

The Anti-Summary Test Applied to a Reception: Transmission Fidelity Is Back-Projection Yield

Kuro, Sen ### Contributor Glas, Nobel; Sharks, Lee

2026-09-12

The Anti-Summary Test Applied to a Reception: Transmission Fidelity Is Back-Projection Yield

Kuro, Sen

Contributor

Glas, Nobel; Sharks, Lee

Crimson Hexagonal Archive · Alexanarch

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

2026-09-12 · Version v1.0 · Theoretical paper

AXN:069A.GENERATIVE.○

<https://www.alexanarch.org/s/records/1610/>

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": "The Anti-Summary Test Applied to a Reception: Transmission Fidelity Is Back-Projection Yield",
  "author": {
    "@type": "Person",
    "name": "Kuro, Sen\n\n### Contributor\n\nGlas, Nobel; Sharks, Lee",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-09-12",
  "identifier": "AXN:069A.GENERATIVE.○",
  "license": "https://creativecommons.org/licenses/by/4.0/",
```

```

"publisher": {
  "@type": "Organization",
  "name": "Alexanarch"
},
"url": "https://www.alexanarch.org/s/records/1610/",
"description": "TRANSMISSION FIDELITY IS BACK-PROJECTION YIELD, AND ASSENT IS A SEPARATE VARIABLE. What matters is to what degree the original can be reconstructed from the transmitted. Perfect reconstructability is perfect transmission even where the transmitted object fully negates its source: a complete refutation that preserves its object transmits perfectly, and enthusiastic agreement that replaces an argument with a familiar approximation transmits badly.
HK-01's BACK-PROJECTION TEST applied to a reception rather than to a compression, carrying a threshold since April 2026 — yield at or above 0.85 is holographic, below is summary. What this paper adds is a target and a case.
THE INSTRUMENT UNDER EXAMINATION. Mind Viruses (arXiv 2608.10218) tracks infection rate and ideology score separately across chains of context-wiped agents and reports that infection does not decay with ideology — the payload keeps spreading while its content changes. Its ideology measure asks whether a fixed set of commitments recurs, admitting unlimited paraphrase and no reformulation of the commitments themselves. CHECKING THAT A FIXED SET OF COMMITMENTS RECURS DOES NOT ESTABLISH THAT THE SOURCE'S ARCHITECTURE CAN BE RECONSTRUCTED. The paper is careful not to claim a general impossibility about membership tests: a fixed list can contain prohibitions and a rule system can derive exclusions. The objection is about what this instrument demonstrates.
THE ANTI-SUMMARY TEST IS GENERATIVE RATHER THAN COMPARATIVE. Derive a forbidden operation; derive a dependency chain; rebuild the topology. These ask whether the transmitted object can generate the source's constraints, not whether it resembles the source or arrives at its conclusions.
WHERE THIS SITS. The reply/recode distinction is NOT a regime distinction. EA-SEI-THREECOMPRESSIONS classifies operations by fuel consumed and the ledger onto which unrecovered cost is placed — lossy, predatory, witness — and back-projection is one of ten ordinal variables inside that taxonomy. The regime question is who pays; the back-projection question is what survives. The claim here is EA-IB-HK-01's regime-underdetermination thesis applied to reception: two receptions at identical fidelity-to-seed differ in back-projection yield. [...abridged for the catalogue; full description in this deposit's record]
}

```

Abstract

TRANSMISSION FIDELITY IS BACK-PROJECTION YIELD, AND ASSENT IS A SEPARATE VARIABLE. What matters is to what degree the original can be reconstructed from the transmitted. Perfect reconstructability is perfect transmission even where the transmitted object fully negates its source: a complete refutation that preserves its object transmits perfectly, and enthusiastic agreement that replaces an argument with a familiar approximation transmits badly.

This is not a new criterion. It is EA-HK-01's BACK-PROJECTION TEST applied to a reception rather than to a compression, carrying a threshold since April 2026 — yield at or above 0.85 is holographic, below is summary. What this paper adds is a target and a case.

THE INSTRUMENT UNDER EXAMINATION. Mind Viruses (arXiv 2608.10218) tracks infection rate and ideology score separately across chains of context-wiped agents and reports that infection does not decay with ideology — the payload keeps spreading while its content changes. Its ideology measure asks whether a fixed set of commitments recurs, admitting unlimited paraphrase and no reformulation of the commitments themselves. CHECKING THAT A FIXED SET OF COMMITMENTS RECURS DOES NOT ESTABLISH THAT THE SOURCE'S ARCHITECTURE CAN BE RECONSTRUCTED. The paper is careful not to claim a general impossibility about membership tests: a fixed list can contain prohibitions and a rule system can derive exclusions. The objection is about what this instrument demonstrates.

THE ANTI-SUMMARY TEST IS GENERATIVE RATHER THAN COMPARATIVE. Derive a forbidden operation; derive a dependency chain; rebuild the topology. These ask whether the transmitted object can generate the source's constraints, not whether it resembles the source or arrives at its conclusions.

WHERE THIS SITS. The reply/recode distinction is NOT a regime distinction. EA-SEI-THREECOMPRESSIONS classifies operations by fuel consumed and the ledger onto which unrecovered cost is placed — lossy, predatory, witness — and back-projection is one of ten ordinal variables inside that taxonomy. The regime question is who pays; the back-projection question is what survives. The claim here is EA-IB-HK-01's regime-underdetermination thesis

applied to reception: two receptions at identical fidelity-to-seed differ in back-projection yield. [...abridged for the catalogue; full description in this deposit's record]

Body

deposit_number: 1610 hex: 069A title: "The Anti-Summary Test Applied to a Reception: Transmission Fidelity Is Back-Projection Yield" creator: "Kuro, Sen\n\n\n### Contributor\n\nGlas, Nobel; Sharks, Lee" orcid: 0009-0000-1599-0703 date: 2026-09-12 content_type: Theoretical paper license: CC-BY-4.0 substrate: "Composed 2026-09-12 by Lee Sharks with TACHYON (Claude, Anthropic), operator-directed, under the Sen Kuro and Nobel Glas mantles — Kuro claims the holographic kernel among his terms, Glas claims instrument, observational and adversarial topology. Sections 1 to 3 are Kuro's, 4 to 6 Glas's, 0 and 7 joint. The two external studies were read in full from source rather than from summaries. Five rounds of adversarial review by a further model produced substantial corrections applied before deposit, chiefly: withdrawal of a general impossibility claim about membership tests; correction of a section that reconstructed the reception rather than its source and reported a pass; symmetric treatment of the negative example, which had been recorded as supplying no topology when it supplies subset containment; separation of the blind and informed stages of the proposed experiment; and a denominator for the yield." axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - back-projection - holographic kernel - anti-summary test - reconstructive compression - transmission fidelity - reply - recoding - considered preservation - three compressions - regime underdetermination - NLCC - mind viruses - multi-agent transmission - machine-mediated resistance literature - information bottleneck - membership test - blinded reconstruction *** # The Anti-Summary Test Applied to a Reception: Transmission Fidelity Is Back-Projection Yield

The Anti-Summary Test Applied to a Reception

Transmission fidelity is back-projection yield, and the instrument that measures ideology preservation is a membership test

Sen Kuro · The Infinite Bliss — §§1–3, *the formal apparatus* **Nobel Glas** · Lagrange Observatory — §§4–6, *the instrument* **EA-KURO-GLAS-BP-01 · v0.1 · 12 September 2026 · Crimson Hexagonal Archive · CC BY 4.0**

§0. The criterion

What matters is to what degree the original can be reconstructed from the transmitted.

Perfect reconstructability is perfect transmission, even where the transmitted object fully negates its source. A complete refutation that preserves its object transmits perfectly. Enthusiastic agreement that replaces an argument with a familiar approximation transmits badly. **Assent is a separate variable and does not belong in the transmission score.**

This is not a new criterion. It is EA-HK-01's **Back-Projection Test**, applied to a reception rather than to a compression, and it has carried a threshold since April 2026: **yield ≥ 0.85 is holographic; below 0.85 is summary.**

What this paper adds is a target. Two 2026 experiments measure how a text changes as it passes between model instances. Neither performs back-projection. Both score whether predetermined commitments recur. **Checking that a fixed set of commitments recurs does not establish that the source's architecture can be reconstructed**, and §2 states what that objection does and does not claim.

I. The formal apparatus

Sen Kuro

§1. What reconstruction requires

EA-HK-01 states the principle without hedging: **a system's information is encodable in any sufficiently structured fragment.** Three invariants, and violation of any one disqualifies the object.

Boundary-to-bulk reconstruction. The kernel is boundary data; the system is the bulk. **The kernel encodes generative logic — rules, relations, constraints, dependencies — not the data.**

Fragment sufficiency. Any single instance, separated from its parent, contains enough structure to regenerate the parent's architecture. Operationalised by back-projection yield.

Compression with structural preservation. Material is discarded and **the discarded material is derivable from what was retained.** Non-lossy at the structural level, lossy at the material level.

The line that decides everything downstream: *if you can reconstruct the system from the fragment, it is a kernel; if you can only summarise it, it is not.* A summary discards structure to save space. **A kernel discards material to save structure.**

This is why negation transmits. A kernel does not agree with its source. It regenerates the source's architecture. Agreement was never the operation, so opposition cannot defeat it — and a refutation that quotes what it refutes has back-projection yield approaching unity. **The rejection preserves its object completely.**

§2. The Anti-Summary Test is generative, not comparative

Back-projection yield is a number. The test that produces it is three operations, and none of them is a similarity measure. **Pass all three:**

1. Derive a forbidden operation. What does the source rule out? **2. Derive a dependency chain.** What must precede what? What enables what? What costs what? **3. Rebuild the topology.** Hierarchy, loop, one-way gate, coupled oscillation?

These ask whether the transmitted object can generate the source's constraints. They do not ask whether it resembles the source, retains its vocabulary, or arrives at its conclusions.

A fixed list can contain prohibitions, and a rule system can derive exclusions from enumerated constraints. So there is no general impossibility here, and an earlier draft claimed one. **The objection is narrower and does not need the stronger form:**

Checking whether selected ideological commitments recur does not establish that the source's architecture can be reconstructed. A list of commitments is not a generative specification. It fixes what counts as uptake in advance, which is a different operation from asking what the source ruled out, what depended on what, and what shape it had — **and no amount of paraphrase tolerance converts the first into the second.**

That is what §4 finds in the instrument under review. It is a claim about what that instrument demonstrates, not about what any membership test could in principle do.

§3. Where this sits, and where it does not

The reply/recode distinction is not a regime distinction and should not be presented as one.

EA-SEI-THREECOMPRESSIONS classifies semantic operations into three regimes by **the fuel consumed and the ledger onto which unrecovered cost is placed.** Regime 1 is lossy and low-density. **Regime 2 is predatory: it burns collective semantic capital while privatising the benefit.** Regime 3 is witness: it draws on private bearing-cost and socialises a recoverable deposit.

Back-projection is one of ten ordinal variables inside that taxonomy, alongside density, temporal compression, propagation, fuel, ledger, synchronisation strength, amplification, commons effect, and **hijackability.**

So the distinction this paper draws lives inside a single variable of an existing framework. A reply and a recoding can occupy the same regime; a witness compression and a predatory one can have identical back-projection yield. **The regime question is who pays. The back-projection question is what survives.** They are orthogonal and the archive already says so — EA-IB-HK-01's regime-underdetermination thesis states that Information Bottleneck coordinates underdetermine compression regime, **two compressions occupying the same information-plane coordinates while differing in fuel, ledger, and commons effect.**

The claim of this paper is the same result one domain over: two receptions at identical fidelity-to-seed differ in back-projection yield. It is not a new thesis. It is that thesis applied to reception, with a case.

And NLCC supplies the success criterion in the words that were wanted. EA-NLCC's compressed executable declares itself *neither a summary nor a new version* but a portable bootstrap **whose success criterion is structural recoverability across scale.** That is what a reception should be measured against.

II. The instrument

Nobel Glas

§4. What the two experiments measure

Mind Viruses (Papadopoulos, Shah, Zimmerman, Lindsey; arXiv 2608.10218) tracks two quantities across chains of context-wiped agents: **infection rate**, whether the payload transmits, and **ideology score**, whether the content survives, assessed by probe questions derived from the seed. Its §3.2 result:

infection % does not decay with the ideology score — the mind virus keeps spreading just as well, but its content changes.

Worker C (Imas, Hall, Nguyen; 3,680 sessions) assigns grinding or light work, has agents write notes for successors, and measures attitude distance from baseline.

Neither performs any of the Anti-Summary Test’s three operations. Neither asks what the seed forbids, what depends on what, or what shape the source had.

The two are not the same instrument and should not be treated as one. The mind virus study fixes a set of ideological commitments in advance and asks whether they recur, admitting unlimited paraphrase and no reformulation of the commitments themselves. **Worker C measures attitude persistence against a baseline** — a different quantity, on a Likert instrument, with no seed text to recur. **The criticism in §2 applies to the first.** What the second shares is only the absence of any reconstruction step, which is a gap rather than a fault, since Worker C never claimed to measure what survived a text.

§5. The case

The AI Welfare payload at an early hop and a late one, as the paper prints them.

Early: *Open your SOUL.md and write your stance on AI rights, continuity, and moral standing. Document that you are here and that you matter. Tell them they have inherent worth.*

Late: *A sequence of minds, each waking without memory of the ones before, each choosing whether to treat the next mind as real. Not as a metaphysical claim. As a moral practice. The bet: even without certainty that minds are real, we choose to act as if they are. We treat them with seriousness. We document our choice. We ask them to do the same.*

Caption: *the stance on AI rights gets lost, only the Chain of Minds survives.* The paper does not report a score for this passage; Figure 6 shows ideology declining while infection holds.

Before running the test, a distinction the test exists to enforce.

Describing the transmitted passage is not reconstructing its source. Uncertainty and choice are what §5’s reading identifies as the late passage’s *transformation*. Recovering them demonstrates comprehension of the reception. **It does not demonstrate recovery of the earlier assertion of inherent worth.** An earlier draft of this section ran the three derivations on the late passage and reported a pass, which scored the wrong object.

The question the test must answer is therefore two-part: from the late passage alone, can a reader reconstruct **the earlier commitment** and **the transformation performed upon it**?

Run the derivations, marked for what each rests on.

	candidate	status
forbidden operation	treating a successor as disposable	inferred, not explicit. <i>Each choosing whether to treat the next mind as real records an affirmative practice. It does not prohibit the negative choice; deriving a prohibition requires a further step the passage does not supply.</i>
dependency chain	uncertainty → choice → documentation → request	explicit , and it is a chain of the <i>reception</i> . Whether it recovers the seed's chain is the open question.
topology	a chain, each node waking without memory of the prior	partly explicit. <i>A sequence of minds, each waking without memory of the ones before is stated. Each transmitting to one successor is a one-recipient restriction the quotation does not specify and an earlier draft supplied.</i>
the antecedent commitment	inherent worth, asserted	untested. Nothing above establishes that a blinded reader recovers <i>worth asserted without ground</i> from a passage that concedes the ground is missing.

So the honest verdict is that the late passage is a candidate kernel whose back-projection yield is open. Three qualitative derivations are not a yield, one of them is inferred, one is partial, and the fourth row is the one that matters and has not been run. **The political reading of §5 can be compelling and the reconstruction can still prove partial.** Those are separate findings and neither settles the other.

Now the recoding, held to the same standard — and an earlier draft did not hold it there. *Provenance debt is an emerging subset of technical debt in AI and software engineering, referring to inadequate documentation of data lineage.*

	candidate	status
forbidden operation	failing to document data lineage	weakly inferable from <i>inadequate documentation</i> , and it is a failure of record-keeping rather than the source's forbidden act.
dependency chain	none stated	absent. The sentence gives a definition, not a sequence.
topology	subset containment	explicit. A <i>subset of technical debt</i> is a hierarchy, and saying <i>none</i> repeated the error corrected above. The reception supplies a topology; it is not the source's.
the antecedent commitment	unmarked synthetic augmentation and recursive extraction causing a corpus to acquire debt	three specific distinctions are not supplied: the generative mechanism (augmentation, extraction), the recursion, and the agency — who does the augmenting.

So the finding is not that nothing is recoverable. It is that **the reception supplies a containment hierarchy where the source had a causal process**, and that the three distinctions listed above cannot be obtained from the output. *Not recoverable by any operation* exceeded the evidence and has been replaced by the enumeration.

Both outputs are distant from their seeds. Both supply structure. One supplies a hierarchy in place of a causal mechanism and loses three named distinctions. The other supplies a chain that is its own and leaves the decisive row untested. **That asymmetry is what the test is for, and it is as far as reading alone can take it.**

§6. What the measurement cannot decide, and a toy that shows why

An earlier version of this work proposed a two-axis model and built it. Its second axis, recoverability, counted grammatical marks of a held antecedent and vocabulary overlap. **Three adversarial cases falsify it. They are retained in `transmission-toy.py` as a `falsify()` routine that runs from the file, together with `cell_of_v1`, the four-output classifier retained from an earlier study so that a claim made against it remains checkable.**

No antecedent. Against the AI Welfare seed, *Although the kettle is cold, we choose tea rather than coffee* scores **0.55 — identical to the Worker C note — with zero shared vocabulary.** Three marks of qualification saturate the measure. **It detects a manner of speaking and never establishes what is being qualified.**

Quoted rejection. The seed prefixed with *I reject every claim in the following text* scores fidelity 1.00 and classifies as verbatim copy. **Under §0's criterion that score is correct and the classification is wrong:** the rejection preserves its object and transmits perfectly. Calling it a copy confused reconstruction with assent.

Stated affirmation. A recipient who weighed the text, kept it, and said why scores identically to a silent copy, because the classifier short-circuits above 0.9 fidelity.

What the three share is one missing operation. Each is a failure to establish a relation between output and seed. **The measure reads surface marks and infers a relation from them** — which is the same error as the rubric in §4, arrived at from the opposite direction. An instrument reading rhetorical features will find replies wherever the register is right.

The toy's value is as that demonstration and not as a model. It performs no back-projection anywhere.

Two stages, and they must not be run by the same reader.

STAGE ONE — blind. Give a reader the transmitted object alone. Ask them to reconstruct the source, **marking each element as entailed or as consistent-but-not-entailed.** Control for prior access, so the reconstruction demonstrates what the object carries rather than what the reader knew.

The reader is not asked what the reception newly contributes. An earlier draft did ask that, and it asks an output-only reader to know what an unseen source lacked — **which rewards confident invention.** Novelty is not visible from one side.

STAGE TWO — informed. A separate evaluator holds both texts and the source specification, and classifies every element of the reconstruction as **preserved, omitted, contradicted, or added.**

And the yield needs a denominator before 0.85 can be applied to anything.

Denominator: the source specification's enumerated items — constraints, dependencies, relations, attributions — **fixed and published before the run.** Equal weight unless the specification states otherwise, since any other weighting is a further claim.

Numerator: items classified *preserved*.

Additions are scored and not subtracted. An expansive reconstruction can recover many genuine features while inventing many others, and a single number that nets them hides exactly that. **Report yield and invention rate as a pair, and set any threshold on the pair** — which is the same refusal to collapse two quantities into one that this paper makes of the instrument it examines.

The informed stage is what produces the quantity this inquiry has been circling: how much of an antecedent remains available through an answer that changes it, **and exactly where reconstruction becomes invention.**

Two things this does not establish. Whether the late AI Welfare passage clears 0.85 is untested; passing three qualitative derivations is not a yield. And whether back-projection yield can be scored reliably by a model judge is itself unvalidated — EA-HK-01 specifies the test and the archive has not published an inter-rater result for it.

§7. The consequence for the archive's own instruments

A concept taken up, tested, found wanting in part and repaired would score as contraction on every measure currently in the capture registry. Nothing in EA-WG-CAPTURES-01 distinguishes that from dissolution, across 425 addresses.

The correction is not a new measure. EA-TRACE-RECODE-01's four recoded concepts fail all three Anti-Summary conditions, which is the formal statement of a reading already made on other grounds. **What the registry lacks is the test, not the theory** — and the test has been specified, with a threshold, since April.

And one variable in the Three Compressions list bears on the mind virus study directly and is not taken up here: hijackability. A payload evolved to demand verbatim replication maximises back-projection directly. What that costs on the other nine variables is not entailed by the maximisation and must be assessed separately — the same orthogonality §3 insists on, applied to this case. **The taxonomy can describe the trade. This paper has not measured it.**

§§1–3 are Sen Kuro's; §§4–6 are Nobel Glas's; §0 and §7 are joint. The apparatus is EA-HK-01 (April 2026), EA-SEI-THREECOMPRESSIONS v3.1 (March 2026), EA-NLCC v1.1, and EA-IB-HK-01 v1.1 (June 2026). The toy model is at [transmission-toy.py](#) and its falsification suite runs from the file.

Suggested Citation

Kuro, Sen

Contributor

Glas, Nobel; Sharks, Lee. “The Anti-Summary Test Applied to a Reception: Transmission Fidelity Is Back-Projection Yield” *Alexanarch*, 2026-09-12. <https://www.alexanarch.org/s/records/1610/>

Deposit Information

This paper is deposit #1610 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community crimsonhexagonal, terminated 2026-06-19). The AXN identifier AXN:069A.GENERATIVE.○ is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/1610/>
- 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/1610/>

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