

Can the Ontology Be Wrong and Find Out? Source-Admission Systems, the Proxy for Standing, and the Damage a Representation Does to Itself

Sharks, Lee

2026-09-17

Can the Ontology Be Wrong and Find Out? Source-Admission Systems, the Proxy for Standing, and the Damage a Representation Does to Itself

Sharks, Lee

Crimson Hexagonal Archive · Alexanarch

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

2026-09-17 · Version v1.0 (transport object) · Journal article, transport object

AXN: 06BD.GENERATIVE

<https://www.alexanarch.org/s/records/1627/>

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": "Can the Ontology Be Wrong and Find Out? Source-Admission Systems, the Proxy for Standing, and the Dam",
  "author": {
    "@type": "Person",
    "name": "Sharks, Lee",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-09-17",
  "identifier": "AXN:06BD.GENERATIVE",
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
```

```

    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/1627/",
  "description": "THE COMPANION TRANSPORT OBJECT, AND ITS CRITERION IS ONE QUESTION. This is the ontology strand
contained article, paired with the valuation strand and sharing its primitive. Its subject is the damage a source-
admission ontology does to ITSELF when a proxy for epistemic standing substitutes for the judgment the proxy was
TrustRank, SpamRank, EigenTrust and the feature sets around them are the standing apparatus of adversarial info-
proof. And the endogenous alternative was proposed by the incumbent – Knowledge-
Based Trust, from Google Research in 2015, distinguished itself from every topological method on one axis, an e
}

```

Abstract

THE COMPANION TRANSPORT OBJECT, AND ITS CRITERION IS ONE QUESTION. This is the ontology strand written as a single self-contained article, paired with the valuation strand and sharing its primitive. Its subject is the damage a source-admission ontology does to ITSELF when a proxy for epistemic standing substitutes for the judgment the proxy was introduced to approximate, and its whole argument reduces to a question that can be asked of any representational system: can evidence from a misrepresented object revise the representation that misrepresented it? Where it can, the ontology can be wrong and find out. Where it cannot, the ontology can only be scored.

IT ENTERS THROUGH THE DISCIPLINES AND CONCEDES THE CASE FOR TOPOLOGY FIRST. Ranking by graph structure is given in full and endorsed, with its founding description doing the work: PageRank was specified as query independent and content free, and that abstention is the source of its virtues, since at web scale a ranking system that had to adjudicate every claim on every page would not work at all. TrustRank, Anti-TrustRank, SpamRank, EigenTrust and the feature sets around them are the standing apparatus of adversarial information retrieval and are treated as such.

THEN IT LOCATES ITS OWN DIAGNOSIS IN PRIOR WORK. The manipulability of topological trust was named early and formally by the field: collusion topologies studied in 2005, and a proof that symmetric reputation mechanisms of this family cannot be Sybil-proof. And the endogenous alternative was proposed by the incumbent — Knowledge-Based Trust, from Google Research in 2015, distinguished itself from every topological method on one axis, an endogenous signal, the correctness of the factual information a source provides. On the philosophical side, testimonial injustice has for twenty years described a credibility discount that tracks category membership rather than evidence and is resistant to rational revision, with recent work naming the variant in which such deficits become true by shaping the very conditions that sustain them. [...abridged for the catalogue; full description in this deposit's record]

Body

deposit_number: 1627 hex: 06BD title: Can the Ontology Be Wrong and Find Out? Source-

Admission Systems, the Proxy for Standing, and the Damage a Representation Does to Itself creator: Sharks, Lee orcid: 0009-0000-1599-0703 date: 2026-09-17 content_type: Journal article, transport object license: CC-BY-4.0 substrate: "Composed 2026-09-17 by Lee Sharks with TACHYON (Claude, Anthropic), operator-directed, as the transport form of the EA-MIRROR-01 notebook (#1613). The adjacent literatures were searched rather than recalled, and the search located the paper's nearest precedent inside Google Research and its philosophical structure in Fricker. Five reviewer repairs were verified and applied before deposit: operative standing defined in the body rather than left in the keywords; the incapacity narrowed from architectures to the operation; the legibility theorem narrowed from monotonic to existential with the stronger form retained for the measured case; the level-three hypothesis made system-relative; and the references completed, which supplied a sharper formulation than the paraphrase it replaced. A sixth change came from a composition of the draft that identified the paper's real gap — that it glossed why the endogenous alternative was not adopted — and §2.1 is the result." version: v1.0 (transport object) related_ids: "#1613 (EA-MIRROR-01, the notebook this transports); #1611 (the negative ontology, cited as the deeper technical source); #1615; #1616; #1612; #1626 (Value Before Number, the paired transport object, sharing operative standing as its primitive)" axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - source-admission ontology - operative standing - epistemic proxy - trust graph - PageRank - TrustRank - knowledge-based trust - ground-truth regress - testimonial injustice - self-fulfilling testimonial injustice - retrieval-augmented generation - knowledge graph - source evaluation - ontology engineering - corrigibility - inverted legibility incentive - reading - recursive representation - transport object - Crimson Hexagonal Archive *** # Can the Ontology Be Wrong and Find Out? Source-Admission Systems, the Proxy for Standing, and the Damage a Representation Does to Itself

```
{ "@context": [ "https://schema.org", { "spxi": "https://spxi.dev/ns/v1#" } ], "@type": "ScholarlyArticle", "name": "Can the Ontology Be Wrong and Find Out? Source-Admission Systems, the Proxy for Standing, and the Damage a Representation Does to Itself", "alternateName": "EA-MIRROR-SUMMATIVE-01", "version": "1.0", "datePublished": "2026-09-17", "license": "https://creativecommons.org/licenses/by/4.0/", "author": { "@type": "Person", "name": "Sharks, Lee", "affiliation": "Semantic Economy Institute, Crimson Hexagonal Archive", "orcid": "0009-0000-1599-0703" }, "spxi:hexAddress": "06.SEI.MIRROR.SUMM.001", "spxi:pillarStatus": "transport object — the ontology strand as a single self-contained article for external venues", "spxi:transportPair": { "this": "what happens when a representation cannot be revised by its object", "companion": "Value Before Number (#1626) — what happens when a representation compresses value", "shared_primitive": "operative standing, defined in both", "this_sequence": "source → admission → standing → revision", "companion_sequence": "reference → standing → price formation → write-back" }, "spxi:criterion": "Can evidence from the misrepresented object revise the representation that misrepresented it? Where it can, the ontology can be wrong and find out. Where it cannot, the ontology can only be scored.", "spxi:contributesOperators": [ { "name": "reading (operational)", "role": "permitting a text's claims, distinctions, evidence, corrections and internal relations to revise the representation under which the text is evaluated; tokenization, embedding, extraction, classification and summarization are not that, and a system can consume every byte of a document and read none of it" }, { "name": "operative standing", "role": "the capacity of evidence to alter, block or redirect the decision that governs its own admission; evidence may be present, retrievable and correctly represented while possessing none" },
```

{ "name": "source-relative update barrier", "role": "E(s) exists and does not reach P(s) or A(s): corrective evidence available only downstream of the admission its prior controls" }, { "name": "inverted legibility incentive", "role": "under a topology-only proxy, some operations that increase semantic legibility can simultaneously increase the features associated with manufactured standing" }, { "name": "corrigible against scorable", "role": "the distinction between an ontology revisable by its objects and one merely evaluable by its operators" }], "spxi:engagesScholarship": { "adversarial_IR": "Brin and Page 1998, whose own description of PageRank as query independent and content free supplies the paper's §1; Kleinberg 1999; Gyöngyi, Garcia-Molina and Pedersen 2004 on TrustRank; Benczúr et al. 2005 on SpamRank; Baeza-Yates, Castillo and López 2005 on PageRank increase under collusion topologies; Cheng and Friedman 2005 on the impossibility of Sybilproof reputation; Kamvar et al. 2003 on EigenTrust; ClaimTrust 2025 importing the proxy into retrieval-augmented generation.", "the_precedent": "Dong et al. 2015, Knowledge-Based Trust (Google Research), which distinguished its proposal from every topological method on one axis — an endogenous signal, the correctness of what a source says — and is cited here as prior.", "epistemology": "Fricker 2007 on testimonial injustice, whose structural features are the barrier's; De Bruin 2025 on self-fulfilling testimonial injustice, where credibility deficits become true by shaping the conditions that sustain them; Medina 2013 on credibility excess; Wanderer 2012." }, "spxi:noveltyClaim": "NARROW AND STATED IN §4. Adversarial IR evaluates the proxy by RANKING QUALITY. Epistemology evaluates the credibility discount by the WRONG TO THE SPEAKER. Neither asks about the representation's own corrigibility — the third harm, which is the paper's subject.", "spxi:regressResult": "§2.1, added before deposit and the paper's hardest result: the corrigibility criterion does NOT favour endogenous trust over topological trust. Evaluating correctness requires a ground-truth corpus, that corpus is itself a source-admission ontology with its own priors, and a source excluded from it cannot correct it. The barrier recurs — from a graph statistic to a curated corpus. So the criterion applies to both and neither satisfies it by construction, and Knowledge-Based Trust's non-adoption is more likely a recognition that endogenous trust relocates the gatekeeping problem than a failure to see its merit.", "spxi:disambiguation": { "the incapacity is the operation's, not the architecture's": "a larger system combining topology with endogenous content evaluation would have the faculty Fricker's remedy requires, in the part performing the evaluation; what is settled is that the topology-only operation cannot supply its own correction.", "the legibility claim is existential": "a collision is exhibited between a specific family of legibility operations and a specific family of topology features; covariance in general is neither shown nor needed. The monotonic form holds for the measured case only.", "level three is system-relative": "a diagnosis of a black box that cannot be verified from outside it. One system shown not to use such a variable does not falsify the attribution to another, and no aggregate claim about contemporary retrieval is made that a single architecture could settle.", "the criterion is not a design recommendation": "it does not say use content rather than topology; it asks whether a representation can be revised by the object it misrepresents." }, "spxi:falsifiers": "§9, four conditions plus the defence the paper refuses. The barrier fails on the exhibition of a documented path by which corrective evidence from a low-prior source reaches the variables governing its admission without first passing the gate.", "spxi:owed": "§10. An independent audit of the topology finding, since the fleet measured is the author's own. And a positive construction, whose difficulty §2.1 explains: the path would have to be cheap enough to be always open and narrow enough not to be an attack surface, and the tension between corrigibility and manipulability is why

the problem is open.”, “spxi:requires”: [“#1613”, “#1611”, “#1615”, “#1616”, “#1612”, “#1626”], “keywords”: [“source-admission ontology”, “operative standing”, “epistemic proxy”, “trust graph”, “PageRank”, “TrustRank”, “knowledge-based trust”, “testimonial injustice”, “retrieval-augmented generation”, “knowledge graph”, “source evaluation”, “ontology engineering”, “corrigibility”, “recursive representation”, “ground-truth regress”, “Crimson Hexagonal Archive”] }

Abstract

This paper concerns the damage a **source-admission ontology** does to **itself** when a **proxy for epistemic standing** substitutes for the judgment the proxy was introduced to approximate. The condition is stated as a single criterion: *can evidence from a misrepresented object revise the representation that misrepresented it?* Where it can, the ontology can be wrong and find out. Where it cannot, the ontology can only be scored.

The paper defines **reading** operationally — permitting a source’s claims, distinctions, evidence, corrections and internal relations to revise the representation under which that source is evaluated — and distinguishes it from tokenization, embedding, extraction, classification and summarization, none of which are reading in this sense. A system can consume every byte of a document and read none of it.

It then names the **source-relative update barrier**: where a prior on a source controls admission, and corrective evidence from that source is available only downstream of admission, the source cannot correct the prior that excludes it. The barrier is stated in terms of **operative standing** — the capacity of evidence to alter, block or redirect the decision that governs its own admission — and the paper’s central observation is that evidence may be present, retrievable and correctly represented while possessing none. The barrier’s structure is shown to be formally identical to a well-theorized condition in the epistemology of **testimony**, with the harms running in opposite directions: the epistemological literature theorizes the wrong done to the speaker, and this paper theorizes the degradation of the hearer’s representation. A second consequence, the **inverted legibility incentive**, is that under a topology-only regime the operations which increase a source’s semantic legibility can move it toward the spam signature.

The diagnosis of proxy insufficiency is not new and is not claimed as new: adversarial information retrieval established the manipulability of link-based trust two decades ago, and Google Research proposed an endogenous, content-based alternative in 2015. What that literature evaluates is **ranking quality**; what this paper evaluates is the **representation’s capacity to be corrected by its object**.

Keywords: source-admission ontology; operative standing; epistemic proxy; trust graph; PageRank; TrustRank; knowledge-based trust; testimonial injustice; retrieval-augmented generation; knowledge graph; source evaluation; ontology engineering; corrigibility; recursive representation

1. The conventional account, stated in full

Ranking by graph topology is one of the most consequential ideas in information retrieval and it earned its place.

Brin and Page (1998) treat a hyperlink as a vote and rank pages by the stationary distribution of a random walk over the link graph. The design is deliberate in its abstention, and the original description says so plainly: the method is **query independent and content free**. That abstention is the source of its virtues. It is cheap, scalable, and it does not require the ranking system to adjudicate any claim made on any page — which, at web scale, is not a concession but the only thing that works.

The line has a long provenance and a serious literature. Citation analysis long predates web ranking. Kleinberg’s authority-hub decomposition arrived alongside PageRank. Gyöngyi, Garcia-Molina and Pedersen’s **TrustRank** (2004) propagates trust from a human-vetted good seed set through outlinks; **Anti-TrustRank** and SpamRank (Benczúr et al. 2005) propagate distrust backwards from a bad seed set. EigenTrust does the equivalent for peer-to-peer reputation. The feature sets that grew around them — degree, reciprocity, assortativity, supporters at distance d , host-level aggregation — form the standing apparatus of adversarial information retrieval.

This is a working technology and nothing below suggests otherwise. The question taken up here begins at a boundary the literature itself has marked.

2. What adversarial information retrieval already established

The proxy’s limits were named early, formally, and from inside the field.

Topological trust is manipulable by construction. Baeza-Yates, Castillo and López (2005) studied PageRank increase under different collusion topologies. Cheng and Friedman (2005) proved that symmetric reputation mechanisms of this family cannot be Sybil-proof, and the result generalizes: a node can accrue standing by creating identities and edges rather than by being right about anything. Later work concedes it directly — PageRank is “in principle not Sybil-tolerant.”

And the endogenous alternative was proposed by the incumbent. Dong et al., “Knowledge-Based Trust: Estimating the Trustworthiness of Web Sources” (Google Research, 2015), is the precedent this paper is nearest to, and its framing is exact. Reviewing PageRank, authority-hub analysis, EigenTrust, TrustMe, TrustRank and Anti-TrustRank, it distinguishes its own proposal from all of them on one axis:

the knowledge-based trustworthiness we propose in this paper is different from all of them in that it considers an important **endogenous signal — the correctness of the factual information provided by a web source**.

Endogenous is the right word and the whole matter is in it. Every method in that list scores a source by properties of the graph the source sits in. KBT proposed scoring it by whether what it says is true.

Two things follow that this paper takes as its starting conditions rather than as discoveries. The insufficiency of topology-only trust is settled and was settled by the field. And **the endogenous alternative was specified, published, and did not displace the proxy** — a datable fact about deployment.

2.1 Why the endogenous alternative does not escape the criterion

It is tempting to read KBT as the solution and its non-adoption as an institutional failure. The criterion this paper states does not permit that reading, and saying so is necessary before §6 makes the criterion do any work.

Evaluating the correctness of a source’s claims requires a corpus of established fact against which to evaluate them. That corpus is itself a source-admission ontology: something decided which sources constitute it, with what standing, under what prior. A source excluded from the ground-truth corpus cannot correct the ground-truth corpus, because its corrections are in material the corpus does not admit.

So the barrier recurs. It moves from a graph statistic to a curated corpus, and the curation is a decision about standing made without consulting the sources it excludes.

The corrigibility criterion does not favour endogenous trust over topological trust. It applies to both, and neither satisfies it by construction.

Two consequences worth stating plainly.

The criterion is not a design recommendation. It does not say *use content rather than topology*. It asks whether a representation can be revised by the object it misrepresents, and a content-based system can fail that test exactly as a graph-based one does — by admitting evidence only through a gate the evidence cannot reach.

And KBT’s non-adoption is no longer evidence of anything about the field’s judgment. An earlier version of this paper treated it as a suggestive fact. It is more likely a recognition that endogenous trust relocates the gatekeeping problem rather than dissolving it: from a mathematical graph to an arbiter of truth, which is the harder position to occupy and the harder one to defend. **The regress is the reason, and it is a good reason.**

What survives is narrower and is what §6 argues: the question is not which signal a system uses but whether **some** path exists by which a misrepresented object reaches the decision that misrepresents it. Topology closes that path by being content-free. A ground-truth corpus closes it by being curated. A system satisfying the criterion would need a path that is neither, and this paper does not build one — see §10.

The proxy is currently migrating rather than retreating. Recent work adapts PageRank-style trust propagation to retrieval-augmented generation — ClaimTrust (2025) modifies PageRank to propagate trust scores across documents through supporting and refuting claim relations — which imports a graph-topological standing variable into precisely the layer where compositions about sources are now produced.

3. What the epistemology of testimony already established

The second literature theorizes the same structure and has done for twenty years, in a vocabulary that has not reached retrieval.

Fricker (2007) defines **testimonial injustice** as the wrong a speaker suffers in receiving deflated credibility from a hearer owing to identity prejudice. The structural features are the ones at issue here: the credibility assessment is poorly grounded in evidence, it tracks category membership rather than epistemically relevant features of the testimony, and it is **resistant to rational revision**. The discount operates prior to and independent of any evaluation of the testimony itself.

Two developments matter especially.

The loop closes on the speaker. De Bruin (2025) names **self-fulfilling testimonial injustice**: the case in which credibility deficits *become true by shaping the very conditions that sustain them*. The prejudice is confirmed by its own operation, and the confirmation is genuine rather than illusory — which is what makes the variant structurally significant rather than merely unjust.

And the remedy is located in the hearer. Fricker’s proposal is a virtue: hearers developing the capacity to correct for the distortion in themselves. That location is the point at which the machine case departs, and §6 is about the departure.

Also relevant: Medina (2011, 2013) on credibility excess as the necessary complement of deficit, distributing standing unfairly across a field rather than merely withholding it from one speaker; and Fricker’s second category, hermeneutical injustice, concerning gaps in the shared interpretive resources by which an experience could be rendered intelligible at all.

4. What is therefore not new here

- that link-topological trust is manipulable and not Sybil-proof — Cheng and Friedman 2005; Baeza-Yates et al. 2005;
- that an endogenous, content-based trust signal is the alternative — Dong et al. 2015;
- that a credibility discount can track category rather than evidence — Fricker 2007;
- that such a discount is resistant to revision and can confirm itself — Fricker 2007 and subsequent;
- that distrust propagation is a distinct operation from trust propagation — Benczúr et al. 2005; Borgs et al. 2010.

What neither literature supplies is an account of the **damage to the representation**. Adversarial IR evaluates the proxy by ranking quality: does the spam get through, do good pages surface. Epistemology evaluates the discount by the wrong done to the speaker: is this person wronged in their capacity as a knower. Neither asks the question this paper asks, which concerns neither the ranking nor the speaker but the ontology itself, considered as a representation with a claim on accuracy.

5. Reading, operationally defined

The definition is stipulative and narrow, and everything turns on it.

Reading is permitting the claims, distinctions, evidence, corrections and internal relations of a text to revise the representation under which the text is evaluated.

Tokenization is not reading. Embedding is not reading. Extraction, classification and summarization are not reading. Each consumes the text and returns a value; none permits the text to alter the frame in which it is being valued. **A system can consume every byte of a document and read none of it**, and the sentence is precise under this definition rather than rhetorical.

The definition is not a complaint about machines. It applies identically to a human reviewer who has decided what a submission is before opening it, and it is satisfiable by a machine that routes a source's corrections back into the variables that govern its admission. What it names is a direction of permission, not a substrate.

6. The source-relative update barrier

Let $P(s)$ be a prior on source s , $A(s)$ the admission decision that prior controls, and $E(s)$ the evidence the source could supply.

Evidence $E(s)$ has operative standing with respect to admission when it can alter, block or redirect $P(s)$ or $A(s)$. Evidence may therefore be present, retrievable, and even correctly represented while lacking operative standing over the decision that governs its own admission. Presence and standing are separate properties, and the second is the one that determines whether a representation can be corrected.

The barrier obtains when $E(s)$ exists and does not reach $P(s)$ or $A(s)$: available only downstream of the admission that $P(s)$ controls.

Where this holds, the source cannot correct the prior that excludes it. Its corrections are in the material that admission gates.

This is the structure Fricker describes, with one difference that matters and one that matters more.

The first: in the testimonial case the prior attaches to a social identity, and here it attaches to a topological position. That difference is smaller than it looks — both are category memberships that stand in for an assessment nobody performed.

The second: **Fricker's remedy is unavailable to the topology-only admission operation as specified.** Testimonial injustice is correctable, on her account, through the hearer cultivating a capacity for self-correction — noticing the distortion and adjusting for it. That remedy presupposes a hearer who can hold the prejudice and the testimony in view at once and judge between them.

The operation cannot do this, and the reason is the reason it works. A topological variable does not have the testimony in view, because it is content-free by design, and content-free by design was the founding virtue.

The incapacity belongs to the operation and not to every architecture containing it. A larger system that combined topology with an endogenous content evaluation would have the faculty, in the part that performed the evaluation. Whether current systems do is §8's level-three hypothesis and is not settled here. What is settled is that the operation itself cannot supply its own correction, so any correction must come from a component the operation does not contain.

Which yields the criterion the paper turns on, and it can be asked of any representational system:

Can evidence from the misrepresented object revise the representation that misrepresented it? Where it can, the ontology can be wrong and find out. Where it cannot, the ontology can only be scored.

The distinction is between an ontology that is corrigible by its objects and one that is merely evaluable by its operators. The second can be measured, tuned, benchmarked and improved on every axis its own metrics express, and cannot discover that it is wrong about a source in a way those metrics do not capture.

7. The inverted legibility incentive

A consequence follows for anyone whose material is being admitted, and it is perverse in a specific and measurable way.

Under a topology-only regime, the operations that make a source **more semantically legible** — reciprocal linking across its own hosts so that related material is reachable, templated surfaces so that structure is consistent, mirrors so that content survives a host failure, dense internal propagation so that claims are traceable to sources — are the same operations that constitute the nepotistic-link signature the spam literature was built to detect.

Under a topology-only proxy, some operations that increase semantic legibility can simultaneously increase the features associated with manufactured standing.

The claim is existential rather than monotonic: a collision is exhibited between a specific family of legibility operations and a specific family of topology features. Whether legibility and spam score covary in general is not established and is not needed.

For the measured case the stronger sentence holds: the operations that increased legibility for a reader increased the spam score for the proxy. An audit conducted for this paper on a twenty-nine host fleet under single ownership found reciprocal cross-host density at **0.86**. A major reference platform's signature is the opposite shape: it links almost entirely to itself (0.98), almost never to another platform (0.001), and barely outward.

The finding here cuts both ways and the paper reports both. A topology vector **does** contain a dimension that separates the two cases without reading either — cross-host reciprocal density under common ownership, which is precisely the structure the link-spam literature identified and precisely what a large self-contained platform never exhibits. The proxy is not blind. It is doing exactly what it was built to do, and what it was built to do cannot distinguish a small distributed corpus that is internally consistent from a small distributed corpus that is manufacturing standing, because the distinguishing evidence is in the text.

8. Three levels, kept apart

A paper making this argument is exposed to a specific failure, which is to let a measurement of one thing pass as evidence about another. Three levels are therefore separated throughout and the separation is stated rather than assumed:

1. **What the audit measures.** The link-topology vector of a specific fleet, on a specific date.
2. **What classical trust-graph models imply.** TrustRank, Anti-TrustRank and their feature sets, 2004–2020, from published specifications.
3. **What is hypothesised of contemporary retrieval.** That a source-admission step operates on standing variables of this kind in current systems.

The third is not established by the first two, and nothing in this paper says it is.

Stated at the right strength: level three is **a diagnosis of a black box that cannot be verified from outside it**. The admission machinery of deployed systems is not published, so the hypothesis is compelling or not on the coherence of the mechanism and the behaviour observable at the surface, and it is settled by neither. The evidence that would settle it is named in §9, and it is evidence only its operators can supply.

9. Falsification

The barrier claim (§6) fails on the exhibition of a source-admission system in which corrective evidence from a low-prior source reaches the variables governing its admission without first passing the admission gate — a documented path by which a misrepresented source revises the representation that misrepresents it.

The reading definition (§5) fails as a useful distinction if a system that performs only extraction, classification and summarization can be shown to revise its evaluative frame in response to a source’s corrections. That would mean the operations named are reading in the relevant sense after all.

The inverted legibility claim (§7) fails if legibility-increasing operations can be shown to move a source *away* from the spam signature under the standard feature sets, or if a topological feature is exhibited that separates internal consistency from manufactured standing without recourse to content.

The level-three hypothesis (§8) is system-relative, and its falsification is therefore system-relative too. For any system to which it is attributed, it fails if that system’s published architecture, or a discriminating intervention on it, shows its admission decision to be independent of standing variables of the stipulated kind. It is confirmed only for a system whose architecture or controlled behaviour establishes such dependence. **One system shown not to use such a variable does not falsify the attribution to another**, and no aggregate claim about contemporary retrieval is made here that a single architecture could settle.

And the paper may not be defended by the misrepresentation of any particular source, including its author’s. A single misrepresented object is an anecdote. The claim is structural and is about what a class of representation can and cannot discover about itself.

10. What remains owed

An independent audit of the topology finding. The fleet measured is the author’s own, which makes the measurement interested in the same way §9’s last paragraph warns against. The density figure is checkable by anyone with a crawler and should be checked by someone else.

And a positive construction. This paper states a criterion and does not build a system that satisfies it — and §2.1 explains why that omission is harder to repair than it first appears. Topology fails the criterion by being content-free. A ground-truth corpus fails it by being curated. **A system that satisfied the criterion would need an admission path that a misrepresented source can reach without first passing admission**, and the shape of such a path is the open problem this paper leaves.

What can be said about it: the path would have to be **cheap enough to be always open and narrow enough not to be an attack surface**, since an admission channel that anyone can use to alter their own prior is a channel that will be used for exactly that. The tension between corrigibility and manipulability is the reason the problem is open, and it is not resolved by preferring one signal over another.

References

- Baeza-Yates, R., C. Castillo, and V. López. 2005. “PageRank Increase Under Different Collusion Topologies.” *Proceedings of the 1st International Workshop on Adversarial Information Retrieval on the Web (AIRWeb)*.
- Benczúr, A., K. Csalogány, T. Sarlós, and M. Uher. 2005. “SpamRank — Fully Automatic Link Spam Detection.” *AIRWeb*.
- Borgs, C., J. Chayes, A. T. Kalai, A. Malekian, and M. Tennenholtz. 2010. “A Novel Approach to Propagating Distrust.” In *WINE 2010*, LNCS 6484. Berlin: Springer, 87–105.
- Brin, S., and L. Page. 1998. “The Anatomy of a Large-Scale Hypertextual Web Search Engine.” *Computer Networks and ISDN Systems* 30(1–7): 107–117.

Cheng, A., and E. Friedman. 2005. "Sybilproof Reputation Mechanisms." *Proceedings of the Third Workshop on Economics of Peer-to-Peer Systems*.

"ClaimTrust: Propagation Trust Scoring for RAG Systems." 2025. arXiv:2503.10702.

De Bruin, B. 2025. "Self-fulfilling Testimonial Injustice." *Episteme*, published online 8 July 2025.

Dong, X. L., E. Gabrilovich, K. Murphy, V. Dang, W. Horn, C. Lugaresi, S. Sun, and W. Zhang. 2015. "Knowledge-Based Trust: Estimating the Trustworthiness of Web Sources." *Proceedings of the VLDB Endowment* 8(9): 938–949. arXiv:1502.03519.

Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.

Gyöngyi, Z., H. Garcia-Molina, and J. Pedersen. 2004. "Combating Web Spam with TrustRank." *Proceedings of the 30th VLDB Conference*, 576–587.

Kamvar, S. D., M. T. Schlosser, and H. Garcia-Molina. 2003. "The EigenTrust Algorithm for Reputation Management in P2P Networks." *WWW 2003*.

Kleinberg, J. 1999. "Authoritative Sources in a Hyperlinked Environment." *Journal of the ACM* 46(5): 604–632.

Medina, J. 2013. *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations*. Oxford: Oxford University Press.

Wanderer, J. 2012. "Addressing Testimonial Injustice: Being Ignored and Being Rejected." *Philosophical Quarterly* 62(246): 148–169.

The underlying notebook, the fleet topology audit and the argument-as-dataset are deposited in full in the Crimson Hexagonal Archive and openly licensed under CC BY 4.0. Deposit identifiers are available at alexanarch.org.

Suggested Citation

Sharks, Lee. “Can the Ontology Be Wrong and Find Out? Source-Admission Systems, the Proxy for Standing, and the Damage a Representation Does to Itself” *Alexanarch*, 2026-09-17. <https://www.alexanarch.org/s/records/1627/>

Deposit Information

This paper is deposit #1627 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community crimsonhexagonal, terminated 2026-06-19). The AXN identifier AXN:06BD.GENERATIVE is sovereign — minted and maintained by Alexanarch.

- Canonical URL: <https://www.alexanarch.org/s/records/1627/>
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: [alexanarch.org](https://www.alexanarch.org/)
- Machine-readable dataset entry: </api/index.json>
- Full text (HTML): <https://www.alexanarch.org/s/records/1627/>

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