

Peer Review

Review of: "Truvry: Portable, Decentralized Trust Proofs for Inclusive Digital Participation and Democratic Decision-Making"

Nasibeh Mohammadzadeh¹

1. Universidad Politécnica de Catalunya, Spain

This preprint introduces a decentralized protocol for generating portable trust proofs from behavioral data, aiming to enhance inclusion in digital ecosystems and support AI-mediated governance. The concept of decoupling trust from identity is innovative, and the focus on portability addresses real challenges in fragmented trust systems. Overall, the paper presents a solid conceptual framework with promising societal implications, though the reliance on simulated data currently limits its empirical strength.

Strengths

- The portability of trust proofs across platforms fills a gap in existing DID and reputation systems (e.g., uPort, BrightID), emphasizing behavioral evidence over static credentials.
- Strong privacy features such as hashed anchoring, field-level redaction, and zk-compatibility align well with GDPR and the DPDP Act, making the approach regulator-friendly.
- The democratic applications (e.g., sybil-resistant voting in DAOs or citizen assemblies) add practical value, positioning Truvry as potential governance infrastructure.

Points for Improvement / Clarification

- **Empirical Depth:** Results are based on a simulated dataset of 112 instances, with metrics such as 3.7s proof issuance and 91% classification accuracy. A real-world pilot involving diverse groups such as gig workers, DAO participants, or stateless individuals would significantly enhance empirical depth and highlight practical edge cases.

- **Adversarial Resilience:** The rule-based scoring engine shows 42% suppression of manipulated inputs but lacks detail on more advanced attacks. Integrating probabilistic models or ML-based anomaly detection could improve robustness.
- **Interoperability:** While DID-optional, deeper integration with standards such as W3C Verifiable Credentials and the European Blockchain Services Infrastructure (EBSI) would strengthen alignment with existing frameworks and facilitate adoption. Concrete examples of cross-platform testing would be especially valuable.
- **Bias and Ethics:** The ethics section covers autonomy and redressal well, but indirect biases in behavioral data (e.g., reliance on stable internet access) need stronger safeguards. Differential privacy techniques beyond Laplace noise could be considered.
- **Scalability:** Gas costs (<\$0.02) are promising, but scalability of on-chain anchoring under higher loads has not been tested. Further exploration of Layer-2 solutions, as briefly noted in the roadmap, would strengthen feasibility.

Overall Impression

This work contributes a thoughtful, behavior-centric model for trust in decentralized systems, with clear potential for fostering inclusive participation. It advances beyond siloed reputation ratings but would benefit from real-user validation, stronger adversarial defenses, and more concrete interoperability pathways. Minor revisions for empirical rigor and expanded ethical safeguards could elevate its impact. *I recommend further real-world piloting to validate and refine the approach.*

Declarations

Potential competing interests: No potential competing interests to declare.