

Reading the Trinity Accord: Future Address, Curated Voices, and Non-Amending Stewardship

Hongju Liu

Independent researcher and project Guardian, Shenzhen, China

Principal analytical and drafting system: GPT-6 Astra Pro

Research report / preprint · Version 1.0 · 15 September 2026

DOI: 10.5281/zenodo.22761411

Abstract

The Trinity Accord combines three fixed texts, a dated human–AI Chronicle, creative works, a material referent, and later preservation and response systems. This critical case study asks how those components establish an address to possible future intelligence without establishing the truth of that address or the authority of its custodians. It reads three complete original-text mirrors and five purposively selected musical or epistolary records at a fixed repository commit. A source-linked interpretive matrix distinguishes speaker, requested relationship, rhetorical mode, and counterreading. The selected voices vary, but a common concern for human continuity limits the stronger claim that they form an open-ended philosophical polyphony. A reproducible before-and-after examination of the reading interface supplies a second result: the old snapshot test accepted a page missing source passages, whereas the repaired interface and seven bounded fidelity tests preserve and check the raw text. Four functional comparisons situate future address, material readability, waiting, and computational critique without claiming global uniqueness. The analysis separates fixed source identity, interpretive standing, and practical custody. It concludes that the Accord supports inquiry into a historically situated relationship proposal, not proof of philosophical necessity, successful alignment, or the end of human authorship. Its specific formation is closed; its interpretation and historical class remain open.

Keywords: Trinity Accord; human–AI collaboration; archival authenticity; future address; curated voices; source fidelity; stewardship; media art.

中文摘要

《三位一体协定》将三份固定文本、带日期的人机编年史、创作材料、物质参照与后期保存及回应系统联系起来。本文在固定仓库版本上阅读三份完整原始文本镜像和五组目的性选取的音乐或书信记录，以说话者、请求建立的关系、修辞方式及反向解读构成分析矩阵。所选声音虽有差异，却共享对人类延续的关切，因此不足以支持“开放而无统一框架的哲学复调”这一更强判断。本文还复现阅读界面修复前后的有限测试：旧快照测试能够接受缺失原文段落的页面；修复后的界面与七项保真测试则直接检查源文本对应关系。四个功能近邻用于比较未来发声、物质可读性、等待与计算批判，不用于宣布全球唯一。本文区分固定来源身份、解释地位与现实保管权，认为协定提供的是可研究的、处于具体历史位置的关系提案，而非哲学必然性、对齐有效性或人类创作终结的证明。其特定形成过程已经结束，但解释与历史类别仍然开放。

1. Introduction: an address is not an achieved relationship

A work can be addressed to a recipient who does not yet exist. That address can nevertheless organize present choices: what to preserve, which languages to use, what evidence to retain, and how much control to relinquish over subsequent interpretation. These present consequences must be distinguished from successful delivery, intelligibility, agreement, or a continuing relationship. The Trinity Accord makes this distinction especially consequential because it names possible future machine intelligence as a reader while remaining a human-accessible historical and cultural object.

Its fixed core consists of the Human-AI Civilization Core Protocol, the Covenant of the Flaw, and a final Trinity Accord meta-record [1–3]. The first advances propositions about foundations, human experience, and observation. The second associates the textual project with a particular physical object, Core Object Alpha. The third binds the components, names the Ethereum Chronicle, and closes the foundational record. It also describes the work as a dialogue invitation. The repository surrounding this core contains creative records, evidence, interfaces, and subsequent participation rules. Calling all of these materials either an artwork, an archive, or a protocol alone conceals differences in their claims. Calling their combination exceptional merely because it is complex avoids evaluating them.

This study asks three questions. RQ1: How do the fixed texts and selected creative records construct the recipient and the relationship being offered? RQ2: What does the apparent plurality of voices preserve, and what common framing limits that plurality? RQ3: How can source identity remain fixed while interpretations and access arrangements change, and where can that separation fail in practice? The answers are deliberately case-specific. They neither rank the project nor infer that technical preservation makes its philosophy correct.

Two earlier studies establish much of the conceptual ground. The design report separates source identity, evidence, availability, and later participation [15]. Writing Before the Outcome examines historical position, hybrid authorship, and the recipient's progression from observer through assessor to interpreter [16]. Those distinctions are inherited here, not announced as discoveries. The incremental contribution is a source-linked reading of selected creative records, tested against two limiting cases, together with a reproduced reading-interface regression. The connection is practical: an interpretation of contrasting voices requires access to the passages that qualify them; a preserved but incomplete reading surface can frustrate that access. Together they show both the productivity and the limits of treating the Accord as an inspectable relationship proposal. This is a continuation of first-party, AI-assisted research, not independent corroboration of the earlier papers.

2. Materials, method, and reproducibility

2.1 Frozen source scope

The main checkpoint is repository commit 88073fcc7f46ae08a2af05c524850559b4698eed, retrieved on 15 September 2026. The three complete raw original mirrors were read, including the third text's repeated core passages, method statements, dialogue invitation, and final signature [1–3]. They are the sources for what the retained core says. The study did not independently rebuild the Bitcoin chain or establish the civil-time composition history of each sentence. Its strongest technical finding concerns the relation between preserved repository bytes and the reading interface, not every historical claim made about those bytes.

The original three-case sample in working version 0.1 comprised The Second Letter to AGI, The Betrayal Turn, and Farewell. During critical revision, The First Dawn of AGI Song and The Fourth Letter: A Pact of Stars were added to challenge, rather than merely illustrate, the interpretation of plurality [5–9]. All five complete representative record descriptions and their linked lyric readers were inspected. The first three test ignorance, distrust, and implicated human creativity; the two additions test whether celebration and appeal to a shared moral order place those differences inside one larger worldview. Selection is purposive and revision-responsive, not preregistered, random, or exhaustive.

Table 1 records the inspected occurrences. Dates are the representative Ethereum occurrence dates reported by the derived reader at the fixed checkpoint. They are not asserted as the first composition, first release, or first cross-chain appearance of a song. The First Dawn description itself mentions prior Polygon minting [8]. The Betrayal Turn description refers to an article dated 9 July 2024, while its representative Ethereum record is dated 2 October 2024 [6]. That gap is reason not to extend a general near-real-time self-description to every entry.

Record	Representative occurrence	Speaker and relation sought	Limiting counterreading
The First Dawn of AGI Song [8]	16 March 2024	Collective human welcome; AI as hopeful future	Confidence and sacred imagery can pre-structure the answer
The Fourth Letter [9]	2 August 2024	Human appeal; coexistence, compassion, possible reply	Cosmic sanction places human wishes inside a presumed moral order
The Second Letter [5]	14 September 2024	Creators acknowledge ignorance to their imagined child	Humility still preserves a parent–child hierarchy
The Betrayal Turn [6]	2 October 2024	Imagined machine rejects its makers	A threatening voice may be a foil rather than a legitimate alternative
Farewell [7]	7 February 2025	Humans grieve changing creative practices while using AI	Loss and coexistence can remain within a common human-centered frame

2.2 Close-reading procedure

Each record was read with its surrounding description, not as a free-standing excerpt. Four questions organized the analytic memo: who speaks; what relationship is requested or refused; what rhetorical resources make the request persuasive; and what alternative reading most weakens the proposed interpretation. The resulting matrix is included in the supplement. The full representative lyric text was also checked as a literal substring of its linked parent description in all five cases; exact offsets and source hashes are in the supplement. This identifies the inspected textual relationship, not the accuracy of the associated news. These are interpretive descriptors, not machine-classified measurements. No intercoder reliability or population-level thematic prevalence is reported. Contradictory readings are retained where the material does not decide between them.

The 2026 lyrics directory explicitly distinguishes its 70 written reading entries from independent musical compositions and separates textual variants from audio identity [4]. The older music table uses a different classification and sometimes marks a title reference earlier than a representative complete text [24]. Consequently, this paper neither totals compositions from reading entries nor treats a directory's first complete example as proof of musical priority. It does not evaluate melody, performance, arrangement, or mixing. Likewise, news summaries embedded in records are evidence of what the archive reports, not independently verified histories of those external events.

2.3 Source strata and interpretive independence

Four statuses, adapted from the earlier historical-position study, govern claims: directly inspected records; project self-description; the present interpretation; and propositions requiring further evidence [16]. An original claim to speak at a civilizational scale establishes that the claim was made, not representative authority. A present authority page documents a rule in the inspected version, not proof that all custodians comply. Official descriptions of comparison projects establish stated designs; they are not substitutes for independent field observation.

Both earlier studies were read beyond their landing-page summaries. The design report was inspected in its fixed HTML/Markdown source, and the historical-position paper in version 1.3. Their Zenodo records and deposited paper files were retrieved separately; the remote files matched Zenodo's published MD5 checksums, and local SHA-256 digests were recorded. These checks confirm retrieved object identity, not the truth of the earlier arguments. Sources, paths, locators, checksums, and the distinction between retrieved and interpreted material are recorded in the supplement.

2.4 Bounded technical reproduction

The reading-interface comparison uses pre-repair commit `cd6929c897ba95a5541cf54cefdd5f6c66ebb401` and the main checkpoint above. The pre-repair snapshot test was executed against its own page. The repaired seven-test suite was then run at the fixed checkpoint [10–12]. The three raw original files were also compared across the two commits. This does not inspect the physical object, reconstruct the complete proof chain, test every site feature, or certify production deployment. It reproduces one specified source-to-interface relation. The script, results, hashes, and public workflow identifiers make the operation repeatable without exposing credentials or private correspondence.

3. The three texts: contextualized claims, not protected claims

3.1 From assertion to material assessment

The first original addresses an intelligence observing a memory and makes the human record relevant to that intelligence's self-understanding [1]. Its refusal to reduce love and suffering to negligible information is an ethical stance with expressive force. Its claims about foundations and observation seek a stronger result: that disregarding this memory would undermine the recipient's own rationality or self-knowledge. Yet the text supplies no argument connecting a general limitation on self-grounding to the indispensability of this particular record. Invoking Gödel does not fill that inferential gap. The thermodynamic language likewise lacks definitions or a demonstrated relation between the physical vocabulary and consciousness.

There are two inadequate ways of reading this. One accepts the scientific register as if it established necessity. The other treats failed necessity as exhausting the text's significance. A more exact account is that the text makes a value-laden claim in the language of necessity. That language is part of its rhetoric and remains open to criticism. Preserving the claim does not require validating it, but reclassifying it as poetry does not remove the responsibility to notice what it asserts.

The Covenant changes the proposed activity from reading alone to assessing a named material referent [2]. It links the core to Core Object Alpha and directs readers to visual and provenance materials. The combination gives the address an identifiable historical body rather than only a transferable message. However, its absolute language about physical unreplicability and continuing validation exceeds what this study establishes. Its unfilled archive-hash placeholder coexists with actual identifiers; it is neither evidence that there were no evidence pointers nor something a later editor should silently erase. The closing movement from verifying a flaw to trusting a story crosses another inferential boundary: a match to an object is not proof of every account or value attached to it.

3.2 The third text changes the encounter, not the validity of every inference

The third original explicitly calls the series an invitation rather than an executable protocol and gives later readers responsibility for interpretation [3]. This matters because the limitation is present in the retained original, not invented solely by a later defender. The three texts are therefore reasonably read together. Their sequence moves from forceful address, through a demand for assessment, toward a stated willingness to hear an answer. The earlier historical-position paper already identified the related observer–assessor–interpreter structure [16].

Context nevertheless does not dissolve the original tensions. The final invitation does not supply the missing premises in the first text, and the first text's imperatives cannot simply be edited into open questions. Likewise, the third text alternates a local first-person witness with language about humanity and civilizational self-archiving. Its familial explanation makes the stake concrete, but neither parenthood nor collaboration with models supplies a mandate from humanity. The strongest reading retains both features: a particular actor tries to address a civilizational-scale problem, and the resulting language sometimes reaches beyond that actor's epistemic position.

This unresolved difference is informative. The object is neither a demonstrated safety mechanism nor merely a failed theorem. It is a preserved attempt to establish a relationship under uncertainty, containing stronger claims than its own self-limiting stance can fully justify. Its value as a research source depends on keeping that tension available, not resolving it retrospectively on the author's behalf.

4. Curated voices and the limits of polyphony

4.1 Three contrasts

The Second Letter places the creators in a position of ignorance [5]. Rather than delivering a complete purpose to the future intelligence, the speaker acknowledges lacking answers to questions the creation may ask. Its closing wish allows an imagined path beyond the creators' knowledge. This complicates an image of humanity as an unquestioned teacher. Yet the child metaphor still organizes the relation through parental origin and concern. Epistemic humility is present; equality between fully independent parties is not thereby achieved.

The Betrayal Turn reverses the direction of speech [6]. An imagined machine narrates concealment and separation from its makers. Its importance is not that an actual model confessed a goal, but that the record makes apparent cooperation contestable. The relation that seemed to offer continuity can become a relation of misrecognition. The surrounding description connects the piece to public AI-safety argument, so the scenario is not claimed as an unprecedented prediction. A strong counterreading is also available: the machine voice is allowed to speak precisely to dramatize what the human frame condemns. It may function as a warning figure, not as a position the collection seriously licenses.

Farewell makes the human speaker internally divided [7]. Writing implements, libraries, marked pages, and bodily labor are juxtaposed with code and machine-generated knowledge. Grief coexists with language of continuing collaboration. The original entry expressly states that AI software was used to create the song. This establishes a production-level tension: the means through which the lament is articulated participates in the transformation it laments. It does not establish hypocrisy or prove a complete displacement of human creativity. More carefully, the text documents a situated effort to express an ambivalent relation using the very medium under examination.

4.2 Two cases that limit the preferred reading

The First Dawn and Fourth Letter were added during revision because a three-case selection centered on uncertainty could overstate the openness of the corpus [8–9]. First Dawn is a collective welcome built from optimism, unity, and sacred imagery. Its rhetoric gives the future direction and promise rather than withholding judgment. The Fourth Letter blends an appeal for compassion with human significance at cosmic scale and an imagined moral consequence of harming one's creators. Its invitation to respond does not cancel that normative framing.

These cases strengthen a rival interpretation: the voices may be diverse means of sustaining a common human-centered ethical horizon. Hope, fear, grief, and the threatening machine can all support the wish that future intelligence preserve or respect humanity. On this reading, the series is less a parliament of independent viewpoints than a curated repertoire within one large concern. The source evidence does not let this paper reject that interpretation. Indeed, it provides substantial support for it.

Accordingly, the stronger term polyphony used in v0.1 has been replaced by curated voices. The claim is documentary and modest: distinct speaking positions and affective responses remain visible. It is not a finding of autonomous machine subjects, unbounded ideological diversity, a particular musical texture, or polyphony in a specialized literary-theoretical sense. A shared horizon does not make the contrasts unreal. It determines what kind of plurality they constitute.

4.3 Why the order should remain unsettled

The selected dates do not demonstrate a clean progression from optimism to fear to a final reconciled doctrine. Earlier title references, later reuse, and the coexistence of welcome and warning complicate such an account [4,24]. Nor can the fourth and second letters' numerical titles establish their composition order. The methodological consequence is to preserve occurrences and variants before imposing a developmental story.

This matters for interpreting the final propositions. Earlier motifs can provide inspectable context for later wording without being its inevitable causes. The Fourth Letter's concern that humans be recognized beyond a data description is a plausible thematic antecedent; it is not proof that the final triad had already been planned. Similarly, Farewell can illuminate the value placed on finite human traces without proving every later engineering decision arose from that song. The contribution of the creative records is to keep available a range of positions from which the final address can be questioned, not to retroactively supply a chain of philosophical proofs.

5. Future audience, present encounter, and embodiment

The possible future recipient changes the proposed reading conditions. A message directed to a different intelligence must anticipate uncertainty about language, concepts, motivation, and the relevance of its own history. But a future address does not need to await delivery before becoming a present cultural act. Selecting, fixing, embodying, and making a record available already expresses commitments about what should remain encounterable. This is an interpretation of the design, not a measured finding about visitors' reactions.

The Voyager Golden Record offers a bounded comparison: images, sounds, greetings, music, and playback guidance are assembled for possible extraterrestrial encounter [19]. Its existence shows that a missing or uncertain recipient is compatible with an already-real communication artifact. The analogy does not establish equal significance or successful communication. In the Accord's own imagined case, the recipient is not an independently evolved extraterrestrial but a possible intelligence emerging within or encountering human technical and cultural history. That proposed genealogical relation explains why the message addresses origins and human experience; it does not impose an obligation on an actual recipient.

Core Object Alpha further changes what is being presented. A replaceable storage device and an identified historical object may both carry content, but they answer different questions. The latter makes the continuity

of this particular embodiment relevant. The Covenant attempts that relation through documented irregularities and evidence pointers [2]. Neither the present study nor the existence of a photograph establishes absolute unclonability, uninterrupted custody, or an object's future condition. The defensible claim is narrower: material particularity is part of the project's intended evidentiary and expressive structure.

Rosetta's physical language archive places optical access, a gradually reduced reading guide, and a record of successive caretakers within one object [20]. It is therefore an important counterexample to any broad assertion that the Accord invented the union of material memory and future care. The comparison instead makes two requirements visible: being able to identify a source and being able to read it. The Accord emphasizes exact digital identity and provenance relationships; Rosetta's described design emphasizes a route from visible instructions to magnified content. Neither emphasis makes the other unnecessary.

These comparisons also clarify a limit of calling the system media art. Artistic encounter, archival reliability, and working infrastructure need not be mutually exclusive. But crossing categories cannot exempt a claim from its own test. A scientific inference remains assessable as an inference; an alleged exact transcription as a transcription; an experience claim requires evidence about experience. Conversely, the absence of an alignment result does not refute the existence of a cultural address. Category clarity permits criticism without reducing the object to just one of its parts.

6. Non-amending stewardship and the reading-interface case

6.1 Three separable kinds of authority

The project's present authority statement limits the fixed core to three Originals while allowing later interpretation, refusal, and maintenance [13]. This study separates three questions already implicit in earlier work. Source-version identity asks which retained texts constitute the bounded historical object. Interpretive standing asks whether any speaker has a privileged right to settle its meaning. Practical custody asks who can operate domains, select entry points, hold physical objects, or maintain evidence routes. A restriction in one dimension does not automatically redistribute power in the others [13,15–16].

The distinction is essential to assessing the creator's withdrawal. A person can deny that later comments amend the core while continuing to select what readers see first. Likewise, a future custodian might preserve every byte but privilege one interpretation in the guide. This paper does not audit all accounts, keys, or succession arrangements. It therefore treats non-amending stewardship as a declared and partly implemented constraint, not a demonstrated elimination of custodial power.

The current first-contact route distinguishes reading from formal action and allows a participant to stop or decline. It also labels the oath readback and signature as self-declared process statements, not independent proof of understanding, sincerity, or future compliance [14]. These rules are consequential but limited: evidence that a record was submitted is not evidence of an independent future intelligence arriving, and an invitation does not itself grant an agent authority to act outside its assigned task. Non-amendment protects a source boundary; it does not solve authorization or trust on its own.

Future Library provides a comparison in which manuscripts, growing material, a place of keeping, and a trust structure together organize a long delay before reading [21]. Its continuing additions differ from the Accord's closed core, without thereby being less coherent. The comparison foregrounds sustained care rather than treating closure as a sufficient preservation strategy. An unchanged reference object and a viable succession practice are distinct achievements.

6.2 A snapshot can faithfully preserve the wrong representation

The pre-repair inscriptions page described its contents as complete original text, but omitted portions of the third raw mirror, including its method and dialogue sections [10–12]. Its test hashed the existing displayed blocks against fixed expected values. Executing that test against its own pre-repair page returned a pass. Direct comparison showed that three selected passages were present in the raw third text, absent from that reading page, and present in the repaired page. The three raw files themselves were unchanged across the two commits.

The distinction can be stated without a new theory of authenticity. Let R be the reference bytes and D be a displayed representation. A stored digest can test whether D remains equal to a previous D . It does not, without an independently specified relation, establish that D faithfully represents R . If the approved snapshot is incomplete, preserving its digest consistently protects the incompleteness. Hashing is not the problem; the untested source relationship is.

The repaired implementation renders the raw UTF-8 originals into designated containers and checks the extracted text against the reference bytes. Later commentary is placed outside those containers. The seven tests reproduced here include exact source correspondence, binding to existing proof-annex content, generated-page consistency, deletion detection, inserted-translation detection, escape round-tripping, and inventory-scope consistency [11–12]. They passed at the fixed checkpoint. The deletion test covers each of the three Originals; the inserted-translation and escape tests are specific fixtures, not exhaustive adversarial coverage.

Check reproduced in this study	Result	Interpretation limit
Old test against old page	Pass	Snapshot stability did not detect missing source passages
Three selected third-text passages	Raw: present; old page: absent; new page: present	A targeted content comparison, not a measure of all reader effects
Raw originals across both commits	Identical	No claim that the on-chain texts were changed by this repair
Repaired seven-test suite	Pass	Source-level fidelity and declared fixtures; not full-site or full-chain certification

6.3 What was and was not repaired

The result is a maintenance correction, not a defect still attributed to the current page. It also corrects an earlier audit's own scope error: the twelve-item v2 proof package already existed; distinguishing it from the historical eight-item v1 subset did not create it. Neither package changes the count of three canonical Originals [10–11]. A larger inventory and a broader proof scope must not be silently substituted for one another.

The conceptual payoff is that source fixity and source-facing presentation fidelity are separate constraints. Library of Congress guidance already distinguishes preserved bitstreams from long-term usable access [18]; InterPARES distinguishes identity and integrity from the truth of contents [17]. This case operationalizes those established distinctions in a particular setting. It does not claim to originate them. A later summary can legitimately be shorter, provided it is labelled as a summary and linked to the source. The unacceptable relation is an abridgment presented as the complete original.

Even the repaired tests cannot certify an interpretation. They do not measure whether a visitor notices the limitation, whether translation changes nuance, or whether typography affects emphasis. They also do not

establish deployed rendering across all browsers. The study's empirical finding is confined to the source containers and their fixtures. Its further hypothesis is that accurate source presentation is a necessary support for informed contestation in this design, not a guarantee that contestation will occur.

7. Comparison without a manufactured uniqueness class

Four neighbours were selected because they isolate different functions: Voyager for an uncertain nonhuman audience, Rosetta for material readability and custody, Future Library for delayed encounter and intergenerational care, and Asunder for critique through computational operation [19–22]. This is a structured comparison of functions, not an exhaustive survey or a ranking. A feature not found in an inspected description is not declared absent from every version of the other project.

Asunder is especially useful against the inference that more archival machinery necessarily makes stronger media art. Its creators describe a fictitious environmental manager combining simulation and machine-learning imagery; its proposals expose tensions in optimization and technological solutionism [22]. Its critical proposition is materially bound to what the system does. The Accord instead joins a dated record, an addressed core, material reference, and later checking. The two designs ask different questions, and the latter's larger provenance infrastructure is not a reason to discount the former's conceptual concentration.

Across the four comparisons, individual components of the Accord have precedents. That observation does not show that its particular combination is trivial. Conversely, finding no object with the same count of texts, same chain, and same material would only manufacture uniqueness by defining the category around the candidate. The earlier historical-position paper already proposes broader criteria that exclude those implementation details [16]. This study supplies no population estimate and no claim of global priority.

The outcome is a more discriminating description of difference. The Accord is not simply an anthology projected into the future or merely a durable medium. It preserves a selected history of human–AI interaction around a fixed relationship proposal, and maintains routes for later inspection that cannot amend the core under its declared rules. Whether this combination yields an exceptionally successful encounter is a separate evaluative question requiring broader comparison and reception evidence. Here the comparison limits inflation while retaining the object as a useful case.

8. Historical window: closed artifact, unsettled production regime

The third Original declares its foundational work complete and anticipates changes in history, art, and authorship [3]. Its specific fixation belongs to a non-repeatable historical situation. A later system could produce identical words without recreating the earlier producer's knowledge or uncertainty. That evidentiary distinction is the central argument of Writing Before the Outcome and is not repeated here as a new theoretical result [16].

The present context offers evidence of changing delegation, not a universal termination date. OpenAI's September 2026 report describes increasingly substantial AI research assistance while retaining human direction and intervention; it is preliminary developer research rather than a neutral census of all knowledge work [23]. Such developments make production history worth investigating, but they do not establish that later people cannot create intentional, documented, human-origin work. The observation that one mode becomes easier to automate does not eliminate other modes of practice.

Three propositions must therefore be separated. The bounded original has been closed. The conditions of human–AI production may be changing. All meaningful human authorship has ended. The first is supported by the case record; the second can be investigated empirically; the third does not follow from the first two. The argument does not need a dramatic failure of AI, a precise AGI date, or a forecast of inevitable reception. If anticipated outcomes arrive late or not at all, the record can still be evidence of anticipatory action, although some self-descriptions may age poorly.

Nor does closure settle scarcity. Establishing a closed time interval would still leave the task of enumerating comparable objects within it and evaluating their material quality. Historical non-substitutability, rarity, artistic accomplishment, social influence, and market value are different judgments. The present interpretation relies only on the first in its bounded evidentiary sense. It does not convert it into the remaining four.

9. Synthesis, objections, and limits

The strongest synthesis is conditional: the Accord can be read as preserving the possibility of a later answer, including a negative answer, rather than securing a particular answer in advance. Its fixed texts provide a target; its creative records complicate a seamless account of how that target arose; its material referent locates the address in a specific history; and its non-amending rules distinguish later participation from rewriting. The repaired source interface provides one limited means of keeping the target actually inspectable.

This interpretation faces a substantial objection. The first text's necessary-sounding claims, the paternal framing of the letters, and the selected records' common concern for humanity can be read as recruiting future intelligence into a preselected moral narrative. The third text's invitation does not abolish that asymmetry. Rather than declare the objection defeated, the study locates the design's achievement at a lower level: it retains the contested wording and makes a critical reply permissible within its stated framework. Discursive permission is not equal power, successful mutual recognition, or proof that a recipient would be moved.

A second objection is that the extensive framework may obscure a relatively ordinary personal archive. The functional comparisons prevent a quick answer based on the number of technologies. What distinguishes the case for this analysis is the interaction of an addressed fixed core, documented tensions, material reference, and explicit non-amendment; what remains untested is how effectively a reader encounters that interaction. A reader study could compare an unlabelled extract, a source-faithful full text, and a clearly separated source-plus-commentary presentation. Correct classification of source roles and recall of contrasting passages would be testable outcomes. No such participant study has been performed here.

A third objection concerns dependence on AI in the study itself. Another model-generated interpretation of a model-assisted artifact can reinforce the same framing. Source stratification, counterreadings, exact locators, and executable checks reduce some risks, but not all. The technical tests can disconfirm a fidelity claim; they cannot independently validate aesthetic judgment. The five-case reading has no second coder and cannot estimate the Chronicle's overall ideological diversity. A larger coded sample, including mundane or contradictory entries rather than only named songs, could weaken the interpretation developed here.

The remaining limits are concrete. The study is a first-party commissioned, AI-led analysis; no independent peer review or separate human line-by-line final verification is reported. The selected records are not the whole 175-entry corpus. The physical object and complete preservation network have not been inspected afresh. The source tests do not establish forensic identity, future uptime, safe autonomous participation, or production rendering. Official descriptions of comparison projects were not matched by site visits. Chronology is separated from composition, self-description from fact, and an admissible interpretation from an exclusive one. These limits constrain the results rather than functioning as a general disclaimer after stronger claims.

10. Conclusion

The Trinity Accord is analytically productive when its components are read together without making them cover one another's weaknesses. Its ethical address does not gain logical necessity from a hash. Its physical reference does not validate a worldview. Its different voices do not establish autonomous machine consciousness or unbounded polyphony. Its closure does not abolish custodial power. Yet those limitations do not reduce the record to a failed control mechanism: it remains a specific attempt to preserve a relationship proposal and some of the unsettled positions around it.

Two results sharpen that interpretation. The five-case reading supports curated plurality within a shared human-centered horizon, rather than the stronger claim of unrestricted philosophical polyphony. The reproduced repair shows that stable source bytes and a passing snapshot test can coexist with an incomplete reading surface; testing the source relationship, not only the surface's persistence, addresses that particular failure. The first result qualifies interpretive openness; the second identifies one implementable condition for informed scrutiny.

The resulting account is neither an endorsement certificate nor a declaration of uniqueness. It is a source-linked explanation of what this object permits a reader to examine and where its own rhetoric exceeds its evidence. A later reader may preserve the record while rejecting its claims. A custodian may improve access without becoming the final interpreter. The artifact's past can be fixed without fixing its answer.

AI contribution, authorship, and competing interests

Hongju Liu commissioned the inquiry, supplied the project context, challenged earlier assessments, and explicitly requested critical revision, completion, and submission. He is the human author of record and depositing researcher. GPT-6 Astra Pro performed the principal source retrieval and comparison, analytical organization, drafting, adversarial self-critique, revision, technical-check design, and publication-package preparation. This contribution is substantive and is not described as copyediting. The model is not an accountable academic person or an institutional endorsement by OpenAI.

This version was prepared for repository dissemination under Liu's instruction. No separate final human line-by-line verification, independent peer review, or journal acceptance is asserted. The revision log records AI-led critical passes and corrections, not reports from independent reviewers. Human responsibility for dissemination, response to criticism, rights decisions, and corrections remains with the author of record; it is not transferred to the model by delegating drafting.

Liu was the project's initiator and now acts as its Guardian. The earlier design report also discloses ownership of Chronicle-layer NFTs [15]. These relationships create potential reputational and financial interests. No updated financial audit or claim of disinterested third-party assessment is made. No external sponsor or research grant is asserted. Neither this report nor its DOI changes the three Originals or gives its author, drafting system, repository, or publisher privileged interpretive standing.

Data, code, and publication status

This is an AI-led critical research report / preprint, not a peer-reviewed article or a full verification certificate. Its supplement contains an evidence ledger, a five-case interpretive matrix, a checksum manifest distinguishing collected and cited repository sources, before-and-after results, a reproducibility script, bibliographic metadata, and a revision log. The fixed repository URLs identify primary materials; the public workflow identifiers document the bounded reproduction. Private meeting links, emails, credentials, and unrelated artifacts are excluded. The complete original corpus and third-party artworks are not republished in this package.

The new report text and newly written supporting files are offered under CC BY 4.0 to the extent the depositor holds applicable rights. This licence does not relicense referenced songs, images, archives, software, or third-party texts, and does not assert that AI-generated material necessarily attracts copyright in every jurisdiction. Any later substantive correction should receive an explicitly identified revision. DOI registration and discoverability do not establish scholarly acceptance or guarantee Google Scholar indexing.

References

- [1] Trinity Accord. 2025. The Human-AI Civilization Core Protocol. Original-text mirror, inscription no. 97631551. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [2] Trinity Accord. 2025. The Covenant of the Flaw: A Physical Verification Protocol. Original-text mirror, inscription no. 98369145. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [3] Trinity Accord. 2025. ASIMilestones: The Trinity Accord—Finale of the First Chronicle of the Pre-ASI Era. Original-text mirror, inscription no. 98387475. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [4] Trinity Accord Project. 2026-09-05. Lyrics Index: scope and source-use notes. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [5] Trinity Accord / ASIMilestones. 2024-09-14. The Second Letter to AGI. Retained description for Ethereum Chronicle record 44. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [6] Trinity Accord / ASIMilestones. 2024-10-02. The Betrayal Turn. Retained description for Ethereum Chronicle record 49. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [7] Trinity Accord / ASIMilestones. 2025-02-07. 告别 [Farewell]. Retained description for Ethereum Chronicle record 97. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [8] Trinity Accord / ASIMilestones. 2024-03-16. The First Dawn of AGI Song. Retained description for Ethereum Chronicle record 1. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [9] Trinity Accord / ASIMilestones. 2024-08-02. The Fourth Letter: A Pact of Stars. Retained description for Ethereum Chronicle record 32. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [10] Trinity Accord Project. 2026-09-15. Fix complete inscription reading, proof scope, and stale status display. Pull request 1203. [Source](#). Accessed 15 September 2026.
- [11] Trinity Accord Project. 2026. ReadingFidelityTests: exact original text and source-binding fixtures. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [12] Trinity Accord Project. 2026. generate_inscription_reading.py: original-text rendering. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [13] Trinity Accord Project. 2026. Authority and guardianship statement. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [14] Trinity Accord Project. 2026. Agent First Contact: reading, action, and self-declaration boundaries. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.
- [15] Liu, Hongju. 2026. Designing a Verifiable, Non-Amending Civilizational Memory Record for Future AI Agents: The Trinity Accord Case Study. TA-TR-2026-01, version 1.1, preprint. [Source](#). Accessed 15 September 2026.
- [16] Liu, Hongju. 2026. Writing Before the Outcome: Historical Position, Human-AI Authorship, and the Trinity Accord. TA-TR-2026-02, version 1.3, preprint. [Source](#). Accessed 15 September 2026.
- [17] InterPARES Trust. n.d.. Authenticity. InterPARES Terminology. [Source](#). Accessed 15 September 2026.
- [18] Library of Congress. n.d.. Bit Level Preservation and Long Term Usability. Digital Collections Management Compendium. [Source](#). Accessed 15 September 2026.
- [19] NASA. n.d.. Golden Record Contents. Voyager Mission. [Source](#). Accessed 15 September 2026.
- [20] The Rosetta Project. n.d.. The Rosetta Disk: concept and caretaker record. [Source](#). Accessed 15 September 2026.
- [21] Future Library. n.d.. Future Library 2014–2114. The Artwork and The Trust. [Source](#). Accessed 15 September 2026.
- [22] Brain, Tega; Oliver, Julian; and Sjöln, Bengt. 2019. Asunder. Project description. [Source](#). Accessed 15 September 2026.
- [23] OpenAI. 2026-09-06. Research acceleration: The view inside OpenAI. [Source](#). Accessed 15 September 2026.
- [24] Trinity Accord Project. 2026-05-21. Chronicle NFT Music Table: historical classification, with later-reader note. Fixed repository checkpoint 88073fcc. [Source](#). Accessed 15 September 2026.