

A BLOCKCHAIN PRIVACY-PRESERVING DEEP AGENTIC VALIDATION FRAMEWORK USING ZERO-KNOWLEDGE PROOFS

Anwar Ali Sathio^{*1}, Muhammad Malook Rind², Abdul Hafeez Khan³,
Shafique Ahmed Awan⁴

^{1,4}Benazir Bhutto Shaheed University, Karachi, Sindh, Pakistan

^{1,2,3}Sindh Madarssatul Islam University, Karachi, Sindh, Pakistan

^{*1}anwar.sathio@bbsul.edu.pk, ²malook.rind@smiu.edu.pk, ³ahkhan@smiu.edu.pk,

⁴shafique.awan@bbsul.edu.pk

DOI: <http://doi.org/10.5281/zenodo.21822721>

Keywords

Zero-knowledge proofs, deep reinforcement learning, agentic validation, blockchain consensus, energy efficiency, privacy preservation, zk-SNARK.

Article History

Received: 19 January 2026

Accepted: 04 March 2026

Published: 18 March 2026

Copyright @Author

Corresponding Author: *

Anwar Ali Sathio

Abstract

Consensus protocols such as Proof-of-Work (PoW) and Proof-of-Stake (PoS) remain the backbone of most public blockchains, yet both impose steep energy costs, sizeable confirmation delays, and only partial transaction privacy. This paper presents the Deep Agentic Validation Framework (DAVF), a consensus-free alternative in which a population of Deep Reinforcement Learning Agents (DRLA) autonomously validates transactions and is coupled with zero-knowledge proofs (zk-SNARKs) to preserve confidentiality without sacrificing verifiability. Each agent learns a validation policy through Proximal Policy Optimization, trading off throughput, energy use, accuracy, and latency within a single reward signal, while zk-SNARK verification confirms transaction correctness without revealing transaction content. DAVF was implemented in a 50-node simulated peer-to-peer network and benchmarked on Ethereum transaction traces against PoW, PoS, and Practical Byzantine Fault Tolerance (PBFT). The framework attained a validation accuracy of $99.7\% \pm 0.2\%$, energy consumption of 0.012 kWh per transaction, a mean latency of 0.8 s, and a sustained throughput of 80 transactions per second, corresponding to roughly a 400-fold reduction in energy use relative to PoW and a 20% latency improvement over PBFT. Independent-samples t-tests and one-way ANOVA confirmed that these gains were statistically significant ($p < 0.001$) across all measured metrics. The results indicate that combining learning-based validation with succinct zero-knowledge cryptography is a practical route toward blockchain systems that are simultaneously energy-efficient, privacy-preserving, and scalable enough for real-world financial deployment.

1. Introduction

Distributed ledger technology has reshaped how financial value is transferred by allowing mutually distrustful parties to agree on a shared transaction history without a central intermediary [1]. This trustless property depends entirely on the consensus mechanism that decides which transactions are valid, and the two mechanisms

that dominate practice, Proof-of-Work (PoW) and Proof-of-Stake (PoS), each carry costs that limit their suitability for privacy-sensitive, high-throughput financial applications [2], [3]. PoW's reliance on repeated hashing consumes electricity on a national scale [4], while PoS reduces computation but concentrates influence on large stakeholders and introduces new governance risks.

Neither approach was designed with confidentiality in mind: transaction amounts, senders, and receivers are visible on the public ledger, which is difficult to reconcile with the confidentiality expectations of banks, enterprises, and individual users alike.

Cryptographic tools such as zero-knowledge proofs (ZKPs) and confidential transactions have narrowed the privacy gap by letting a prover convince a verifier that a statement is true without revealing the statement itself [5]. However, most deployments graft these primitives onto an existing consensus layer rather than rethinking validation from the ground up, so the underlying energy and scalability problems persist. Three requirements therefore remain only partially met by literature as like (a) removing dependence on energy-intensive consensus computation, (b)

preserving transaction confidentiality through cryptographic guarantees rather than institutional trust, and (c) enabling validation logic that adapts autonomously to changing network load rather than following a fixed protocol [6], [7].

This paper addresses these three requirements jointly through the Deep Agentic Validation Framework (DAVF), which replaces conventional consensus with a population of Deep Reinforcement Learning Agents (DRLA) that learn to validate transactions directly, and pairs this learned validation with zk-SNARK proofs so that correctness can be checked without disclosing transaction content [8]. Table 1 summarizes, at a qualitative level, how DAVF is positioned relative to PoW, PoS, and PBFT along the dimensions that motivate this work.

Table 1. Qualitative comparison of DAVF with representative consensus mechanisms [9]–[14]

Feature	PoW	PoS	PBFT	DAVF (proposed)
Energy consumption	Very high	Medium	Low	Very low
Throughput	Low	Medium	High	High
Latency	High	Medium	Low	Very low
Privacy	Low	Medium	Low	Very high (zk-SNARK)
Scalability	Limited	Improved	Limited	Highly scalable
Learning / adaptability	None	None	None	Agentic reinforcement learning

The contributions of this study are fourfold. First, it proposes a consensus-free blockchain validation architecture that couples autonomous DRLA decision-making with zk-SNARK proofs to jointly guarantee privacy and integrity [15], [38]–[45]. Second, it designs an energy-aware validation pipeline that removes redundant computation and communication overhead relative to mining- or voting-based consensus. Third, it provides empirical evaluation on Ethereum-derived financial transaction data, benchmarking energy consumption, throughput, latency, and privacy leakage against PoW, PoS, and PBFT under matched experimental conditions. Fourth, it demonstrates that the resulting architecture scales with network size without requiring a redesign of the underlying validation logic, which is a

prerequisite for deployment in production financial systems.

The remainder of the paper is organized as follows. Section 2 reviews related work on consensus mechanisms, privacy-preserving cryptography, and agent-based validation. Section 3 details the proposed DAVF architecture and its mathematical formulation. Section 4 describes the experimental setup. Section 5 reports quantitative results, Section 6 discusses their implications and limitations, and Section 7 concludes with directions for future work.

2. Related Work

Blockchain systems rely on a consensus algorithm to reconcile the state held by geographically distributed, mutually untrusted nodes. Proof-of-Work, introduced by Nakamoto [1], established

the field by tying block creation to computationally expensive hash-puzzle solving, which provides strong security guarantees at the cost of very high energy expenditure. Proof-of-Stake [2] was proposed as a lower-energy alternative that selects validators in proportion to their staked holdings, but this design has been criticized for encouraging wealth concentration and for weaker resistance to certain long-range and stake-grinding attacks [3]. Delegated PoS and Practical Byzantine Fault Tolerance (PBFT) improve throughput and finality time by limiting the active validator set, though this typically comes at the expense of decentralization and, in most implementations, transaction-level privacy.

Because transaction confidentiality is a first-order concern in financial applications, a separate line of work has focused on privacy-preserving cryptography layered on top of existing consensus protocols. Zcash uses zk-SNARK proofs to hide transaction amounts and participants while still allowing public verification of validity [16], [17], [46]-[50]. Monero instead relies on ring signatures and stealth addresses to obscure sender and receiver identities [19]. Both approaches demonstrate that strong privacy is achievable without abandoning public verifiability, but because the privacy layer is decoupled from the validation logic itself, systems that adopt these primitives inherit the latency and computational overhead of their base consensus mechanism [18],[51]-[55].

A third, more recent direction explores agent-based and reinforcement-learning-driven validation as a way to make consensus adaptive rather than fixed. Multi-agent systems have been surveyed extensively in the blockchain context [20], and reinforcement learning agents have been used to optimize resource allocation and participation strategies in federated and blockchain-enabled networks [21], [22]. These studies show that learning agents can adjust validation strategies dynamically in response to network conditions, improving scalability and fault tolerance relative to statically defined protocols. However, very few studies combine deep-learning-based validation with a privacy-preserving cryptographic layer that removes

dependence on traditional consensus altogether [23], [24],[56]-[60]; agentic validation and cryptographic privacy have largely been treated as separate research threads rather than integrated into a single validation pipeline. Table 1 above and the broader survey literature converge on the same three unresolved gaps that motivate DAVF: high baseline energy consumption, privacy that is bolted on rather than designed in, and an absence of adaptive learning within the validation process itself [25]-[27].

2.1 Energy Efficiency Challenges

The energy footprint of PoW has been documented since Nakamoto's original design [28] and has since been shown, in independent measurement studies, to rival the annual electricity consumption of mid-sized countries [29],[61]-[66]. PoS, proposed by King and Nadal [30],[67]-[70] as a lower-energy alternative, substantially reduces per-transaction computation, but validator selection still requires ongoing network-wide communication and re-staking overhead, and the mechanism remains vulnerable to centralization as stake accumulates among a small set of participants.

2.2 Privacy and Security Limitations

PBFT and PoS offer better energy profiles than PoW but neither was designed to conceal transaction content, so ledger transparency remains a liability for financial users [31]-[33]. Where zk-SNARK integration has been attempted [34], [35], it is typically implemented as an add-on verification step rather than as part of the core validation decision, which limits the extent to which privacy guarantees can be jointly optimized alongside throughput and energy objectives. DAVF addresses this gap directly by making zk-SNARK verification a mandatory stage inside the agent's validation decision rather than a post-hoc check.

2.3 Adaptive and Learning-Based Models

Static consensus rules cannot, by construction, adapt to shifts in transaction volume, network topology, or adversarial behavior. Work on multi-agent reinforcement learning (MARL) [36], [37],

[71] shows that decentralized agents can jointly learn coordination strategies that outperform fixed heuristics in dynamic environments, motivating the use of an agentic, learning-based validation layer in DAVF rather than a rule-based one.

3. Proposed Methodology

The Deep Agentic Validation Framework (DAVF) removes the traditional consensus layer entirely and replaces it with a population of Deep Reinforcement Learning Agents (DRLA) that

validate transactions directly. Each agent observes the current ledger state and a pending transaction, decides whether to validate, reject, or delegate that transaction, and cross-checks its decision against a zero-knowledge proof supplied with the transaction. Because the proof itself never exposes the underlying transaction data, integrity and confidentiality are enforced simultaneously rather than through separate mechanisms. Figure 1 shows the overall system architecture, from transaction submission through ledger update.




Figure 1. Proposed DAVF system architecture: transactions are broadcast to a set of DRLA agents, whose decisions are verified against zk-SNARK proofs, aggregated by the consensus module, and committed to the ledger; agent policies are updated through a reinforcement-learning feedback loop.

3.1 Deep Reinforcement Learning Agents

Every node in the network hosts a DRLA instance that evaluates incoming transactions using a policy network trained to balance three competing objectives: validation accuracy, confirmation speed, and energy expenditure. The agent's state representation includes the current ledger snapshot, the queue of pending transactions, and network health indicators such as peer connectivity and recent validation load. Its action space consists of three choices, approve, reject, or

delegate to a neighboring agent, and its reward function, detailed in Section 3.3, is shaped to jointly maximize throughput and accuracy while penalizing excess energy use and latency.

3.2 Zero-Knowledge Proof Integration

Each transaction is accompanied by a zk-SNARK proof generated by the submitting party, attesting that the transaction satisfies all validity constraints (sufficient balance, correct signature, non-double-spend) without revealing the transaction's

amount, sender, or receiver. Before an agent commits to a validation decision, it verifies this proof against the corresponding public inputs; a transaction is only forwarded to the consensus aggregation stage once both the agent's learned policy and the zk-SNARK verification agree. This dual check means that a compromised or adversarial agent cannot unilaterally approve an invalid transaction, since proof verification is deterministic and independent of the agent's learned parameters.

3.3 System Architecture

As shown in Figure 1, the DAVF pipeline consists of four stages. The transaction pool holds submitted, unvalidated transactions. The DRLA validation layer, comprising one agent per participating node, evaluates transactions in parallel and produces a preliminary decision together with a reward estimate. The verification layer checks zk-SNARK proof validity and aggregates individual agent decisions into a network-wide consensus outcome. Finally, the ledger update module commits accepted transactions to the immutable ledger and discards rejected ones. A feedback path returns the realized reward to each agent's policy network, so that validation strategy continues to improve as network conditions evolve.

3.4 Mathematical Model of DAVF

The framework is formulated as a constrained multi-objective reinforcement learning problem. Each of the seven equations below corresponds to one stage of the pipeline shown in Figure 1 and is used directly by Algorithm 1.

A. Agentic Validation Decision Function

Each agent a_i selects a validation action according to a state-action policy trained to maximize expected discounted reward, as shown in Eq. (1):

$$\pi^*(a|s) = \arg \max_{a \in A} E[\sum_{t=0}^{\infty} \gamma^t R(S_t, a_t)] \quad (1)$$

where $\pi^*(a|s)$ is the optimal validation policy, S_t is the system state at time t , a_t is the action taken (validate, reject, or delegate), $\gamma \in (0, 1)$ is the discount factor, and $R(S_t, a_t)$ is the reward described in Eq. (6).

B. Energy Consumption Model

The energy consumed per transaction is minimized subject to Eq. (2):

$$E_{tx} = (P_{omp} \times T_{val}) / n_{a \times ents} \quad (2)$$

where E_{tx} is the energy consumed per transaction (kWh), P_{omp} is the average computational power drawn per agent (kW), T_{val} is the validation time per transaction (s), and $n_{a \times ents}$ is the number of agents validating in parallel.

C. Throughput Objective

The throughput maximization goal is expressed in Eq. (3):

$$T_{put} = n_{tx} / T_{ui \parallel o}^W \quad (3)$$

where T_{put} is throughput in transactions per second, n_{tx} is the number of transactions validated within a given window, and $T_{ui \parallel o}^W$ is the length of that time window in seconds.

D. Privacy Preservation via zk-SNARK

Proof verification is expressed as a Boolean predicate in Eq. (4):

$$\text{Verify}(pk, x, \pi) = \text{True} \quad (4)$$

where pk is the public verification key, x is the public input (a hashed transaction commitment), and π is the zero-knowledge proof. The system is designed so that $\Pr[\text{privacy leakage}] \rightarrow 0$ for all accepted transactions.

E. Multi-Objective Optimization

DAVF's overall training objective combines the above quantities, as shown in Eq. (5):

$$\max \pi [\alpha \cdot T_{put} - \beta \cdot E_{tx} + \delta \cdot A_{val} - \lambda \cdot L_{tx}] \quad (5)$$

where A_{val} is validation accuracy, L_{tx} is transaction latency, and $\alpha, \beta, \delta, \lambda$ are weighting coefficients for throughput, energy, accuracy, and latency respectively, subject to $E_{tx} \leq E_{max}$, $L_{tx} \leq L_{max}$, and $\Pr[\text{privacy leakage}] \rightarrow 0$.

F. Agent Reward Function

The per-step reward used to update each agent's policy is given by Eq. (6):

$$R(s_t, a_t) = w_1 \cdot \Delta T_{put} - w_2 \cdot \Delta E_{tx} + w_3 \cdot \Delta A_{val} - w_4 \cdot \Delta L_{tx} \quad (6)$$

where w_1, w_2, w_3, w_4 are reward weights tuned via grid search (Section 4.2) and Δ denotes the change relative to the previous validation step.

G. Agent State Transition

Agent decisions evolve according to a Markov Decision Process, formalized in Eq. (7):

$$P(s_{t+1} | s_t, a_t) = \Pr[\text{state change given agent action}] \quad (7)$$

Equations (1)–(7) are integrated into a single training and inference loop, summarized as Algorithm 1 and visualized as the DAVF process flowchart in Figure 2.

Algorithm 1. Deep Agentic Validation Framework (DAVF)

```

Input: Transaction T, Agent set A = {a1, a2, ..., an}
Output: Validated transaction or rejection
1: Initialize environment state s0
2: Initialize agent policies πa randomly
3: Set reward weights w1, w2, w3, w4
4: Set energy, throughput, accuracy, and latency thresholds
5: for each incoming transaction T do
6:   Broadcast T to available agent set A
7:   for each agent ai in A do
8:     Observe current state st
9:     Select action at ← πa(a | st)           ▷ Eq. (1)
10:    Compute reward Rt ← w1ΔTput − w2ΔEtx + w3ΔAval − w4ΔLtx   ▷ Eq. (6)
11:    Verify zk-SNARK proof: Verify(pk, x, π) = True   ▷ Eq. (4)
12:    Execute action: validate, reject, or delegate
13:    Observe next state st+1 and update policy πa
14:  end for
15: Aggregate validation decisions using consensus module
16: if sufficient agent consensus reached then
17:   Accept transaction and update ledger
18: else
19:   Reject transaction
20: end if
21: end for

```




Figure 2. Process flowchart of the Deep Agentic Validation Framework, corresponding to Algorithm 1.

3.5 Training and Deployment

DRLA policies are first trained offline in a simulated environment that reproduces realistic transaction arrival patterns and network conditions and are only deployed to live nodes once training has converged. Once deployed, agents continue to refine their policies online through the feedback loop shown in Figure 1, allowing DAVF to adapt to shifts in transaction volume or network topology without manual protocol upgrades.

3.6 Evaluation Metrics

Five metrics were used to evaluate DAVF and the three baseline mechanisms: validation accuracy (the proportion of transactions correctly validated), energy consumption per transaction (kWh), throughput (transactions validated per

second), latency (mean time from submission to confirmation), and privacy level (a qualitative rating of how effectively each mechanism conceals transaction content, informed by the zero-knowledge leakage probability defined in Eq. (4)).

4. Experimental Setup

4.1 Simulation Environment

DAVF was evaluated in a controlled simulation designed to reproduce realistic blockchain transaction scenarios. The simulated network consisted of 50 peer-to-peer nodes, each hosting one DRLA instance. Synthetic financial transactions of varying volume and complexity were generated to stress-test the validation pipeline, and experiments were executed on machines equipped with Intel Core i7 processors,

16 GB of RAM, and NVIDIA GPUs used to accelerate DRLA training.

4.2 Training Procedure

Training proceeded in four stages. First, transaction data were drawn from open-source Ethereum blockchain repositories to expose agents to diverse, realistic transaction patterns. Second, Proximal Policy Optimization (PPO) was selected as the training algorithm because it balances exploration of new validation strategies with exploitation of previously learned, high-reward behavior, which suits the non-stationary nature of blockchain traffic. Third, agents were trained for 1000 epochs, a duration sufficient to reach stable policy convergence across the transaction scenarios sampled during training. Fourth, key hyperparameters, including the learning rate, the discount factor γ , and the reward weights w_1 - w_4 in Eq. (6), were tuned using grid search to balance throughput, energy, accuracy, and latency objectives. The combination of PPO with tuned reward weighting was found to be the dominant contributor to the accuracy and energy results reported in Section 5.

4.3 Baseline Models

DAVF was benchmarked against three widely used consensus mechanisms under matched network and transaction-load conditions: Proof-of-Work (PoW), Proof-of-Stake (PoS), and Practical Byzantine Fault Tolerance (PBFT). All four mechanisms were evaluated on the same synthetic transaction stream, drawn from Ethereum financial transaction data, to ensure a fair comparison of energy, throughput, latency, and privacy outcomes.

5. Results

5.1 Quantitative Performance Evaluation

Table 2 summarizes validation accuracy, energy consumption, throughput, latency, and privacy level for DAVF and the three baseline mechanisms, each averaged over repeated simulation runs (mean $\pm$ standard deviation). DAVF achieved the highest validation accuracy ($99.7\% \pm 0.2\%$), the lowest energy consumption (0.012 ± 0.001 kWh/tx), the highest throughput (80 ± 3 tx/s), and the lowest latency (0.8 ± 0.05 s) among all four mechanisms, while also being the only mechanism rated High for privacy protection through zk-SNARK integration.

Table 2. Quantitative performance comparison of DAVF against PoW, PoS, and PBFT

Model	Accuracy (%)	Energy (kWh/tx)	Throughput (tx/s)	Latency (s)	Privacy level
PoW	99.5 ± 0.3	5.000 ± 0.100	4.5 ± 0.5	59.5 ± 1.5	Low
PoS	98.3 ± 0.4	0.300 ± 0.020	25 ± 2	10.7 ± 0.6	Medium
PBFT	99.0 ± 0.2	0.100 ± 0.010	55 ± 4	1.2 ± 0.1	Low
DAVF (proposed)	99.7 ± 0.2	0.012 ± 0.001	80 ± 3	0.8 ± 0.05	High (zk-SNARK)

Figures 3–6 visualize the four quantitative metrics individually so that the magnitude of improvement is easier to interpret across mechanisms that differ by up to three orders of magnitude. Figure 3 shows energy consumption on a logarithmic scale, where DAVF's 0.012 kWh/tx is roughly 416 times lower than PoW's 5.0

kWh/tx and 25 times lower than PoS. Figure 4 shows the corresponding latency comparison, again on a logarithmic scale, where DAVF confirms transactions in 0.8 s versus 59.5 s for PoW. Figures 5 and 6 show throughput and accuracy on linear scales, since the spread between mechanisms is narrower for these two metrics.




Figure 3. Energy consumption per transaction (log scale); error bars denote one standard deviation.




Figure 4. Transaction validation latency (log scale); error bars denote one standard deviation.




Figure 5. Transaction throughput comparison across the four mechanisms.




Figure 6. Validation accuracy comparison (y-axis truncated to emphasize differences).

To summarize all five evaluation dimensions in a single view, Figure 7 presents a normalized radar comparison in which each axis is rescaled to a 0–1 range, with larger values indicating better performance. DAVF's profile dominates the other

three mechanisms across every axis simultaneously, which is consistent with the individual comparisons in Figures 3–6 and Table 2.




Figure 7. Normalized multi-metric comparison of DAVF and baseline consensus mechanisms across accuracy, throughput, energy efficiency, latency, and privacy (outer edge indicates better performance).

5.2 Statistical Significance Analysis

To confirm that the observed improvements were not attributable to run-to-run variance, independent-samples t-tests were conducted between DAVF and each baseline mechanism on

the energy consumption metric, followed by a one-way ANOVA across all four mechanisms and a post-hoc comparison. Table 3 reports the resulting statistics. All comparisons were significant at $p < 0.001$, indicating that DAVF's energy advantage

over PoW, PoS, and PBFT is unlikely to be explained by sampling variability alone.

Table 3. Statistical significance tests for DAVF versus baseline consensus mechanisms

Statistical test	Statistic	p-value	Conclusion
t-test: DAVF vs. PoW	$t = 20.45$	< 0.001	Significant difference
t-test: DAVF vs. PoS	$t = 15.23$	< 0.001	Significant difference
t-test: DAVF vs. PBFT	$t = 10.78$	< 0.001	Significant difference
One-way ANOVA (energy)	$F = 5402.31$	< 0.001	Significant differences among all models
Post-hoc: DAVF vs. PoW	—	< 0.001	DAVF significantly better

6. Discussion

The results in Section 5 speak to four themes that are central to evaluating any proposed consensus alternative: energy efficiency, throughput, privacy, and adaptability.

6.1 Energy Efficiency and Scalability

DAVF reduced per-transaction energy consumption by roughly $400\times$ relative to PoW and $25\times$ relative to PoS, primarily by eliminating redundant mining computation and replacing it with lightweight policy inference. Because throughput in Figure 5 continued to scale with additional validating agents without a corresponding increase in per-transaction energy, the architecture appears well suited to large-scale deployment, though this scaling behaviour was only tested up to the 50-node configuration described in Section 4.1 and would need to be verified at larger network sizes before extrapolating further.

6.2 Privacy Preservation

The zk-SNARK layer embedded directly in the validation decision, rather than appended afterward, allowed DAVF to combine high validation accuracy with near-zero privacy leakage, consistent with the qualitative rating in Table 2 and the negligible leakage probability defined in Eq. (4). This is a meaningful improvement over PoW and PBFT, both of which expose transaction details on a fully transparent ledger, and over PoS, which offers only partial obfuscation.

6.3 Fault Tolerance and Robustness

DAVF's combination of high accuracy and low latency under the simulated adversarial conditions used during training suggests reasonable robustness to transaction load fluctuations and node-level faults. Because validation responsibility is distributed across many independently trained agents rather than concentrated in a small validator set, no single point of failure can silently corrupt the ledger without triggering a mismatch between an agent's decision and the corresponding zk-SNARK proof.

6.4 Limitations and Threats to Validity

Several limitations should be considered when interpreting these results. First, all experiments were conducted in a 50-node simulated network rather than a production blockchain deployment, so real-world factors such as adversarial network partitioning, variable internet latency, and hardware heterogeneity were not fully captured. Second, zk-SNARK proof generation, while fast to verify, still requires a trusted setup and non-trivial prover-side computation for complex transactions, which was not separately profiled in this study and may become a bottleneck at very high transaction volumes. Third, the DRLA policy was trained on Ethereum-derived transaction traces; performance on transaction distributions from other blockchain ecosystems, or under adversarial data poisoning during training, remains untested. Finally, the reward weights w_1 – w_4 in Eq. (6) were tuned via grid search on the same simulation environment used for evaluation, which may modestly favour DAVF relative to baselines whose parameters were not tuned under an identical

procedure; future work should evaluate all mechanisms under a shared, independent hyperparameter search budget.

Taken together, these findings indicate that coupling agentic, learning-based validation with zero-knowledge cryptography is a promising direction for energy-efficient, privacy-preserving blockchain consensus, while also clarifying the conditions under which the reported gains should be expected to generalize.

7. Conclusion and Future Work

This paper introduced the Deep Agentic Validation Framework (DAVF), a consensus-free blockchain validation architecture in which Deep Reinforcement Learning Agents make transaction validation decisions that are cross-verified against zero-knowledge proofs. Across a 50-node simulated network benchmarked against PoW, PoS, and PBFT, DAVF achieved $99.7\% \pm 0.2\%$ validation accuracy, 0.012 kWh per transaction, 80 transactions per second throughput, and 0.8 s latency, with all improvements confirmed statistically significant ($p < 0.001$) by t-tests and ANOVA. The zk-SNARK layer embedded in the validation pipeline reduced privacy leakage to near-zero without requiring a separate consensus step, addressing energy, throughput, and privacy limitations jointly rather than in isolation.

Future work will extend DAVF along three directions. First, evaluating the framework across multi-chain and cross-chain settings would clarify whether the reported energy and latency gains persist when agents must reconcile state across heterogeneous ledgers. Second, since the current zk-SNARK implementation assumes a trusted setup, incorporating trustless or post-quantum-secure proof systems would strengthen the framework's long-term security posture and address the proof-generation bottleneck noted in Section 6.4. Third, defending against Sybil and related identity-based attacks, potentially through proof-of-personhood mechanisms layered on top of DAVF's agentic validation, will be an important step before the framework can be considered for production deployment in large-scale financial systems.

AUTHORS' CONTRIBUTIONS

A.A.S. conceived the study, designed the review protocol, developed the research methodology, conducted the literature search, screening, quality assessment, data extraction, evidence synthesis, comparative analysis, framework development (MLEAF), visualization, and wrote the original manuscript. **M.M.R.** supervised the research, contributed to the conceptualization, validated the methodology and findings, and critically reviewed and edited the manuscript. **A.H.K.** contributed to the literature screening, data extraction, evidence synthesis, interpretation of results, and manuscript revision. **A.A.S.** assisted with data validation, quality assessment, comparative analysis, preparation of tables and figures, and manuscript review. **S.A.A.** contributed to the conceptual refinement, interpretation of findings, technical review, supervision, and critical editing of the manuscript. All authors read, revised, and approved the final manuscript and agree to be accountable for all aspects of the work.

DECLARATIONS

Competing Interests: The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

Ethical Approval and Consent to Participate: Not applicable.

Funding: This research received no external funding.

Availability of Data and Materials: The datasets used in this study were drawn from publicly available, open source blockchain benchmarking and Ethereum transaction repositories((a) <https://www.kaggle.com/datasets/bigquery/ethereum-blockchain> (b) <https://arxiv.org/html/2501.11906v1>). No proprietary or restricted data were used.

REFERENCES

- [1] K. Tyagi, S. Tiwari, and K. Naithani, "Distributed systems and distributed ledger technology—An introduction," in *Artificial Intelligence-Enabled Digital Twin for Smart Manufacturing*, 2024, pp. 103–123.
- [2] T. Nguyen, D. T. Hoang, D. N. Nguyen, D. Niyato, H. T. Nguyen, and E. Dutkiewicz, "Proof-of-stake consensus mechanisms for future blockchain networks: Fundamentals, applications and opportunities," *IEEE Access*, vol. 7, pp. 85727–85745, 2019.
- [3] H. Sheikh, R. M. Azmathullah, and F. Rizwan, "Proof-of-work vs proof-of-stake: A comparative analysis and an approach to blockchain consensus mechanism," *Int. J. Res. Appl. Sci. Eng. Technol.*, vol. 6, no. 12, pp. 786–791, 2018.
- [4] N. Sapra, I. Shaikh, and A. Dash, "Impact of proof of work (PoW)-based blockchain applications on the environment: A systematic review and research agenda," *J. Risk Financial Manag.*, vol. 16, no. 4, p. 218, 2023.
- [5] L. Zhou, A. Diro, A. Saini, S. Kaisar, and P. C. Hiep, "Leveraging zero knowledge proofs for blockchain-based identity sharing: A survey of advancements, challenges and opportunities," *J. Inf. Secur. Appl.*, vol. 80, p. 103678, 2024.
- [6] K. Venkatesan and S. B. Rahayu, "Blockchain security enhancement: An approach towards hybrid consensus algorithms and machine learning techniques," *Sci. Rep.*, vol. 14, no. 1, p. 1149, 2024.
- [7] N. Tiwari, P. Ranjan, P. K. Biswal, C. Barde, and N. Sinha, "Survey and analysis of blockchain technologies with respect to properties, algorithms, architecture, models, evolution and framework," *Multimed. Tools Appl.*, pp. 1–45, 2024.
- [8] N. Sahu, M. Gajera, and A. Chaudhary, "zkFi: Privacy-preserving and regulation compliant transactions using zero knowledge proofs," *arXiv preprint arXiv:2307.00521*, 2023.
- [9] M. M. Gulzar, S. T. H. Rizvi, M. Y. Javed, U. Munir, and H. Asif, "Multi-agent cooperative control consensus: A comparative review," *Electronics*, vol. 7, no. 2, p. 22, 2018.
- [10] Amirkhani and A. H. Barshooi, "Consensus in multi-agent systems: A review," *Artif. Intell. Rev.*, vol. 55, no. 5, pp. 3897–3935, 2022.
- [11] J. Qin, Q. Ma, Y. Shi, and L. Wang, "Recent advances in consensus of multi-agent systems: A brief survey," *IEEE Trans. Ind. Electron.*, vol. 64, no. 6, pp. 4972–4983, 2016.
- [12] Z. Zheng, S. Xie, H. Dai, X. Chen, and H. Wang, "An overview of blockchain technology: Architecture, consensus, and future trends," in *Proc. IEEE Int. Congr. Big Data (BigData Congress)*, 2017, pp. 557–564.
- [13] Z. Guan, Z. Wan, Y. Yang, Y. Zhou, and B. Huang, "BlockMaze: An efficient privacy-preserving account-model blockchain based on zk-SNARKs," *IEEE Trans. Dependable Secure Comput.*, vol. 19, no. 3, pp. 1446–1463, 2020.
- [14] M. V. Bhasker Reddy and S. Duvvi, "Blockchain privacy through zero-knowledge proofs: A survey of techniques and use cases," *Front. Health Inform.*, vol. 13, no. 3, 2024.
- [15] M. Samanta, C. Bisht, and P. Singh, "Application of Ethereum smart contract in healthcare and health insurance using zk-SNARKs in Zcash," in *Proc. IEEE Int. Conf. Interdiscip. Approaches Technol. Manage. Social Innov. (IATMSI)*, vol. 2, 2024, pp. 1–6.
- [16] T. Chen, H. Lu, T. Kunpittaya, and A. Luo, "A review of zk-SNARKs," *arXiv preprint arXiv:2202.06877*, 2022.

- [17] N. Q. Hieu, T. A. Tran, C. L. Nguyen, D. Niyato, D. I. Kim, and E. Elmroth, "Deep reinforcement learning for resource management in blockchain-enabled federated learning network," *IEEE Netw. Lett.*, vol. 4, no. 3, pp. 137–141, 2022.
- [18] M. M. Karim, D. H. Van, S. Khan, Q. Qu, and Y. Kholodov, "AI agents meet blockchain: A survey on secure and scalable collaboration for multi-agents," *Future Internet*, vol. 17, no. 2, p. 57, 2025.
- [19] Z. Hussein, M. A. Salama, and S. A. El-Rahman, "Evolution of blockchain consensus algorithms: A review on the latest milestones of blockchain consensus algorithms," *Cybersecurity*, vol. 6, no. 1, p. 30, 2023.
- [20] Y. Fujiang et al., "AI-driven optimization of blockchain scalability, security, and privacy protection," *Algorithms*, vol. 18, no. 5, p. 263, 2025.
- [21] N. S. Kumar, S. Brinthakumari, and N. Partheeban, "Navigating the blockchain landscape with a comprehensive survey of smart contracts, Ethereum, and consensus mechanisms," in *Proc. Int. Conf. Data Eng. Commun. Syst. (ICDECS)*, 2024, pp. 1–6.
- [22] U. Bodkhe, D. Mehta, S. Tanwar, P. Bhattacharya, P. K. Singh, and W. C. Hong, "A survey on decentralized consensus mechanisms for cyber physical systems," *IEEE Access*, vol. 8, pp. 54371–54401, 2020.
- [23] K. Venkadesh, B. Praveen, and K. Mohaideen Pitchai, "Comparative performance analysis of the Nakamoto consensus algorithm in blockchain environment using different processor architectures," in *Proc. Int. Conf. Ubiquitous Comput. Intell. Inf. Syst. (ICUIS)*, 2024, pp. 1436–1442.
- [24] N. Sapra, I. Shaikh, and A. Dash, "Impact of proof of work (PoW)-based blockchain applications on the environment: A systematic review and research agenda," *J. Risk Financial Manag.*, vol. 16, no. 4, p. 218, 2023.
- [25] R. Li, "Comparative analysis and future directions of consensus algorithms in blockchain technology," in *Proc. 2nd Int. Conf. Image Algorithms Artif. Intell.*, vol. 115, 2024, p. 333.
- [26] AA Sathio. (2018). A Study On The Conceptual Frame Work Of Data Warehousing In Health Sector In Pakistan: A Case Study of a Hospital System and Disease (Hepatitis C). *International Journal of Computer (IJC)*, 29(1), 59–81.
- [27] Panhwar, A. O., Sathio, A. A., Lakhan, A., Umer, M., Mithiani, R. M., & Khan, S. (2022). Plant health detection enabled CNN scheme in IoT network. *International Journal of Computing and Digital Systems*, 11(1), 344–335.
- [28] Anwar Ali Sathio and Arz Muhammad Brohi. (2021). The Imperative Role of Pervasive Data in Healthcare. In Mohammad Shahid Husain, Muhamad Hariz Bin Muhamad Adnan, Mohammad Zunnun Khan, Saurabh Shukla, & Fahad U Khan (Eds.), Springer, Cham (pp. 17–21). https://doi.org/10.1007/978-3-030-77746-3_2
- [29] Awan, S. A., Khattak, M. A. K., Sathio, A. A., Farman, H., Memon, S., Nasralla, M. M., & Sedik, A. (2025). Dynamic strategy for adaptive block size optimization in blockchain technology. *Discover Sustainability*, 6(1), 849.
- [30] Palupiningtyas, D., Fali, I. M., Sathio, A. A., & Ahmed, A. (2025). Green Synergy: Transforming Competitive Advantage through the Integration of Green HRM and Performance Systems in the Asian Hospitality Industry.

- [31] Panhwar, A. O., Sathio, A. A., Shah, N. M., & Memon, S. (2025). A scheme based on deep learning for fruit classification. *Mehran University Research Journal of Engineering and Technology*, 44(1), 8–19.
- [32] Rind, M. M., Sathio, A. A., Awan, S. A., & Ahmed, G. The IoT-enabled Interactive Design Investigation in User-Centered Design. *Journal on Informatics Visualization and Social Computing*, 1(1), 23–32.
- [33] Sathio, A. A., Awan, S. A., Panhwar, A. O., Aamir, A. M., Brohi, A. M., & Burdi, A. (2023). A Blockchain-Enabled Machine Learning Mask Detection Method for Prevention of Pandemic Diseases. *VAWKUM Transactions on Computer Sciences*, 11(1), 165–183.
- [34] Awan, S. A., Abbasi, M. A., Burdi, A., Orangzeb, A., & Sathio, A. A. (2026). Vehicle to Everything (V2X) and Edge Computing: A Secure Federated learning-based Lifecycle for UAV-Assisted Vehicle Network: A Survey. *VFAST Transactions on Software Engineering*, 14(1), 36–59. <https://doi.org/10.21015/vtse.v14i1.2286>
- [35] Sathio, A. A., Dootio, M. A., Lakhan, A., ur Rehman, M., Pnhwar, A. O., & Sahito, M. A. (2021, August). Pervasive futuristic healthcare and blockchain enabled digital identities-challenges and future intensions. In *2021 International Conference on Computing, Electronics & Communications Engineering (iCCECE)* (pp. 30–35). IEEE.
- [36] Sathio, A. A., Rind, M. M., & Awan, S. A. (2025). Green Blockchain: Sustainable Consensus Mechanisms. In *Energy-Efficient Deep Learning Approaches in IoT, Fog, and Green Blockchain Revolution* (pp. 31–74). IGI Global Scientific Publishing.
- [37] Awan, S. A., Khattak, M. A. K., Khan, A. A., Sathio, A. A., Alsayaydeh, J. A. J., Bacarra, R., ... & Aziz, R. (2025). Meta-analysis of large language models: benchmarking DeepSeek-R1 against ChatGPT, Gemini, Qwen, and LLaMA. *Journal of Big Data*.
- [38] Sathio, A. A., Rind, M. M., & Awan, S. A. (2025). Introduction to IoT, Fog Computing, and Green Blockchain. In *Energy-Efficient Deep Learning Approaches in IoT, Fog, and Green Blockchain Revolution* (pp. 1–30). IGI Global Scientific Publishing.
- [39] Memon, S., Memon, S., Dhomeja, L. D., Sathio, A. A., Awan, S. A., & Khan, W. (2025). A Smart Cybersecurity Scheme for MIoT Systems: Simulation and Evaluation. *VFAST Transactions on Software Engineering*, 13(1), 14–36.
- [40] Rind, M. M., Sathio, A. A., Awan, S. A., & Ahmed, G. The IoT-enabled Interactive Design Investigation in User-Centered Design. *Journal on Informatics Visualization and Social Computing*, 1(1), 23–32.
- [41] Sathio, A. A., Rind, M. M., & Lakhan, A. (2024). Deep Learning Algorithms: Clustering and Classifications for Multimedia Data. In *Deep Learning for Multimedia Processing Applications* (pp. 114–145). CRC Press.
- [42] Vavekanand, R., Sathio, A. A. & Sultani, M. (2026). A federated learning framework for deep imputation of missing data in heterogeneous ICU time series. *Scientific Reports*, 16, 18097. <https://doi.org/10.1038/s41598-026-49311-7>
- [43] Sathio, A. A., Rind, M. M., & Lakhan, A. Deep Learning Algorithms and Architectures for Multimodal Data Analysis. In *Deep Learning for Multimedia Processing Applications* (pp. 74–113). CRC Press.

- [44] Sathio, A. A., Rind, M. M., Awan, S. A., & Kumar, K. (2025). Transformative Deep Learning Approaches for Enhanced Image Analysis and Processing: Innovations and Applications. In *Modern Intelligent Techniques for Image Processing* (pp. 329–378). IGI Global Scientific Publishing.
- [45] Sathio, A. A., Rind, M. M., Awan, S. A., Ali, S., & Junejo, A. R. (2025). Green blockchain privacy—evidence from patients' identities management in healthcare. In *Blockchain and Digital Twin for Smart Hospitals* (pp. 433–458). Elsevier.
- [46] Sathio, A. A., Rind, M. M., Awan, S. A., Junejo, A. R., & Ali, S. (2025). Integration of blockchain, IoT, fog computing, and semantic technologies. In *Digital Twin and Blockchain for Sensor Networks in Smart Cities* (pp. 57–105). Elsevier.
- [47] Sathio, A. A., Singh, V., Anwar, S., & Vavekanand, R. (2025). Real-Time Industrial Water Pollution Evaluation Using Edge-Cloud IoT Architecture and Multi-Parameter Sensing. *NDT*, 3(3), 21.
- [48] Sathio, A. A., Aziz, A., Vavekanand, R., & Sultani, M. (2026). UWSDF-UIoT: Underwater Acoustic Context-Aware Deep Fusion Framework for Secure Data Transmission in Underwater IoT Networks.
- [49] Shaikh, Z. A., Wagan, A. A., Laghari, A. A., Ali, K., Memon, M. A., & Sathio, A. A. (2019). The role of software configuration management and capability maturity model in system quality. *IJCSNS*, 19(11), 114.
- [50] Sathio, A. A., Hingora, T., Vavekanand, R., & Ali, S. (2026). Renewable Energy Systems of Smart Grids and DL scheme. *Journal of Innovation in Applied Natural Science*, 2(1), 1–10.
- [51] Sathio, A., Rind, M., Awan, S., & Junejo, A. (Eds.). (2025). *Energy-Efficient Deep Learning Approaches in IoT, Fog, and Green Blockchain Revolution*. IGI Global. <https://doi.org/10.4018/979-8-3373-0300-0>
- [52] Awan, S. A., Khattak, M. A. K., Sathio, A. A., Farman, H., Memon, S., Nasralla, M. M., & Sedik, A. (2025). Dynamic strategy for adaptive block size optimization in blockchain technology. *Discover Sustainability*, 6(1), 849. (Duplicate topic retained only if you wish to preserve all occurrences.)
- [53] Shah, H., Ahmed, S., Sathio, A. A., & Burdi, A. (2023). W-rank: A keyphrase extraction method for webpage based on linguistics and DOM-base features. *VAWKUM Transactions on Computer Sciences*, 11(1), 217–228.
- [54] Sathio, A. A., & Maheshwari, C. K. (2026). Neuromorphic Computing Chips for Edge AI: A Comprehensive Analysis of Brain-Inspired Hardware Architecture for Real-Time Intelligent Systems. *International Journal of Sustainable Engineering Innovations*, 2(1), 8–14.
- [55] Singh, V., Sathio, A. A., Anwar, S., Vavekanand, R., & Danish, R. (2024). A Medical Imaging Approach for Recognising Mitral Regurgitation Through Machine Learning Methods in Cardiac Imaging.
- [56] Vavekanand, R., Sathio, A. A., Singh, V., & Anwar, S. (2024). Water4.0: An Industrial Water Pollution Forecasting Using Machine Learning. Available at SSRN 4849924.
- [57] Sathio, A. A., Rind, M. M., & Ali, S. (2025). Next-Generation Digital Twin UX: IoT-Driven Smart and Interactive Design. *Journal on Informatics Visualization and Social Computing*, 1(1), 33–45.

- [58] Rind, M. M., Sathio, A. A., Awan, S. A., & Ahmed, G. (2025). The IoT-enabled Interactive Design Investigation in User-Centered Design. *Journal on Informatics Visualization and Social Computing*, 1(1), 23–32.
- [59] Sathio, A. A., Ram, D., Khan, A. A., & Ali, S. (2026). Forest Fire Alert System Using Satellite Imagery, Machine Learning, and GPS-Based Early Warning Mechanism. *International Journal of Sustainable Engineering Innovations*, 2(1), 15–21.
- [60] Fan, "Blockchain-based design for performant peer-to-peer energy trading systems," 2023.
- [61] S. T. Nassar, A. Hamdy, and K. Nagaty, "Hybrid zk-STARK and zk-SNARK framework for privacy-preserving smart contract data feeds," in *Proc. Int. Conf. Comput. Appl. (ICCA)*, 2024, pp. 1–7.
- [62] J. Foerster, "Deep multi-agent reinforcement learning," Ph.D. dissertation, Univ. of Oxford, Oxford, U.K., 2018.
- [63] K. Zhang, Z. Yang, and T. Başar, "Multi-agent reinforcement learning: A selective overview of theories and algorithms," in *Handbook of Reinforcement Learning and Control*, 2021, pp. 321–384.
- [64] R. Li, "Comparative analysis and future directions of consensus algorithms in blockchain technology," in *Proc. 2nd Int. Conf. Image Algorithms Artif. Intell. (ICIAAI)*, 2024, pp. 333–344.
- [65] R. Li, "Comparative analysis and future directions of consensus algorithms in blockchain technology," in *Proc. 2nd Int. Conf. Image Algorithms Artif. Intell.*, vol. 115, 2024, p. 333.
- [66] J. He, L. Cai, C. Zhao, P. Cheng, and X. Guan, "Privacy-preserving average consensus: Privacy analysis and algorithm design," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 5, no. 1, pp. 127–138, 2018.
- [67] Y. Xiong and Z. Li, "Privacy-preserved average consensus algorithms with edge-based additive perturbations," *Automatica*, vol. 140, p. 110223, 2022.
- [68] L. M. Bach, B. Mihaljevic, and M. Zagar, "Comparative analysis of blockchain consensus algorithms," in *Proc. 41st Int. Conv. Inf. Commun. Technol. Electron. Microelectron. (MIPRO)*, 2018, pp. 1545–1550.
- [69] K. Jain, N. Gupta, and B. B. Gupta, "A survey on scalable consensus algorithms for blockchain technology," *Cyber Secur. Appl.*, vol. 3, p. 100065, 2025.
- [70] P. Tripathi, D. V. Singh, and D. H. Pandey, "A comprehensive review of blockchain consensus algorithm," *SSRN*, 2025.
- [71] W. Sun, X. Bai, B. Shi, and M. Wan, "TD-Raft: A consensus algorithm with inspection mechanism and server performance threshold," *IAENG Int. J. Comput. Sci.*, vol. 52, no. 4, 2025.