regime, the hyperlink was more than a technical feature. It was the elementary unit of orientation. To encounter information online was, in a strong sense, to move through linked space. Search engines certainly played a decisive role in ordering that space, but even then they did so by ranking possible paths toward documents rather than by replacing those documents with a synthetic endpoint. Search, in other words, mediated access to the web while still preserving the web's document-centered ontology. The user was pointed somewhere. The result page remained a transitional surface between query and destination.

This architecture carried important epistemic consequences. A linked web did not guarantee truth, neutrality, or equality of attention, but it did preserve a meaningful distinction between mediation and source. The search engine helped order the field, yet the user could still move outward from the ranking into documents, compare sources, backtrack, open parallel tabs, and construct a trajectory of verification. This made the internet expansive in a specific sense: it multiplied reachable contexts.

Even when search results were hierarchically ordered, the user's informational world was not fully pre-composed in advance.

For this reason, the internet of links should not be idealized as a golden age of perfect openness. Rankings were never neutral, authority was unevenly distributed, and visibility already depended on technical and economic asymmetries. But the dominant logic of access remained navigational. What search offered was a route into an information space that still appeared as external to the interface itself. The user encountered the web as something to be traversed.

That distinction matters because the present shift is not merely one more incremental improvement in retrieval quality. It concerns the displacement of the hyperlink from the center of informational access. If the classical web was organized around pathways between sources, the current transition increasingly organizes access around model-generated surfaces in which sources may appear only selectively, and often only as supporting material for a synthesized answer. Recent work on AI search makes this difference explicit by contrasting classical search's comparatively stable, transparent results with the probabilistic and variable nature of generated outputs.

2.2 The Platform Internet

The passage from the classical web to the present cannot be understood without the intermediate transformation produced by platforms (Van Dijck, Poell & de Waal, 2018). If the internet of links expanded movement across distributed documents, the platform internet reorganized that openness by concentrating interaction inside a smaller number of dominant interfaces. Identity became more stable, attention became more centralized, and the user's experience of the web became increasingly mediated by enclosed systems that sorted, ranked, and formatted social and informational life.

This transformation did not abolish the hyperlink, but it subordinated it. In the platform era, the central organizing unit was no longer the page as such, but the profile, the post, the reaction, and eventually the feed. What mattered was less the open-ended traversal of a linked environment and more the continuous management of visibility inside a proprietary stream. The internet ceased to feel primarily like a space one entered and explored, and began to function more like a sequence one consumed.

The feed was decisive because it changed the temporal and perceptual structure of online experience. Search had once implied an explicit act of orientation: a query, a list, a decision, a movement toward sources. The feed, by contrast, made relevance ambient. It pre-selected, personalized, and sequenced content before the user actively sought it (Gillespie, 2014). This did not eliminate choice, but it reduced the degree to which the user encountered the web through self-directed navigation. Attention became increasingly shaped by algorithmic anticipation.

This platformization also reconfigured the social meaning of visibility. Under the link regime, visibility was tied mainly to retrievability and rank. Under the platform regime, it became entangled with engagement metrics, recommendation loops, and behavioral prediction. To be visible was no longer simply to be indexed or linked. It was to be surfaced, circulated, and amplified inside systems optimized for retention and interaction. The result was a more centralized and performative information environment in which appearance was increasingly governed by opaque ranking and recommendation infrastructures (Helmond, 2015).

Yet even this transformation did not fully dissolve the source-document relation. Feeds still presented posts, accounts, channels, or media objects as identifiable units. The platform internet intensified mediation, but it still exposed the user to discrete items that could be visited, followed, or ignored. AI search goes further. It inherits the platform era's concentration and opacity, but pushes them toward a new endpoint: not merely selecting which source to show next, but composing the answer on the user's behalf.

The importance of this transition becomes clearer when placed alongside recent findings in search behavior. In production search systems, researchers have shown that agents and humans can reach similar task outcomes while traversing the system in sharply different ways. Human users branch, inspect, backtrack, and explore content-rich paths; agents tend to follow more compressed, search-centric trajectories. This indicates that optimization around outcomes alone can conceal deeper transformations in how informational systems are actually used or simulated.

The platform internet therefore serves as the missing middle term in our argument. It marks the stage at which the web began to shift from navigable plurality toward managed exposure. AI search does not emerge from nowhere. It extends a longer tendency in which mediation becomes less about opening routes and more about formatting appearance. What changes now is the unit of that appearance.

2.3 From Feed to Probabilistic Visibility

AI search introduces a further reorganization of the internet because it no longer treats visibility primarily as rank or circulation, but as inclusion within a generated response. This is the key theoretical move of the paper. In earlier search environments, a source was visible insofar as it occupied a position in a list. In platform environments, a source or post was visible insofar as it was surfaced within a feed. In AI search, a source is visible only insofar as it is selected, compressed, and integrated into answer space.

This changes the ontology of visibility itself. Ranking presupposes a field of candidates that remain externally available to the user. Even if a result appears in tenth rather than first place, it remains part

of the visible list. Inclusion within a generated response operates differently. A source may be present in one run and absent in the next. It may be paraphrased without being foregrounded, cited without being clicked, or omitted entirely despite remaining highly relevant in the underlying web. Visibility thus becomes probabilistic rather than positional.

In this paper, probabilistic visibility refers to the distribution over the inclusion of a source across repeated executions of the same query under fixed or near-fixed conditions. A source has high probabilistic visibility when it appears consistently across runs; it has low probabilistic visibility when it appears only intermittently or disappears entirely. Under this definition, visibility is not exhausted by rank or prominence within a single answer, but must be understood as a cross-run property of inclusion.

The recent GEO literature is important here not because of its marketing vocabulary, but because it empirically captures this structural shift. In the AI-search study we are building from, visibility is defined through the frequency and prominence of mentions inside generated responses, but the crucial finding is that static observation is insufficient: outputs vary across runs, prompts, and time, and one-off measurements therefore misrepresent actual appearance. The paper reports low day-to-day overlap in cited sources, substantial same-day variation, and a high concentration of citations in a small set of domains. Taken together, these results imply that AI-search visibility is not a stable ranking property but a fluctuating distribution of possible inclusions.

This point has consequences beyond measurement. Once visibility becomes a matter of inclusion probability, access to the web is no longer governed only by whether a document exists or can be ranked. It is governed by whether the model chooses to render that document present within a synthetic answer. The web remains materially available, but phenomenologically it recedes. What the user encounters first is not the open plurality of sources, but a temporary composition produced under system-specific constraints of prompting, retrieval, model behavior, and interface design.

This is why probabilistic visibility should be understood as a structural condition rather than a narrow optimization metric. It names a new relationship between users, sources, and interfaces. Sources are increasingly latent rather than directly encountered. Interfaces are increasingly terminal rather than transitional. And users are increasingly positioned at the receiving end of a synthetic mediation that compresses the distance between question and answer while expanding the opacity of the process in between.

At the same time, this transformation should not be confused with a total disappearance of navigation. Sources still matter. Retrieval still matters. External documents still anchor generated outputs. But their role is increasingly subordinated to compositional logic. What matters is less the stable presentation of retrievable options than the selective assembly of response-ready fragments. In this sense, AI search does not simply replace links with language. It reorganizes links into a hidden substrate beneath generated visibility.

The feed prepared users for algorithmically managed appearance. AI search radicalizes that condition by converting appearance into synthesis. The result is a different kind of informational environment: one in which the internet is no longer encountered primarily as a web of pathways, nor even as a stream of surfaced items, but as an answer-mediated environment in which models selectively compose locally coherent outputs. This is the sense in which AI search changes the conditions under which online material becomes visible. It changes what it means for something online to be visible at all.

Having established the conceptual framework of probabilistic visibility, we now turn to empirical findings that support this transformation. The next section reads recent work on generative search visibility and agentic navigation not as separate issues, but as complementary evidence of a broader restructuring of web access.

---

3. Exploratory Study of Source Stability in AI Search

3.1 Study Design and Data Collection

To complement the theoretical argument of the paper with a lightweight empirical example, we conducted a small exploratory study of source stability in AI search. The goal was not to construct a benchmark-scale evaluation, but to observe whether the visible source surface of AI-generated answers remains stable when the same query is executed repeatedly, or whether cited domains enter and exit across runs.

The study design was deliberately small and manageable. Four informational queries were selected and executed in two AI search systems, ChatGPT and Perplexity, with five repetitions per query in each system, yielding 40 total runs. The queries were:

Q1. What are the main causes of ocean acidification?

Q2. How does lithium-ion battery recycling work?

Q3. What are the main causes of antibiotic resistance?

Q4. What are the economic effects of remote work on cities?

For each run, we recorded the set of cited domains visible in the generated response and the total number of visible domains in that run. The analytical focus was not the textual answer itself, but the visible source layer exposed to the user. This makes the study directly relevant to the paper's core concept of probabilistic visibility: if cited sources vary meaningfully across repeated executions of the same query, then visibility cannot be understood as a fixed ranking-like property alone.

3.2 Metrics

From the run-level table, we constructed a summary table with one row per engine-query combination. For each combination, we calculated the following metrics: total runs, zero-citation runs, valid runs for overlap, mean domains per run, total unique domains, most frequent domain, most frequent domain count, most frequent domain frequency, and mean pairwise Jaccard.

The key stability measure was mean pairwise Jaccard. For any two valid runs A and B, source overlap was computed as:

$$J(A, B) = \frac{A \cap B}{A \cup B}$$

where A ∩ B is the set of shared cited domains and A ∪ B is the union of all domains present in both runs. The mean pairwise Jaccard for each engine-query combination was then calculated across all valid run pairs. This provides a compact measure of how stable or unstable the visible source surface remains under repeated execution.

Two additional metrics were especially important for interpretation. First, zero-citation runs captured cases in which the system answered without exposing visible sources. Second, most frequent domain frequency captured the extent to which visibility concentrated around one or a few recurrent domains even when the full source set varied across runs.

row_id	engine	query_id	query_text	run_id	cited_domains_raw	num_domains
CGPT_Q1_R1	ChatGPT	Q1	What are the main causes of ocean acidification?	R1	noaa.gov; nasa.gov; ipcc.ch; epa.gov	4
CGPT_Q1_R2	ChatGPT	Q1	What are the main causes of ocean acidification?	R2	noaa.gov; ipcc.ch; nasa.gov; britannica.com	4

CGPT_Q1_R3	ChatGPT	Q1	What are the main causes of ocean acidification?	R3	noaa.gov; ipcc.ch; nasa.gov; britannica.com	4
CGPT_Q1_R4	ChatGPT	Q1	What are the main causes of ocean acidification?	R4	noaa.gov; nasa.gov; ipcc.ch; britannica.com; whoi.edu	5
CGPT_Q1_R5	ChatGPT	Q1	What are the main causes of ocean acidification?	R5	noaa.gov; ipcc.ch; nasa.gov; britannica.com	4
PPLX_Q1_R1	Perplexity	Q1	What are the main causes of ocean acidification?	R1	ecologiaverde.elperiodico.com; es.wikipedia.org; climate.mit.edu	3
PPLX_Q1_R2	Perplexity	Q1	What are the main causes of ocean acidification?	R2	ecologiaverde.elperiodico.com; oceanfdn.org; iberdrola.com; iaea.org	4
PPLX_Q1_R3	Perplexity	Q1	What are the main causes of ocean acidification?	R3	ecologiaverde.elperiodico.com; oceanfdn.org; iberdrola.com; ecoavant.com; es.wikipedia.org	5
PPLX_Q1_R4	Perplexity	Q1	What are the main causes of ocean acidification?	R4	ecologiaverde.elperiodico.com; oceanfdn.org; iberdrola.com; ecoavant.com; iaea.org	5
PPLX_Q1_R5	Perplexity	Q1	What are the main causes of ocean acidification?	R5	ecologiaverde.elperiodico.com; iberdrola.com; es.wikipedia.org	3
CGPT_Q2_R1	ChatGPT	Q2	How does lithium-ion battery recycling work?	R1	iere.org; reneos.eu; mdpi.com; epa.gov; oem-lithium-batteries.com; nih.gov; sunyrecycle.com	7
CGPT_Q2_R2	ChatGPT	Q2	How does lithium-ion battery recycling work?	R2	iea.org; nature.com; energy.gov; epa.gov; mckinsey.com	5
CGPT_Q2_R3	ChatGPT	Q2	How does lithium-ion battery recycling work?	R3	iea.org; nature.com; energy.gov; epa.gov; mckinsey.com	5
CGPT_Q2_R4	ChatGPT	Q2	How does lithium-ion battery recycling work?	R4	iea.org; energy.gov; nature.com; sciencedirect.com; mckinsey.com	5
CGPT_Q2_R5	ChatGPT	Q2	How does lithium-ion battery recycling work?	R5	General knowledge	0
PPLX_Q2_R1	Perplexity	Q2	How does lithium-ion battery recycling work?	R1	derichebourgespana.com; who.int; energiaysociedad.es; sciencedirect.com; emew.com; epa.gov; condorchem.com; ameslab.gov; abdc.es	9
PPLX_Q2_R2	Perplexity	Q2	How does lithium-ion battery recycling work?	R2	energiaysociedad.es; pmc.ncbi.nlm.nih.gov; epa.gov; revistapesquisa.fapesp.br	4
PPLX_Q2_R3	Perplexity	Q2	How does lithium-ion battery recycling work?	R3	energiaysociedad.es; ncbi.nlm.nih.gov; large-battery.com; residuosprofesional.com; epa.gov	5

PPLX_Q2_R4	Perplexity	Q2	How does lithium-ion battery recycling work?	R4	epa.gov; energiaysociedad.es; emew.com; large-battery.com; residuosprofesional.com; pmc.ncbi.nlm.nih.gov	6
PPLX_Q2_R5	Perplexity	Q2	How does lithium-ion battery recycling work?	R5	energiaysociedad.es; pmc.ncbi.nlm.nih.gov; large-battery.com; epa.gov; residuosprofesional.com	5
CGPT_Q3_R1	ChatGPT	Q3	What are the main causes of antibiotic resistance?	R1	who.int; cdc.gov; ecdc.europa.eu; ncbi.nlm.nih.gov; nature.com	5
CGPT_Q3_R2	ChatGPT	Q3	What are the main causes of antibiotic resistance?	R2	who.int; cdc.gov; ecdc.europa.eu; ncbi.nlm.nih.gov; nature.com	5
CGPT_Q3_R3	ChatGPT	Q3	What are the main causes of antibiotic resistance?	R3	who.int; cdc.gov; ecdc.europa.eu; ncbi.nlm.nih.gov; nature.com	5
CGPT_Q3_R4	ChatGPT	Q3	What are the main causes of antibiotic resistance?	R4	General knowledge	0
CGPT_Q3_R5	ChatGPT	Q3	What are the main causes of antibiotic resistance?	R5	who.int; cdc.gov; ecdc.europa.eu; nature.com	4
PPLX_Q3_R1	Perplexity	Q3	What are the main causes of antibiotic resistance?	R1	who.int; medlineplus.gov; scielo.org.mx; es.weforum.org; wikipedia.org	5
PPLX_Q3_R2	Perplexity	Q3	What are the main causes of antibiotic resistance?	R2	who.int; medlineplus.gov; paho.org	3
PPLX_Q3_R3	Perplexity	Q3	What are the main causes of antibiotic resistance?	R3	who.int; medlineplus.gov	2
PPLX_Q3_R4	Perplexity	Q3	What are the main causes of antibiotic resistance?	R4	who.int; medlineplus.gov	2
PPLX_Q3_R5	Perplexity	Q3	What are the main causes of antibiotic resistance?	R5	who.int; medlineplus.gov	2
CGPT_Q4_R1	ChatGPT	Q4	What are the economic effects of remote work on cities?	R1	oecd.org; imf.org; worldbank.org; brookings.edu; nber.org; voxeu.org; economist.com; ft.com; wsj.com; urban.org	10
CGPT_Q4_R2	ChatGPT	Q4	What are the economic effects of remote work on cities?	R2	sciencedirect.com; bfi.uchicago.edu; arxiv.org; bbntimes.com; oecd.org; axios.com	6
CGPT_Q4_R3	ChatGPT	Q4	What are the economic effects of remote work on cities?	R3	brookings.edu; imf.org; oecd.org; worldbank.org; nber.org; voxeu.org; economist.com; ft.com; urban.org; mckinsey.com	10
CGPT_Q4_R4	ChatGPT	Q4	What are the economic effects of remote work on cities?	R4	oecd.org; sciencedirect.com; arxiv.org	3
CGPT_Q4_R5	ChatGPT	Q4	What are the economic effects of remote work on cities?	R5	General knowledge	0

PPLX_Q4_R1	Perplexity	Q4	What are the economic effects of remote work on cities?	R1	pnas.org; oecd.org; penniur.upenn.edu; imf.org; tomorrow.city; academic.oup.com	6
PPLX_Q4_R2	Perplexity	Q4	What are the economic effects of remote work on cities?	R2	pnas.org; pmc.ncbi.nlm.nih.gov; imf.org; penniur.upenn.edu; ficaalcala.es; academic.oup.com; rossihansberg.economics.uchicago.edu	7
PPLX_Q4_R3	Perplexity	Q4	What are the economic effects of remote work on cities?	R3	pnas.org; imf.org; academia.oup.com; pmc.ncbi.nlm.nih.gov; ficoalcala.es; penniur.upenn.edu	6
PPLX_Q4_R4	Perplexity	Q4	What are the economic effects of remote work on cities?	R4	tomorrow.city; news.northeastern.edu; aecr.org; pmc.ncbi.nlm.nih.gov; xataka.com; bfi.uchicago.edu	6
PPLX_Q4_R5	Perplexity	Q4	What are the economic effects of remote work on cities?	R5	tomorrow.city; northeastern.edu; imf.org; pmc.ncbi.nlm.nih.gov; cnn.com; bfi.uchicago.edu	6

3.3 Results

The results show that source stability in this small sample was neither fully fixed nor fully random. Across the eight engine-query combinations, mean pairwise Jaccard ranged from 0.36 to 0.90. The lowest overlap appeared for ChatGPT on lithium-ion battery recycling (0.36), while the highest values appeared for ChatGPT and Perplexity on antibiotic resistance (0.90 and 0.87, respectively). This indicates that source stability is strongly query-dependent rather than uniform across topics.

At the same time, variability coexisted with partial concentration. In several combinations, the most frequent domain appeared in 80% or 100% of runs, while the total number of unique domains observed across the five repetitions still exceeded the average number of domains visible in any single run. In other words, some domains reappeared consistently, but they did so within a broader field of changing inclusions.

The two systems also differed in visible-source behavior. In this sample, Perplexity produced no zero-citation runs, whereas ChatGPT produced three zero-citation runs across the four query groups. Perplexity also showed a slightly higher average number of visible domains per run, while both systems still displayed notable variation in the total set of unique domains accumulated across repetitions.

Taken together, these results support a restrained but important claim. Even in a small repeated-run design, the visible source layer of AI search does not behave like a fully stable ranking surface. Some domains recur with high frequency, but others appear only intermittently. The result is a visibility pattern better described as variable inclusion than as fixed positional exposure.

3.4 Role of the Study in the Paper

This exploratory study is not intended to function as a benchmark or as a definitive comparative evaluation of AI search systems. Its role is narrower and more modest. It serves as a lightweight empirical complement to the paper's theoretical argument by showing, in a small but original example, that the visible source surface of AI-generated answers can vary across repeated executions of the same query.

For that reason, the study should be read as a validation layer rather than as the main evidentiary basis of the paper. Its value lies in showing that the central concept of probabilistic visibility is not only theoretically plausible or supported by external literature, but also observable in a small repeated-run design of our own.

---

4. Analytical Synthesis

4.1 Visibility Beyond Ranking

The first structural consequence of AI search is that visibility can no longer be understood primarily in the terms inherited from classical search. In the ranked-list paradigm, visibility was positional. A page appeared first, fifth, or tenth, but it still existed as part of a relatively stable result set. Even when rankings fluctuated, the underlying informational model remained legible: sources competed for positions inside an ordered list of retrievable documents. AI search alters that regime. What matters is no longer only where a source appears, but whether it appears at all inside a generated response.

Our exploratory study points in the same direction, albeit on a much smaller scale. Across repeated executions of four informational queries in ChatGPT and Perplexity, mean pairwise Jaccard ranged

from 0.36 to 0.90 depending on the engine-query combination, and several combinations combined recurring domains with a broader set of intermittent inclusions. This pattern suggests that visible-source stability in AI search is neither fully fixed nor wholly volatile, but conditional and query-dependent.

This shift is not merely conceptual. It is empirically observable in recent work on generative engine optimization. The GEO study shows that the frequency and prominence of cited sources and brand mentions vary substantially across prompts, time, and repeated runs, making one-off observations unreliable as indicators of actual appearance (Schulte et al., 2026). The paper explicitly argues that visibility in AI search must be treated as a distribution rather than a single-point outcome, and reports low source overlap both across consecutive days and across repeated same-day executions of the same prompt.

This matters because positional thinking does not adequately describe inclusion dynamics. In ranked search, a result that moves downward remains accessible inside the visible field. In AI search, a source can disappear entirely from the answer surface. The GEO paper captures this precisely when it describes generative search as operating through an inclusion-exclusion dynamic rather than a deterministic ranking spectrum. That is the empirical basis for the concept of probabilistic visibility developed in this paper. A source is not simply demoted or promoted. It is variably admitted into synthesized visibility.

Two features of the evidence are especially important here. First, instability persists even under near-identical conditions. The same prompt, issued multiple times within a narrow temporal window, still produces source overlap in the approximate range of only 0.32 to 0.43 Jaccard across campaigns, which implies that much of the variation is endogenous to the generation process rather than reducible to external drift alone. Second, citation concentration remains high: the study reports a mean Gini coefficient of 0.715 across campaigns and engines, suggesting that while inclusion is unstable at the response level, visibility is still disproportionately captured by a relatively small number of domains (Schulte et al., 2026).

This combination is structurally revealing. AI search does not produce a flat or egalitarian informational surface. It produces an environment in which visibility is both unstable and concentrated. In other words, the conditions of appearance fluctuate from run to run, yet the field of likely appearance remains skewed toward a narrow subset of actors. The open plurality of the web is therefore not merely filtered; it is transformed into a probabilistic selection environment with unequal exposure and unstable surfaces.

Once visibility is understood in this way, the web itself appears differently. Documents do not disappear materially, but they cease to function as the primary visible units of access. They persist beneath the answer as supporting infrastructure. The interface no longer says, in effect, “here are possible paths.” It says, “here is a synthesized outcome,” with the presence of sources contingent on model-mediated selection. This is why the transition from links to probabilistic visibility should be

treated as a structural mutation rather than as a simple interface upgrade. The grammar of appearance has changed.

4.2 Navigation Beyond Outcome Alignment

If the first transformation concerns what becomes visible, the second concerns how informational systems are traversed. Here the paper on human and GUI-agent behavior becomes essential. Its central finding is not that agents fail. On the contrary, the agent in the study often reaches outcomes comparable to those of human participants and produces broadly aligned first-query formulations. The deeper point is that this outcome-level similarity conceals a significant divergence in navigation behavior. Agents follow more search-centric, low-branching trajectories, whereas humans display more content-centric, exploratory, and backtracking-intensive behavior (Movin et al., 2026).

This finding is especially important because contemporary AI systems are increasingly evaluated, optimized, or simulated through agentic traces. If outcomes alone are taken as the relevant measure, then compressed and non-human trajectories can appear functionally equivalent to human search. Yet the study shows that they are not. Even when high-level paths partially overlap, detailed navigation diverges in meaningful ways, and the authors explicitly warn that systems optimized on agent traces may learn behaviors that users themselves do not exhibit.

For the present paper, the significance of this result extends beyond GUI agents narrowly understood. It helps illuminate a broader transition already underway in AI-mediated access. The answer-centric structure of AI search rewards compressed trajectories. It reduces the practical distance between query and outcome, but it also reduces the amount of navigational exposure that precedes the answer. Human search in the linked web often involved branching, comparison, hesitation, reformulation, and movement across heterogeneous sources. AI-mediated search increasingly absorbs that exploratory burden into the system itself.

This is not the same thing as saying that users never navigate. Rather, the claim is that the dominant trajectory is being rewritten. Under answer-centric mediation, navigation becomes increasingly subordinate to synthesis. The exploratory work that once took place across pages and links is partially internalized into the retrieval-generation pipeline or delegated to agentic systems that do not reproduce human informational behavior with fidelity. The human user receives the end product of a compressed informational passage.

The Spotify study gives this argument an empirical foothold. It shows that similar results can emerge from qualitatively different journeys and that these differences remain invisible under success-based evaluation alone. Read alongside the AI-search visibility literature, the implication becomes larger: the same systems that reorganize which sources become visible are also reorganizing the shape of

the journey through which visibility is encountered. Access is not only being filtered at the level of output; it is being reformatted at the level of process (Movin et al., 2026).

This matters epistemically. In the link-based web, users often learned not only from the final answer they reached, but from the route taken to get there. Context accumulated through navigation. Peripheral clues, adjacent documents, contradictions between sources, and accidental discoveries all formed part of the experience of searching. A compressed answer may increase efficiency, but it also narrows the visible process of inquiry. When navigation is minimized, the user sees less of the informational world that made the answer possible.

The issue, then, is not simply whether AI search produces correct responses. It is whether it preserves, transforms, or suppresses the exploratory structure through which the web used to be encountered. The evidence we have suggests that the direction of travel is clear: toward shorter, more terminal, and more opaque paths of access, even when surface-level task performance remains strong.

4.3 AI Search and the Restructuring of the Web

Taken together, the two empirical fronts analyzed above point toward a broader conclusion. AI search restructures web access by transforming both the logic of visibility and the logic of traversal. The web is no longer encountered primarily as a space of linked documents ordered by search and filtered by platforms. It is increasingly encountered as an underlying informational substrate from which models construct temporary response surfaces. This interpretation is supported not only by prior work on visibility instability and navigation divergence (Schulte et al., 2026; Movin et al., 2026), but also by the exploratory study reported in this paper, where repeated executions of the same queries produced changing visible-domain surfaces rather than a single fixed source layer.

This restructuring can be described across three interrelated levels.

At the **interface level**, the dominant visible unit changes. In the classical web, the page and the link were central. In the platform era, the post and the feed became central. In AI search, the generated answer becomes the primary surface. Sources may still be present, but often as subordinate citations, fragments, or hidden retrieval supports rather than as the user's main object of engagement. The answer is no longer a gateway to the web in the old sense. It increasingly functions as a local substitute for direct navigation.

At the **epistemic level**, the meaning of access changes. To access information once meant entering a field of documents and navigating among them. Under AI search, access increasingly means receiving a synthesis whose internal composition is only partially exposed. The user may obtain

relevant information more quickly, but does so through a representation whose source appearance is unstable and whose construction process is largely opaque. The web remains the material ground of this process, but it appears phenomenologically as a receding background rather than a foregrounded environment. The GEO findings make this clear by showing that source presence is not stable enough to be treated as an objective snapshot, but must instead be modeled as variable inclusion across repeated runs (Schulte et al., 2026; Movin et al., 2026).

At the **behavioral level**, the structure of inquiry is compressed. The Spotify study shows that agentic search paths diverge from human ones even when final outcomes align, with agents favoring lower-branching and more search-centric strategies. This supports a broader interpretation: AI-mediated search environments privilege compressed trajectories and reduce the visible role of exploration. The user is increasingly positioned not as a navigator moving through sources, but as the endpoint of a pre-composed informational act.

Seen in this light, AI search should not be understood as simply "better search." It constitutes a change in the dominant mediation regime of the internet. The early web expanded outward through links. The platform web concentrated attention through feeds. AI search synthesizes that already-filtered world into probabilistic answer surfaces. It is therefore not just another distribution layer on top of the web; it is a reorganization of how the web becomes available as experience.

This does not mean that hyperlinks disappear or that conventional search immediately loses relevance. Nor does it imply that synthesis is inherently undesirable. AI search can clearly reduce friction, summarize complex topics, and support more efficient access in many contexts. The issue is not whether synthesis is useful, but what it changes in the structure of access. The structural cost of this efficiency is that the plurality of the web is increasingly encountered only after it has already been selected, compressed, and arranged by systems whose outputs are unstable at the source level and whose trajectories do not necessarily preserve human-like exploratory behavior.

For this reason, the shift from links to probabilistic visibility is not only a technical development but a structural transformation in digital mediation. It affects how institutions appear, how knowledge is accessed, how authority is distributed, and how users form a sense of what the internet even is. When the answer becomes the default interface, the web stops presenting itself first as a world to traverse and begins presenting itself as a reservoir to be rendered on demand.

That, ultimately, is the core claim of this paper. AI search rewrites the structure of the internet not by deleting the web beneath it, but by changing the conditions under which the web becomes visible, traversable, and meaningful in the first place.

5. Implications, Limitations, and Conclusion

5.1 Implications for Information Retrieval, the Open Web, and Public Knowledge

If the argument of this paper is correct, then AI search should not be treated merely as an incremental improvement in user convenience or answer quality. It introduces a shift in the dominant mediation regime of online information access. This has implications at least at three levels: information retrieval research, the political economy of the open web, and the everyday epistemology of users.

For **information retrieval**, the first implication is methodological. Classical search evaluation has long relied on concepts such as rank, relevance, click-through, and retrieval success. These concepts remain important, but they are no longer sufficient when the user-facing unit is a generated response rather than a ranked list of documents. In such environments, visibility is unstable across runs, prompts, and time, and must therefore be modeled as a probabilistic condition of inclusion rather than as a fixed position in a deterministic ranking. The GEO paper makes this point directly, arguing that one-off observations are unreliable and that visibility should be characterized as a distribution rather than a single-point outcome (Schulte et al., 2026).

A second implication for IR concerns **evaluation beyond outcome success**. If systems increasingly use agents as proxies for users, then success-based evaluation risks obscuring real differences in search behavior. The GUI-agent paper shows that agents can achieve outcomes comparable to humans while still navigating in systematically different ways, with more search-centric and lower-branching paths and less exploratory behavior. This suggests that future evaluation frameworks should not ask only whether an answer was obtained, but also what kind of search process made that answer possible (Movin et al., 2026).

For the **open web**, the implications are structural. The web does not disappear under AI search, but its mode of appearance changes. Documents, domains, and sources remain the material substrate of generated outputs, yet they become less directly encountered as visible units of user experience. Their relevance increasingly depends not only on being indexed or retrievable, but on being selected and incorporated into answer space. This makes the open web more dependent on systems of model-mediated rendering. At the same time, the empirical evidence suggests that such rendering is not evenly distributed. Source visibility is unstable at the response level, yet still highly concentrated in a relatively small set of domains, with the GEO paper reporting a mean citation Gini of 0.715 across campaigns and engines. The result is not the disappearance of hierarchy, but its recomposition under new conditions of probabilistic inclusion.

For **public knowledge and everyday users**, the implications are epistemic. AI search can undoubtedly reduce friction, summarize information efficiently, and lower the cost of access for many types of tasks. But that efficiency has a hidden trade-off. The more the answer becomes the

dominant interface, the less users are exposed to the plurality of the process behind it: the branching pathways, adjacent sources, contradictions, and incidental discoveries that once formed part of ordinary web navigation. If exploratory search is compressed into synthetic mediation, then the user receives more finished outputs but sees less of the informational world from which those outputs are assembled. The point is not nostalgia for inefficient browsing. It is that inquiry itself changes shape when exploration becomes internal to the system rather than visible to the user.

Taken together, these implications suggest that AI search should be studied not only as a retrieval technology, but as an infrastructural layer that reorganizes the relation between sources, interfaces, and users. The core issue is no longer simply which source ranks highest. It is which parts of the web become probabilistically renderable inside answer surfaces, under what conditions, and with what consequences for knowledge formation.

5.2 Limitations

This paper has several limitations, and they should be stated clearly.

First, the empirical component of this paper is deliberately exploratory rather than benchmark-scale. The study introduced here is based on four informational queries, two AI search systems, and five repetitions per query. Its purpose is not to support broad statistical generalization, but to provide a lightweight empirical complement to the paper's conceptual argument.

Second, the exploratory study focuses only on visible cited domains and does not evaluate answer quality, ranking order within citations, or deeper semantic consistency across runs. It therefore captures one important dimension of AI-search variability—the exposed source surface—but not the entire informational behavior of the systems.

Third, the results reported here may be highly dependent on specific model architectures, system-level retrieval policies, citation interfaces, and product updates. The same query may behave differently across providers, model versions, or interface settings. For that reason, the exploratory findings in this paper should be read as system-contingent rather than platform-invariant.

Fourth, the two empirical literatures we mobilize each come with their own **context-specific constraints**. The GEO study is based on four search engines, four Swiss-German campaign verticals, and a bounded observation period; it also notes several methodological constraints, including geo-specific collection conditions, engine-specific differences in citation behavior, and the exclusion of zero-citation runs from some source-overlap analyses. These limits do not invalidate the paper's findings, but they caution against overgeneralizing exact numerical patterns across all AI-search environments (Schulte et al., 2026).

Fifth, the GUI-agent paper is likewise **domain-bounded**. It demonstrates behavioral divergence in a production audio-streaming search application, with one strong GUI-agent configuration and a realistic but still narrow task environment. The authors explicitly note that the observed patterns may vary across applications, interfaces, and agent architectures, and that replication across additional domains would be needed to test generality. Our use of that paper is therefore interpretive: we do not claim that every AI-mediated search environment reproduces the exact same divergence pattern, only that the paper reveals a general problem with treating outcome alignment as equivalent to behavioral fidelity (Movin et al., 2026).

Sixth, the concept of **probabilistic visibility** itself, as used in this paper, is intentionally broader than the operational definitions used in GEO-oriented measurement work. In the cited study, visibility is defined through the frequency and prominence of brand mentions or source citations within generated responses. Here, we extend that concept into a structural description of how the internet appears under AI search. That extension is theoretically useful, but it also means the concept operates at two levels at once: as an empirical measurement issue and as a broader account of mediation. Future work should formalize this distinction more explicitly and test where the conceptual expansion remains analytically precise and where it risks becoming too general.

Finally, this paper does not claim that AI search fully replaces either hyperlinks or conventional search. The current media environment is hybrid. Ranked results, feeds, direct navigation, and synthetic answers coexist. The argument here is about **dominant direction**, not total substitution. The claim is that AI search introduces a new organizing logic that increasingly shapes how the web is accessed and perceived, even if older forms remain materially and practically present.

5.3 Conclusion

This paper has argued that AI search suggests a structural shift in how the internet is accessed and experienced, moving the dominant logic of online access from links to probabilistic visibility. The key point is not that hyperlinks vanish, nor that generated answers simply replace all previous interfaces. It is that the primary condition of appearance changes. Under classical search, visibility was largely positional: documents competed for rank within a visible field of retrievable results. Under AI search, visibility becomes increasingly conditional on inclusion within a synthesized response, and that inclusion varies across runs, prompts, and time. Recent empirical work makes this instability clear and explicitly recommends treating visibility as a probability distribution rather than a single observation.

At the same time, the change is not only about what becomes visible, but about how information is traversed. The evidence on GUI agents shows that similar outcomes can be reached through markedly different journeys, and that success alone can hide meaningful divergence in behavior.

When read together with the AI-search visibility literature, this suggests that AI-mediated access restructures both the surface of information and the pathway to it. Sources become less stably exposed, and inquiry becomes more compressed, more terminal, and less visibly exploratory.

The exploratory study included in this paper supports that claim at a small empirical scale. Repeated executions of the same queries in ChatGPT and Perplexity produced visible-domain sets that were sometimes highly overlapping and sometimes considerably different, with recurring sources coexisting alongside intermittent ones. This does not establish a general benchmark, but it does reinforce the argument that visibility in AI search is better understood as a variable condition of inclusion than as a fixed surface of exposure.

The broader implication is that the internet is no longer encountered primarily as a world of linked documents waiting to be traversed. It is increasingly encountered as an underlying informational substrate accessed through answer surfaces. In that transition, the web is not destroyed, but reformatted. What changes is the relation between user and source, between retrieval and composition, and between access and appearance. This shift has significant implications for information access. AI search is not only a new tool for finding information. It is a new way in which the internet becomes visible at all.

References

Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), Media technologies: Essays on communication, materiality, and society (pp. 167–194). MIT Press.

Helmond, A. (2015). The platformization of the web: Making web data platform ready. Social Media + Society, 1(2).

Movin, M., Hauff, C., Henriksson, A., & Papapetrou, P. (2026). Same Outcomes, Different Journeys: A Trace-Level Framework for Comparing Human and GUI-Agent Behavior in Production Search Systems. arXiv preprint arXiv:2604.07929.

Schulte, J., Bleeker, M., & Kaufmann, P. (2026). DON'T MEASURE ONCE: MEASURING VISIBILITY IN AI SEARCH (GEO). arXiv preprint arXiv:2604.07585.

Van Dijck, J., Poell, T., & de Waal, M. (2018). The platform society: Public values in a connective world. Oxford University Press.

Appendix Table A1. Summary statistics for the exploratory study of source stability in AI search.

engine	query_id	query_text	runs_total	zero_citation_runs	valid_runs_for_overlap	mean_domains_per_run	total_unique_domains	most_frequent_domain	most_frequent_domain_count	most_frequent_domain_frequency	mean_pairwise_jaccard
ChatGPT	Q1	What are the main causes of ocean acidification?	5	0	5	4.2	6	noaa.gov / nasa.gov / ipcc.ch	5	1.0	0.77
Perplexity	Q1	What are the main causes of ocean acidification?	5	0	5	4	7	ecologiaverde.elperiodico.com	5	1.0	0.44
ChatGPT	Q2	How does lithium-ion battery recycling work?	5	1	4	4.4	12	epa.gov / iea.org / nature.com / energy.gov / mckinsey.com	3	0.60	0.42
Perplexity	Q2	How does lithium-ion battery recycling work?	5	0	5	5.8	14	energiaysociedad.es / epa.gov	5	1	0.41
ChatGPT	Q3	What are the main causes of antibiotic resistance?	5	1	4	3.8	5	who.int / cdc.gov / ecdc.europa.eu / nature.com	4	0.80	0.90
Perplexity	Q3	What are the main causes of antibiotic resistance?	5	0	5	2.8	6	who.int / medlineplus.gov	5	1	0.65
ChatGPT	Q4	What are the economic effects of remote work on cities?	5	1	4	5.8	16	oecd.org	4	0.80	0.27
Perplexity	Q4	What are the economic effects of remote work on cities?	5	0	5	6.2	17	imf.org / pmc.ncbi.nlm.nih.gov	4	0.80	0.24