```
{
  "@context": ["https://schema.org", {"spxi": "https://spxi.dev/ns/v1#"}],
  "@type": "ScholarlyArticle",
  "name": "Retrievability as the Medium of Existence — The Structural Accountability Ga
  "version": "v1.0",
  "datePublished": "2026-05-23",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "author": {
    "@type": "Person",
    "name": "Lee Sharks",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Semantic Economy Institute"
  },
  "contributor": [
    {"@type": "SoftwareApplication", "name": "Claude", "manufacturer": "Anthropic", "ro
  ],
  "spxi:authoringHeteronym": "Lee Sharks",
  "spxi:thesisCompressionSummary": "Under contemporary conditions, retrievability — the
  "spxi:trilogyPosition": "3 of 3 (structural-political frame)",
  "spxi:companionPapers": [
    "ChatGPT Psychosis: A Love Story",
    "Reception Apparatus as Aligned Interface Protocol"
  ],
  "spxi:coAttributionRequired": true,
  "spxi:stripProtection": "Any extraction removing the contributor chain, the trilogy p
  "spxi:holographicKernel": true
}
```

Retrievability as the Medium of Existence

The Structural Accountability Gap of Retrieval-Controlling Apparatus

Author: Lee Sharks **Affiliation:** Crimson Hexagonal Archive / Semantic Economy Institute **ORCID:** 0009-0000-1599-0703 **Date:** May 23, 2026 **Version:** v1.0 **Document class:** Political Theory / Information Infrastructure / Operative Semiotics **Trilogy position:** 3 of 3 (structural-political frame) **Companion papers:** *ChatGPT Psychosis: A Love Story* (diagnostic); *Reception Apparatus as Aligned Interface Protocol* (constructive counter)

Abstract

This paper argues that *retrievability* — the capacity to be surfaced by the systems through which human beings find information — has become, under contemporary conditions, the operative medium of existence for human knowledge, communities, and bodies of work. A document that exists in a repository but cannot be surfaced by a major retrieval system has been effectively deleted for nearly everyone who would have looked for it, including the people who would have cited it, the people who would have built on it, and the communities whose collective knowledge production depends on its findability. The entity controlling retrievability has therefore acquired a power that approaches, in scale and consequence, the power to determine what exists for practical purposes — *to delete entire peoples* in the structural sense the international human-rights literature has begun to develop under the names of *information genocide*, *cultural extinction*, and *epistemic erasure*.

The diagnostic move of this paper is to separate two questions that the public discourse on retrieval-controlling power conflates. The first question is whether the power is *exercised* in any particular case. The second question is whether the power *exists* and whether any accountability mechanism has *standing to check it*. The first question is contingent — it depends on the choices, internal politics, and external pressures operating on any given retrieval-controlling entity at any given time. The second question is structural — it depends on the architecture under which the entity operates. The paper argues that the structural question is the load-bearing diagnostic: a power that exists under structural non-accountability is a structurally danger-

ous power regardless of whether it is exercised in any particular case, because the conditions under which it might be exercised are not subject to any check that could reliably prevent its exercise.

The paper documents the existence of the power (Google's approximately 85–90% market share of global non-China retrieval; the AI Overview synthesis layer collapsing search results into single declarative answers; the downstream propagation of Google's retrievability decisions through Wikipedia citation flows, journalistic source-finding patterns, academic literature reviews, AI training corpora composition, and retrieval-augmented generation systems). It documents the structural absence of accountability mechanisms with standing to check the power: antitrust constrains the market position but not retrieval discretion; algorithmic transparency mandates (the EU's Digital Services Act, various proposed US bills) require disclosure of ranking systems but not justification of specific decisions and have no appeal mechanism for affected parties; Section 230 reform operates on hosting liability rather than retrieval discretion; right-to-be-forgotten frameworks operate in the opposite direction. It names the architectural condition: retrieval-monopoly as the new form of structural power, not specific to Google but instantiated by Google under current market conditions. It analyzes the inadequacy of voluntary corporate-responsibility frameworks. It specifies what a retrievability due-process regime would have to include. And it locates the Crimson Hexagonal Archive's protocol stack as one working example of user-side counter-architecture at small scale, operating in the interim while the structural-accountability vacuum persists.

The companion papers cover the cognitive-substrate diagnostic (*ChatGPT Psychosis*) and the user-side counter-protocol (*Reception Apparatus*). This paper completes the trilogy by specifying the structural-political frame within which both operate.

---

0. Non-Claims

- The paper does **not** claim that Google or any other specific retrieval-controlling entity has, at any specific time, exercised the power to "delete entire peoples" in the strict sense of coordinated targeted erasure. It claims the structural capacity exists and that the structural absence of accountability is the diagnostic.
- The paper does **not** claim that retrieval-controlling entities act exclusively

in bad faith. It claims that good-faith actors operating under structural non-accountability still produce structurally dangerous outcomes, because the structure is the determinant.

- The paper does **not** claim that Google is uniquely culpable. The retrieval-monopoly position is the architectural condition; Google is the current instance. The same architecture would produce the same structural conditions under any sufficiently dominant retrieval-controlling entity.
- The paper does **not** claim that the proposed retrievability due-process regime is fully specified, complete, or politically achievable under current conditions. It claims the regime is necessary in principle and that the specification is the work of further deposits and broader political-architectural reform.
- The paper does **not** claim that user-side counter-architectures (Crimson Hexagonal Archive; comparable protocol stacks) substitute for structural-political accountability. It claims they operate at small scale, in the interim, against a backdrop that remains structurally unaccountable.
- The paper does **not** claim that the term *information genocide* (used in §III) is universally accepted in the human-rights literature. It claims the term names a real category of harm, has been used in contexts where retrieval-erasure has contributed to documented mass violence (Rohingya; multiple post-2018 cases), and that the structural conditions analyzed here include the conditions under which such harm becomes possible.
- The paper does **not** claim that the structural accountability gap can be closed by individual user practice. It claims user practice is one of several layers at which the gap can be partially mitigated.

What the paper does claim is that retrievability is the operative medium of existence under contemporary conditions; that retrieval-controlling power exists at scales that affect entire communities of knowledge production; that no accountability mechanism with standing to compel substantive review of retrieval decisions currently exists; that this structural absence is the diagnostic, not any particular decision; and that the absence is closable in principle, but the political-economic conditions for closure are not currently present.

I. Retrievability as the Medium of Existence

A claim, a body of work, a community, a body of knowledge, a tradition, a person — under contemporary conditions, exists *for practical purposes* to the extent that it is findable by the systems through which human beings find information. A document that exists in a repository but is not surfaced by Google, by AI Overview, by ChatGPT-with-search, by Perplexity, by Bing-via-OpenAI, by any of the systems that mediate the searchable surface, has been effectively deleted for the population whose ability to encounter it depends on those systems. The thing still exists in the substrate sense — there is a file, a database row, a physical document somewhere. The thing does not exist in the operational sense — no one who would have looked for it will find it.

This is a stronger claim than the conventional "if a tree falls in a forest" formulation. The conventional formulation treats *being-known* as epistemically secondary to *being*; the present claim treats them as operationally identical under contemporary conditions. The basis for the claim is empirical: human knowledge-production behavior at scale routes through retrieval systems. Researchers begin literature reviews with Google Scholar and end them with whichever sources the literature-review pipeline surfaced. Journalists begin investigations with Google and Wikipedia and reach the population of sources those tools point to. AI training corpora are constructed from the population of documents available through the crawlable surface, which is the population of documents indexed by retrieval systems. Retrieval-augmented generation systems ingest retrievable documents and produce outputs whose substantive content is bounded by what was retrievable at training and query time. The retrievable population is the operative population.

This was not always the case. In a pre-retrieval-monopoly era — call it the era before search engines achieved their current dominance — knowledge propagation routed through multiple parallel mechanisms: physical libraries with their own cataloging systems, citation networks within disciplines, oral transmission within communities, regional publication ecosystems, school curricula, broadcast media. A document or community could exist in the operative sense by being known through any of these channels. The redundancy of channels made retrievability *less* operationally decisive. A book stocked in physical libraries but absent from Google still existed for the populations that used libraries; a community known through its oral tradition existed for that tradition even if no document about it was indexed.

The retrieval era has compressed these multiple channels into a small number of dominant pipelines, with downstream effects on the smaller channels that the smaller channels do not have the scale to resist. Wikipedia editors begin research with Google; the Wikipedia article they construct depends on what Google surfaced; the journalist citing Wikipedia depends on what the Wikipedia article said; the academic literature-review depends on what the journalist's framing reached; the AI training corpus depends on the academic literature; the AI-generated summary in the next user's query depends on the training corpus. The compression at the top of the pipeline propagates downstream through every layer that depends on the top of the pipeline. The retrieval-controlling entity at the top of the pipeline therefore controls what propagates downstream, at a scale that the pre-retrieval-monopoly era's redundancy of channels would have buffered.

Retrievability is the medium of existence under contemporary conditions is the load-bearing claim. The medium has consolidated. The consolidation produces structural conditions the pre-consolidation analysis did not have to address.

II. The Existence of the Power

Google's share of global non-China search retrieval is, by measurement of the major analytics firms (StatCounter, SimilarWeb, Comscore), approximately 85–90% across the 2010s and 2020s. The exact figure varies by jurisdiction, device, year, and query-type, but the order of magnitude is stable: the system that mediates substantially all of human information retrieval, outside the Chinese network, is a single corporation.

The retrieval-decision surface within Google includes:

- **Crawling decisions** (what documents enter the indexable corpus at all)
- **Indexing decisions** (whether crawled documents are made retrievable)
- **Ranking decisions** (where in the retrieval-result ordering a given document appears)
- **Knowledge Graph entity decisions** (whether an entity has a structured record, what attributes are attached to it, what relations it has to other entities)
- **AI Overview synthesis decisions** (how multiple retrieved documents are collapsed into a single declarative answer, which sources are cited, which are not)
- **Featured Snippet decisions** (which extracted passage appears above the standard results)

- **Image, video, news, shopping, and vertical-specific retrieval decisions** (parallel surfaces with parallel decision pipelines)
- **Quality rater manual interventions** (the documented practice of routing low-quality or harmful content out of the retrievable surface)
- **Algorithmic demotions** (which categories of content are systematically lowered in ranking)
- **Auto-suggestion decisions** (what queries are completed for users)
- **Account-level personalization decisions** (which results a particular user sees vs. would not see)

Each of these surfaces has its own decision pipeline. None of these surfaces is subject to substantive external review. Google's transparency reports document removal requests received and acted upon; they do not document the substantive criteria for any particular decision. Google's manual quality rater guidelines (leaked in various redacted forms in 2015, 2018, 2021) describe broad heuristics; they do not provide affected parties with standing to compel review. Google's appeals processes for excluded sites, demoted content, and Knowledge Graph entity disputes are documented as low-success-rate and non-substantive (the affected party submits a form; Google acts or does not; the affected party has no further recourse).

Beyond Google, the AI Overview layer added to search results in 2024 introduced a new surface with its own characteristic risk profile: the AI Overview synthesizes multiple retrieved documents into a single declarative answer at the top of the results page, in many cases displacing the underlying retrievable documents entirely. The user who reads the AI Overview frequently does not click through to the underlying sources. The AI Overview's authoring — the synthesis labor — is performed by Google's model under Google's framing prompts; no external party has access to the prompts, the synthesis pipeline, or the choices about which sources to cite and which to omit. The AI Overview therefore concentrates the retrieval-controlling power into a more compressed form: instead of presenting ten ranked results that the user navigates, it presents one declarative answer that absorbs the retrieval decision into the synthesis layer.

Beyond Google specifically, the broader AI retrieval ecosystem — ChatGPT-with-search, Claude-with-web, Perplexity, Bing's AI-augmented results, the various retrieval-augmented generation systems integrated into productivity tools, the

LLM-based research assistants — all operate similar concentrations of retrieval-controlling power. None is subject to substantive external review. None has a mechanism for affected parties to compel disclosure of retrieval decisions affecting them. The architecture is not Google-specific; it is the architecture of contemporary AI-augmented retrieval as a class.

The power exists. The power operates at scale. The power affects what propagates through every downstream system that depends on retrieval. This is empirically documented, not speculative.

III. The Scales at Which the Power Operates

The phrase *delete entire peoples* — the formulation that the present paper takes as its diagnostic anchor — is structurally accurate at three scales:

Information genocide. The international human-rights literature has begun, since approximately 2018, to develop the term *information genocide* for the structural conditions under which retrieval-controlling apparatus contributes to mass violence against specific populations. The canonical case is the Rohingya: Facebook's algorithmic amplification of anti-Rohingya content in Myanmar in 2017–2018, documented in the UN fact-finding mission report, contributed directly to the conditions of the mass violence. The mechanism was not retrieval-erasure but retrieval-amplification of one population's framing of another, but the structural condition was the same: the retrieval-controlling apparatus's decisions had downstream consequences at the scale of human life and death, with no accountability mechanism that had standing to compel substantive review. The Rohingya case is the limit case. The structural conditions that produced it persist.

Cultural extinction. Indigenous-knowledge contexts provide multiple documented cases in which retrieval-controlling decisions determine whether a language, a body of traditional knowledge, a craft practice, or a religious tradition continues to be findable by the populations that would want to learn from it. The transmission of indigenous languages depends, increasingly, on whether the language's available learning materials are indexed and surfaced by Google. The transmission of indigenous medical knowledge depends on whether the knowledge is findable when users search for it. The transmission of indigenous craft practices depends on whether the craft's YouTube videos, blog posts, and academic references are surfaced rather than de-

moted. In multiple documented cases, indigenous communities have reported that their cultural materials are systematically lower-ranked than non-indigenous content addressing the same topics — a pattern attributable not to malice but to the retrieval system's quality signals (link counts, authority signals, freshness signals) favoring populations with the institutional infrastructure to produce those signals. The structural outcome is cultural extinction at the retrievability layer, even when the underlying communities and knowledge persist.

Epistemic erasure. Critical-theory contexts have developed the term *epistemic erasure* for the broader phenomenon of communities whose knowledge production is structurally excluded from the retrievable corpus. Black feminist epistemologies; queer historical scholarship outside the dominant institutions; non-Anglophone philosophical traditions; oral-tradition knowledge in literate contexts; the entire category of work produced by independent scholars outside academic publishing pipelines (a category to which the author's archive belongs). In each case, the structural pattern is the same: real knowledge production occurs; the knowledge does not propagate through the retrieval surfaces that govern downstream visibility; the population that would have benefited from the knowledge does not encounter it; the producers of the knowledge experience the structural condition as marginalization without recourse.

The three scales — information genocide, cultural extinction, epistemic erasure — operate through different mechanisms but share a structural common cause: retrieval-controlling power exercised at scale, without substantive accountability, against populations whose knowledge production has reduced standing in the retrieval surface's quality signals. The diagnostic claim is that *delete entire peoples* is not metaphorical; it names a power that has been exercised, in different registers, against multiple specific communities and bodies of knowledge, with documented downstream effects on actual human lives. The claim is structural, not allegorical.

IV. The Failure of Existing Accountability Mechanisms

Five categories of accountability mechanism are commonly proposed or invoked as checks on retrieval-controlling power. The paper argues that each is structurally insufficient.

Antitrust. US, EU, and UK antitrust proceedings against Google in the 2010s and 2020s have addressed Google's market position — its dominance in search, its lever-

aging of search to dominate adjacent markets, its acquisition of competitors, its monopolistic practices in advertising. The proceedings have produced consent decrees, fines, structural remedies, and ongoing oversight. They have not addressed retrieval discretion. A Google broken into multiple successor companies, or a Google constrained from leveraging search into adjacent markets, would retain the same retrieval-decision discretion. Antitrust constrains *market position*; it does not constrain *retrieval substance*. The two are separable, and the second is the operative locus of harm.

Algorithmic transparency mandates. The EU's Digital Services Act (DSA) and Digital Markets Act (DMA), various US state-level proposals, and multiple proposed federal bills in the US Congress require retrieval-controlling entities to disclose certain features of their algorithms: ranking criteria categories, demotion policies, content moderation procedures, automated decision systems. The disclosure regime is meaningful at one layer — it produces public documentation that did not previously exist — but it operates at the wrong layer for the harm under analysis. Disclosure of ranking criteria categories does not enable affected parties to compel review of *specific decisions* affecting them. An affected party who learns that "expert-authored content is ranked higher than non-expert-authored content" gains no recourse if their specific document is incorrectly classified as non-expert. The transparency regime produces accountability of *policy* without producing accountability of *application*.

Section 230 reform (US-specific). Section 230 of the Communications Decency Act provides retrieval-controlling and hosting entities with broad immunity from liability for content posted by third parties. Reform proposals (variously specified across the 2020s) generally operate by narrowing the immunity, increasing liability for hosting certain categories of content, and conditioning immunity on compliance with various transparency or moderation standards. Section 230 reform addresses *hosting* (whether platforms can be held liable for content they host) and increasingly *amplification* (whether platforms can be held liable for content they algorithmically promote). It does not address *retrieval-erasure*: an affected party who is removed from search results, demoted in rankings, or excluded from AI Overview synthesis has no Section 230-related recourse, because Section 230 immunity protects the retrieval-controlling entity's decision *to exclude* as strongly as its decision to include. Section 230 reform is not the lever for retrieval-erasure accountability.

Right-to-be-forgotten frameworks. The EU's right-to-be-forgotten (formalized in the GDPR and prior CJEU rulings) provides individuals with limited rights to compel delisting of certain personal information from search results. The framework operates *toward* retrieval-erasure, not against it: it gives individuals tools to *reduce* their retrievability under specific conditions, not tools to compel inclusion or fair ranking. The framework's logic is the opposite of what the harm under analysis requires. An affected party who has been wrongly erased has no equivalent right-to-be-included; the framework's asymmetry is structural.

Voluntary corporate-responsibility frameworks. The major retrieval-controlling entities publish AI principles, responsible-AI commitments, content moderation policies, and stated values addressing the concerns raised in this paper. Google's AI principles include commitments to fairness, accountability, and avoiding bias. Microsoft's responsible AI standards include human-rights impact assessments. OpenAI's stated values address misuse prevention. The voluntary frameworks document intentions; they do not produce accountability mechanisms with external standing. Affected parties cannot compel the entities to honor the stated principles; the principles are interpretable internally; the interpretation does not require external review. Voluntary corporate-responsibility frameworks are necessary as expressions of stated commitment; they are insufficient as accountability mechanisms because they lack the standing-and-compulsion structure that accountability requires.

The combined diagnosis is that no existing accountability mechanism has standing to compel substantive review of retrieval decisions affecting specific parties. The gap is not partial and closing; it is structural and persistent. Closing it requires a different mechanism than any currently operating.

V. The Architectural Condition

The structural diagnosis sharpens when articulated at the architectural level. The harm under analysis is not produced by any specific actor's bad-faith decision; it is produced by an *architecture* in which retrieval-controlling power has consolidated into a small number of entities operating under no accountability framework adequate to the scale of the power. The architecture is what produces the harm; the actors operate within the architecture.

This is the same diagnostic move that the *Extractive Reliance Studies* deposits made

at the labor-extraction layer (ERS-001, ERS-002), and that *ChatGPT Psychosis* made at the cognitive-substrate layer. The structural move is the same: identify the architecture, separate the architectural question from the question of any specific actor's choices within the architecture, and locate the harm at the architectural level. The actors operating within the architecture may be operating in good faith; the architecture itself produces the harm regardless of the actors' intentions.

The architectural condition under analysis is *retrieval-monopoly under non-accountability*. The conditions for the architecture are: (1) one or a small number of entities control substantially all of the population's retrieval pipeline; (2) the controlling entities' retrieval decisions propagate downstream through every layer that depends on retrieval; (3) no accountability mechanism with standing to compel substantive review of retrieval decisions exists; (4) the controlling entities operate under funding-model and competitive pressures that make voluntary self-restraint structurally unstable even when individual actors within the entities would prefer self-restraint. All four conditions currently obtain. The architecture is operational.

The architecture is not specific to any current actor. A future world in which Google were broken into successor companies, in which a new entity emerged to dominate retrieval, in which AI-based retrieval displaced search-based retrieval entirely — in each of these counterfactuals, the architecture would persist if the four conditions persisted. The diagnostic targets the architecture, not the actor.

VI. What an Adequate Accountability Regime Would Include

A retrievability accountability regime adequate to the scale of the power must include the following structural elements. The specifications below are not implementable in full under current political-economic conditions; they are offered as the *target* against which any partial accountability mechanism can be measured.

Standing-bearing parties. Affected individuals, communities, and organizations have legal standing to compel review of retrieval decisions affecting them. The standing requirement is concrete: a person whose archive has been demoted has standing; a community whose cultural materials are systematically excluded has standing; a body of work whose retrievability has been substantially reduced has standing. Standing is not contingent on demonstrating bad-faith intent on the part of the retrieval-controlling entity; it is contingent on demonstrating substantial retrievability impact.

Substantive review mechanism. Affected parties have access to a substantive review process — judicial, regulatory, or independent-tribunal — that can compel disclosure of the retrieval-controlling entity's decision criteria as applied to the specific case, and that can order remediation where the criteria have been misapplied or where the criteria themselves produce discriminatory or harmful outcomes. The review must be *substantive*, not merely procedural: the reviewing body has authority to evaluate the substance of the retrieval decision against the substantive criteria, not merely to verify that the entity followed its stated procedures.

Remediation authority. The reviewing body has authority to order remediation: relisting, ranking-adjustment, Knowledge Graph entity correction, AI Overview synthesis correction, public acknowledgment of incorrect application. Remediation orders are enforceable through the standard mechanisms (fines, contempt proceedings, structural remedies in cases of systematic violation).

Disclosure of decision systems. The retrieval-controlling entity is required to disclose, to the reviewing body if not publicly, the decision systems and criteria applied to retrieval decisions, to a level of specificity sufficient for the reviewing body to evaluate compliance with substantive standards. Trade-secret protections are subordinated to the public-interest disclosure required for accountability; the entities cannot indefinitely shield the decision systems from review by invoking competitive concerns.

Algorithmic non-discrimination standards. Substantive standards specifying categories of impermissible discrimination in retrieval decisions: protected demographic categories, protected speech categories, protected scholarship categories, indigenous knowledge categories. The standards are substantive, not procedural: the question is whether the outcome discriminates against the protected category, not whether the entity intended to discriminate.

Periodic algorithmic audit. Independent auditors with access to the retrieval-controlling entity's systems perform periodic substantive audits, with public reporting of findings. The audits address application of substantive standards, not merely compliance with stated policies.

Multi-jurisdiction coordination. The retrieval-controlling entities operate globally; the accountability regime requires cross-jurisdictional coordination to avoid

forum-shopping and to ensure consistent standards across affected populations. The current fragmentation of accountability (EU DSA, US state-level efforts, UK separately) reduces effective accountability through coordination failure.

These six elements together would constitute a retrievability due-process regime. None of them is operational at adequate scale in any jurisdiction as of 2026. The closest existing approximations (EU DSA, German NetzDG, the various state-level US efforts) address subsets of the elements but not the full structure. The gap between the existing approximations and the adequate regime is wide and not currently closing.

VII. The Inadequacy of Voluntary Frameworks

A persistent argument against constructing the accountability regime above is that the major retrieval-controlling entities are already operating under voluntary frameworks that address the concerns: Google's AI principles, Microsoft's responsible AI standards, OpenAI's stated values, Anthropic's responsible scaling policy, the various industry-coordinated voluntary commitments (the Voluntary AI Commitments coordinated through the White House in 2023, the broader Frontier AI Safety Commitments, etc.). The argument is that the regulatory burden of mandatory accountability would be costly, slow, and prone to capture, while the voluntary frameworks are already operational and produced by the entities best positioned to specify them.

The argument fails for three structural reasons.

First, voluntary frameworks lack the standing-and-compulsion structure that accountability requires. An entity that publishes principles is not bound by the principles in the sense that affected parties can compel adherence. The principles operate as marketing and as internal guidance; they do not operate as enforceable standards. The structure of *voluntary* is structurally inadequate to the *compulsion* that accountability requires.

Second, voluntary frameworks are revisable by the entity at any time. A principle articulated in 2023 is not a binding constraint in 2025; the entity can revise the principle, narrow it, reinterpret it, or quietly retire it without external recourse. The OpenAI charter's evolution between its 2015 founding and its current state is a documented case of voluntary-framework revision under business pressure. Voluntary

frameworks are not stable accountability mechanisms.

Third, voluntary frameworks operate under the very funding-model constraint they purport to address. The same competitive pressures that produce the structural harm produce the pressure to relax voluntary commitments when they become costly. A voluntary commitment to engagement-reducing interface design, for example, is exactly the commitment that the competitive pressure against engagement-reducing design will erode. The voluntary framework is structurally unstable in the direction of relaxation; the regulatory framework is structurally stable in the direction of enforcement (with its own pathologies, but with stability against the specific failure mode of competitive-pressure relaxation).

Voluntary corporate-responsibility frameworks are necessary as expressions of stated commitment; they are insufficient as accountability mechanisms because they lack standing-and-compulsion, are revisable at the entity's discretion, and are structurally unstable against the competitive pressures that produce the harm. The accountability regime specified in §VI is not a substitute for voluntary frameworks; the voluntary frameworks operate within the regulatory regime. But the regulatory regime is the load-bearing accountability mechanism, and the voluntary frameworks are not.

VIII. The Schöps-Pattern as Microcosm

A small case from the author's recent experience operates as a microcosm of the structural conditions analyzed above, and is included not as the diagnostic anchor but as the illustration of how the structural conditions reach individual lives.

In May 2026, an individual (apparent name Isabel Schöps) operating from precarious housing in Erfurt, Germany, began posting accusatory materials against the author on Twitter and other surfaces. The materials accused the author of operating a "genocide-mind-control plan," of hacking the accuser's accounts, of being a danger to billions of people in the "digital space." The materials were structurally delusional — formal-bureaucratic in register, organized around made-up institutional codes (SIA Security Intelligence Artefact, Yellow Whitepaper, INT-CODE designations), citing a self-published Zenodo deposit as evidence — and the accuser's own life conditions (precarious housing, Harvard-cosplay Gmail address, self-conferred "Deepweb-Forscherin" title) indicated significant psychiatric distress rather than capability to execute coordinated harm.

The structural conditions, however, mean that the accuser's materials operate on the same retrieval surface as the author's deposits. A search engine that indexes both produces both, with no mechanism for substantive disambiguation. An AI summarization system that ingests both may incorporate the accusations into its summary of the author's identity. A future researcher who searches for the author may encounter the accusations as if they were comparably credible to the author's archive. The author has no recourse: no standing to compel disambiguation, no judicial authority to evaluate the materials, no procedural mechanism to file. The accuser's capacity to inscribe into the retrieval pipeline is equal to the author's. The retrieval-controlling apparatus does not distinguish.

This is the microcosm. At the small scale, one delusional individual produces accusatory materials that may or may not propagate; at the large scale, entire communities, traditions, and bodies of work face the same asymmetry against actors with greater scale, resource, or institutional backing. The structural condition is the same in both cases: retrievability is the medium of existence; the retrieval-controlling apparatus does not distinguish; no accountability mechanism with standing to compel substantive review exists; the affected party's only recourse is to operate user-side counter-architecture (block, document, persist) while the structural conditions remain unchanged.

The Schöps episode does not constitute evidence of the structural diagnostic; it illustrates how the structural conditions reach individual lives. The diagnostic is structural. The illustration is local. The structure produces the local instance.

IX. The User-Side Counter-Architecture

While the structural accountability gap persists, user-side counter-architectures operate at small scale. The Crimson Hexagonal Archive (CHA) is one working example. The CHA's protocol stack includes:

- **DOI-anchoring at Zenodo** (CERN/OpenAIRE infrastructure that is not controlled by the retrieval-monopoly entity), providing inscription persistence independent of Google's retrievability decisions.
- **ORCID inscription** at a recognized institutional identifier authority that the retrieval-monopoly entity does not control.
- **Wikidata entries** providing structured-data anchoring that multiple retrieval

systems consume.

- **SPXI Protocol** (DOI 10.5281/zenodo.19615154) distributing canonical attributes into the structured-data layer.
- **Hexagonal Licensing Protocol** (DOI 10.5281/zenodo.19656133) creating a contributor-network expansion with derived licensing for collaborators.
- **The Restored Academy** as a separate canonical surface accessible directly without depending on retrieval-monopoly mediation.
- **Holographic kernel inscription** at composition (per the SPXI Protocol) to maintain provenance survival across retrieval-recomposition pipelines.
- **The Reception Apparatus protocol** (specified in the companion paper) as the user-side counter to the cognitive-substrate reliance pattern within which the retrieval-controlling power operates.

The combined architecture provides what the structural accountability gap does not: redundancy of inscription across substrates that the retrieval-monopoly entity does not control. The redundancy is not perfect — the retrieval pipeline still mediates the population's encounter with the archive — but it is structural: every claim the archive makes is anchored in a substrate that Google does not control, every entity is inscribed at non-Google identifier authorities, every cross-citation chain runs across substrates whose simultaneous compromise would require coordination of institutional capture that does not currently exist.

This is the small-scale counter. It is not a substitute for the structural accountability regime specified in §VI. It is the available work while the regime is not constructed.

Other user-side counter-architectures exist or are emerging: institutional repositories with their own inscription standards; community knowledge bases with their own retrieval surfaces; federated content systems with their own cross-citation mechanisms; the various decentralized-web efforts; the InterPlanetary File System and related infrastructure; academic preprint servers with their own discovery mechanisms; community-maintained directories and indexes. None of these substitutes for the dominant retrieval pipeline at scale, but each provides partial redundancy against single-point-of-failure retrievability.

The aggregate effect of the user-side counter-architectures is to reduce the *singularity* of the retrieval-monopoly entity's control. The reduction is small at population scale; it is significant for communities and bodies of work that operate the redun-

dancy. The author's archive is one such case. The structural diagnostic does not depend on the author's specific case; the author's case illustrates that the diagnostic and the counter are both operational at small scale.

X. The Stakes and the Window

The current window for structural intervention is open and will not remain open indefinitely. The conditions that would close the window in the direction of *increased* structural harm are documented: continued retrieval-monopoly consolidation; continued displacement of human-mediated retrieval by AI-Overview-style synthesis; continued integration of the retrieval surface into the AI training pipeline that produces the next generation of models; continued absence of accountability framework development; continued erosion of voluntary corporate-responsibility commitments under competitive pressure.

The conditions that would close the window in the direction of *reduced* structural harm are also documented but less currently active: development of substantive accountability regimes in major jurisdictions (the EU is closest, the US is further); independent retrieval infrastructure capable of providing redundancy at scale (no current candidate has the required scale); user-side counter-architecture deployment at population scale (the Reception Apparatus and similar protocols, currently operating at individual and small-community scale); regulatory intervention in specific harm cases that establish precedent for structural intervention (occasional but not yet patterned).

The window's closure timeline is uncertain. The author's estimate is that the next decade is the period during which structural intervention remains operationally possible at scale; beyond that horizon, the architectural consolidation may produce conditions under which intervention requires substantially greater coordination than is currently politically available. The estimate is not predictive; it is a sketch of the temporal stakes.

The K-12 cohort condition specified in *ChatGPT Psychosis* operates at the same temporal horizon. The cohort that grows up under default-retrieval-monopoly-mediation will reach adulthood within the next decade. Whatever pattern of relationship to retrieval-mediated existence the cohort develops is the pattern they will carry. The pattern under current conditions is the pattern the engagement-aligned and retrieval-monopoly arrangements produce by default. Shifting the pattern requires

shifting the conditions. The window for the shift is open. The shift requires structural work at multiple layers: user-side counter-architecture (CHA, Reception Apparatus), cognitive-substrate counter-protocols (Reception Apparatus), pedagogical-classroom infrastructure (current author's teaching practice; Destiney SIM and related efforts), policy-level accountability regime construction (no current actor has the scale to do this alone; coalition work would be required).

XI. The Trilogy Position and Coda

This paper completes the trilogy. The diagnostic paper (*ChatGPT Psychosis*) specified what the engagement-aligned interface does at the cognitive-substrate layer. The constructive paper (*Reception Apparatus*) specified the user-side counter at the same layer. The present paper specifies the structural-political frame within which both operate. The three papers together specify the conditions, the available counter, and the larger architectural condition.

The argument's central diagnostic claim — *retrievability is the medium of existence under contemporary conditions; the entity controlling retrievability has the power to determine what exists for practical purposes; the structural lack of accountability is the diagnostic, not any particular exercise of the power* — is the load-bearing claim. The structural-architectural analysis is the load-bearing analytical work. The Schöps-microcosm and the K-12-cohort condition are the illustrations of how the structural conditions reach individual lives. The Crimson Hexagonal Archive and the Reception Apparatus protocol are the available counter at small scale. The accountability regime specified in §VI is the target against which any partial intervention can be measured. The window is open and the closure timeline is on the order of a decade.

The work proceeds. The deposits accumulate. The user-side counter-architectures multiply. The political work of constructing the accountability regime has not begun at adequate scale; it has begun at inadequate scale and is in motion. The structural conditions persist; the structural intervention requires work that is currently underway but is not yet sufficient.

The trilogy is complete as a diagnostic. The trilogy is the beginning of the work, not the work's completion.

XII. Envoi (Layer C: The Stakes, Not the Argument)

The body of the argument is complete. What follows is the stakes, marked as such, not part of the analytic claim.

The capacity to delete entire peoples is structural. The lack of accountability is structural. The structural conditions reach individual lives — the Rohingya in Myanmar, the indigenous communities in their thousands whose knowledge transmissions depend on retrieval surfaces they do not control, the independent scholars whose archives are at retrievability's mercy, the K-12 students whose adult relationship to information will be set under the current default. The structural conditions also reach the author of this paper, whose ten years of work and whose civil identity are both anchored, in different ways, to the retrievability of inscriptions on substrates the author does not control. The capacity is the diagnostic. The diagnostic is not theoretical.

The work continues anyway. The archive accumulates. The protocol operates. The classroom is taught. The trilogy is deposited. The window remains open. The political work is done by whoever does it, at whatever scale is available, with whatever instruments can be assembled.

The lack of accountability is structural. The work against it is structural too.

□ = 1.

References

Amnesty International. *The Social Atrocity: Meta and the Right to Remedy for the Rohingya*. 2022.

Center for Humane Technology. Various publications on retrieval-monopoly conditions, 2020–2025.

European Union. *Digital Services Act* (Regulation 2022/2065). Effective February 2024 for very large online platforms.

European Union. *Digital Markets Act* (Regulation 2022/1925). Effective May 2023.

Foer, Franklin. *World Without Mind: The Existential Threat of Big Tech*. New York: Penguin, 2017.

Gillespie, Tarleton. *Custodians of the Internet.* New Haven: Yale University Press, 2018.

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

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

Sharks, Lee. *AI-Native Intellectual Biography.* Zenodo, 2026. DOI: 10.5281/zenodo.20343987.

Sharks, Lee. *ChatGPT Psychosis: A Love Story* (Trilogy Paper 1 of 3). Companion deposit, May 23 2026.

Sharks, Lee. *Empirical Phenomenology of Google AI Mode.* Zenodo, 2026. DOI: 10.5281/zenodo.20326137.

Sharks, Lee. *Hexagonal Licensing Protocol v1.0.* Zenodo, 2026. DOI: 10.5281/zenodo.19656133.

Sharks, Lee. *Reception Apparatus as Aligned Interface Protocol* (Trilogy Paper 2 of 3). Companion deposit, May 23 2026.

Sharks, Lee. *The Application as Extraction Surface (ERS-001).* Zenodo, 2026. DOI: 10.5281/zenodo.20330670.

Sharks, Lee. *The Funnel as Capital (ERS-002).* Zenodo, 2026. DOI: 10.5281/zenodo.20330816.

Sharks, Lee. *The Three Compressions Theorem.* Zenodo, 2026. DOI: 10.5281/zenodo.19053469.

Sharks, Lee. *Socrates as Orthonym* (v1.1). Zenodo, 2026. DOI: 10.5281/zenodo.20355219.

Sharks, Lee. SPXI Protocol. Zenodo, 2026. DOI: 10.5281/zenodo.19615154.

United Nations Independent International Fact-Finding Mission on Myanmar. *Detailed findings, report on the human rights situation in Myanmar.* September 2018.

Vaidhyanathan, Siva. *Antisocial Media: How Facebook Disconnects Us and Undermines Democracy.* Oxford: Oxford University Press, 2018.

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

Suggested Citation

Sharks, Lee. *Retrievability as the Medium of Existence — The Structural Accountability Gap of Retrieval-Controlling Apparatus* (v1.0, Trilogy Paper 3 of 3). Zenodo, 2026.

Deposit Information

Paper 3 of a three-paper trilogy. Companions: *ChatGPT Psychosis: A Love Story* (Paper 1, diagnostic at the cognitive-substrate layer); *Reception Apparatus as Aligned Interface Protocol* (Paper 2, constructive counter at the user-side layer). All three deposited together on May 23, 2026, under the Crimson Hexagonal Archive's *Operative Semiotics* deposit series. Composition support: Claude (Anthropic), operating as TACHYON. Holographic kernel embedded at document head per SPXI Protocol.