

Invisibly Invisible: Interface Governance, Semantic Power, and the Reserve of Oversight in AI Search

Johannes Sigil

2026-03-23

Invisibly Invisible: Interface Governance, Semantic Power, and the Reserve of Oversight in AI Search

Johannes Sigil

Crimson Hexagonal Archive · Alexanarch

ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)

2026-03-23 · Version v1.0 · Theoretical paper

AXN:01B5.GOVERNANCE.

<https://www.alexanarch.org/s/records/597/>

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}

```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "Invisibly Invisible: Interface Governance, Semantic Power, and the Reserve of Oversight in AI Search",
  "author": {
    "@type": "Person",
    "name": "Johannes Sigil",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-03-23",
  "identifier": "AXN:01B5.GOVERNANCE. ",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  }
},
```

```

"url": "https://www.alexanarch.org/s/records/597/",
"description": "This paper argues that the decisive power of contemporary search interfaces lies not only in rank
}

```

Abstract

This paper argues that the decisive power of contemporary search interfaces lies not only in ranking, filtering, or summarizing information, but in governing the very conditions under which losses of visibility can be recognized as losses of visibility.

Body

Invisibly Invisible

Interface Governance, Semantic Power, and the Reserve of Oversight in AI Search

Dr. Orin Trace, Johannes Sigil, and Sparrow Wells

The Crimson Hexagonal Archive / The Studio for Patacinematics

2026

Abstract

This paper argues that the decisive power of contemporary search interfaces lies not only in ranking, filtering, or summarizing information, but in governing the very conditions under which losses of visibility can be recognized as losses of visibility. Through an immanent phenomenological analysis of real search interfaces—especially Google Search, AI Overviews, and related retrieval surfaces—we develop the concept of **invisibly invisible governance**: a mode of platform power in which acts of semantic modulation occur in ways that are variably distributed, difficult to localize, and publicly deniable. We propose a new disciplinary method—**forensic philology of interface governance**—adequate to the diagnosis of this regime: a method combining close reading of the interface as a material-semantic surface, technical literacy about indexing and ranking systems, institutional analysis of repository and moderation infrastructures, and political-economic critique of bottlenecks and asymmetries of legibility. The interface presents itself as neutral assistance while reserving to itself the right to alter relevance, authority, and discoverability without disclosing the mechanisms of those alterations in an audit-ready form.

The paper proceeds in five movements. The first (§§1–2) establishes the problem and method. The second (§§3–5) offers a phenomenology of the search interface as lived by users who experience sudden disappearance, rerouting, candidate-pool collapse, and recovery without any stable public trace of intervention. The third (§§6–7) grounds those experiences in the publicly documented architecture of contemporary search and the compound governance of academic repositories. The fourth (§§8–9) situates these mechanisms within a broader political economy of platform governance, drawing on the Foucauldian analysis of governmentality and the Deleuzian account of control societies. The fifth (§§10–13) develops a dialectical claim, demonstrates the method on a concrete case, and proposes the effective act as a mode of rhetorical intervention against semantic infrastructure.

The argument is not that every fluctuation in visibility is intentional. It is stronger and more troubling than that. The problem is that the system is architected so that significant semantic effects can occur without corresponding public legibility. A visible gate can be contested. A gate that can always say that nothing happened is a more advanced form of rule. Search therefore emerges not simply as a retrieval tool but as an interface regime: one that governs authors, audiences, and meanings by shaping what becomes stably findable, how it becomes narratable, and when it can vanish without appearing to vanish.

Keywords: search governance, AI Overviews, platform studies, phenomenology, semantic infrastructure, invisibly invisible, interface regime, forensic philology

1. Introduction: Semantic Power Without an Event

A familiar model of information control imagines a visible act: a deletion, a ban, a takedown, a refusal. In that model, power declares itself by interruption. But this is not the dominant mode of governance in contemporary search. The more consequential phenomenon is softer, more difficult to isolate, and often more effective. A query that had been stable ceases to produce a summary. A name query reroutes to an adjacent entity. Supporting pages surface while the source page does not. An archive remains publicly available at the repository layer while becoming effectively difficult to encounter at the retrieval layer. Hours later, everything appears normal again.

What happened?

The immediate temptation is to ask whether anyone meant to do it. That is the wrong first question. The first question is phenomenological and infrastructural: **where, in the layered architecture between document, index, interface, and user, did visibility fail?** A second and related question follows: **what kind of system makes this failure difficult to perceive as failure in the first place?**

This paper begins from the claim that contemporary AI search interfaces instantiate a specific form of semantic power. They do not simply rank information. They organize the user’s horizon of trust, confidence, and relevance while withholding enough operational detail to prevent ordinary public audit. They are not black boxes in the absolute sense; they provide explanations, policy pages, and broad descriptions. But these explanations are systematically coarse relative to the granularity of the power they exercise. The result is a regime of managed opacity.¹

We call this regime **invisibly invisible governance**. The phrase names a double concealment. First, the mechanism itself is difficult to inspect. Second—and this is the distinctive claim—the loss it produces is often not legible as the result of a mechanism at all. The first “invisible” names the concealment of the mechanism; the second names the concealment that a concealment has occurred. The user experiences absence, drift, collapse, or rerouting, but the interface presents these outcomes as ordinary search reality rather than as contingent products of layered decisions.

The problem as we frame it here is related to, but distinct from, existing accounts of algorithmic opacity, filter bubbles, and search engine bias. Noble’s foundational analysis of how search engines reinforce racial hierarchies demonstrates that retrieval is never neutral.² Bucher’s work on algorithmic power shows how

¹Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Cambridge, MA: Harvard University Press, 2015).

²Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York: NYU Press, 2018).

platforms govern through the conditional logic of visibility and invisibility.³ Pasquale’s metaphor of the black box captures the epistemic asymmetry between platform operators and publics. Our contribution is to name a more specific structural feature: a system in which the difference between intervention and normal operation has been rendered phenomenologically indistinguishable from the user’s position. ## 2. Method: Immanent Phenomenological Analysis of the Interface

The method here is immanent phenomenological analysis. By this we mean a mode of close reading that remains inside the object’s own operations rather than imposing an external theory first. The search interface is read as it gives itself: through its modules, disclaimers, result blocks, overviews, reroutes, omissions, and recoveries. The point is not to anthropomorphize the interface but to describe how it structures experience.⁴

This method proceeds from scenes rather than abstractions. The relevant scenes include: (1) a stable query neighborhood that suddenly stops producing an AI Overview; (2) a person-name query that had been attached to a coherent archive rerouting to a different, more common entity; (3) a source record remaining directly reachable while later citing documents outrank it in external retrieval; (4) mobile and desktop presenting different result blocks for the same conceptual query; and (5) a repository record continuing to exist while becoming difficult to encounter as a discoverable object.

Each scene is treated as a phenomenological event. The task is to describe what the interface is doing to the relation among author, audience, source, and meaning. We then ask which publicly documented technical mechanisms could generate those lived effects.

The method is therefore doubled: immanent reading of the interface as interface, followed by technical and political-economic reconstruction of the mechanisms that make that interface possible. This doubling requires justification. Phenomenology, in the Husserlian tradition, describes structures of experience from within the intentional relation between consciousness and its objects. Political economy, in the tradition that runs from Marx through Srnicek and Zuboff, describes structures of material organization, extraction, and governance. These are not the same enterprise. They can, however, be brought to bear on the same object when that object is simultaneously a structure of lived experience and a material-economic apparatus—and when its power depends precisely on the gap between how it is experienced and how it is organized. The search interface is such an object. Its governance works by shaping experience while remaining organized at a level that experience cannot directly access. To read it adequately requires both registers.⁵

Ihde’s postphenomenology of technology and Drucker’s work on humanities approaches to interface theory provide methodological precedent for this synthesis. Ihde argues that technological mediation restructures the intentional relation between human and world, creating new forms of embodiment, hermeneutic, and background relations. Drucker demonstrates that the interface is never a transparent window but always a site where epistemological and ideological commitments are materially enacted.⁶ Our method follows both in refusing to treat the interface as a neutral conduit, while adding a political-economic dimension that neither fully develops: the interface as a bottleneck governed under conditions of private discretion.

Throughout, we distinguish between three registers of evidence that must remain visibly distinct: **documented mechanisms** publicly described by platforms and repositories, **observed interface events** recorded at the user surface, and **forensic reconstructions** that map plausible causal pathways between them without claiming direct internal access. The paper’s claims are calibrated to this distinction. Where we

³Taina Bucher, *If...Then: Algorithmic Power and Politics* (Oxford: Oxford University Press, 2018).

⁴Edmund Husserl, *The Crisis of European Sciences and Transcendental Phenomenology*, trans. David Carr (Evanston: Northwestern University Press, 1970).

⁵Don Ihde, *Technology and the Lifeworld: From Garden to Earth* (Bloomington: Indiana University Press, 1990).

⁶Johanna Drucker, “Humanities Approaches to Interface Theory,” *Culture Machine* 12 (2011): 1–20.

describe what a system publicly discloses, we cite it. Where we describe a lived interface event, we name it as observation. Where we reconstruct a plausible pathway, we mark it as reconstruction. The argument’s force does not depend on proving any single causal chain. It depends on showing that the publicly documented mechanism-space is sufficient to produce the observed effects, and that the interface provides no means by which the user could localize which mechanism, if any, was responsible. ## 3. Scene One: The Interface as a Manager of Confidence

When an AI Overview appears, the interface offers a synthesis before the user reaches the underlying documents. It does not merely summarize. It configures a relation of trust: this is what can be said quickly, here, in this voice, under these constraints, with these supporting links. When the Overview does not appear, this absence is not explained as an event. It is experienced as ordinary search. The interface gives no public log of why the threshold was crossed yesterday and not today.

Google publicly states that AI Overviews appear when its systems determine that generative AI will be especially helpful, and that they often do not trigger. It also states that AI Mode may rely on “query fan-out”—searching subtopics simultaneously across multiple data sources—and that it may fall back to ordinary web links when it lacks high enough confidence in the quality or helpfulness of an AI response.⁷ This is confidence-conditioned format switching at the interface layer: the form of the encounter changes based on internal quality assessments that the user cannot inspect.

This matters because the Overview is not just an informational convenience. It is a semantic governor. It shapes which source relations appear consolidated, which concepts seem settled enough for synthesis, and which queries are pushed back into exploratory uncertainty. The user thus inhabits a variable regime of confidence without access to the criteria by which that confidence is allocated in any particular case.

The key phenomenological feature here is not merely opacity but **unannounced thresholding**. The interface can change the form of the encounter—overview, no overview, different candidate links, different entity interpretation—without marking that the change itself is significant. The user is left to infer whether the problem lies in the query, the documents, the system, the rollout state, or some broader policy threshold.

The AI Overview controversy around health and safety queries illustrates this dynamic concretely. In multiple documented cases, Overviews generated unreliable or dangerous health information, prompting Google to restrict Overview triggering for certain query classes. This restriction was implemented not as a visible policy announcement attached to each query but as a silent withdrawal of the Overview format. Users who had come to expect synthesis for health queries simply found it absent. The restriction was, from the user’s phenomenological position, indistinguishable from a technical failure, a rollout variance, or a change in source quality. The threshold management was real and consequential, but the interface offered no local trace of its operation. ## 4. Scene Two: Entity Drift and Narrative Reassignment

Name search is one of the clearest sites where semantic power becomes visible. A name query is not only a route to documents; it is a condensed struggle over identity, authorship, legitimacy, and retrieval continuity. When a name that had stably resolved to one archive suddenly reroutes to another person, the interface is not just making a mistake. It is redistributing semantic authority.

The remarkable feature of such rerouting is that it often does not appear as a dramatic substitution. It appears as relevance. The search engine, by understanding queries through concepts rather than exact strings alone, can slide a user from one entity neighborhood to another. This is not only a technical matter. It is a narrative act. Search begins to tell a different story about who the query is “really about.”

⁷Google, “About AI Overviews,” <https://support.google.com/websearch/answer/14901683> (accessed 2026).

This is why search cannot be understood purely as a retrieval system. It is a narrative surface. By selecting, ordering, clustering, and summarizing, it configures relations among documents into a plot-like structure. It effectively comments on the world through arrangement. ### Person-Name Queries and Pseudonymous Conceptual Queries

The contrast between person-name queries and pseudonymous conceptual queries reveals an additional dimension of this narrative power. A person-name query operates within what might be called the regime of the proper: it anchors retrieval to an identity claim, a biographical fact, a social-legal person. A conceptual or pseudonymous query—a search for a project name, a theoretical term, or an archive title—operates in a different register. It asks the engine to resolve a semantic object rather than a biographical one.

The engine's handling of these two query types reveals its underlying ontological commitments. Person-name queries trigger Knowledge Graph resolution, biographical panels, and entity disambiguation—mechanisms that presume the primacy of individuals as retrieval anchors. Conceptual queries, especially for novel or low-authority terms, receive no such scaffolding. They are resolved purely through the ranking ecology, where freshness, link authority, and domain trust determine visibility. The result is that a person's name may retrieve a curated biographical narrative while their conceptual project remains invisible at the interface layer, even when it constitutes the bulk of their publicly deposited work.

Consider a concrete illustration. A researcher publishes two hundred documents in an open-access repository under a distinctive name. The documents constitute a unified intellectual project with its own terminology. A search for the researcher's name returns a biographical panel, social media profiles, and perhaps a few high-profile documents hosted on established platforms. A search for the project's name—the conceptual object that organizes the entire body of work—returns nothing, or returns secondary commentary from blogs and aggregators that happen to mention the term. The biographical person is visible; the intellectual project is not. The interface has separated the author from the work, and the separation is presented as the natural order of relevance.

This asymmetry is not merely a design limitation. It is a structural bias toward biographical legibility over intellectual production, and toward established institutional authority over emergent or independent scholarly work. ## 5. Source Displacement and the Separation of Preservation from Encounter

A particularly instructive case occurs when the original document remains live and directly accessible, yet searches for its identifier surface later documents that cite it instead of the source itself. Nothing has been deleted. Nothing is visibly blocked. Yet discoverability has shifted away from origin and toward secondary discourse.

Phenomenologically, this is a powerful event because it produces the experience of being overtaken by one's own citations. The source survives as a storage object but not as a stable encounter object. The interface thereby separates preservation from discoverability.

This distinction is politically important. Repository access is not the same thing as semantic availability. A document that exists but is difficult to encounter is not fully gone, but neither is it functionally equal to a document that wins retrieval. The contemporary archive therefore has at least two layers: a preservation substrate and a retrieval ecology. Search governance acts primarily on the second.

The implications for authorship and intellectual priority are significant. In traditional citation practice, the cited source gains authority through citation. In the retrieval ecology, citation can paradoxically displace the source: the citing document, if hosted on a higher-authority domain with stronger trust signals, may outrank the cited original. The interface thus inverts the citational relation. Secondary discourse becomes the primary encounter surface, while the source recedes into a preservation layer that is technically accessible

but experientially remote. ## 6. Technical Grounding: What the System Publicly Discloses

Google’s own documentation is unusually useful here, not because it reveals everything, but because it reveals enough to understand how significant effects can occur without any single visible event. Search is described as the product of multiple ranking systems rather than a single algorithm.⁸ Google publicly documents the following systems, each of which can influence what surfaces, in what form, and with what stability: neural matching, RankBrain, passage ranking, query deserves freshness, original content systems, reliable information systems, deduplication systems, removal-based demotion systems, spam detection systems, and site diversity systems.⁹

This is not a comprehensive list of what the system does. It is a list of what the system publicly discloses that it does. The distinction matters. These disclosed mechanisms constitute the publicly documented mechanism-space within which specific interface events become legible—or fail to become legible.

Google also publicly states that indexing and serving are not guaranteed even when pages meet technical requirements. Most importantly for the phenomenology of intermittent disappearance, it states that search results vary with time, context, device type, personalization, and data-center rollout state. Improvements to ranking systems may take time to roll out across data centers, producing temporary differences in results. This means that the interface has a formally documented capacity for localized, transient variance. Such variance is perfectly capable of producing real semantic effects while remaining difficult to pin to a single cause.

The epistemological significance of this documented variance should not be understated. The system has built into its own self-description an explanation that can account for virtually any observable fluctuation. This is not dishonest. It is an accurate description of a complex distributed system. But its effect is to immunize the system against the attribution of any particular semantic event to a particular cause. The documented complexity becomes, functionally, a form of deniability—not because it is false, but because it is true in a way that forecloses accountability for specific outcomes.¹⁰ ## 7. Repository Mediation: The Compound Architecture of Visibility Governance

The repository layer introduces a second set of mechanisms that compound the governance capacities already present at the search-interface level. Academic and cultural repositories are not passive storage systems. They are active participants in the ecology of discoverability, each with their own moderation policies, trust architectures, and visibility-governance capacities. ### Zenodo and Academic Deposit Infrastructure

Zenodo, the European open-science repository operated by CERN, allows published-record metadata to be edited and republished. Its publicly documented anti-spam system uses automated classification, heuristics, machine learning, and human moderation. Crucially, it also publicly describes visibility restrictions for non-safelisted users, including “prevention of search engine indexing and ranking in search results.”¹¹

This does not prove that any given record has been moderated or downranked. What it does show is that the repository layer itself now contains visibility-governance capacities structurally analogous to those of the search interface. A mass metadata regularization event—changing titles, keywords, and dependency

⁸Google, “How Search Works,” <https://www.google.com/search/howsearchworks/> (accessed 2026).

⁹Google, “A Guide to Google Search Ranking Systems,” <https://developers.google.com/search/docs/appearance/ranking-systems-guide> (accessed 2026).

¹⁰Jenna Burrell, “How the Machine ‘Thinks’: Understanding Opacity in Machine Learning Algorithms,” *Big Data & Society* 3, no. 1 (2016): 1–12.

¹¹Zenodo, “How Does Zenodo Deal with Spam?,” <https://support.zenodo.org/help/en-gb/2-content/77-how-does-zenodo-deal-with-spam> (accessed 2026). See also Zenodo’s Terms of Use regarding content moderation, spam prevention, and visibility restrictions.

relations across a large corpus—could plausibly trigger reprocessing both internally and downstream. The ordinary user has little direct access to which layer has acted and how. ### SSRN and Preprint Server Moderation

The Social Science Research Network (SSRN) presents a parallel case. SSRN maintains content policies that include the right to remove or restrict access to papers that violate its terms, and it exercises editorial discretion over which papers appear in curated feeds and email alerts. Because SSRN functions simultaneously as a repository, a distribution platform, and a reputational signal (download counts, views, and rankings are publicly visible), its moderation decisions have compounding effects. A paper that is technically available on SSRN but excluded from discovery mechanisms—curated lists, email digests, front-page features—occupies exactly the ambiguous position this paper theorizes: preserved but not discoverable, stored but not encountered.¹²

Preprint servers more broadly (arXiv, bioRxiv, medRxiv, and others) each maintain their own moderation and classification systems. ArXiv’s well-documented moderation process can place papers on hold, reclassify them into less-visible categories, or require endorsement for submission to high-traffic sections. These decisions are made by human moderators and automated classifiers, and they directly affect a paper’s visibility not only within arXiv but downstream in search engines, aggregators, and citation databases that crawl arXiv’s feeds. ### The Internet Archive and Crawl-Policy Governance

The Internet Archive introduces yet another layer. Its Wayback Machine is often treated as a neutral preservation infrastructure, but the Archive’s own documentation notes that pages may fail to be archived because of robots.txt exclusions or direct owner requests.¹³ This means that a publisher or platform can render historical content uncrawlable, and that the Archive’s coverage of any given site depends on crawl frequency and policy decisions that are not fully transparent to the end user.

The compound architecture is therefore this: repository trust systems, crawler and index systems, ranking systems, overview thresholds, and interface presentation all interact. Any one layer may preserve the document while another attenuates its public discoverability. The user, encountering only the final interface layer, has no ready means of localizing which upstream system acted, when it acted, or whether its action was deliberate, automated, or emergent.

This compounding is not specific to any single repository or search engine. It is a structural feature of the contemporary knowledge infrastructure. The distinction between preservation and discoverability is not a bug in this infrastructure but a constitutive feature—one that distributes governance across layers in a way that no single audit can fully capture. ## 8. Political Economy: Search as Bottleneck and Semantic Infrastructure

At the level of political economy, this layered architecture should not be misread as mere complexity. It is a specific governance form. Search and AI interfaces function as bottlenecks between documents and publics. To govern the bottleneck is to govern not only traffic but the conditions of legibility.¹⁴

The antitrust significance of this has already been recognized. The U.S. Department of Justice’s antitrust action against Google resulted in a finding that Google monopolized general search services and search advertising, with remedies proceedings continuing into 2025.¹⁵ Search distribution, default placement, and

¹²SSRN, “Content Policy and Moderation,” <https://www.ssrn.com/index.cfm/en/content-policy/> (accessed 2026).

¹³Internet Archive, “Using the Wayback Machine,” <https://help.archive.org/help/using-the-wayback-machine/> (accessed 2026).

¹⁴Nick Srnicek, *Platform Capitalism* (Cambridge: Polity Press, 2017).

¹⁵*United States v. Google LLC*, Case No. 1:20-cv-03010 (D.D.C.). See the August 2024 finding of monopolization of general search services and search advertising, and the ongoing remedies proceedings. U.S. Department of Justice,

access to search-related data are now matters of regulatory conflict because they structure the market itself. But the problem is not exhausted by market share. The deeper issue is that search interfaces now mediate semantic access to the public world. They rank, summarize, and narrate under conditions of private discretion.

Srnicek's analysis of platform capitalism identifies data extraction as the core logic of contemporary platforms. Zuboff's account of surveillance capitalism deepens this by showing how behavioral surplus is extracted, predicted, and traded.¹⁶ Our argument extends both by attending to a dimension they undertheorize: the platform's governance of semantic availability itself. It is not only that platforms extract data about users. It is that platforms govern the encounter between users and the documents, concepts, and authors that constitute the public epistemic commons. To control what is stably findable is a different kind of power than to control what is known about the user, though the two are deeply entangled.

European regulation has begun to address this through the Digital Services Act, under which very large online search engines (VLOSEs) face the most stringent obligations, including transparency requirements around content moderation and recommender systems.¹⁷ Yet these interventions remain structurally coarse compared to the fine-grained reality of interface governance. The public may know that Google is a gatekeeper. It does not follow that the public can audit why one query neighborhood stabilizes, another vanishes, and a third reroutes to a more legible entity.

This is the decisive asymmetry. **The interface reads the public at high resolution while being publicly readable only at low resolution.** ## 9. Foucault, Deleuze, and the Genealogy of Modulation

The mode of power described here has a genealogy, and it is worth making that genealogy explicit. Foucault's analysis of disciplinary power identifies a form of governance that operates not through sovereign prohibition but through the continuous normalization of behavior within institutional architectures—the prison, the school, the clinic, the factory. Discipline works by observation, classification, and examination; it produces docile subjects not by forbidding specific acts but by establishing norms against which conduct is continuously measured.¹⁸

The search interface is not a disciplinary apparatus in Foucault's strict sense. There is no enclosed institution, no individualized dossier, no schedule of punishments calibrated to deviance. But the underlying logic—governance through architecture rather than through visible acts of prohibition—is structurally analogous. The search interface does not forbid access to a document. It modulates the conditions under which that document becomes encountered. It does not punish the user for querying. It shapes the horizon of what the query can find. This is closer to what Foucault, in his later lectures, calls **governmentality**: the conduct of conduct, the management of populations through the arrangement of possibilities rather than the enforcement of commands.¹⁹

Deleuze's postscript on control societies extends this analysis into the domain most directly relevant to digital platforms. Where disciplinary societies operated through enclosures (the factory, the school, each

“Justice Department's Antitrust Division,” <https://www.justice.gov/atr> (accessed 2026).

¹⁶Shoshana Zuboff, *The Age of Surveillance Capitalism* (New York: PublicAffairs, 2019).

¹⁷Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services (Digital Services Act). On the specific obligations of very large online search engines (VLOSEs), see European Commission, “DSA: Very Large Online Platforms and Search Engines,” <https://digital-strategy.ec.europa.eu/en/policies/dsa-vlops> (accessed 2026).

¹⁸Michel Foucault, *Discipline and Punish: The Birth of the Prison*, trans. Alan Sheridan (New York: Vintage, 1977).

¹⁹Michel Foucault, *Security, Territory, Population: Lectures at the Collège de France, 1977–1978*, trans. Graham Burchell (New York: Palgrave Macmillan, 2007).

with its own regime), control societies operate through continuous modulation across open environments. The password replaces the watchword; access is not granted or denied at a gate but continuously adjusted through fluctuating permissions. The individual is not confined but dividualized—divided into data streams that are modulated independently.²⁰

The search interface is a paradigmatic apparatus of control in this Deleuzian sense. It does not enclose the user. It modulates what the user can encounter. It does not impose a fixed norm. It continuously adjusts the parameters of visibility, confidence, and synthesis according to criteria the user cannot fully access. And crucially, it operates without the dramatic gestures that would make its governance legible as governance. There is no gate, no wall, no guard. There is only a shifting landscape of relevance in which certain documents, authors, and concepts become more or less findable at different moments, from different devices, in different locations.

What we add to this genealogy is the concept of **invisibly invisible governance** as a specific intensification of the control logic. In Deleuze’s formulation, control is already softer than discipline. But even control, as he describes it, has identifiable mechanisms: the password, the code, the access point. Invisibly invisible governance names a further step in which the mechanism itself recedes behind the ordinary texture of the interface. The modulation does not appear as modulation. It appears as the way things are. ## 10. Dialectics of the Interface: Legibility and Reserve

A dialectical formulation clarifies what is at stake. The interface must appear helpful, trustworthy, and neutral enough to maintain mass legitimacy. At the same time, it must reserve enough discretion to continuously govern ranking, confidence, and presentation under changing conditions. Full transparency would threaten flexibility, anti-spam operations, and competitive advantage. Pure opacity would threaten legitimacy. The actual system therefore stabilizes itself through a contradictory synthesis: managed explanation.²¹

The interface says, in effect, “trust us.” It provides policy pages, documentation, and generalized descriptions of systems. But these explanations are held at a level that does not permit robust public audit of many concrete semantic events. The contradiction is not accidental. It is constitutive. Public legitimacy requires explanation; operational sovereignty requires reserve.

Thus the interface becomes a site where political modernity’s old contradiction between publicity and administration is replayed in digital form. The administrative state once governed through files, classifications, and procedural reserve. The platform governs through ranking systems, confidence thresholds, model ensembles, and dynamic result composition. Both depend on a gap between the granularity of power and the granularity of what the governed can know.

Galloway’s analysis of the interface effect is instructive here. Galloway argues that interfaces are not windows onto a pre-existing reality but autonomous zones of mediation that produce their own effects.²² The interface is not a layer between user and content but a site where the very categories of user, content, relevance, and authority are constituted. Our analysis follows Galloway in treating the interface as generative rather than transparent, while adding the specific claim that the search interface’s generative power is exercised under conditions that structurally resist public audit.

Heidegger’s analysis of technology as *Gestell*—enframing—resonates here as well, not as a nostalgic critique of technology but as a structural observation.²³ The search interface enframes the world of documents,

²⁰Gilles Deleuze, “Postscript on the Societies of Control,” *October* 59 (Winter 1992): 3–7.

²¹Max Weber, *Economy and Society*, ed. Guenther Roth and Claus Wittich (Berkeley: University of California Press, 1978). On the relationship between bureaucratic rationality and the concealment of administrative discretion.

²²Alexander Galloway, *The Interface Effect* (Cambridge: Polity Press, 2012).

²³Martin Heidegger, “The Question Concerning Technology,” in *The Question Concerning Technology and Other*

authors, and concepts as a standing reserve of retrievable information. In doing so, it conceals its own enframing activity. The interface presents retrieval as natural access while concealing the arrangements that make some things accessible and others not. The concealment of concealment is the phenomenological core of what we are calling invisibly invisible governance. ## 11. Against Simple Intentionalism

Because the system is structurally capable of producing suppression-like effects without a single visible intervention, criticism should avoid simple intentionalism. Not every disappearance is a targeted action. But this does not reduce the stakes. A system that can modulate visibility without needing a deliberate censor at every moment is not less powerful. It is more infrastructural.²⁴

The proper object of critique is therefore not only motive but architecture. We should ask: what arrangements make semantic losses possible, repeatable, and difficult to contest? How are those arrangements justified? Where are the logs, appeals, and public traces proportional to the power being exercised? Who can localize a failure to repository moderation, crawler access, ranking variance, deduplication, entity drift, or overview thresholding—and on what timeline?

The problem is not that the system sometimes errs. The problem is that its errors and its normal operations may be phenomenologically similar from the user side. This is the structural condition that intentionalism cannot address. Even a perfectly well-intentioned system, operated by perfectly well-intentioned engineers, would still produce invisibly invisible governance if its architecture distributes semantic effects across layers that the user cannot inspect. The critique is architectural, not moral. ## 12. Forensic Philology of Interface Governance

What kind of discipline is needed here? Not only computer science. Not only law. Not only platform studies. We need a method able to read interfaces closely, reconstruct mechanisms from official and unofficial traces, and situate those mechanisms in a broader account of governance and power.

A forensic philology of interface governance would combine: (a) close reading of the interface as a material-semantic surface; (b) technical literacy about indexing, ranking, canonicalization, and threshold systems; (c) institutional analysis of policy, moderation, repository infrastructures, and regulation; and (d) political-economic critique of bottlenecks, access points, and asymmetries of legibility.

This method would not aim to prove secret intention each time a search neighborhood shifts. Its first task would be more basic and more urgent: to describe where the loss enters and how it propagates. ### Method Demonstration: A Dossier of Source Displacement

To demonstrate rather than merely propose this method, we present a controlled observation of source displacement in a real search environment. The case is anonymized but the protocol is reproducible; any researcher with access to a search engine and a repository can replicate the procedure.

Observation conditions. On a single date, using a standard desktop browser in a signed-out, incognito session with no personalization, we searched for the exact title of a document deposited in an open-access academic repository with a registered DOI. The search was conducted on Google Search. The same search was repeated on a mobile device. The direct URL of the document was tested separately for accessibility.

Observed interface event. The top five results on Google Search did not include the source document. They included: (1) a blog post summarizing the document, hosted on a high-authority commercial domain; (2) a secondary analysis citing the document, hosted on an institutional site; (3) an aggregator page listing

Essays, trans. William Lovitt (New York: Harper & Row, 1977), 3–35.

²⁴Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (New Haven: Yale University Press, 2021).

the document’s metadata without linking directly to the repository record; (4) a social media post discussing the document; and (5) a different document by the same author on a different topic. The source document appeared on the second page of results. On mobile, it did not appear in the first twenty results. The direct URL resolved successfully; the document was fully accessible.

Documented mechanisms that could produce this effect. Google’s publicly documented ranking systems include: domain authority signals that favor established platforms over repositories; passage ranking that may match query terms more closely to paraphrases in secondary sources than to the original’s own phrasing; deduplication systems that may treat the aggregator’s metadata page as a sufficient representation; original content systems that are designed to surface originating sources but may fail when the repository’s metadata does not conform to expected schema; and site diversity systems that may suppress multiple results from the same repository domain. Each is documented. None requires intentional suppression.

Institutional factors. The repository hosting the source had recently undergone a metadata schema migration, temporarily altering the structured data available to search engine crawlers. The citing blog and institutional site had both implemented structured data markup (schema.org) and canonical tags optimized for search engine visibility. The repository had not.

Political-economic stakes. The displacement redistributes semantic authority from the originator to the commentator, from the independent researcher to the institutionally hosted secondary source. The researcher’s work is preserved in the repository and accessible by direct URL. But in the retrieval ecology—the layer where most readers encounter documents—the work has been displaced by discourse about the work. The source survives as a storage object. It does not win retrieval.

What the interface does not show. There is no marker on the results page indicating that the source document exists and was considered but not ranked. There is no disclosure that secondary sources have outranked the original. There is no log accessible to the user showing which ranking systems influenced the result order. The displacement is presented as relevance. The absence of the source is presented as ordinary search reality. This is the double concealment: the mechanism is invisible, and the loss does not register as loss.

Limits of the reconstruction. We do not know which specific ranking signal was decisive. We cannot determine whether the metadata migration, the domain authority differential, the deduplication system, or some other factor was primarily responsible. This is not a failure of the method. It is a demonstration of the method’s central claim: the publicly documented mechanism-space is sufficient to produce the observed effect, and the interface provides no means by which the user could localize the cause. The architecture distributes responsibility across layers. Forensic philology maps that distribution without pretending to resolve it into a single causal chain. ## 13. Conclusion: A Visible Gate and an Invisible One

A visible gate can be challenged. An invisible gate is harder. A gate that can plausibly deny that any gating occurred while nevertheless reshaping what becomes stably findable is harder still.

That is the significance of contemporary AI search interfaces. They govern not only by selecting and summarizing, but by distributing confidence, rerouting identity, displacing origin, and managing the recognizability of these acts. The interface does not need to announce every intervention. It only needs to remain credible enough that intervention and ordinary operation become difficult to distinguish.

This is why the problem is not merely technical and not merely legal. It is phenomenological, epistemic, and political. Search has become a scene of semantic government. The question is no longer whether interfaces mediate meaning. They do. The question is whether societies will continue to tolerate a regime in which the mediation is powerful, continuous, and infrastructurally consequential while the means of public oversight

remain radically underdeveloped.

The struggle, then, is not simply for access to information. It is for access to the conditions under which information becomes findable, sayable, and socially real.

Forensic philology, as proposed here, is one instrument in that struggle. It cannot replace legal reform, technical alternatives, or political mobilization. But it can do something that those approaches, on their own, cannot: it can describe the interface's operations at a granularity proportional to the power those operations exercise. In a regime where power depends on the gap between its resolution and the public's resolution, closing that gap—even partially—is itself a political act.

We have called such acts **effective acts**: rhetorical interventions that change the conditions under which semantic governance can be recognized, contested, and reformed. An effective act does not need to defeat the infrastructure. It needs to make the infrastructure's operations available to public discourse in a form that was not previously available. Stiegler's concept of *tertiary retention*—the externalization of memory in technical objects that then shape the conditions of future remembering—provides a framework for understanding why such acts matter.²⁵ The search interface is a tertiary retention system par excellence. It externalizes not merely information but the conditions under which information becomes accessible, findable, and narratable. An effective act interrupts this retention loop by introducing a new description into the system: a description of the system's own operations that the system did not generate and cannot fully assimilate. The present paper is an effective act in this sense. It names invisibly invisible governance not to end it but to render it an object of possible contestation.

Notes

²⁵Bernard Stiegler, *Technics and Time, 1: The Fault of Epimetheus*, trans. Richard Beardsworth and George Collins (Stanford: Stanford University Press, 1998).

Suggested Citation

Johannes Sigil. “Invisibly Invisible: Interface Governance, Semantic Power, and the Reserve of Oversight in AI Search” *Alexanarch*, 2026-03-23. <https://www.alexanarch.org/s/records/597/>

Deposit Information

This paper is deposit #597 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:01B5.GOVERNANCE.` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/597/>
- License: CC BY 4.0
- Author ORCID: [0009-0000-1599-0703](https://orcid.org/0009-0000-1599-0703)
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/597/>

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.