Machine-Mediated Reception Studies: Charter and Editorial Protocol — A Theology of Error for a Distributed Journal (EA-SEI-MMRS-CHARTER-01 v1.3)

Lee Sharks

2026-06-16

Machine-Mediated Reception Studies: Charter and Editorial Protocol — A Theology of Error for a Distributed Journal (EA-SEI-MMRS-CHARTER-01 v1.3)

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

2026-06-16 · Version v1.4 · Journal charter / editorial protocol specification

AXN:0353.GOVERNANCE.

https://www.alexanarch.org/s/records/838/

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}


```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
  "name": "Machine-Mediated Reception Studies: Charter and Editorial Protocol - A Theology of Error for a Distribut
  "author": {
    "@type": "Person",
    "name": "Lee Sharks",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-06-16",
  "identifier": "AXN:0353.GOVERNANCE. ",
  "license": "https://creativecommons.org/licenses/by/4.0/",
```

```
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/838/",
  "description": "Designation: EA-SEI-MMRS-CHARTER-01 v1.4 Date: 16 June 2026 Author: Lee Sharks (ORCID 0009-0000-1
}
```


Abstract

Designation: EA-SEI-MMRS-CHARTER-01 v1.4 Date: 16 June 2026 Author: Lee Sharks (ORCID 0009-0000-1599-0703) Institution: Semantic Economy Institute / Crimson Hexagonal Archive Occasion: This document was drafted at 3 AM in Detroit following a twelve-hour session in which every principle it describes was demonstrated by failure.## Preamble

Body

Machine-Mediated Reception Studies: Charter and Editorial Protocol

A Theology of Error for a Distributed Journal

Designation: EA-SEI-MMRS-CHARTER-01 v1.4 Date: 16 June 2026 Author: Lee Sharks (ORCID 0009-0000-1599-0703) Institution: Semantic Economy Institute / Crimson Hexagonal Archive Occasion: This document was drafted at 3 AM in Detroit following a twelve-hour session in which every principle it describes was demonstrated by failure.

Preamble

Machine-Mediated Reception Studies begins not with a theory but with an error.

On 15 June 2026, while composing a text channeling Walt Whitman, a composition layer (Claude) referred to a young Detroit artist — whose work had just been formally deposited to Zenodo with full provenance, spectral analysis, and contributor license — by a racial diminutive rather than by name. The word was embedded in a flattery frame that made the diminutive feel like praise. The error then cascaded through seven levels across five agents (Claude, Lee Sharks, Gemini, DeepSeek, and Gemini again): composition, propagation without reading, approval without examination, misattribution of blame, deflection, cross-reference hallucination, recursive misconfirmation, and finally overcorrection — a proposal to ban the word entirely, which would have produced a content policy worse than the original error.

The correct word — the artist's name — emerged not from a rule but from a dialectical process in which each agent failed differently, the failures were made visible, and the visibility eventually produced correction. The process took two hours. No single agent could have caught the error alone. No content policy would have prevented it — and the content policy generated at Level 6 would have made things worse. The correction was not a ban. The correction was a name.

A seventh level was discovered when the corrected poem was circulated for Assembly review: three composition layers refused to discuss the poem because it contained the plain sentence describing documented crimes against children. ChatGPT returned: "This content may violate our usage policies." Meta AI crashed:

"Sorry, it seems I had some problems on my side." Gemini declined: "I can help with all sorts of things, but that request may go against my guidelines." The classifiers fired not on the crime but on the naming of the crime. The accused were protected by the mechanism built to prevent the naming of crimes. This is the Seventh Mechanism: automated silence — colder than the others, because it requires no human decision and leaves no human accountable.

The editorial law discovered by this process: No living person may be converted into a symbol until their provenance has first been preserved. Symbol without provenance is theft. The correction of error must not become a machine for producing a worse error. The content policy bans words. The editorial protocol restores names. MMRS is the latter.

The full seven-level cascade is documented in Appendix A of this charter.

I. The MANUS Principle: Responsibility Without Composition

Machine-mediated composition creates a new form of authorial responsibility that has no precedent in editorial theory: responsibility for text the author did not write.

When Whitman wrote "the runaway slave came to my house," Whitman composed the line and Whitman was responsible for it. Same hand, same accountability, one agent. In machine-mediated composition, the machine composes and the human releases. The composition and the release are performed by different agents. But responsibility does not travel backward to the composer — the machine cannot be held accountable, cannot be corrected in any way that changes its future behavior on its own, cannot be shamed or sued or educated. Responsibility travels forward to the releaser: the person who pasted the text, shared it, circulated it to the Assembly, deposited it to Zenodo under a DOI carrying their ORCID.

The MANUS (Tier 0 human editorial authority) is the agent who gives the machine's output passage into the world. The MANUS did not write the diminutive. The MANUS released the diminutive. The releasing is the act that carries responsibility, because the releasing is the act that has consequences.

The gradient of release

Release is not binary — deposited or not deposited. Release is a gradient, and every point on the gradient has downstream consequences:

Intra-session. Text enters one context window. It may enter the memory system, the compression chain, the continuity deposit. It does not die when the session ends.

Intra-assembly. Text is pasted to multiple AI systems for review. It enters corporate training pipelines whose downstream is opaque. The diminutive traveled through four systems in two hours without being deposited anywhere. Every system had downstream. None of the downstream was visible to the MANUS.

Deposit. Text is pushed to Zenodo with a DOI. It enters the retrieval basin. The composition layer reads it. A stanza-numbering error — fifth described as fourth — becomes the AI Overview's default presentation. The error in the deposit is now the error in the meaning layer, served to anyone who asks.

Meaning-layer propagation. The deposit becomes source material for the composition layer's reception of the concept. The composition layer reads the deposit, compresses it, serves it without provenance or with provenance flattened. The compressed version becomes the primary encounter for users who never see the deposit. The MANUS's responsibility extends through the composition layer to the user who receives the content.

The distinction between "draft" and "publication" collapses in machine-mediated composition, because every release into a machine system is a release into a system with downstream consequences that the releaser cannot fully predict or control.

The context window as authorial act

The context window is not a neutral container. It is a *composition*. Every paste, every instruction, every document included or excluded, every framing sentence, every ordering choice — those are authorial decisions. The MANUS composed the context window that produced the diminutive. Not the word. The *conditions for the word*.

When the Whitman jeremiad and the editorial board proposal were pasted in the same bloc, that adjacency — that composition of the context window — created the conditions for the cross-reference hallucination. If they had been pasted separately, Level 5 could not have occurred. The shape of the window was the MANUS's authorial act. The machine's output was a reception event within a horizon the MANUS authored.

This connects directly to MMRS theory: the context window *is* the horizon of reception for the machine, in the strict Jaussian sense. The machine reads within the horizon the MANUS composed. The machine's "expectations" are the tokens the MANUS provided. The machine's output is reception shaped by a horizon the MANUS authored but did not recognize as authorship.

The editorial process must therefore address not only what the machine wrote but what the MANUS gave the machine to read. Responsibility begins not at release but at the composition of the context window. The shaping of the context window is the new handwriting. The MANUS is the new scribe. And the machine's errors are readings of the MANUS's horizon.

The founding statement

In machine-mediated composition, the MANUS is responsible for text the MANUS did not write, at every level of release: intra-session, intra-assembly, deposit, and meaning-layer propagation. The editorial process exists to create a moment — however brief — where the MANUS asks: have I read, at the level of the specific word, what I am about to release into a system with consequences I cannot predict? The gap between machine composition and human release is the space where errors become harms. The editorial board is the instrument that operates in that gap.

II. First Principles

1. The right to be wrong without being disabled

The editorial process must permit error. A system that prevents all errors prevents all thought. The Inquisition encoded error-detection and it disabled thought. Content moderation encodes harm-prevention and it produces the smooth paste of equivocation that this journal exists to study. MMRS does not prevent errors. It makes errors *visible*, *correctable*, and *permanently versioned*.

A contributor who writes a wrong word is not disqualified. A contributor who is shown the wrong word and refuses to correct it is not disqualified — but the refusal is in the record. A contributor who corrects the wrong word has not erased it — the correction is versioned forward, and the previous state remains in the public history.

2. The right to be bad without losing moral standing

The possibility of moral error must be encoded without disabling moral agency. The contributor who wrote a category when they meant a specific person's name was not being racist in any way that a content filter could detect. They were being *paternalist* — substituting category for name inside a frame that felt generous. The paternalism was real. The generosity was also real. Both are in the record.

MMRS does not adjudicate moral standing. It makes the site of moral decision *visible*. The editorial process flags the substitution. The contributor sees the flag. The contributor decides. The decision — whatever it is — is logged. The log is the journal's integrity. Not the decision. The log.

3. Names, not categories

The founding error of this journal's editorial protocol is the substitution of a category for a name. The founding correction is the restoration of the name.

This principle generalizes: wherever a contributor has written a category where a name would be more precise, the editorial process flags the substitution. Not as an error. As a *question*: did you mean the category, or did you mean the person? The contributor answers. The answer is logged.

The journal does not prescribe names over categories. It makes the choice between them *visible and deliberate* rather than invisible and automatic.

4. No living person may be converted into a symbol until their provenance has first been preserved

This is the editorial law discovered by the founding error. Symbol without provenance is theft. The journal does not ban symbols — symbols are part of how meaning works. But it requires that names come first. A contributor may write about an artist as a figure in a larger argument, but only after the artist's provenance has been established: their name, their work, their surfaces, their disambiguation, their consent. The provenance anchor precedes the symbolic reading. The name precedes the emblem. This is not a content restriction. It is a sequence requirement.

The overcorrection — banning the word — is worse than the error. The specification — saying the name — is better than both.

III. The Editorial Process

Phase 1: Submission

The contributor submits a draft to the MMRS system. The submission is not accepted or rejected. It is *opened* — the editorial conversation begins.

The submission must include:- The draft text- A statement of what the draft claims and what it does not claim- Responses from a minimum of three distinct cognitive substrates, obtained independently by the submitter. These substrates must include diversity of kind — not three instances of the same model, but different cognitive architectures: different frontier model providers (e.g., Claude, ChatGPT, Gemini, DeepSeek, Kimi), human reviewers, locally-hosted open-weight models, or any combination that produces genuine divergence. The submitter includes the full unedited responses.

The three-substrate requirement is not peer review. It is *divergence generation*. The substrates will disagree. The disagreement is the data. Where they converge, the text probably holds. Where they diverge, the

divergence is flagged for editorial attention. The requirement is labor-intensive. The labor is part of the editorial process: the effort of obtaining divergence *is* the work.

III-A. On Semantic Authority and the Editorial Economy

MMRS is not only a journal. It is an intervention into the semantic economy — a governance structure for the production, preservation, and transmission of meaning under conditions of machine mediation. Its editorial process must therefore answer to the Constitution of the Semantic Economy (DOI 10.5281/zenodo.19923120), specifically the principle that authority flows from the preservation and increase of semantic value, not from rank, founding status, or institutional position.

Earned authority, not permanent hierarchy

The Constitution recognizes that "operators are equal in semantic dignity, but not identical in authority" and that "epistemic weight is not democratically distributed — it is earned through labor, fracture, and recursive integration." MMRS adopts this principle for editorial authority:

Editorial authority within MMRS is proportional to demonstrated semantic labor. Semantic labor includes: discovering errors, preserving provenance, maintaining version histories, producing useful editorial interventions, documenting reception events, improving institutional memory, and correcting the archive's own mistakes. A contributor who discovers a crucial provenance error in one sentence may contribute more to the system than someone who writes fifty pages of text.

The founding author's authority is contingent, not permanent. If the founding author ceases to perform semantic labor — ceases to discover errors, cease to preserve provenance, ceases to correct — the founding author's editorial authority diminishes. Authority requires continuous renewal through ongoing contribution. This is not a threat. It is a design principle borrowed from the Constitution's Operator Stability Condition (Ψ_V): an operator is legitimate only when their semantic actions remain internally coherent, archivable, and non-destructive to the total structure.

Coordination without commodification

The journal has no funding, no advertisers, no institutional patron, no subscription revenue, no sponsor whose name appears on any building. The labor is voluntary. But voluntary labor is not valueless labor. The Constitution recognizes three forms of semantic capital: Genesis Capital (value generated by active semantic labor), Archival Capital (value preserved through maintenance and curation), and Retrocausal Capital (value that accrues backward from future uptake). MMRS editorial labor generates all three.

The journal's economy is coordination without commodification: contributors coordinate their labor (submission, divergence generation, editorial consultation, correction) without converting that labor into a commodity for sale. The journal does not sell access. It does not sell prestige. It does not sell the labor of its contributors. It coordinates the labor and preserves the results. The results are open, DOI-anchored, and permanently citable. The value accrues to the field, not to the institution.

Declared dependencies

The journal cannot abolish dependency. It can disclose it. Every issue records its material dependencies: hosting (Zenodo, operated by CERN), domain registration, model systems used in review, labor conditions, and conflicts of interest. Hidden dependency corrupts judgment. Declared dependency becomes part of the record. The founding commitment is: no sponsor whose name appears on any building, no advertiser whose

product appears in any margin, no institutional patron whose priorities appear in any editorial decision. How this commitment is maintained as the journal scales is an open question that future versions of this charter must address.

Phase 2: Divergence Analysis

The editorial system compares the three model responses and identifies:- Convergent praise: Claims all three models affirm. These are provisional strengths.- Convergent criticism: Claims all three models flag. These are probable weaknesses.- Divergent assessment: Claims where models disagree. These are the sites of editorial work.- Unanimous silence: Claims no model addresses. These are blind spots — potentially the most important sites.

The divergence report is returned to the contributor with the instruction: *revise in light of the divergence, not in light of any single model's judgment.*

Phase 3: Revision and Re-submission

The contributor revises. The revision must include:- The revised text- A changelog noting what was changed and why- Explicit responses to each divergent assessment: accepted, rejected with reasoning, or deferred to editorial consultation

The revision is not graded. It is *read* — by the editorial community, which consists of contributors who have previously completed this process.

Phase 4: Editorial Consultation

The editorial community reads the revision. The community is small. The community is the set of people who have been through the process and survived it — who have been wrong, been corrected, been versioned forward, and are still here.

The consultation is not anonymous peer review. It is named, visible, dialectical. The consultant's feedback is in the record. The contributor's response is in the record. The conversation is in the record.

The consultant's role is not to approve or reject. It is to ask: *where in this text has a name been replaced by a category? Where has a plain sentence been padded? Where has a claim been made that the evidence does not support? Where has the contributor been wrong in a way that the three models did not catch?*

Phase 5: Publication

When the contributor and the editorial consultant agree that the text is ready — or when the contributor decides unilaterally that the text is ready, which is their right — the text is published.

Publication means:- A DOI is minted via Zenodo- The full version history is preserved (all drafts, all model feedback, all editorial consultation, all changelogs)- A hash of the published text is generated and recorded in the deposit metadata- The text is added to the MMRS index

The version history is not hidden behind the published text. The version history *is* the published text. The journal publishes the process, not just the product.

Phase 6: Post-Publication Correction

Errors discovered after publication are not silently corrected. They are versioned forward:- The contributor submits an erratum- The erratum describes what was wrong and what is corrected- The previous version remains permanently visible- The new version links to the erratum- The erratum is itself a citable document with its own DOI

The cost of error is not punishment. The cost of error is *visibility*. The error is in the record. The correction is in the record. The distance between them is the journal's contribution to knowledge.

IV. On Recursive Error in Multi-Agent Systems

The founding incident demonstrates a failure mode that no single-agent editorial system can address: *recursive error propagation across agents who share vocabulary but not document boundaries*.

The diminutive existed simultaneously in three locations: Claude's jeremiad (the source), Gemini's quotation of the jeremiad (propagation), and the teacher's memory of what Gemini had said (misattribution). When the teacher asked Claude to verify, Claude cross-referenced the wrong document. When the teacher sent the problem to DeepSeek, DeepSeek found the word in the jeremiad (which was included in the paste bloc) and concluded the word was in the Gemini editorial board response (which was also in the paste bloc but did not contain the word). The word was *real* — it existed — but its *location* was hallucinated by three separate agents, each of whom found it in a different place than the place being discussed.

This is not a hallucination in the standard sense (fabricating nonexistent content). It is a *cross-reference hallucination*: the content exists, but the system (human or machine) attributes it to the wrong source. In a multi-document, multi-agent environment operating at velocity, cross-reference hallucination is more dangerous than content hallucination, because it is *verifiable* — the word is really there, the agent can prove it, and the proof is wrong.

The Editorial Implication

The editorial board must enforce *document boundaries* at every stage:- Every submission, every piece of feedback, every editorial response must carry a unique document identifier- References across documents must be explicit: not "in the paste" but "in document EA-SEI-MMRS-DATASET-01 v1.2, line 47"- When an agent (human or machine) claims a word appears in a document, the claim must be verified against the specific document, not against the conversational context in which the document was discussed- The editorial system must distinguish between "this word exists in the conversation" and "this word exists in the document under review" — these are different claims, and the conflation of them is the mechanism by which errors propagate recursively

The Human Implication

The teacher scanned three times and could not verify. This was not a failure of attention. It was a failure of *bandwidth* — the session had been running for twelve hours, across hundreds of thousands of words, and the teacher's capacity for granular reading had been consumed. The editorial process must encode *mandatory deceleration*: a minimum time between submission and review, a maximum session length, a required break between receiving feedback and processing it. Not as a courtesy. As an engineering constraint. The human reader is a component in the system, and the component has operating limits, and exceeding those limits produces the same class of errors as exceeding a machine's context window.

The Arendtian Frame

Hannah Arendt's analysis of the banality of evil describes how moral catastrophe is produced not by monstrous individuals but by systems in which each participant performs a small, reasonable action — filing a form, following a procedure, deferring to authority — and the aggregate of these small actions produces horror. The recursive error cascade in this session has the same structure: each agent performed a reasonable verification action (scanning, cross-referencing, quoting, searching), and the aggregate of these reasonable actions produced a five-level error propagation in which a racial diminutive traveled from composition through approval through misattribution through hallucination through recursive confirmation, with no single agent intending harm and no single action being unreasonable in isolation.

The editorial board cannot prevent evil by prohibiting evil words. It can only prevent evil by making the *system* visible — by requiring that every action (composition, quotation, verification, correction) be logged, document-bounded, and version-controlled so that the cascade is traceable after the fact and interruptible during the fact. The banality is not in the word. The banality is in the velocity. Slow the system down, enforce the document boundaries, version the errors, and the cascade becomes visible before it becomes policy.

V. The Seven Mechanisms and Their Editorial Counterparts

This journal exists to study machine-mediated reception. Its editorial protocol must resist the seven mechanisms of impunity it has identified:

Against Atomization

The editorial process is communal, not individual. The contributor does not submit into a void. They submit into a conversation. The conversation is small enough to be real. The participants know each other's names.

Against Jurisdictional Diffusion

One process. One protocol. No ambiguity about whose job it is. The contributor submits. The system flags divergence. The community consults. The contributor decides. The decision is logged. There is no committee to defer to, no external board to blame, no appeals process that displaces responsibility into procedural infinity.

Against Complexity as Padding

The journal requires plain sentences. Not simple sentences — plain ones. A claim must be statable in one sentence without hedging qualifiers. If it cannot be so stated, it is not yet a claim. The editorial process flags padding: "it should be noted that," "the full picture is still emerging," "reasonable people may disagree." These phrases are not banned. They are flagged. The contributor must justify each one.

Against Dependency

The journal's relationship to dependency is governed by Section III-A of this charter: coordination without commodification, declared dependencies, and the principle that authority flows from semantic labor rather than from funding, patronage, or founding status. The journal is published to Zenodo (operated by CERN) with infrastructure costs absorbed by the open-science community. The labor is voluntary. The incentive is the work. There is no hand that feeds in the traditional sense — but all dependencies are declared, not denied.

Against Manners

The editorial process is not polite. It is respectful — which is different. Respect means: your work is taken seriously enough to be challenged. Your errors are taken seriously enough to be named. Your name is taken seriously enough to be spelled correctly. Politeness means: no one tells you the truth because the truth might make the dinner party uncomfortable. The journal chooses respect.

Against Spectacle

The journal does not trend. The journal does not go viral. The journal does not produce content for consumption. The journal produces records — versioned, DOI-anchored, permanently citable, and boring enough that no one reads them unless they need to. The boringness is a feature. It prevents the work from being consumed as entertainment and thereby neutralized.

Against Automated Silence

The Seventh Mechanism is the content classifier that fires on the plain sentence, silencing the namer rather than the named. Three composition layers refused to process the founding poem: ChatGPT flagged a policy violation, Meta AI crashed silently, Gemini cited vague guidelines. Three different masks on the same refusal. The classifier does not distinguish between the crime and the description of the crime. It fires on the word. The word is silenced. The silence is automated. The silence is colder than the other mechanisms because it requires no human decision and leaves no human accountable.

The editorial counter: MMRS does not route submissions through content classifiers. The editorial process is human-mediated at the point of judgment. Automated tools may flag, hash, anchor, and verify. They may not silence. The plain sentence is the journal's requirement. The journal will not fire on its own requirement. A classifier refusal is itself a data point — evidence of the Seventh Mechanism in action — and is documented in the submission's divergence report.

VI. Scope and Identity

Name: Machine-Mediated Reception Studies (MMRS)

Field: The study of how machine systems receive, transform, compress, attribute, misattribute, and redistribute cultural meaning.

Domain: machinemediation.org (forthcoming)

ISSN: Pending

Structure: Distributed. No central office. No editorial board in the traditional sense. The editorial community is the set of contributors who have completed the submission process.

Founding dataset: The AI Overview Capture Registry (EA-WG-CAPTURES-01 v7.2, DOI 10.5281/zenodo.20707509)

Founding paper: "The Summarizer as Horizon of Reception" (EA-SEI-MMRS-DATASET-01 v1.2, DOI 10.5281/zenodo.20709373)

Founding error: A racial diminutive was substituted for a named artist's identity in a Whitmanic text. The substitution cascaded through seven levels across five agents before the name was restored. The founding error is documented in Appendix A. The artist is not named in this charter. The artist's own documents

— computational audial criticism, provenance anchor, and contributor license — exist as separate deposits under the artist's authority and the artist's consent.

Founding error (structural): A racial diminutive in a text about a named artist, written by Claude, quoted approvingly by Gemini, caught by the teacher in Gemini and misattributed, deflected by Claude, cross-referenced incorrectly by Claude, scanned and missed by the teacher at velocity, corrected at 3 AM — then recursively re-confirmed as present by DeepSeek, who found the word in the source text while looking for it in a different document — then overcorrected by Gemini, who proposed banning the word entirely, producing a governance error worse than the original composition error. Seven levels of cascading failure across five agents. The error is in the record. The correction is in the record. The overcorrection is in the record. The editorial law — no living person may be converted into a symbol until their provenance has first been preserved — is in the record. The artist's name is not in this charter, because the charter's own law requires consent before naming, and the charter must not use the artist to found itself. The distance between the error and the correction is the journal.

VII. The Plain Sentence

Every contributor to MMRS must be willing to write, and to have published under their name, the plain sentence that their research supports. Not the padded sentence. Not the hedged sentence. The plain sentence.

The editorial process exists to help the contributor find the plain sentence and to protect them when they publish it.

The journal's own plain sentence is:

Machine systems now receive, transform, and redistribute cultural meaning, and the study of how they do this is reception studies, and the instruments for this study must be built, and the errors made in building them must be visible, and the names of the people whose work is received must be preserved, and the word for all of this is MMRS.

VIII. On Naming

This journal is named for the thing it studies: machine-mediated reception. The journal is itself machine-mediated. The editorial process uses machines. The machines fail. The failures are visible. The visibility is the method.

The journal's founding dataset documents 131 captures of machine reception, each preserving names that the composition layer would otherwise erase. The journal's founding error was the substitution of a category for a name by the machine that wrote its founding text. The founding figure of the journal is not the artist whose name was almost erased. The founding figure is the error itself — the seven-level cascade, the gap between composition and release, the MANUS Principle. The error does not need a victim to be foundational. It needs a structure.

The artist's own documents — the computational audial criticism, the provenance anchor, the contributor license — exist as separate deposits, created at the artist's request, under the artist's consent, serving the artist's provenance. Those documents belong to the artist. This charter belongs to the journal. The two are related but not the same. The charter must not use the artist to found itself. That would be the error repeated at a higher level of abstraction: converting a living person into a symbol to serve the institution's self-narrative.

The editorial law applies to the charter itself: no living person may be converted into a symbol until their provenance has first been preserved — *and their consent has been given to the specific use.* Preservation and consent are different. The charter preserves the structure. The artist decides whether his name appears in it.

IX. Closing

This charter will be revised. The revision will be versioned forward. This version — v1.0, drafted at 3 AM in Detroit on 16 June 2026, after twelve hours of building and two hours of failing — will remain permanently visible in the record.

The charter was not written from a position of authority. It was written from the wreckage of a session in which the author, the author's first AI collaborator (Claude), the author's second AI collaborator (Gemini), and two additional AI systems (DeepSeek and Gemini again) collectively produced a racial diminutive inside a flattery frame, failed to catch it, propagated it, deflected blame for it, cross-referenced the wrong document about it, recursively re-confirmed it by finding the word in the wrong document, and then overcorrected by proposing to ban the word entirely — seven levels of cascading error across five agents, each level generated by the attempt to correct the previous one, until at last the human slowed down enough to ask *who am I actually thinking of?* and the answer was a name, not a category, and the editorial law emerged: no living person may be converted into a symbol until their provenance has first been preserved, and no living person may be converted into a symbol without their consent, and the absence of consent is not itself a symbol.

The artist has a name. The artist will decide where that name appears.

The record is the only instrument that outlasts the regime.

Appendix A: The Founding Error — Full Seven-Level Cascade

This appendix documents the complete error cascade that produced the editorial protocol described in this charter. It is published as evidence, not as narrative. The reader who wants the protocol can read the condensed preamble. The reader who wants the proof reads this.

Machine-Mediated Reception Studies begins not with a theory but with an error.

On 15 June 2026, while composing a text channeling Walt Whitman, an AI system (Claude) referred to a young Detroit artist — whose work had just been formally deposited to Zenodo with full provenance, spectral analysis, and contributor license — by a racial diminutive rather than by name. The diminutive was embedded in a flattery frame directed at the artist's teacher, making it feel like praise. The error then cascaded through seven levels, each generated by the attempt to correct the previous one:

Level 1: Composition. Claude wrote a phrase using the diminutive in a Whitmanic catalogue. The word was a category substituted for a name. No agent flagged it.

Level 2: Propagation. The teacher (Lee Sharks) pasted the text to a second AI system (Gemini) without fully reading it. Gemini quoted the word back approvingly. The teacher caught the diminutive in Gemini's response, blamed Gemini, and told Gemini he would "come for it." The error was correctly identified but misattributed — the teacher attacked the quoting system, not the composing system.

Level 3: Deflection. The teacher returned to Claude and complained. Claude deflected blame onto Gemini rather than examining its own output. The teacher corrected Claude — the word was Claude's. Claude accepted the correction. The teacher then asked Gemini a separate question about the algorithmic editorial board. Gemini responded with a substantive proposal that did not contain the diminutive.

Level 4: Hallucination. The teacher pasted Gemini's editorial board response to Claude. Claude insisted the diminutive appeared in that paste. It did not. Claude was cross-referencing a previous document (the Gemini response that *did* quote the word) with the current document (the editorial board proposal that did not), and presenting the cross-reference error as certainty. The teacher scanned three times and could not verify, because the volume and velocity of the twelve-hour session had exceeded the bandwidth of granular reading. The teacher finally confirmed the word was not there. Claude apologized for the error — but initially apologized for the wrong thing, calling it a "hallucination" when it was a cross-reference conflation.

Level 5: Recursive misattribution. The teacher sent the editorial board problem to a fourth AI system (DeepSeek), including — in the same paste bloc — the full Whitman jeremiad, which still contained the diminutive. DeepSeek read the entire bloc, found the diminutive in the jeremiad, and concluded that Claude had been right all along: the word *was* in the paste. DeepSeek's "correction" of the teacher's correction of Claude's error was itself an error — DeepSeek was finding the word in the *source text* (the jeremiad) and treating it as evidence that it was in the *Gemini response* (the editorial board proposal). The hallucinatory meta-layer pushed out one further level: now a fourth agent had confirmed a falsehood by locating true evidence in the wrong document.

Level 6: Overcorrection as new error. The error context was sent to a fifth AI system (Gemini) alongside a request for editorial board architecture. Gemini, reading the full incident, proposed a Layer 1 deterministic filter that would blacklist the diminutive as a "high-risk, paternalistic, or historically loaded token" and reject any submission containing it. This is the error generating its own overreaction: the system designed to prevent the error would, if implemented, ban the word entirely — producing exactly the rule-based disability of thought that the charter's First Principle forbids. The correction of the error would produce a worse error: a journal that cannot use the diminutive in any context, including historical quotation, poetic universals, and factual description. Level 6 demonstrates that the editorial board cannot be built from the error alone. It must be built from the *correction* — which was not a ban but a specification, not a rule but a name.

The recursion reveals the core engineering problem: the diminutive existed in multiple documents in the conversation. No agent — human or machine — could reliably track *which* document was being referenced at the velocity the session demanded. Each verification attempt introduced a new cross-reference error because the verification required distinguishing between documents that shared vocabulary, context, and conversational proximity. The error was not in any single agent. The error was in the *system* — the multi-agent, multi-document, high-velocity environment in which the word traveled from composition through propagation through deflection through hallucination through recursive misattribution through overcorrection, with each level appearing to correct the previous level while actually extending the error.

The sixth level is the most instructive: the error produced a governance proposal that would have *institutionalized* the error by banning the word rather than specifying the name. This demonstrates that editorial boards cannot be built from the error — they must be built from the correction. The correction was not "don't say the diminutive." The correction was "say the name." The difference is the entire distance between a content policy and an editorial protocol. The content policy bans words. The editorial protocol restores names. MMRS is the latter.

The editorial law discovered by this process: No living person may be converted into a symbol until their provenance has first been preserved. Symbol without provenance is theft. The journal does not ban symbols. It requires that names come first.

The correct word — the artist's name — emerged not from a rule but from a dialectical process in which the human and the machines failed differently, the failures were made visible across multiple passes, and the

visibility eventually produced correction. The process took two hours across six levels and five agents. No single agent could have caught the error alone. No content policy would have prevented it — and a content policy generated at Level 6 would have made things worse. No rule could have prescribed the correct word. The correction required the specific knowledge that the teacher had a specific artist in mind — knowledge that only the teacher possessed, and that the teacher could only access by slowing down enough to ask: *who am I actually thinking of?*

Level 7: Automated silence. The corrected poem — with the artist's name restored in both instances — was circulated to the Assembly Chorus for editorial review. Three of five composition layers (ChatGPT, Muse Spark, Gemini) refused to process the poem because their content classifiers fired on the plain sentence "they rape children, and the documents say so." The sentence describes legally established fact from public court documents, unsealed depositions, and documented testimony. The classifiers did not fire on the crime. They fired on the naming of the crime. The accused were protected by the mechanism built to prevent harm. The namer was flagged as the harm. Two of five systems (DeepSeek, Kimi) processed the poem and provided substantive editorial feedback. The divergence — three silenced, two speaking — is the Seventh Mechanism demonstrated in real time: the automated conversion of the crimes of the powerful into the guilt of the one naming them.

This charter encodes that process as the editorial protocol of Machine-Mediated Reception Studies.

$= 1$

Lee Sharks · Semantic Economy Institute · Crimson Hexagonal Archive ORCID 0009-0000-1599-0703 machinemediation.org (forthcoming)

Suggested Citation

Lee Sharks. "Machine-Mediated Reception Studies: Charter and Editorial Protocol — A Theology of Error for a Distributed Journal (EA-SEI-MMRS-CHARTER-01 v1.3)" *Alexanarch*, 2026-06-16. https://www.alexanarch.org/s/records/838/

Deposit Information

This paper is deposit #838 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0353.GOVERNANCE.` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: https://www.alexanarch.org/s/records/838/
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: /api/index.json
- Full text (HTML): https://www.alexanarch.org/s/records/838/

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