Evidence-Based Subjective Logic in ZKML Identity Systems: Towards Verifiable Epistemic Trust

Oliver C. Hirst [Steake]
Shadowgraph Labs

October 2025

Abstract

This paper formalizes the integration of **Evidence-Based Subjective Logic (EBSL)** into **Zero-Knowledge Machine Learning (ZKML)** frameworks for decentralised identity. EBSL extends classical subjective logic by introducing direct mappings between observed evidence, belief entropy, and proof verifiability, enabling self-sovereign identities to express uncertainty as a first-class cryptographic primitive. When embedded into ZKML systems, EBSL provides the mathematical substrate for privacy-preserving, evidential reasoning — allowing trust computations to be both explainable and verifiable without revealing underlying data.

Contents

1 Introduction

The question of how to mathematically model *trust* lies at the heart of all decentralised identity systems. Traditional models reduce trust to binary verification: a credential is valid or invalid. However, in social, financial or epistemic systems, belief is not binary — it exists as a gradient of certainty informed by limited evidence.

Evidence-Based Subjective Logic (EBSL) bridges this gap. It allows systems to reason over uncertain, incomplete, and conflicting data using probabilistic opinions rather than fixed truth values.

This becomes particularly powerful in the context of **Zero-Knowledge Machine Learning (ZKML)** — systems that can compute over encrypted data and prove the correctness of inference without revealing inputs or model parameters. By combining EBSL and ZKML, we arrive at a new class of identity architecture: *verifiable epistemic trust systems* — networks capable of reasoning about belief, evidence, and doubt in a mathematically grounded and cryptographically secure manner.

2 Conceptual Framework

2.1 From Probability to Belief

Subjective Logic generalizes classical probability by decomposing knowledge about a proposition into three components:

$$b = \text{belief}, \quad d = \text{disbelief}, \quad u = \text{uncertainty},$$

subject to

$$b + d + u = 1, \quad b, d, u \in [0, 1].$$

The expected probability of a proposition X given an opinion ω is:

$$E(\omega) = b + a \cdot u,$$

where a is the base rate — a prior expectation in the absence of evidence.

2.2 Evidence Incorporation

EBSL extends this model by linking belief components directly to observed evidence counts:

$$b = \frac{r}{r + s + 2}, \quad d = \frac{s}{r + s + 2}, \quad u = \frac{2}{r + s + 2},$$

where r and s denote positive and negative evidence supporting a claim. This allows a seamless Bayesian update mechanism as new data arrive.

3 Mathematical Formulation

3.1 Opinion Space

Define the opinion space:

$$\Omega = \{(b, d, u, a) \in [0, 1]^4 \mid b + d + u = 1\}.$$

Each opinion represents a subjective state on the proposition X.

3.2 Fusion Operator

Two opinions ω_1 and ω_2 can be fused under independence as:

$$\omega_1 \oplus \omega_2 = \left(\frac{b_1 u_2 + b_2 u_1}{k},\ \frac{d_1 u_2 + d_2 u_1}{k},\ \frac{u_1 u_2}{k},\ a \right),$$

where $k = u_1 + u_2 - u_1 u_2$. This operator allows for the propagation of trust through a network of peers.

3.3 Discounting Operator

To propagate trust through transitive relationships:

$$\omega_{ik} = \omega_{ij} \otimes \omega_{jk},$$

with

$$b_{ik} = b_{ij} b_{jk}, \quad d_{ik} = d_{ij} + u_{ij} d_{jk}, \quad u_{ik} = u_{ij} u_{jk}.$$

4 Integration into ZKML

4.1 Zero-Knowledge Constraints

Each opinion update is represented as a circuit:

$$\mathcal{C}_{EBSL}(\omega_1, \omega_2) = \omega_3,$$

and verified under zero-knowledge constraint:

$$\text{VerifyZK}(\pi) \Rightarrow b + d + u = 1,\ 0 \leq b, d, u \leq 1.$$

Thus, trust propagation can be verified cryptographically without exposing the contributing evidence.

4.2 ZKML Training and Inference

Machine learning models operating in zero-knowledge can utilize belief vectors ω_i as input features. The model computes:

$$\hat{y}_i = f_\theta(\omega_i)$$

where f_θ is a parameterized function (e.g. neural network). Proof of correct inference is produced by a ZKML prover:

$$\text{Proof} = \text{ProveZK}(f_\theta, \omega_i, \hat{y}_i).$$

This enables explainable predictions over uncertain inputs without revealing raw evidence or model parameters.

5 Identity and Reputation Modeling

5.1 Belief Tensor for Identity

Each identity I_i is defined as a tensor of opinions:

$$\mathcal{T}_i = [\omega_{i,1}, \omega_{i,2}, \ldots, \omega_{i,n}],$$

representing beliefs across different contexts — e.g. reliability, honesty, performance. Identity thus becomes an evidential field, evolving as opinions are fused over time.

5.2 Temporal Update Rule

Reputation updates follow an exponential decay model:

$$R_i^{(t+1)} = \beta R_i^{(t)} + (1 - \beta) E(\omega_i^{(t)}),$$

where β encodes memory persistence. New evidence adjusts the identity's belief distribution without revealing raw transactions.

6 Algorithmic Implementation

Algorithm 1 EBSL-Based Reputation Update in ZKML System

Require: Prior reputation $R_i^{(t)}$, evidence set $E_i^{(t)}$

1: Compute opinion from evidence: $\omega_i \leftarrow f_{EBSL}(E_i^{(t)})$
2: Fuse neighbor opinions: $\omega_i' \leftarrow \bigoplus_{j \in N(i)} \omega_j$
3: Generate proof: $\pi \leftarrow \text{ProveZK}(\mathcal{C}_{EBSL}, \omega_i, \omega_i')$
4: Verify proof: $\text{VerifyZK}(\pi) = 1$
5: Update reputation: $R_i^{(t+1)} \leftarrow \beta R_i^{(t)} + (1 - \beta) E(\omega_i')$

This process allows decentralised identity networks to maintain verifiable reputation states even under partial observability and privacy constraints.

7 Entropy and Uncertainty

Uncertainty u measures informational entropy:

$$H(\omega) = -b \log b - d \log d - u \log u.$$

This entropy directly informs the model's confidence calibration, allowing agents to evaluate trustworthiness as a function of evidential entropy rather than deterministic assertions.

Entropy regularization in ZKML training can be achieved via:

$$\mathcal{L}_{trust} = \|E(\omega) - \hat{y}\|^2 + \lambda H(\omega),$$

balancing belief fidelity with epistemic humility.

8 Applications

8.1 Verifiable Reputation Graphs

Reputation is computed through EBSL fusion and proven in ZK, forming a decentralised, privacy-preserving trust graph suitable for social, financial, or DAO governance applications.

8.2 Evidential Credential Verification

Institutions issue attestations in the form of opinions rather than static credentials. A verifier checks only the ZK proof that the opinion satisfies belief normalization and threshold conditions.

8.3 Sybil Resistance

EBSL-based identities possess measurable uncertainty; newly created identities exhibit high u, reducing their influence until evidence accumulates.

9 Discussion

EBSL provides a rigorous bridge between epistemology and cryptography. When embedded into ZKML frameworks, it enables a new class of self-sovereign systems where knowledge, uncertainty, and prediction coexist under cryptographic guarantees.

Trust is no longer an assumption or a score — it is a dynamic, evidential waveform verified by mathematics.

10 Conclusion

Evidence-Based Subjective Logic in ZKML Identity Systems establishes a foundation for verifiable epistemic reasoning — a world where digital identity evolves through evidence, where belief becomes a measurable construct, and where privacy and transparency are no longer opposites.

By uniting evidential logic, probabilistic reasoning, and zero-knowledge cryptography, EBSL represents not just a framework for decentralised identity, but the blueprint of a new digital epistemology.

References

References

[1] A. Jøsang, "A Logic for Uncertain Probabilities," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 9, no. 3, pp. 279–311, 2001.

[2] EZKL Team, "Zero-Knowledge Machine Learning Framework," GitHub Repository, 2024.

[3] O. C. Hirst, "Evidence-Based Subjective Logic in decentralised Identity Systems," Shadowgraph Research, Technical Report, 2025.