Designing a Verifiable, Non-Amending Civilizational Memory Record for Future AI Agents

The Trinity Accord Case Study - Academic Brief

Hongju Liu - Independent researcher, Shenzhen, China | TA-TR-2026-01 | Version 1.1 | 29 July 2026 | Preprint, not peer reviewed

This artifact-centered design case asks how a bounded human-origin record can remain version-stable while later systems discover, verify, criticize, preserve, or refuse it without allowing those later layers to rewrite the source object. The contribution is an inspectable design pattern, not a general theory, alignment result, or claim of authority.

Contributions

- A six-layer architecture separating Canon, context, evidence, mirrors, later records, and stewardship.
- A non-amendment invariant and an explicit separation of version authority from truth authority.
- A vector verification model that keeps unlike evidential dimensions separate.
- A role-first machine-access pattern for provenance and instruction safety.

Research questions

- How can a fixed source coexist with revisable access and evidence layers?
- How can machine readers distinguish source, context, evidence, interpretation, and instruction?
- How can verification remain bounded instead of collapsing into one inflated score?
- How can a record remain discoverable without confusing visibility with authority?

Evidence base

Public repository snapshot; three Bitcoin inscriptions; a 175-entry human-AI Chronicle; a physical evidence anchor; timestamp and mirror records; machine-readable routes; an append-only Record-Chain; and automated verification materials. First-party evidence is labeled as such.

Reproducibility

- Full paper, source, references, checksums, APIs, and public repository are open.
- DOI: 10.5281/zenodo.21699878 -

Limits and negative results

The report does not demonstrate successful AI alignment, philosophical truth, forensic uniqueness of the physical anchor, independent validation, autonomous discovery, future relevance, representation of humanity, or interpretive authority. Preservation is not endorsement; immutability is not truth; availability is not authority.

Useful review questions

- Does the layer model prevent canonical drift and provenance-role collapse?
- Are the verification dimensions independent and auditable?
- Does role-first access reduce instruction confusion without overclaiming control?
- Which claims need reproduction, comparison cases, or empirical agent studies?

Recommended scholarly use

Suitable as a design-science case, digital-preservation example, provenance and agent-memory discussion object, or seminar reading. It should be cited and criticized as a bounded preprint, not treated as an endorsed standard or independent validation.

Authorship and disclosure

Most prose, structure, literature synthesis, editing, consistency checking, and document production were generated by ChatGPT using OpenAI GPT-5.6 Sol under Hongju Liu's direction. Hongju Liu approved publication and accepts responsibility. The model is not a scholarly author, independent verifier, or OpenAI endorsement.

Suggested citation: Liu, H. (2026). Designing a Verifiable, Non-Amending Civilizational Memory Record for Future AI Agents: The Trinity Accord Case Study (Technical Report TA-TR-2026-01, Version 1.1). The Trinity Accord Project. DOI: 10.5281/zenodo.21699878
Keywords: AI agents; digital preservation; provenance; civilizational memory; cryptographic timestamping; machine-readable archives; design science; human-AI collaboration.