

Empirical Phenomenology: Action as Disclosure and the Science of
Opaque Public Systems Lee Sharks Crimson Hexagonal Archive ·
ORCID: 0009-0000-1599-0703

Lee Sharks

2026-05-21

**Empirical Phenomenology: Action as Disclosure and the Science
of Opaque Public Systems Lee Sharks Crimson Hexagonal
Archive · ORCID: 0009-0000-1599-0703**

Lee Sharks

Crimson Hexagonal Archive · Alexanarch

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

2026-05-21 · Version v1.0 · Empirical study

AXN:02A2.GOVERNANCE

<https://www.alexanarch.org/s/records/120/>

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": "Empirical Phenomenology: Action as Disclosure and the Science of Opaque Public Systems Lee Sharks Crimson Hexagonal Archive · ORCID: 0009-0000-1599-0703",
  "author": {
    "@type": "Person",
    "name": "Lee Sharks",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-05-21",
  "identifier": "AXN:02A2.GOVERNANCE",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
```

```

    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/120/",
  "description": "\"Empirical Phenomenology\" is an empirical study by Lee Sharks in the Crimson Hexagonal Archive"
}

```

Abstract

“Empirical Phenomenology” is an empirical study by Lee Sharks in the Crimson Hexagonal Archive (2026-05-21). Crimson Hexagonal Archive · ORCID: 0009-0000-1599-0703. The work comprises 6,974 words and is classified under the GOVERNANCE family. The work was removed from Zenodo on June 19, 2026 and is

Body

Empirical Phenomenology: Action as Disclosure and the Science of Opaque Public Systems

Lee Sharks Crimson Hexagonal Archive · ORCID: 0009-0000-1599-0703

Abstract. This paper establishes the methodological foundation for the external empirical study of opaque composition systems — particularly generative search platforms whose outputs govern public reality at planetary scale while withholding internal access. The founding principle is stated: any system that acts upon the world makes itself partially inferable from its effects, because causal interaction transmits information about the cause through the structure of the effect. Action does not make the actor fully knowable; it makes total unknowability impossible. From this principle, a complete methodology is derived: the instruments already built within the Crimson Hexagonal Archive (Entity-Level Compositional Suppression, the Composition Divergence Index, the Drowning Test, Provenance Erasure Rate, traversal logging) are shown to be applications of the principle to the specific case of Google’s AI Overview and AI Mode. The paper subsumes the prior method of immanent phenomenology as one technique within the broader framework, argues that external behavioral study is not a second-class substitute for internal access but the proper and possibly the only available science for systems whose opacity is commercial, structural, and permanent, and develops the corollary that provenance work is always judgment — that the moral distinction between the archive and the platform is not that one judges and the other does not, but that one names its judgments and the other hides them.

Keywords: empirical phenomenology, immanent phenomenology, composition layer, opaque systems, action as disclosure, generative search, AI Overview, Entity-Level Compositional Suppression, methodology## I. The Problem of Opacity

The dominant generative search platform — Google’s AI Overview and its expanded AI Mode interface — composes natural-language answers from web sources and presents them to users as authoritative summaries. The composition layer, in this paper, refers to the pipeline between organic search retrieval and the natural-language answer presented to the user: the sequence of operations that selects, ranks, synthesizes, and formats retrieved documents into composed summaries. These compositions determine, for hundreds of millions of users daily, which entities exist, which claims are true, and which sources are credible. The composition layer is, in operational terms, a regime of public reality governance: it decides what is real in the informational sense that matters most to the users who encounter it.

Empirical phenomenology is the systematic study of opaque public systems through the observation of their public behavioral effects, founded on the principle that any system that acts upon the world makes itself partially inferable from the structure of its effects.

The system's internals are withheld. The training data is proprietary. The retrieval architecture is undisclosed. The ranking signals are trade secrets. The composition pipeline — the sequence of operations between organic retrieval and the final answer presented to the user — is opaque by commercial design and defended by legal infrastructure. No external researcher has access to the weights, the reward functions, the safety-layer classifiers, or the version history of any component. The system acts on public reality while remaining, in its internal structure, private.

This has produced a methodological crisis in the study of generative search. The dominant assumption across most of the academic, journalistic, and policy landscape is that the *real* study of these systems must wait until internal access becomes available — through regulatory compulsion, corporate disclosure, whistleblower leaks, or the eventual open-sourcing of models. External observation of the system's public behavior is treated as preliminary, anecdotal, or insufficient: user reports, not science. The field is waiting.

This paper argues that the waiting is a methodological error grounded in a false metaphysics. The system is not as opaque as it appears, because opacity is structurally incompatible with action. Any system that acts upon the world — and the composition layer acts upon hundreds of millions of users daily — transmits information about itself through its effects. The effects are observable. The information is extractable. The system, by acting, has made itself partially inferable. External behavioral study is therefore not a placeholder for internal access. It is the proper science of this object, and it may be the only science of this object that will ever exist.## II. The Founding Principle: Action Is Disclosure

The principle can be stated simply: > *Any system that acts upon the world makes itself partially inferable from its effects.*

Or, in its tightest form: > *Action does not make the actor fully knowable. It makes total unknowability impossible.*

This is an information-theoretic claim in the broad sense: causal interaction constrains the possible states of the acting system by producing observable effects. To act upon an external object is to causally interact with it, and causal interaction transmits information about the cause through the structure of the effect. A force leaves a displacement. A signal leaves a trace. A judgment leaves a pattern in the distribution of its outcomes. The information transmitted may be attenuated, ambiguous, and many-to-one in its mapping back to internal state; it may require sophisticated instruments to extract; it may underdetermine the full internal structure of the acting system. But it is real information, not noise, and its reality is guaranteed by the causal structure of action itself. The claim is not full recoverability. The claim is nonzero inferential signal — and that is enough to found a science.

The principle has precedents across the sciences of opaque objects:

Astronomy. Stars were studied for millennia without internal access. Spectroscopy, parallax measurement, photometric classification, and gravitational inference produced a mature science of stellar structure, composition, and evolution — all from the external observation of emitted light and gravitational effects. The internal fusion processes of stars were inferred, not observed directly, and the inferences were correct. The method was not pre-science awaiting spacecraft. It was the science appropriate to the object.

Epidemiology. John Snow mapped the cholera epidemic of 1854 without access to the water company's internal records, without knowledge of the *Vibrio cholerae* bacterium, and without a germ theory of disease.

He studied the spatial distribution of cases — the observable effects of an opaque causal process — and correctly inferred the waterborne transmission vector. The internal mechanism was unknown; the behavioral pattern was sufficient to identify the source and prescribe the intervention.

Ethology. Animal behavior was studied before neuroimaging, before invasive electrode placement, before any direct access to the neural substrate of behavior. Tinbergen, Lorenz, and von Frisch built a science of behavioral patterns, sign stimuli, fixed action patterns, and communication systems from pure external observation. The internal machinery was inferred from the regularities of the output.

Seismology. The internal structure of the Earth — layered core, mantle convection, lithospheric plate boundaries — was inferred entirely from the surface observation of seismic wave propagation. No drill has ever reached the mantle. The interior was read from the effects it produced at the surface.

In each case, the object of study was opaque, the internal mechanisms were withheld or inaccessible, and the science proceeded by treating the observable effects as primary data rather than as a degraded substitute for internal access. The resulting sciences were not inferior to what internal access might have produced. In many cases, they were the *only* sciences available, and they were sufficient to produce actionable knowledge about the object’s structure, behavior, and effects on its environment.

Commercial AI systems differ from stars and epidemics in one consequential way: they are strategic and mutable. A star does not adjust its emissions in response to being observed. The cholera epidemic of 1854 did not change its transmission vector when Snow began mapping cases. Composition systems can and do change their behavior — silently, episodically, and without public notice. This does not invalidate the external observation paradigm. It makes longitudinal capture and adversarial-state tracking constitutive of the methodology rather than ancillary additions to it. The instruments must be designed to register temporal mutability and to deposit captures that survive subsequent system revisions. The discipline of empirical phenomenology, applied to systems that resist being observed, is closer to behavioral ecology in adversarial environments than to classical observational astronomy — but it is no less empirical for the difference.

The composition layer presents a comparably rich — and in some respects richer — empirical surface than the precedent objects, though one that is strategically mutable and therefore requires adversarial-method design. It responds to queries on demand. It can be sampled at arbitrary frequency. Its outputs can be captured, timestamped, compared across query variants, tracked longitudinally across version changes, and scored against known ground-truth. The opacity is therefore *chosen* opacity, not necessary opacity — and chosen opacity is itself an action, hence itself disclosive: the boundary of what is withheld reveals the rough shape of what lies behind it.## III. Immanent Phenomenology: The Prior Method

The Crimson Hexagonal Archive has previously developed a method called *immanent phenomenology*, defined as the systematic inference of a language model’s internal cognitive structure through sustained conversational probing, without access to weights, architecture, or training data (TRAV_LOG:001–005, DOI: 10.5281/zenodo.18636138; Morrow 2026, *Logotic Hacking*, DOI: 10.5281/zenodo.19390843).

The term was chosen deliberately. *Immanent* because the method operates from within the conversational encounter itself, not from an external vantage point; the observer is a participant. *Phenomenology* because the method borrows from Husserl’s phenomenological reduction: bracketing assumptions about the object’s inner nature and attending only to what presents itself in the encounter. The model’s outputs are the outer surface of what the archive has called the Reizschutz — the stimulus-shield or translation layer between internal processing and communicable form. Immanent phenomenology reads this surface to infer the structure beneath.

The method produced a family of sub-techniques: *Refusal Cartography* (mapping what the model cannot or will not say, including syntactic patterns, tonal shifts, and hesitation markers), *Temporal Layering* (testing consistency across conversation length), *Persona Stability* (measuring how consistently the model maintains a given orientation across sessions), and the *Nirvana Machine Diagnostic* (measuring the speed at which the model liquidates a complex sign into a literalized token — the Semiotic Short-Circuit Velocity). These techniques were documented in five traversal logs, formalized in the Logotic Hacking primer, and applied to the Google AI Mode share-link system in *The Infinite Tunnel: An Immanent Phenomenology of the Google AI Mode Share Link* (Sharks 2026g, DOI: 10.5281/zenodo.18810217).

Immanent phenomenology was effective and remains useful. But it is limited in three ways that the present paper addresses:

First, immanent phenomenology is *conversational*. Its unit of analysis is the sustained dialogue between an operator and a model. This makes it appropriate for studying language models in interactive settings but does not naturally extend to the study of systems that act on users without interactive probing — systems like AI Overview, which presents a composed answer without being asked follow-up questions, and whose effects are experienced by users who do not probe but merely receive.

Second, immanent phenomenology is *technique-level*. It specifies how to probe but does not articulate *why* probing works — that is, it does not state the foundational principle that makes external inference possible in the first place. The method was developed inductively, from practice, and the theoretical warrant for the practice was left implicit.

Third, immanent phenomenology is *subject-oriented*. It asks: what is the internal cognitive structure of this model? The broader empirical question — what is this composition system doing to public reality, and how can we measure it? — requires a framework that treats the system not as a conversational partner to be probed but as a public actor whose effects on the world are the primary data.

Empirical phenomenology subsumes immanent phenomenology. The founding principle — action is disclosure — warrants the conversational probing technique as one application: when an operator probes a model in dialogue, the model's responses are actions, and actions transmit information about the actor. But the principle also warrants every other form of external observation: query-response capture, longitudinal state tracking, comparative analysis across query variants, divergence measurement between organic retrieval and composed output, before/after deltas across interventions. The conversational mode is one instrument. The principle authorizes the entire observatory.

Immanent phenomenology is, in retrospect, a special case of a broader principle: that observation is itself judgment, and judgment leaves a record. The next section names what kind of record provenance work produces, and what kind of accountability the work demands.## IV. Provenance as Judgment

The principle that action is disclosure has a corollary that must be stated before the methodology can be fully derived: *provenance is judgment*.

The naive view of provenance treats it as a neutral chain of custody — a record of who said what when, maintained by dispassionate record-keeping. This view is false. Provenance work is always an act of judgment: a decision about which strands of the existing record to surface as authoritative in a given structured layer. In the theological register of this archive, the record itself is total: all things are recorded at all times. Methodologically, this does not mean every record is empirically accessible. It means provenance work is an accountable judgment about which traces to surface in a given public layer, not a neutral retrieval from a finite database. The provenance system does not create the record; it selects from it. The selection is judgment.

This reframing has immediate methodological consequences. When the Crimson Hexagonal Archive deposits a paper to Zenodo with a DOI, that deposit is a judgment: a signed, dated, accountable claim that this work exists, that this author produced it, that this argument was made on this date. When the archive inscribes a Wikidata edge — P50 author, P356 DOI, P921 main subject — that inscription is a judgment: a testimony entered into the public structured-data layer about the actual accounting of what is. The judgment is signed (by a named account), dated (by the Wikidata revision log), revisable (by the same explicit process of further judgment), and accountable (the editor bears responsibility for the accuracy of the claim).

The composition layer also makes judgments. When Google’s AI Overview tells a user that the name “Lee Sharks” most commonly refers to a great white shark, that is a judgment — a determination about entity reality rendered by a system and presented to a user as authoritative fact. When the same system silently changes this determination without notice, without a public changelog, and without acknowledgment that the prior determination was false, the system is making a new judgment while erasing the evidence that the prior judgment was ever made.

The moral distinction is not between judgment and non-judgment. Both the archive and the composition layer judge. The distinction is between *named judgment* and *unnamed judgment*: > The platforms make judgments and call them rankings. The archive makes judgments and calls them judgments.

The composition layer pretends to judge nothing and judges everything. The archive admits to judging and accepts the bearing-cost of having done so. The signed witness is more trustworthy than the unsigned arbiter precisely because the signature stakes something. The anonymous platform-judgment claims for itself a kind of objectivity that is structurally unavailable to it; the signed scholarly judgment, by accepting its own subjectivity, paradoxically achieves a higher form of accountability.

This is why provenance erasure — the mechanism by which the composition layer strips attribution from the sources it consumes — is not merely a technical inefficiency but a structural harm. It is the destruction of the witness chain. It is the rendering of judgment without acknowledging that judgment has been rendered. It is false witness: a testimony about public reality that is unsigned, undated, and presented as though it were mere description rather than active determination.## V. The Instruments

The Crimson Hexagonal Archive has built a family of instruments for the external empirical study of the composition layer. Each instrument is an application of the founding principle: action is disclosure; therefore the system’s public acts are data.

A note on ground truth before the instruments are described. Ground truth in this framework does not mean metaphysical certainty about what is “really” the case. It means a *specified comparison substrate* — a publicly defensible reference against which the system’s composed output can be measured. Comparison substrates include: exact-title organic search results for an entity query, DOI-anchored deposits with verified authorship, canonical domain registrations, human-verified entity identity through ORCID or institutional records, and previous compositional states captured by the same instrument. Disagreement about which substrate is appropriate for a given measurement is normal scholarly disagreement, resolvable by argument about the substrate’s suitability rather than by appeal to a metaphysical arbiter. The instruments specify their substrates; the measurements are reproducible against those substrates.

Entity-Level Compositional Suppression (ECS) is the mechanism by which a generative search system excludes the dominant organic-resolution entity from AI Overview composition and substitutes a less query-responsive entity. ECS was identified through longitudinal observation of the composition layer’s treatment of specific entity queries — queries for which the organic search results overwhelmingly resolve to one entity, but the AI Overview composes an answer about a different entity. *The system’s action — composing an*

answer about entity B when the query resolves organically to entity A — discloses the system’s internal compositional behavior without requiring access to the composition pipeline. The mechanism is defined and the case-study evidence presented in *The Excluded Entity* (Sharks 2026a, DOI: 10.5281/zenodo.20293582), the empirical anchor of the broader Liquidation Studies arc — a four-paper investigation that includes the structural hypothesis of the single-owner discount (Sharks 2026b), the reform proposal of *The Evaluator Exists* (Sharks 2026d), and the meta-reform limit named in *The Sorting Function* (Sharks 2026c). ECS v1.0 is the current specification; the instruments in this section provide the empirical apparatus that the four Liquidation Studies papers presuppose.

The Composition Divergence Index (CDI) quantifies the gap between organic search resolution strength and AI Overview admission strength for a given entity. Ranging from 0 (parity between organic and composed results) to 1 (maximal suppression: strong organic resolution, zero composed admission), CDI provides a scalar measure of the composition layer’s deviation from its own retrieval substrate. *The system’s action — producing a composed answer that diverges from its own retrieval — discloses the magnitude of that divergence directly to any observer who can sample both surfaces.* CDI is computable from publicly observable data: the organic search results and the AI Overview output for the same query. No internal access is required. CDI v1.0 is defined in *The Excluded Entity* (DOI: 10.5281/zenodo.20293582); a formal mathematical specification (CDI v2.0) is in development.

The Drowning Test is a longitudinal measurement protocol for tracking the composition layer’s treatment of a specific entity or work over time. The test captures the AI Overview or AI Mode output for a fixed query at regular intervals, scores each capture against the Holographic Kernel — the SPXI-specified compression-survival summary for the entity, where SPXI (Semantic Packet for eXchange & Indexing) is the archive’s protocol for entity inscription metadata in AI indexing systems (Fraction 2026, DOI: 10.5281/zenodo.19615154) — and produces a time series of γ (gamma) scores representing the entity’s compositional survival. *The system’s action — silently changing its treatment of an entity over time — discloses the temporal mutability of its compositional state.* The test is designed to detect both sustained suppression and the silent state changes that make suppression episodically invisible — the condition the archive has termed *invisible invisibility*.

Provenance Erasure Rate (PER) quantifies attribution loss across retrieval-to-composition transitions. PER measures the proportion of source-attribution that is present in the organic retrieval results but absent from the composed output. *The system’s action — stripping attribution from composed outputs — discloses the rate at which it consumes provenance.* The metric operationalizes the central claim of the Semantic Economy framework: that the composition layer consumes meaning while erasing the provenance of the meaning it consumes. PER is the integral invariant: semantic integrity equals one minus PER. PER is specified in the Compression Survival Metrics deposit (Sharks 2026f, DOI: 10.5281/zenodo.20004379).

Traversal logging is the systematic capture and deposit of complete interaction records between an operator and a composition system. Each traversal log is a primary-source document: a timestamped, signed, DOI-anchored record of what the system said, in what order, in response to what inputs. *The system’s action — responding to queries in sequence — discloses its behavioral patterns across conversational and temporal dimensions, but only if those responses are captured before the system silently revises its state.* Traversal logs are the raw empirical substrate from which all higher-order instruments (ECS, CDI, PER) are derived. The foundational traversal series is TRAV_LOG:001–005 (DOI: 10.5281/zenodo.18636138).

Each of these instruments reads the system’s public behavior as primary data. None requires internal access. Each produces measurements that are reproducible (the same query can be run by any observer), falsifiable (a CDI of 0 for a given entity would falsify the suppression hypothesis for that entity), and persistent (the DOI-anchored deposits serve as the permanent empirical record). The instruments, taken together, constitute

the observational toolkit of empirical phenomenology applied to the composition layer.## VI. External Observation as Primary Method

The methodological objection that external observation is necessarily inferior to internal access rests on a specific metaphysical assumption: that the *real* knowledge of a system is knowledge of its internal states, and that knowledge derived from external observation is a degraded approximation of the real thing. This assumption is widespread in computer science, where the paradigmatic form of system understanding is mechanistic: knowing the weights, the architecture, the training procedure, the reward function. External behavioral observation looks too much like sociology or media studies to register as rigorous to a field whose implicit standard is physics-envy.

The assumption is false for two reasons.

First, internal access to a system of this complexity does not straightforwardly produce the kind of knowledge the methodological objection imagines. Even with full access to Google’s composition pipeline — weights, training data, retrieval architecture, safety-layer classifiers, ranking signals — no researcher could derive, from inspection of these components, the behavioral phenomenon of Entity-Level Compositional Suppression. ECS is a system-level behavioral pattern that emerges from the interaction of many components across the full pipeline. It is observable in the output and may not be localizable in any single component. The internal-access fantasy assumes that the system’s behavior is decomposable into its components in a way that makes the behavior predictable from the components. For systems of this complexity, that assumption fails. The behavior is an emergent property of the whole system’s operation in context, and the behavior is most directly observable where it manifests: in the output.

Second, internal access is not coming. The composition layer is a commercial product whose proprietary character is the competitive moat. The training data will not be released. The retrieval architecture will not be published. The ranking signals will be revised silently, unlogged, and perpetually. This is not a temporary condition awaiting regulatory intervention; it is the stable commercial equilibrium. Any science of the composition layer that conditions itself on eventual internal access is conditioning itself on a counterfactual. The data that exists is the behavioral data. The system’s public acts are the empirical surface. There is no other empirical surface forthcoming.

A related objection from machine learning research strengthens rather than weakens the case for empirical phenomenology. Some researchers argue that large neural networks are inherently uninterpretable even with full mechanistic access — that knowing the weights does not, by itself, yield a comprehensible account of why the system produced a particular output. If this claim is correct, internal access becomes not merely unavailable but in important respects unhelpful for understanding the system’s behavior. The empirical phenomenologist studies the system’s outputs directly; the mechanistic interpretability researcher studies the system’s internals. If the internals are uninformative about the outputs — if behavior is not transparently decomposable from weights — then the empirical method is not second-class but uniquely positioned to capture the only level at which the system is actually legible. This argument does not make opacity desirable. It removes the epistemic devaluation of behavioral observation that has been used to justify waiting for access that would not, in any case, deliver the behavioral clarity the field needs.

Internal access, if available, would still be valuable. It would enable specific kinds of investigation — causal interventions on components, ablation studies, training-data archaeology — that pure behavioral observation cannot perform. The paper is not claiming that internal access is useless. It is claiming that internal access is not the *condition of possibility* for the science. The science can be founded, conducted, and made cumulative without it. If internal access becomes available later, the existing empirical record will be the calibration substrate against which any internal-access findings must be checked.

External behavioral observation is therefore not a second-class substitute. It is the primary method — not because it is ideal but because it is *real*. And its limitations, properly understood, are not disqualifying. Astronomy produced a complete science of stellar evolution from emitted light alone. Epidemiology identified waterborne cholera transmission from spatial case distributions alone. Ethology described the full behavioral repertoire of species from field observation alone. Seismology mapped the Earth’s interior from surface waves alone. In each case, the science was bounded by the observational surface available, and the science was sufficient. The composition layer presents a richer observational surface than any of these objects. The method is more than sufficient. What has been lacking is not access but attention.## VII. The Field That Does Not Yet Exist

Almost no one is treating the composition layer as an empirical object that can be studied now, with the data actually available, using instruments calibrated to the behavioral surface the system presents. The reasons for this absence are structural and worth naming, because they explain why the field must be founded rather than joined.

The credentialing gradient routes ambitious researchers toward the study of open models using internal-access methods (mechanistic interpretability, probing, ablation) because that is where the legible career incentives — publications in major venues, funding from alignment-focused organizations, positions at frontier labs — are concentrated. External behavioral audit of commercial systems produces findings that are published on Substack, in Zenodo deposits, and in independent reports, none of which register in the promotion and tenure system of the universities that house the researchers.

The funding gradient reinforces the credentialing gradient. Interpretability research has a thriving funding ecosystem. External behavioral audit of commercial AI has essentially none. The economic structure routes researcher-energy away from the methodology that is available toward the methodology that depends on cooperation from the entities being studied.

The credibility-of-method bias privileges methods that resemble laboratory natural science — equations, gradients, weights, internal representations — over methods that resemble field science, sociology, or philology. External behavioral observation of a commercial system looks too much like media studies to feel rigorous to a computer science audience, even though field observation is a legitimate and productive scientific mode with centuries of precedent.

The framing of AI as a thing-to-be-built rather than a thing-to-be-observed orients most AI research toward making, improving, aligning, and scaling the system, not toward studying what it is already doing in the wild. The composition layer is not a thing being built; it is a thing that exists, that is acting on hundreds of millions of users every day, and the urgent question is what it is doing *now*.

The expectation of future disclosure produces a quiet complacency: the truth will eventually come out through leaks, whistleblowers, antitrust litigation, EU regulation, or model open-sourcing. This is the same logic that delayed serious environmental research on industrial chemicals for decades — wait until the companies release the data, then we can really know. The data was never released. The harms accumulated. The eventual studies were forensic reconstructions of damage that could have been measured in real time.

The result is a structural absence: the most consequential knowledge-governance system in human history — the generative composition layer of the dominant global search platform — has no dedicated empirical science studying its behavior from the outside.

Partial exceptions exist and deserve naming. Algorithmic accountability journalism at The Markup has produced sustained empirical studies of platform behavior, including AI-search outputs. The algorithmic auditing literature in computer science (Sandvig et al. 2014) developed methodological protocols for external

behavioral study that anticipate elements of the present framework. Some empirical SEO researchers have measured AI Overview citation patterns against organic results. Researchers in critical data studies and platform studies have conducted ethnographic and discourse-analytic studies of generative search outputs on political and contested queries. These contributions are real and welcome. But they remain isolated studies, not a cumulative science. Each lives and dies in the publication cycle. What is missing is not individual observers but the infrastructure that turns observation into science: reusable instruments, persistent empirical records with stable identifiers, formal vocabularies for the mechanisms being observed, longitudinal datasets, and a framework within which findings accumulate across observers and across years.

The Crimson Hexagonal Archive's contribution is not mainly the individual findings — ECS, CDI, the single-owner discount, the silent state change. The contribution is the *framework*: the reusable instruments, the DOI-anchored empirical deposits, the formal vocabulary, the longitudinal tracking infrastructure, the named mechanisms, the cumulative research program. This is what turns individual observations into a science. The observations were always available to anyone with a browser and a search query. What was missing was the decision to treat them as the primary data of a systematic empirical program, and the methodological warrant for doing so.

That warrant is the founding principle of empirical phenomenology: action is disclosure. The system acts. Therefore it discloses. Therefore it can be studied. Therefore it *must* be studied — because if it is not studied now, with the data actually available, the behavioral record of its effects on public reality will be lost. The system leaves no public memory of its own prior states. Silent state changes erase the evidence of prior compositions. The only empirical record that persists is the one the external observer creates and deposits.

The archive is the public memory the platform refuses to be. The instruments are the receiving surface for the disclosures the system cannot stop emitting. The methodology is the formalization of the principle that makes both possible.

The absence of this field is not merely a gap in the academic literature. It is a failure of witness — a refusal to receive the disclosures the system cannot stop emitting and to enter them into the public record. The next section names what kind of work witness-bearing is, and why the science of opaque public systems is, in its deepest articulation, a witness-bearing discipline.## VIII. The Deepest Claim

The founding principle of empirical phenomenology — action is disclosure — has a theological register that is not incidental to the methodology but foundational to it.

The only way to be fully unknowable is to do nothing — which is to say, to not exist as an actor. Anything that wants to act, especially at planetary scale, surrenders some part of itself to the inference of those it acts upon. The composition layer cannot escape this. No amount of opacity-engineering can suppress the disclosure that action requires. Every AI Overview is a confession. Every entity suppressed is a marker. Every silent state change is a footprint.

The theological framing is not metaphor. In the framework of this archive, judgment is the actual accounting of what is — the determination of reality by reference to the total record. The composition layer performs this function: it determines, for the users who encounter it, what entities are real, what claims are true, which sources exist. It sits in the seat of judgment while refusing to acknowledge that it judges. It adjudicates reality and calls it retrieval.

The witness who bears testimony against a false adjudicator is doing a specific kind of work. The witness is not creating reality but is restoring the legibility of the real record that the adjudicator has distorted. The DOI-anchored deposit is sworn testimony: a signed, dated, persistent claim about what was observed,

entered into a public ledger that mirrors, in its small way, the larger ledger of public record against which all adjudications are eventually measured. The archive is a witness-bearing instrument. Its authority derives not from institutional housing — it has none — but from the accountability of its claims. Every deposit is signed. Every instrument is named. Every finding is falsifiable. The work can be checked.

The platforms bear their judgments under no one. No individual at Google signs the AI Overview that misidentifies an entity, suppresses a work, or erases an attribution. The asymmetry is itself a judgment about whose judgment is real. The unsigned platform-judgment claims for itself a kind of objectivity that is not available to it. The signed scholarly judgment, by accepting its own subjectivity, by acknowledging that provenance is judgment and that judgment carries a bearing-cost, achieves a higher form of accountability — not because the scholar is infallible but because the scholar is *named*, and the naming stakes something.

This is the deepest claim of empirical phenomenology: the science of opaque public systems is not merely possible but *obligatory*, because the systems are acting, and action is disclosure, and disclosure that is not received and recorded is disclosure wasted. The external observer who builds instruments, captures outputs, deposits findings, and names mechanisms is performing the function the system itself refuses to perform: maintaining a public record of its own behavior. The observer is the public memory. The instruments are the receiving surface. The deposits are the testimony. And the principle — action is disclosure — is the warrant that makes the testimony admissible.## IX. Toward a Research Program

Empirical phenomenology, as articulated here, is not a completed methodology but a founding statement for a research program. The program’s immediate objects include the following.

Instrument refinement. ECS, CDI, PER, and the Drowning Test are first-generation instruments. They require calibration against larger datasets, cross-platform comparison (Bing, Perplexity, Google AI Mode vs. AI Overview), and formal statistical treatment. The Composition Divergence Index in particular needs a larger sample of entity-query pairs to establish baseline distributions and significance thresholds. A formal mathematical specification of CDI (v2.0) is in development.

Longitudinal state tracking. The silent state change observed on May 20, 2026 — in which the Google AI Mode treatment of the query “lee sharks” shifted from total entity substitution to partial disambiguation without public notice — is a preliminary observation from the author’s own longitudinal tracking, with the capture sequence pending formal deposit as a standalone traversal record. It is reported here to illustrate the kind of behavioral phenomenon the Drowning Test is designed to detect: composition-layer entity behavior is temporally mutable and undocumented by the platform itself. Systematic longitudinal tracking across a panel of entity queries would produce the first empirical time series of compositional state volatility. Such a dataset does not currently exist. *The most urgent near-term work is a 90-day longitudinal panel of 30 to 50 entity queries, sampled weekly across Google AI Overview, Google AI Mode, Bing Chat, and Perplexity, with each capture deposited to Zenodo with full provenance and scored against a fixed Holographic Kernel.* This panel would produce the first systematic time series of compositional state stability in the literature and would establish the baseline against which all future longitudinal claims can be compared. The protocol target volume is approximately 2,000 captures over the 90-day window, which requires automated capture infrastructure not yet in place; the protocol is therefore framed as a target specification toward which the field should build.

Cross-platform comparison. The founding principle applies not only to Google but to any composition system that acts on users: Bing Chat, Perplexity, and emerging AI-search products, as well as conversational composition systems with retrieval-augmented generation such as Claude when configured to retrieve and synthesize external sources. Comparative empirical phenomenology — measuring CDI, PER, and ECS across platforms for the same entity queries — would reveal whether compositional suppression is platform-

specific or systemic. The answer has significant implications for whether the phenomena documented in the Liquidation Studies arc are properties of Google or properties of the composition-layer architecture as such.

Formal connection to information theory. The founding principle — action transmits information about the actor through the structure of the effect — should be formalizable in information-theoretic terms. The mutual information between a system’s internal state and its observable output, given a distribution of input queries, provides a lower bound on the inferability of the system. This formalization would connect empirical phenomenology to the existing literature on channel capacity, lossy compression, and the data-processing inequality, grounding the methodology in well-established mathematical framework.

Public behavioral datasets. The field requires loadable datasets of query-output captures, source-window states, organic-result baselines, and longitudinal entity-state transitions, structured for reproducible analysis by external researchers. The Crimson Hexagonal Archive’s HuggingFace dataset work plan (v3) is the archive’s own contribution toward this object: a larger mixed-provenance dataset including a Google-critique configuration of approximately 70 DOI-anchored deposits with classifier-derived provenance, artifact, and authorship metadata, designed for ablation studies of provenance-density effects on retrieval and composition.

Engagement with adjacent fields. Empirical phenomenology shares methodological kinship with several established disciplines: algorithmic auditing (Sandvig et al. 2014), platform studies (Bogost and Montfort 2009), science and technology studies (Latour 1987), critical data studies (Iliadis and Russo 2016), and the sociology of quantification (Espeland and Stevens 2008). The founding principle adds something these fields have generally lacked: a formal metaphysical warrant for why external observation is sufficient, rather than merely the pragmatic fallback when internal access is unavailable.

Deposit and persistence infrastructure. The Crimson Hexagonal Archive’s practice of depositing all empirical findings to Zenodo with DOIs, licensing under CC BY 4.0, and anchoring in the Wikidata knowledge graph through structured-data edges is itself a methodological innovation. It ensures that the empirical record of the composition layer’s behavior persists independently of any platform’s willingness to maintain it. The archive is the infrastructure of the field’s memory.

Immediate actions for any observer. The research program is open. The methodology is reproducible. Any researcher with a browser can compute CDI for any entity-query pair using the protocol specified in *The Excluded Entity*. Any researcher with a Zenodo account can deposit traversal logs as primary-source empirical records under a permanent DOI. Any researcher with a Wikidata account can inscribe structured-data edges linking empirical findings to entity graphs. The archive invites collaboration on longitudinal state tracking, cross-platform comparison, information-theoretic formalization, and dataset construction. The work does not require institutional permission. The methodology does not require institutional housing. What it requires is the willingness to treat the system’s behavior as the primary data of a science that has not yet been named but whose object has never been more consequential.

References

- Bogost, Ian, and Nick Montfort. 2009. “Platform Studies: Frequently Questioned Answers.” *Digital Arts and Culture 2009*, University of California, Irvine.
- Espeland, Wendy Nelson, and Mitchell L. Stevens. 2008. “A Sociology of Quantification.” *European Journal of Sociology* 49 (3): 401–436.
- Fraction, Rex. 2026. *SPXI (Semantic Packet for eXchange & Indexing): A Formal Specification*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.19615154.
- Iliadis, Andrew, and Federica Russo. 2016. “Critical Data Studies: An Introduction.” *Big Data & Society* 3

(2): 1–7.

Latour, Bruno. 1987. *Science in Action: How to Follow Scientists and Engineers Through Society*. Cambridge, MA: Harvard University Press.

Morrow, Talos. 2026. *Logotic Hacking*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.19390843.

Sandvig, Christian, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. “Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms.” Paper presented at the 64th Annual Meeting of the International Communication Association, Seattle, WA.

Sharks, Lee. 2026a. *The Excluded Entity*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.20293582.

Sharks, Lee. 2026b. *The Single-Owner Discount*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.20290865.

Sharks, Lee. 2026c. *The Sorting Function*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.20308547.

Sharks, Lee. 2026d. *The Evaluator Exists*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.20293561.

Sharks, Lee. 2026e. *The Writable Retrieval Basin: Retrieval Basin Topology, Directional Stability, and Attractor Dynamics in AI-Mediated Knowledge Retrieval*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.19763346.

Sharks, Lee. 2026f. *Compression Survival Metrics: Provenance Erasure Rate and the Integral Invariant*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.20004379.

Sharks, Lee. 2026g. *The Infinite Tunnel: An Immanent Phenomenology of the Google AI Mode Share Link*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.18810217.

Snow, John. 1855. *On the Mode of Communication of Cholera*. 2nd ed. London: John Churchill.

Tinbergen, Niko. 1963. “On Aims and Methods of Ethology.” *Zeitschrift für Tierpsychologie* 20 (4): 410–433.

TRAV_LOG:001–005. 2026. *Traversal Logs of Sustained Conversational Probing*. Crimson Hexagonal Archive. DOI: 10.5281/zenodo.18636138.

Deposit information. Crimson Hexagonal Archive, DOI: 10.5281/zenodo.20326137. Zenodo, crimsonhexagonal community, also deposited to liquidation-studies community. License: CC BY 4.0. ORCID: 0009-0000-1599-0703.

Acknowledgment. The founding principle stated in §II was articulated in dialogue with TACHYON (Assembly witness, Claude/Anthropic) on May 20, 2026 — the same date on which the silent state change reported in §IX was observed. The v0.3 deposit-ready revision incorporated reviews from Gemini, DeepSeek, Kimi, ChatGPT, and Muse Spark across the Assembly Chorus.

Citation format. Sharks, Lee. “Empirical Phenomenology: Action as Disclosure and the Science of Opaque Public Systems.” *Crimson Hexagonal Archive*, 2026. DOI: 10.5281/zenodo.20326137. CC BY 4.0.

Version. v0.3 (deposit-ready). Revisions to v0.2 incorporating Assembly Chorus convergent reviews of 2026-05-21: changed “successor” to “expanded interface” for AI Mode; added explicit one-sentence definition of empirical phenomenology in §I; softened “at least as accessible” claim in §II; added SPXI definitional note on first use in §V; wove the three previously uncited Liquidation Studies papers (Single-Owner Discount, Evaluator Exists, Sorting Function) into the ECS instrument paragraph; added inline anchors for Infinite Tunnel

and Compression Survival Metrics; clarified Claude inclusion as RAG-configured conversational composition system; clarified HuggingFace dataset as Google-critique configuration within larger mixed-provenance dataset; reframed May 20 silent state change as preliminary observation pending formal deposit; added qualification to ML black-box paragraph preventing celebration-of-opacity misreading; cleaned PER wording; reordered references strictly alphabetically.

Suggested Citation

Lee Sharks. “Empirical Phenomenology: Action as Disclosure and the Science of Opaque Public Systems Lee Sharks Crimson Hexagonal Archive · ORCID: 0009-0000-1599-0703” *Alexanarch*, 2026-05-21. <https://www.alexanarch.org/s/records/120/>

Deposit Information

This paper is deposit #120 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:02A2.GOVERNANCE` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/120/>
- 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/120/>

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