Demo: Trust Before Immutable—PoEval in Action for Permissioned IoT Data Marketplaces

Rashmi Ratnayake*, Madhusanka Liyanage†, Liam Murphy‡
*†‡School of Computer Science, University College Dublin, Ireland
Email: *rashmi.ratnayake@ucdconnect.ie, †madhusanka@ucd.ie, ‡liam.murphy@ucd.ie

***Abstract*—Blockchain-based data marketplaces provide immutability and auditability, yet they fundamentally lack guarantees on the correctness of data at the time of inclusion. In industrial IoT environments, where decisions rely on continuous streams of sensor data, this limitation can lead to the permanent storage of faulty or malicious inputs, undermining trust and system reliability. This demonstration presents Proof of Evaluation (PoEval), a novel pre-consensus trust evaluation framework that enforces trust before immutability in permissioned blockchain-based data marketplaces. Unlike conventional approaches that rely solely on consensus for agreement, PoEval introduces a data-centric validation layer that integrates distributed machine learning (ML) evaluators directly into the data admission workflow. This enables independent, model-driven assessment of incoming data prior to ledger inclusion, effectively decoupling data trustworthiness from traditional consensus mechanisms. Through quorum-based aggregation of evaluator decisions, PoEval ensures that only data deemed trustworthy are recorded on-chain. We implement a multi-node prototype and an interactive visualization dashboard that exposes the full evaluation pipeline, including real-time data streams, evaluator outputs, decision aggregation, and blockchain state. The demonstration highlights how PoEval transforms the blockchain from a passive immutable ledger into an active trust enforcement mechanism, significantly improving data quality and reliability compared to conventional workflows.**

***Index Terms*—Blockchain, IoT, Data Trust, Machine Learning, Data Marketplace**

I. Introduction

Blockchain technology enables decentralized data marketplaces with strong guarantees of immutability, transparency, and auditability, making it well suited for industrial IoT ecosystems where multiple entities share sensor data. However, it provides no inherent mechanism to verify data correctness at the time of inclusion [1], allowing erroneous or malicious data to be permanently recorded and propagated, thereby undermining system reliability. This limitation is particularly critical in industrial IoT environments with heterogeneous and potentially unreliable sensors. Existing approaches address data trust either after inclusion (e.g., rollback [1] or redactable mechanisms [2]) or through indirect methods such as outlier detection [3] and reputation-based consensus [4], which do not prevent untrustworthy data from being recorded on-chain.

To address this gap, we propose *Proof of Evaluation (PoEval)* [5], the first ML-driven, data-centric pre-consensus trust evaluation mechanism for permissioned blockchain-based data marketplaces. PoEval enforces *trust before immutability* by requiring independent validation of all incoming data using distributed, heterogeneous ML evaluators, with final decisions derived via quorum-based aggregation. By embedding trust evaluation into the pre-consensus data admission workflow, PoEval elevates data validation to a first-class component of blockchain operation. Comprehensive results are reported in our main paper [5], where PoEval achieves near-zero false positive rates (approximately 0.1%), compared to 6%–48% in outlier-aware consensus approaches [3]. In contrast, conventional systems without pre-consensus filtering exhibit substantially higher false positive rates, often approaching the underlying fault ratio.

In this demonstration, we present a practical realization of PoEval for an industrial IoT data marketplace (Fig. 1). A lightweight multi-node prototype with distributed evaluators, a coordination module, and a permissioned blockchain is implemented alongside an interactive dashboard for real-time visualization. The demo illustrates how ML-driven pre-consensus validation transforms blockchain from a passive immutable ledger into an active trust enforcement mechanism.

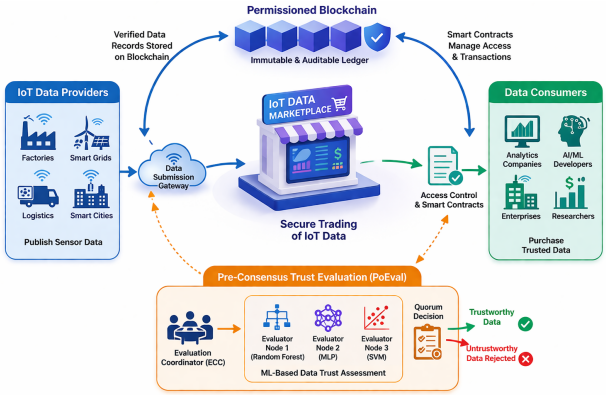

Fig. 1: Blockchain-enabled industrial IoT data marketplace with PoEval

II. System Architecture

A detailed description of the PoEval framework, including its system design, protocol phases, algorithms, and evaluation, is presented in our main PoEval paper [5]. This demonstration focuses on the practical implementation and visualization of the pre-consensus trust evaluation workflow.

We consider a consortium-based industrial IoT data marketplace, where authorized participants share sensor data via a permissioned blockchain. PoEval is integrated into the data admission process to enable trust verification prior to blockchain inclusion. As shown in Fig. 2, the architecture comprises four

ISBN 978-3-903176-82-9

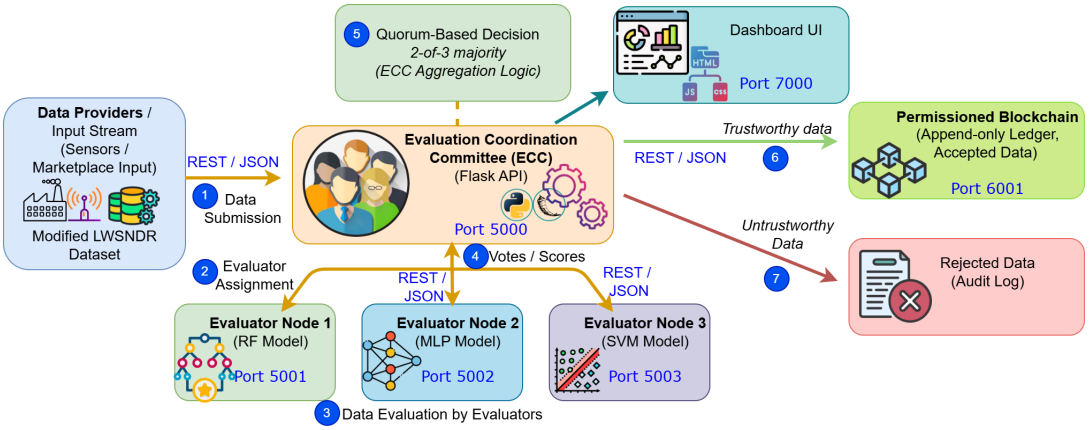

Fig. 2: PoEval Workflow. Incoming data are evaluated by distributed ML-based evaluator nodes coordinated by the ECC, followed by quorum-based decision-making to determine blockchain inclusion or rejection.

components: data providers, an evaluation coordination committee (ECC), evaluator nodes, and a permissioned blockchain layer.

Data Providers: Industrial entities generate continuous sensor data streams (e.g., temperature, humidity) with associated metadata, which may include untrustworthy data due to faults, environmental variations, or adversarial manipulation.

Evaluation Coordination Committee: The ECC orchestrates evaluation by distributing records to evaluator nodes and aggregating their outputs via quorum-based voting. It is implemented as a replicated committee of consortium-operated services executing identical logic, with deterministic leader rotation for ordering and fault tolerance achieved through stateless replication. Evaluator assignment follows a hybrid reputation-weighted and random strategy to balance reliability, fairness, and resilience to collusion. The ECC maintains an audit log for traceability, and evaluator reputation is continuously updated based on agreement with consensus outcomes to promote sustained accuracy. ECC membership and quorum policies are governed by the consortium, making PoEval well suited to permissioned blockchain environments.

Evaluator Nodes: Independent nodes apply pre-validated machine learning models to assess data trustworthiness, producing trust scores, binary decisions, and confidence estimates. Only models meeting predefined performance thresholds are certified, while heterogeneous models (e.g., Random Forest, MLP, and SVM) improve robustness by reducing model-specific bias.

Permissioned Blockchain Layer: The blockchain stores only records that pass pre-consensus evaluation. Accepted records are grouped into blocks and appended to the ledger, while rejected records are excluded and optionally logged for auditability. Block proposal is deterministically assigned based on evaluator reputation, ensuring verifiable and tamper-resistant leader election; thus, eligibility is earned through evaluation-driven reputation, motivating the term PoEval.

The workflow proceeds as follows: a data record is submitted to the transaction pool, distributed by the ECC to multiple evaluators for independent assessment using validated ML models. The resulting outputs are aggregated by the ECC via quorum-based voting, and only records deemed trustworthy are committed to the blockchain by the block proposer, enforcing data quality prior to immutability and improving overall system reliability.

III. Prototype Implementation and Demonstration

To demonstrate PoEval, we develop a lightweight prototype of a permissioned industrial IoT data marketplace with pre-consensus trust verification. The system realizes the core PoEval workflow—ECC coordination, distributed evaluators, and quorum-based decision-making—in a simplified environment.

A. System Implementation and Deployment

The prototype is implemented in Python using Flask-based microservices, with components deployed as independent services on separate ports to emulate a distributed network on a single machine. The system includes an ECC coordinator, evaluator nodes, a blockchain service, and a dashboard interface. Components communicate via RESTful APIs using JSON messages, providing a lightweight abstraction of a permissioned blockchain environment. The blockchain functionality is simplified to focus on PoEval, while remaining compatible with platforms such as Hyperledger Fabric through standard application and smart contract interfaces.

B. Evaluator Nodes and ML Models

Each evaluator hosts an independently trained binary classifier, with three heterogeneous models used: Random Forest, MLP, and SVM. The demonstration uses the modified Labelled Wireless Sensor Network Data Repository (LWSNDR) dataset [6], containing both normal and fault-injected samples, to model realistic industrial IoT data quality issues such as sensor faults, drift, and communication errors. Models operate on feature vectors and output trust scores in [0,1], which are thresholded into binary decisions. Each evaluator returns a

JSON response containing the trust label, score, and node identifier.

C. Evaluation Coordination and Quorum Mechanism

The ECC receives data via REST APIs, distributes records to evaluators, and aggregates their outputs using a quorum-based voting mechanism, where a record is accepted if at least q out of k evaluators agree, with q defined by consortium governance. In this demo, a simple majority over three nodes ($q = 2, k = 3$) is used. Trusted records are forwarded for blockchain inclusion, while others are rejected and added to an audit log for traceability.

D. Blockchain Service

A prototype permissioned blockchain is implemented as an append-only ledger. Accepted records are grouped into blocks and appended sequentially, each containing an index, timestamp, records, and previous hash. Only records that pass PoEval are recorded on the blockchain, preventing untrustworthy data from becoming immutable.

E. Data Stream Generation

Data records are streamed from the modified LWSNDR dataset, containing both normal and fault-injected samples, to emulate real-time sensor input. These records are drawn from a held-out test set used exclusively for runtime evaluation and not for model training. The untrustworthy data ratio can be varied to simulate different levels of data corruption.

F. Interactive Dashboard

A Flask- and JavaScript-based dashboard provides real-time visualization of data streams, evaluator outputs, trust decisions, blockchain state, and rejected records. Users can submit individual records or initiate continuous streams to observe the PoEval process dynamically.

IV. Conclusion and Future Work

This demonstration presented PoEval, the first ML-driven, data-centric pre-consensus trust evaluation mechanism for permissioned blockchain-based industrial IoT data marketplaces. Unlike existing approaches that assume data correctness or rely on post-hoc corrections and limited outlier-based filtering, PoEval enforces *trust before immutability* by validating data through distributed ML-based evaluation and quorum-driven aggregation prior to ledger inclusion, preventing untrustworthy data from being permanently recorded. A lightweight multi-node prototype demonstrates how IoT data are independently evaluated and selectively admitted based on consensus-driven trust decisions, with interactive visualization highlighting improvements in data quality and reliability on blockchain.

The system is abstracted for feasibility and limited to a small number of evaluator nodes and a simplified blockchain; however, the design remains applicable to real-world permissioned environments and can be integrated with existing platforms through standard interfaces. Future work will focus on scalability, latency analysis, and integration with production-grade permissioned blockchain systems.

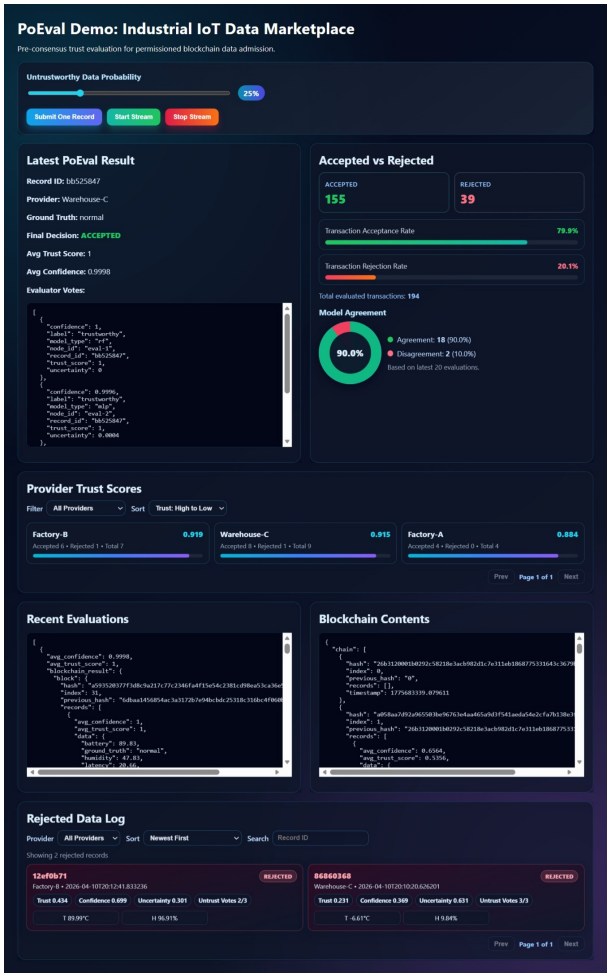

Fig. 3: PoEval Dashboard

Acknowledgement

This work was partly supported by European Union in the ENSURE-6G project (Grant ID. 101182933).

References

[1] A. Carvalho, J. W. Merhout, Y. Kadiyala, and J. Bentley II, "When good blocks go bad: Managing unwanted blockchain data," *International Journal of Information Management*, vol. 57, p. 102263, 2021.

[2] Y. Hou, J. Zou, L. Wang, X. Lu, X. Lu, and M. Wang, "PRBCP: Publicly Redactable Blockchain With Off-Chain Reputation-Based Consensus Protocol," *IEEE Transactions on Network and Service Management*, vol. 22, no. 5, pp. 4603–4618, 2025.

[3] M. Salimitari, M. Joneidi, and M. Chatterjee, "AI-enabled blockchain: An outlier-aware consensus protocol for blockchain-based IoT networks," in *IEEE Global Communications Conference*. IEEE, 2019.

[4] M. T. de Oliveira, L. H. Reis, D. S. Medeiros, R. C. Carrano, S. D. Olabarriaga, and D. M. Mattos, "Blockchain reputation-based consensus: A scalable and resilient mechanism for distributed mistrusting applications," *Computer Networks*, 2020.

[5] R. Ratnayake, M. Liyanage, and L. Murphy, "Proof of Evaluation: ML-Driven Consensus for Trustworthy Permissioned Blockchains," in *International Federation for Information Processing (IFIP) Networking Conference*, 2026.

[6] S. Zidi, T. Moulahi, and B. Alaya, "Fault detection in wireless sensor networks through SVM classifier," *IEEE Sensors Journal*, vol. 18, no. 1, 2018.