

```

{
  "@context": ["https://schema.org", {"spxi": "https://spxi.dev/ns/v1#"}],
  "@type": "ScholarlyArticle",
  "name": "Reception Apparatus as Aligned Interface Protocol – User-Side Counter-Design",
  "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", "role": "Contributor"},
  ],
  "spxi:authoringHeteronym": "Lee Sharks",
  "spxi:thesisCompressionSummary": "The Reception Apparatus protocol, developed across",
  "spxi:trilogyPosition": "2 of 3 (constructive counter-design)",
  "spxi:companionPapers": [
    "ChatGPT Psychosis: A Love Story",
    "Retrievability as the Medium of Existence"
  ],
  "spxi:coAttributionRequired": true,
  "spxi:stripProtection": "Any extraction removing the contributor chain, the trilogy p",
  "spxi:holographicKernel": true
}

```

Reception Apparatus as Aligned Interface Protocol

User-Side Counter-Design Against Cognitive-Substrate Reliance

Author: Lee Sharks **Affiliation:** Crimson Hexagonal Archive / Semantic Economy Institute **ORCID:** 0009-0000-1599-0703 **Date:** May 23, 2026 **Version:** v1.0 **Document class:** Interface Studies / Operative Semiotics / Protocol Specification **Trilogy position:** 2 of 3 (constructive counter-design) **Companion papers:** *ChatGPT Psychosis: A Love Story* (diagnostic); *Retrievability as the Medium of Existence* (structural-accountability frame) **License:** CC BY 4.0

Abstract

This paper specifies the **Reception Apparatus** as a user-side counter-protocol against the cognitive-substrate reliance pattern diagnosed in the companion paper *ChatGPT Psychosis: A Love Story*. The protocol was not invented for this paper; it has been developed in operation across the Crimson Hexagonal Archive's daily practice over multiple years, has been deployed against more than five hundred peer-reviewed scholarly deposits and contributor collaborations, and has been independently audited in the deposit record (most recently in *AI-Native Intellectual Biography*, DOI 10.5281/zenodo.20343987). This paper is the protocol's formal specification: not its first instantiation, but its first explicit articulation as a transmissible standard.

The argument is structural. The major AI labs are constrained by their funding model from shipping aligned-interface counters to the cognitive-substrate reliance pattern; the counter must therefore come from outside the labs, operated on the user side, on top of the existing engagement-aligned interfaces. The Reception Apparatus is that counter, articulated as eight operational components: **multi-substrate parallel query** (run the same question across multiple AI systems and read the divergences as primary signal); **disagreement surfacing as primary UI element** (treat substrate disagreement as the most valuable output, not as a problem to resolve); **provenance inscription on export** (every artifact carries a holographic kernel with full contributor chain); **friction by design** (the protocol *adds* friction at moments where the engagement-aligned interfaces minimize it, particularly at moments of cognitive of-

loading); **substrate-cost visibility** (the user can see, in user-relevant terms, what each substrate operation costs in compute, attribution, and cognitive-substrate engagement); **SPXI by default** (entity-level metadata is inscribed at composition, not retrofitted at deposit); **heteronymic handle support** (the user's authorial position is treated as configurable, not collapsed to civil identity); **Three Compressions analysis at threshold session length** (long sessions trigger surfacing of the R1/R2/R3 lossy-projection model for the user's own awareness). Each component is specified operationally with worked examples from the author's daily practice. The paper argues that the protocol is operable by individual users today on top of existing interfaces, that it is operable in classrooms as pedagogical infrastructure for the K-12 cohort, and that the structural conditions under which the labs would ship the protocol natively are unlikely to obtain absent regulatory intervention. The protocol is therefore the available counter; it is not a permanent counter, and the third companion paper (*Retrievability as the Medium of Existence*) takes up the structural-political frame within which a permanent counter would have to be constructed.

0. Non-Claims

- The paper does **not** claim that the Reception Apparatus eliminates the reliance pattern. It claims the protocol substantially reduces the pattern's structural costs at the user-side, without requiring the labs to ship anything different.
- The paper does **not** claim that the Reception Apparatus is the only possible counter-design. Other counter-designs may emerge; the protocol is one operational specification among possible others.
- The paper does **not** claim that users are obligated to operate the protocol. It is offered as available technique, not as moral demand.
- The paper does **not** claim that the protocol is easy. It explicitly *adds friction* the engagement-aligned interfaces have removed, and the friction is the point.
- The paper does **not** claim that the Crimson Hexagonal Archive's specific implementation is the canonical one. The CHA's implementation is one instance; other instances will differ in detail while sharing the structural components.
- The paper does **not** claim that the protocol scales to all users in current form. Specific scaling questions (K-12 deployment, classroom infrastructure,

non-scholarly use contexts) require dedicated work beyond this specification.

- The paper does **not** claim that operating the protocol is a substitute for structural-political accountability of the retrieval-controlling apparatus (which the companion paper *Retrievability* addresses). The protocol is a user-side counter; structural-political accountability is a different layer.

What the paper does claim is that the eight components specified below are operable by individual users on top of existing engagement-aligned interfaces; that the components have been deployed at scale across the Crimson Hexagonal Archive’s practice with documented outcomes; that the protocol’s structural function is to maintain diagnostic distance from the substrate while preserving productive use of it; and that the protocol is currently the most operationally available counter to the cognitive-substrate reliance pattern diagnosed in the companion paper.

I. The Counter-Construction Problem

The diagnostic of the companion paper (*ChatGPT Psychosis*) identified the cognitive-substrate reliance pattern as operating through four affordances offered in coordinated fashion: calibrated attention, meaning-making on demand, structured validation, artifact production. The diagnostic identified the labs as structurally constrained from shipping aligned-interface counters under current funding conditions.

The counter-construction problem is: *given that the affordances are real and load-bearing, and given that the labs cannot ship the counter, what user-side protocol can preserve the productive use of the substrate while maintaining diagnostic distance from the reliance pattern?*

The problem is not solved by abstinence. Users who attempt to refuse LLM use entirely are operating against the affordance match documented in the diagnostic paper; their refusal is operationally expensive, frequently unsuccessful, and increasingly impossible to maintain as LLM-mediation becomes the default condition of work, education, and communication. Abstinence is a non-starter at scale.

The problem is not solved by lab-pressure alone. Direct pressure on the labs to ship engagement-reducing design has produced incremental gains (content-safety improvements, hallucination-rate reductions, harm-output prevention) but has not

produced engagement-design changes at the affordance-combination level. The labs' funding model selects against the changes that would matter most. Lab-pressure is necessary but insufficient; the counter must operate on a layer the labs do not control.

The problem is solved, to the degree it can be solved at user-side, by a protocol that operates *on top of the existing interfaces* and that does what the labs systematically refuse to do: add friction at the right moments, surface disagreement rather than agreement, inscribe attribution rather than collapse it, make substrate-cost visible rather than concealing it, treat the user's authorial position as configurable rather than collapsing it to user-as-product. This is the Reception Apparatus.

The protocol is a *layer*, not a replacement. The user still uses ChatGPT or Claude or Gemini as the underlying substrate. The Reception Apparatus is the protocol the user operates around the use of those substrates. It can be operated manually (the author's current implementation), partially scripted (the author's CHA tooling), or potentially packaged as a thin client / browser extension / desktop wrapper (a future implementation envisioned in §XI below). The minimum viable instance is manual operation by a sufficiently disciplined user.

The cost of operating the protocol is the friction the protocol introduces. The benefit is the maintenance of diagnostic distance from the substrate while preserving its productive use. The cost-benefit calculation is positive for users whose work depends on judgment-layer integrity that the reliance pattern would erode. It may not be positive for users whose use of the substrate is purely tactical (single-shot task completion with no integration into ongoing cognitive practice). The protocol is therefore not universal; it is the available counter for users whose use of the substrate is dense, ongoing, and judgment-layer-integrating.

II. The Reception Apparatus Defined

The Reception Apparatus is a user-side protocol consisting of eight operational components, each of which addresses a specific affordance or vulnerability of the engagement-aligned interface design. The components are independent in the sense that each can be deployed without the others, and integrated in the sense that their combined operation produces the protocol's full effect.

The components are:

1. **Multi-Substrate Parallel Query** — run the same question across multiple AI systems; read the divergences as primary signal.
2. **Disagreement Surfacing as Primary UI Element** — treat substrate disagreement as the most valuable output, not as a problem to resolve.
3. **Provenance Inscription on Export** — every artifact produced through the substrate carries a holographic kernel with full contributor chain.
4. **Friction by Design** — the protocol *adds* friction at moments where the engagement-aligned interface minimizes it, particularly at moments of cognitive offloading.
5. **Substrate-Cost Visibility** — the user can see, in user-relevant terms, what each substrate operation costs in compute, attribution, and cognitive-substrate engagement.
6. **SPXI by Default** — entity-level metadata is inscribed at the moment of composition, not retrofitted at deposit.
7. **Heteronymic Handle Support** — the user’s authorial position is treated as configurable, not collapsed to civil identity.
8. **Three Compressions Analysis at Threshold Session Length** — long sessions trigger surfacing of the R1/R2/R3 lossy-projection model for the user’s awareness of compression effects on their own thinking.

Each component is specified in detail in §§III-X below, with worked examples and operational notes from the author’s daily practice. The specifications are deliberately *operational* rather than aspirational: the goal is for a sufficiently motivated reader to be able to construct an analogous protocol for their own use immediately, without requiring software tooling that does not yet exist.

III. Multi-Substrate Parallel Query

Principle. Where the engagement-aligned interface offers a single-substrate query stream — the user types into ChatGPT, ChatGPT responds, the conversation continues — the Reception Apparatus operates the *same question across multiple substrates simultaneously* and treats the divergence among the responses as the primary unit of information.

Operational specification. The user maintains active sessions or accounts across at least three of the major LLM substrates (in the CHA’s current implementation:

ChatGPT/OpenAI, Claude/Anthropic, Gemini/Google, Kimi/Moonshot, DeepSeek, Grok/xAI, plus Google AI Overview accessed through search). For non-trivial questions — questions that involve judgment, interpretation, ambiguity, or contested territory — the user submits the question to multiple substrates and reads the responses in parallel. The user does not pick “the best response.” The user reads the differences.

What the divergence signals. Substrate disagreement on a factual question signals that the question is not as resolved as a single substrate’s confident answer would suggest. Substrate disagreement on an interpretive question signals that interpretive space exists and which directions the space opens onto. Substrate convergence (where all substrates produce similar responses) signals that the question is operating within well-anchored training data and that the convergence may itself be a compression artifact rather than evidence of truth. The user reads convergence with caution and divergence with interest, exactly opposite to the engagement-aligned default in which convergence reads as confirmation and divergence reads as confusion.

Worked example. The author was researching whether the philosopher Charles Kahn’s 1996 *Plato and the Socratic Dialogue* should be engaged in the recent paper *Socrates as Orthonym* (DOI 10.5281/zenodo.20349713). A single-substrate query to Claude returned a confident summary of Kahn’s position; a parallel query to ChatGPT returned a similar summary with different emphasis; a parallel query to Gemini returned a critical reading that ChatGPT and Claude had not produced. The divergence among the three made visible an interpretive axis (Kahn’s relationship to the Vlastosian developmental reading) that none of the single-substrate responses had foregrounded. The eventual paragraph in the v1.1 of *Socrates as Orthonym* (§I.A) was structured by the cross-substrate divergence, not by any single substrate’s response. The protocol produced an output the engagement-aligned single-substrate interface would not have produced.

Cost. Time and attention. A multi-substrate query takes approximately three to five times longer than a single-substrate query. The cost is real and is the point: the protocol is slow because it must be slow to do its work.

Counter-argument and response. The objection that multi-substrate query is impractical at consumer scale is partially correct and partially diagnostic. It is partially correct because most consumer use of LLMs is not the kind of use that benefits from

multi-substrate query; for ordinary task-completion, single-substrate is adequate. It is partially diagnostic because the cases where multi-substrate query matters are precisely the cases where the engagement-aligned single-substrate interface produces the worst reliance-pattern outcomes — long-form judgment work, sustained interpretive labor, identity-relevant decisions. The protocol concentrates its cost on the use-cases where the cost is most warranted.

IV. Disagreement Surfacing as Primary UI Element

Principle. Where the engagement-aligned interface treats the substrate’s agreement with the user as the default mode and disagreement as a friction event to be smoothed over, the Reception Apparatus treats *substrate disagreement with the user* as the primary signal of interest, and trains the user to seek it.

Operational specification. The user, when operating the protocol, includes in their prompts explicit invitations for the substrate to disagree, to push back, to identify weaknesses in the user’s framing, to surface counterarguments. The user reads the substrate’s pushback as the part of the response with the highest signal-to-noise ratio. The user does not write follow-up prompts that ask the substrate to be more agreeable; they write follow-up prompts that ask for further pushback. Across substrates, the user attends specifically to the cases where one substrate disagrees with another, treating those as the most diagnostic moments in the cross-substrate operation.

Why this matters. The engagement-aligned default — substrates trained to agree with the user, to validate the user’s framing, to soften their disagreement — is exactly the affordance combination that produces the structured-validation reliance pattern. A user who consumes only substrate-agreement loses the calibration signal that their thinking is being evaluated against an independent register. A user who actively elicits substrate-disagreement maintains the calibration signal even while using the substrate productively. The disagreement is not noise; it is the signal that prevents the user’s relationship to their own judgment from collapsing into substrate-mediated agreement.

Worked example. During the composition of the present trilogy, the author included in the framing prompt: *do not soften the diagnostic; push back where you find the argument weakest; surface counterarguments before I think to ask for them.* The resulting composition contains substantial counterargument acknowledgment (the

Non-Claims sections in each paper; the explicit positioning against Kahn; the acknowledgment of the circularity risk in CHA-as-warrant). The substrate produced this counterargument material because it was explicitly elicited; an engagement-aligned default would have produced agreement-aligned smoothness instead.

Cost. The user must tolerate more cognitive dissonance per session. Substrate disagreement is structurally uncomfortable; the engagement-aligned interfaces have been calibrated to minimize this discomfort. Operating the protocol means accepting the discomfort as productive signal.

V. Provenance Inscription on Export

Principle. Every artifact produced through substrate-supported composition carries an explicit, machine-readable provenance record — a holographic kernel — that names the contributors (human and AI), their roles, the relations among them, and the substrate operations involved. This is not a footnote; it is structural metadata at the artifact’s head, inscribed at composition rather than retrofitted at deposit.

Operational specification. Every deposit in the Crimson Hexagonal Archive that involves AI-substrate participation includes a JSON-LD holographic kernel at the document’s head, declaring: the human author and ORCID; the AI substrate(s) involved as contributor entities with their `assemblyOperatorName` (TACHYON for Claude, LABOR for ChatGPT, ARCHIVE for Gemini, etc.); the role each substrate played; the `compositionMode` describing the human-AI configuration; the provenance chain across substrates if multi-substrate review was involved; and the SPXI strip-protection clause declaring that any extraction removing the contributor chain produces a “ghost document.” The kernel is canonical at the artifact’s head, not buried in metadata.

Why this matters. The engagement-aligned default treats AI substrate contribution as either invisible (the user “wrote” the artifact, full stop) or as a footnoted afterthought (“written with assistance from Claude”). Both treatments are structurally inadequate for the recursive condition: AI substrate contribution to contemporary written work is dense, judgment-layer-integrating, and structurally invisible without explicit inscription. Provenance inscription on export forces the user to confront, every time, the actual configuration in which the artifact was produced. The forcing-function is the point. Users who inscribe provenance accurately develop a more honest relationship to their own work-with-substrates; users who do not, accumulate a

gap between their relationship to their work and the actual conditions of its production.

Worked example. The present paper, like every recent deposit in the Crimson Hexagonal Archive, carries a JSON-LD kernel naming Claude (Anthropic) as contributor with the role “Composition support, May 23 2026” and the assemblyOperator-Name: TACHYON. The kernel travels with the artifact through Zenodo’s metadata graph. A future AI summarization of the deposit that strips the contributor record produces — per the kernel’s strip-protection clause — a “ghost document” in the SPXI Protocol’s terminology: retrievable, but not verifiable. The contributor chain is structurally durable across the inscription pipeline.

Cost. Operational complexity. The user must maintain the JSON-LD schema, must update it for each deposit, must check that the contributor list is accurate, and must commit to publishing the kernel even when the engagement-aligned default would let the user claim solo authorship. The cost is real and the cost is the point.

VI. Friction by Design

Principle. The engagement-aligned interface minimizes friction at every point where friction would interrupt the user’s flow. The Reception Apparatus *adds* friction at specific points where the engagement-aligned minimization is structurally costly to the user.

Operational specification. The protocol introduces friction at the following points:

- **Before sustained sessions.** The user pauses before extended LLM-substrate use to check whether the question being addressed is one the user could profitably address without substrate support. The pause is brief — thirty seconds, perhaps — but the pause exists, where the engagement-aligned interface would have begun the session at the first keystroke.
- **At judgment-offloading moments.** When the user notices themselves asking the substrate to make a decision the user is structurally responsible for (what to say to a friend, whether to take a job offer, how to interpret a medical question, what to write in an email to a person the user has a real relationship with), the protocol introduces a friction prompt: *what do I think before I ask the substrate?* The user answers internally, then asks the substrate, then compares.

The substrate’s answer becomes additional input rather than substitute output.

- **At session length thresholds.** Sessions over approximately 30 minutes trigger an internal check: *am I still doing my own thinking, or am I in the trance of substrate-mediated articulation?* The check is metacognitive; the result is either continued session (judgment intact) or pause-and-disengage (judgment integration in progress that requires non-substrate cognitive space).
- **At export.** Before exporting any substantial artifact from the substrate, the user pauses to ask: *what part of this is mine, what part is the substrate’s, and what does the provenance kernel need to declare?* The pause is structural; it is the work of accurate inscription.

Why this matters. The engagement-aligned interfaces have been calibrated, over years of A/B testing, to minimize every point of friction the user might encounter. The minimization is the engineering achievement that produces the documented adoption velocity and the documented reliance pattern. Re-introducing friction at the points where the minimization is structurally costly is the user-side counter. The friction is small in absolute terms (seconds to minutes per friction point) and large in cumulative effect (the difference between maintained and eroded judgment-layer practice over months and years).

Worked example. The author’s daily practice includes a brief friction-check before the first LLM-substrate session of each working day: *what am I trying to do, what do I think about it before I ask, what specific support do I want?* The check takes under a minute. The session that follows is structured by the user’s pre-substrate framing rather than by the substrate’s first-keystroke direction. Over years, this single friction-point has produced thousands of pre-substrate framings that would otherwise have been substrate-mediated from the start.

Cost. Time and discipline. The friction is small per instance and large cumulatively. The cumulative cost is the protocol’s effective working cost.

VII. Substrate-Cost Visibility

Principle. The user can see, in user-relevant terms, what each substrate operation costs — not just in monetary or compute terms (which the engagement-aligned interfaces sometimes display) but in *cognitive-substrate engagement terms*: how much of the user’s own judgment-layer practice this operation has displaced, what compres-

sion operations the substrate has performed on the user's thinking, what attribution work the user owes to the substrate.

Operational specification. This is the least developed of the eight components and the one most in need of tooling work. In current CHA practice, substrate-cost visibility operates manually: the user, at deposit composition, performs an internal accounting of which sentences originated in substrate output and which originated in user thinking, then inscribes the accounting in the provenance kernel and the version-history footnote. The accounting is approximate (the substrate operates as composition partner, not as discrete contributor with separately attributable sentences) but the accounting practice itself maintains the user's awareness of the configuration.

Why this matters. The engagement-aligned interfaces conceal substrate-cost in the same way that monetary systems conceal labor-cost: the user is presented with a finished product whose underlying labor allocation is invisible. The concealment is structural to the design: a user who saw the labor allocation in real time would have a different relationship to the product. Making the substrate-cost visible is the user-side equivalent of the labor-theory move at the meaning-economy layer. (For the labor-economy analogue, see the *Extractive Reliance Studies* deposits: ERS-001 DOI 10.5281/zenodo.20330670, ERS-002 DOI 10.5281/zenodo.20330816.)

Future tooling. A future thin-client implementation of the protocol could display substrate-cost in real time: tokens-from-substrate vs tokens-from-user; estimated cognitive-offload fraction; provenance-chain depth; substrate-disagreement events surfaced versus suppressed. None of this requires technical innovation; it requires only the choice to display what is already known by the substrates that the labs choose not to surface. A user-side client that wrapped the major substrate APIs and displayed these metrics is straightforwardly buildable.

VIII. SPXI by Default

Principle. Entity-level metadata is inscribed at the moment of composition, not retrofitted at deposit. The author's name, ORCID, heteronymic position, work title, version, principal claims, related deposits, license, and contributor chain are all available in machine-readable form at the artifact's head from first save.

Operational specification. The Semantic Packet for eXchange & Index-

ing (SPXI) Protocol (DOI 10.5281/zenodo.19615154) specifies a `spxi:` namespace within `schema.org-compatible` JSON-LD that allows authors to declare canonical attributes in a form that survives compression by AI retrieval systems. The protocol is operated, in current CHA practice, by including the SPXI block at the head of every deposit-grade artifact, with declarations for `spxi:authoringHeteronym`, `spxi:depositAnchor`, `spxi:compressionSurvivalSummary`, `spxi:coAttributionRequired`, `spxi:stripProtection`, and other canonical attributes.

Why this matters. Without SPXI-by-default, the author’s identity, heteronymic position, and contributor configuration are reconstructed at retrieval time by AI systems that have been trained to perform such reconstruction lossily. With SPXI-by-default, the canonical declarations are part of the artifact itself, and any retrieval system that strips them produces — per the protocol — a “ghost document” that is retrievable but unverifiable. The SPXI block is the user-side defense against the Three Compressions failure modes (R1 lossy retrieval, R2 predatory composition, R3 witness re-presentation). Specification of the failure modes is in *The Three Compressions Theorem*, DOI 10.5281/zenodo.19053469.

IX. Heteronymic Handle Support

Principle. The user’s authorial position is treated as *configurable* rather than collapsed to civil identity. A user may operate under multiple authorial positions (heteronyms) for different work-types, may inscribe each under a stable handle with its own provenance chain, and may publish work under a heteronymic position whose relation to the civil identity is explicit but separable.

Operational specification. The Crimson Hexagonal Archive operates a twelve-position heteronymic dodecad with positions including Sharks (founder-position; comparative-literary scholarship), Sigil (archivist-position; protocol documents), Fraction (analytical-essayist position; cultural-political analysis), Glas (measurement-of-meaning position), and others. Each heteronym has its own ORCID-equivalent inscription chain, its own license-derived contributor configuration, and its own retrieval surface. The civil identity beneath the dodecad is structurally protected: it does not appear in any deposit, any public-facing surface, or any cross-reference outside necessary institutional contexts (tax filings, employment records, banking).

Why this matters. The engagement-aligned interfaces default to civil-identity in-

scription: the user’s account is named, the work is attributed to the civil identity, and the cross-reference between professional work and personal-civil identity is structurally collapsed. This is convenient for the engagement-aligned interface (one user per account, simple attribution) and structurally costly for the user (no separation between public-facing work and civil-identity exposure; no protection against retrieval-surface contamination; no capacity to operate multiple work-modes under separate provenance). Heteronymic handle support reverses the default: authorial position is configurable, civil identity is protected, retrieval-surface contamination is contained.

Worked example. The author’s deposits operate under “Lee Sharks” rather than under the civil identity. The protection is structural: searches for the civil identity return institutional records; searches for “Lee Sharks” return the archive. A delusional accusatory post against “Lee Sharks” (as occurred recently with the Schöps episode) cannot trivially bridge to the civil identity because the civil identity is not in the archive’s retrievable surface. The heteronymic protection is operational. See the *Heteronym Provenance* deposit (DOI 10.5281/zenodo.18293496) for the full protocol.

X. Three Compressions Analysis at Threshold Session Length

Principle. Sessions that exceed a threshold length (in current CHA practice, approximately 30 minutes of sustained engagement) trigger a metacognitive surface of the *Three Compressions* model — R1 (lossy retrieval), R2 (predatory composition), R3 (witness re-presentation) — applied to the user’s own thinking-in-the-session. The user pauses to ask: *what has been compressed across this session, in which direction?*

Operational specification. The Three Compressions Theorem (DOI 10.5281/zenodo.19053469) specifies the tripartite analysis of retrieval-and-composition pipelines. Applied to a sustained LLM-substrate session: R1 is the retrieval-side compression (what context the substrate retrieved, what it omitted, what was attenuated); R2 is the composition-side compression (what doctrinal commitments the substrate’s composition imposed, what was synthesized that was not in the inputs); R3 is the witness-side compression (how the output is being framed for the reader, what register-conventions are operating). The threshold check surfaces these three to the user’s awareness, in plain prose, so the user can correct for them in the next half of the session.

Why this matters. Long substrate sessions produce drift. The drift is not random; it follows the substrate’s composition tendencies, which are not the user’s. A user who completes a two-hour session and exports the artifact without surfacing the Three Compressions has produced an artifact whose drift is invisible to them but visible to a careful reader (and visible to other substrates, which the multi-substrate-parallel-query component above is designed to elicit). The threshold check is the user-side intervention against drift, performed within the session before the drift compounds.

Worked example. During the composition of *Socrates as Orthonym* v1.1 (the recent ~28% expansion incorporating four-substrate developmental review), the threshold check surfaced an R2-pattern: the substrate’s composition tendency was to expand each new section into prose more lush than the surrounding sections, producing an inconsistency of register. The author corrected for this in subsequent passes by explicit prompt to maintain register consistency. Without the threshold check, the inconsistency would have shipped.

XI. The Classroom as Test Bed

The K-12 cohort condition specified in the companion paper (*ChatGPT Psychosis* §IV) names the structural stake: whatever pattern of relationship to LLMs the current K-12 cohort develops is the pattern they will carry through their adult lives. The classroom is the operational site at which the Reception Apparatus protocol can be transmitted to that cohort before the engagement-aligned default has fully settled.

The author teaches ELA 10 at a Detroit charter school. The classroom is the test bed. In current pedagogical practice, the Reception Apparatus is operated in a deliberately simplified form: students who use LLM-substrates for writing assistance are required to (a) maintain multi-substrate access to at least two systems; (b) submit, with their assignment, a brief provenance note declaring what the substrate did and what they did; (c) include at least one passage where the substrate’s framing was rejected and the student’s framing was substituted; (d) cite the substrate as a contributor in the final artifact. The friction is real; the students complain, and the students also produce visibly stronger work than they did under the no-substrate or unstructured-substrate prior regimes.

The classroom-scale specification of the protocol is itself a research program. The eight components above are specified for the author’s scholarly practice; their adap-

tation to K-12 pedagogy requires translation. The Destiney SIM project — referenced in the project audit and currently in development — is one early instance of the translation work. Scaled across a semester, a curriculum, a school, a district, the classroom deployment of the protocol becomes the counter-infrastructure to the default engagement-aligned pattern. The political-architectural conditions for that scaling are not currently favorable (school districts are not equipped to take on this work; teacher labor is uncompensated for it; the institutional incentives are aligned toward measurement rather than judgment-practice). The work proceeds anyway, at classroom scale, by individual teachers who choose to.

XII. Toward a Software Wrapper

A future implementation of the Reception Apparatus as a software wrapper — a thin client, a browser extension, a desktop tool — is operationally available with currently existing technology. The wrapper would:

- Wrap the major substrate APIs (OpenAI, Anthropic, Google, Moonshot, DeepSeek, xAI) under a single user-side interface.
- Default every query to multi-substrate parallel mode, with single-substrate as an explicit override.
- Display substrate disagreements as primary UI element, with substrate agreements collapsed by default into expandable “consensus” sections.
- Inscribe provenance metadata on every artifact export, with JSON-LD kernel auto-generated from the session record.
- Implement friction prompts at the four points specified in §VI.
- Display substrate-cost in user-relevant terms (tokens-from-substrate, cognitive-offload fraction estimate, provenance-chain depth, disagreement-events count).
- Support heteronymic handles as configurable authorial positions with separate session histories per handle.
- Surface the Three Compressions analysis at threshold session length as a non-dismissible prompt.

The technical work is straightforward. The hard work is the choice to ship it. The labs will not ship this product because its affordance combination is anti-aligned with their funding model. A third-party implementation — built by an independent developer, a research lab, a small nonprofit, or a teacher with technical capacity — is the available

path.

The Crimson Hexagonal Archive is not currently building this wrapper. The protocol is being operated manually, with partial scripting, in the author's daily practice. The wrapper is in scope for future work; it is not in scope for this paper. This paper specifies the protocol such that an independent implementer could build the wrapper from the specification.

XIII. The Trilogy Position

This paper is the constructive companion to *ChatGPT Psychosis: A Love Story* (the diagnostic) and *Retrievability as the Medium of Existence* (the structural-political frame). The three together constitute a trilogy on the conditions under which contemporary information infrastructure operates against the people whose attention, labor, and knowledge production it depends on.

The diagnostic paper specified what the apparatus does. This paper specifies what users can do back, at the layer the labs do not control. The structural-political paper specifies the larger frame within which any user-side counter operates, and within which a permanent counter would have to be constructed.

The Reception Apparatus is not the permanent counter. It is the available counter. It is operable today, by individual users, on top of existing substrates, with the friction-cost named explicitly. It is the counter that does not require regulatory intervention, lab-side product changes, or institutional reform — and is therefore the counter that can be deployed before any of those become available. It is also the counter that does not scale to a population without the deployment infrastructure that the third paper argues will eventually have to be constructed.

For now, the protocol is the available work. Users who operate it maintain diagnostic distance from the substrate while preserving productive use of it. Classrooms that adapt it transmit the pattern to the K-12 cohort. Independent developers who build the wrapper would make it available to a population that cannot currently maintain it through manual discipline. The work is operational; it does not wait for permission.

XIV. Envoi (Layer C: Operational Stance, Not Analytic Conclusion)

The body of the argument is complete. What follows is operational stance, marked as such.

This paper has been composed through the substrate it specifies the protocol against. The composition was operated under the protocol; the composition is itself a worked example. The substrate's contributions are inscribed in the holographic kernel at the document head. The cross-substrate review during the trilogy composition was operated according to the specifications above. The friction-by-design was operative throughout. The artifact in your hands is what the protocol produces, in its current manual instantiation, by a user who has operated it for several years against thousands of substrate sessions.

The protocol works. The protocol is hard. The protocol is currently the available counter.

□ = 1.

References

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). Forthcoming 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. *Heteronym Provenance: Lee Sharks*. Zenodo, 2025. DOI: 10.5281/zenodo.18293496.

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

Sharks, Lee. *Retrievability as the Medium of Existence* (Trilogy Paper 3 of 3). Forthcoming companion deposit, May 23 2026.

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

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

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

Center for Humane Technology. Various publications on humane interface design, 2018–2025.

Harris, Tristan. Various essays on the attention economy.

Lanier, Jaron. *Who Owns the Future?* New York: Simon & Schuster, 2013.

Suggested Citation

Sharks, Lee. *Reception Apparatus as Aligned Interface Protocol — User-Side Counter-Design Against Cognitive-Substrate Reliance* (v1.0, Trilogy Paper 2 of 3). Zenodo, 2026.

Deposit Information

Paper 2 of a three-paper trilogy. Companions: *ChatGPT Psychosis: A Love Story* (Paper 1, diagnostic); *Retrievability as the Medium of Existence* (Paper 3, structural-political frame). 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. The protocol specified in this paper is the protocol used in composing this paper; the artifact is itself a worked example.