World Journal of Advanced Research and Reviews
eISSN: 2581-9615 CODEN (USA): WJARAI
Cross Ref DOI: 10.30574/wjarr
Journal homepage: https://wjarr.com/

(REVIEW ARTICLE)

Check for updates

Energy-aware blockchain consensus enhanced by graph neural networks for sustainable, scalable transaction verification across heterogeneous IoT networks

Oyegoke Oyebode *

Visa Inc. USA.

World Journal of Advanced Research and Reviews, 2023, 20(03), 2354-2373

Publication history: Received on 20 November 2023; revised on 27 December 2023; accepted on 29 December 2023

Article DOI: https://doi.org/10.30574/wjarr.2023.20.3.2678

Abstract

The exponential growth of heterogeneous Internet of Things (IoT) networks has amplified demands for secure, scalable, and sustainable transaction verification mechanisms. Traditional blockchain consensus protocols, such as Proof-of-Work (PoW), offer robust security but impose prohibitive energy costs, limiting their viability for resource-constrained IoT environments. Proof-of-Stake (Po's) and lightweight consensus schemes improve efficiency but often compromise scalability or fairness. To address this trade-off, this study introduces an energy-aware blockchain consensus framework enhanced by graph neural networks (GNNs) for sustainable, scalable verification across heterogeneous IoT ecosystems. In this approach, GNNs are applied to dynamically model IoT device interconnections, enabling efficient clustering, adaptive leader election, and optimized consensus pathways. By learning the structural and temporal patterns of IoT networks, GNNs reduce redundant computations and allocate verification tasks intelligently, minimizing energy consumption while maintaining security. The consensus framework integrates energy profiling of devices with predictive workload balancing, ensuring equitable participation across diverse hardware capacities. Blockchain provides the foundation for immutable, decentralized trust, while the GNN-enhanced consensus mechanism improves throughput, latency, and energy efficiency in large-scale deployments. Simulation studies of smart grids, industrial IoT, and urban sensor networks demonstrate measurable improvements in energy savings, scalability, and fault tolerance. The proposed architecture contributes to the vision of sustainable blockchain systems that can operate effectively in energy-sensitive, heterogeneous IoT contexts. By fusing blockchain's decentralized trust with GNN-based intelligence, the framework offers a pathway toward greener, more scalable transaction verification tailored for next-generation IoT infrastructures.

Keywords: Energy-Aware Blockchain; Graph Neural Networks; Sustainable Consensus; Iott Scalability; Transaction Verification; Heterogeneous Networks

1. Introduction

1.1. Background: Growth of IoT and decentralized trust

The Internet of Things (IoT) has emerged as a transformative paradigm, connecting billions of devices across domains such as healthcare, energy, logistics, and smart cities. These devices generate and exchange vast volumes of data, enabling automation, predictive analytics, and responsive infrastructures [1]. However, the distributed nature of IoT also creates new trust challenges. Traditional centralized models of authentication and control are ill-suited for ecosystems where devices operate autonomously and often with minimal human oversight [2].

Blockchain technology has been proposed as a trust-enabling infrastructure for IoT. Its distributed ledger provides immutability, provenance, and consensus-driven validation of device transactions [3]. By eliminating reliance on a

* Corresponding author: Oyegoke Oyebode

single authority, blockchain aligns with the decentralized character of IoT networks. For example, in supply chain applications, blockchain can ensure that data recorded by sensors is tamper-resistant and auditable, improving accountability across stakeholders [4].

The convergence of IoT and blockchain has therefore attracted significant research interest. On the one hand, IoT requires secure, scalable, and verifiable data management. On the other, blockchain benefits from the proliferation of edge devices capable of participating in distributed consensus [5]. Yet this convergence introduces technical and sustainability challenges that must be carefully addressed. These include the energy intensity of consensus protocols and the computational limitations of IoT devices. Thus, while blockchain offers a promising framework for decentralized trust, new mechanisms are needed to ensure it remains efficient and adaptable for large-scale IoT deployments [6].

1.2. Energy and scalability challenges in blockchain-based IoT

Despite its promise, blockchain integration into IoT raises significant concerns regarding energy consumption and scalability. Traditional consensus mechanisms, such as Proof-of-Work (PoW), are computationally intensive and consume large amounts of energy, making them unsuitable for resource-constrained IoT environments [3]. Lightweight devices, such as sensors and wearables, lack the processing and battery capacity to sustain PoW operations. This mismatch creates barriers to widespread adoption in IoT ecosystems.

Scalability compounds these concerns. IoT networks may involve thousands or even millions of devices producing continuous streams of data. Consensus protocols must validate and record these transactions efficiently, yet many blockchains struggle with throughput limitations. For instance, while some networks process only a few dozen transactions per second, IoT applications in smart grids or industrial automation may demand near-real-time processing [6]. The discrepancy between IoT data intensity and blockchain throughput highlights a structural limitation.

Furthermore, latency introduced by consensus delays can hinder IoT applications that rely on rapid responsiveness, such as autonomous vehicles or emergency healthcare monitoring. Without addressing scalability, blockchain-enabled IoT risks undermining the very efficiencies it aims to deliver.

These energy and scalability challenges indicate the need for adaptive consensus protocols tailored to IoT's constraints. Innovative approaches must balance the immutability and trust guarantees of blockchain with the efficiency requirements of large-scale device networks. Researchers are increasingly exploring solutions that leverage advanced computational techniques, such as graph-based optimization, to reduce overhead while maintaining reliability [1].

1.3. Role of graph neural networks (GNNs) in optimizing consensus

Graph neural networks (GNNs) offer a promising pathway for addressing energy and scalability challenges in blockchain-enabled IoT systems. IoT networks can be naturally modeled as graphs, where nodes represent devices and edges capture communication or transactional relationships [7]. GNNs leverage this structure to learn representations that optimize consensus dynamics, identifying efficient validation pathways while minimizing redundant computations.

By analyzing network topology, GNNs can prioritize nodes with stronger connectivity or higher trust scores, reducing the number of participants required for consensus at any given time [4]. This targeted participation lowers energy consumption and accelerates validation, making consensus more suitable for IoT environments. Moreover, GNNs adapt dynamically as device connections evolve, maintaining efficiency even in highly dynamic networks such as vehicular IoT or smart energy grids [5].

Importantly, GNN-enhanced consensus maintains blockchain's core guarantees of immutability and verifiability. Instead of undermining trust, optimization strengthens it by ensuring resources are allocated more intelligently. As such, GNNs represent a bridge between machine learning advances and distributed systems design, aligning computational efficiency with decentralized trust. Their role in blockchain-IoT convergence highlights the potential of hybrid approaches to overcome the limitations of conventional protocols [2].

1.4. Research objectives and scope

This study investigates how graph neural networks can enhance consensus mechanisms for blockchain-enabled IoT systems by addressing energy and scalability challenges. The objectives are threefold: (i) to analyze the limitations of existing blockchain consensus protocols in IoT contexts, (ii) to demonstrate how GNNs can optimize resource allocation and validation processes, and (iii) to evaluate implications for secure, sustainable IoT deployments [6]. The scope focuses on IoT domains requiring high trust and responsiveness, including healthcare, energy, and transportation.

Transitioning from these general challenges, the discussion now turns to the proposed energy-aware, GNN-enhanced consensus model for decentralized IoT systems [3].

2. IOT ecosystems and blockchain integration

2.1. Characteristics of heterogeneous IoT networks

IoT networks are inherently heterogeneous, encompassing a wide range of devices with varying computational, storage, and communication capabilities. At one end of the spectrum are lightweight sensors and wearable devices, constrained by limited energy and processing power. At the other are more capable nodes, such as gateways and edge servers, which provide higher bandwidth and computational resources [9]. This diversity creates challenges in maintaining reliable connectivity and consistent performance across the network.

The heterogeneity extends to communication protocols and data formats. IoT ecosystems rely on multiple wireless technologies Wi-Fi, Bluetooth, Zigbee, Lora WAN each optimized for different coverage and energy trade-offs [7]. As devices interact through these protocols, interoperability becomes a central concern. Without standardized frameworks, heterogeneous networks struggle to support seamless integration, limiting their effectiveness in distributed applications.

Security vulnerabilities also stem from heterogeneity. Resource-constrained devices often lack strong cryptographic protections, leaving them exposed to attacks such as spoofing, eavesdropping, or denial of service [10]. When such devices participate in distributed infrastructures, they can become weak links that compromise system-wide security.

As illustrated in Figure 1, IoT heterogeneity complicates blockchain integration by introducing asymmetries in device capacity, protocol compatibility, and trustworthiness. Forecasting, decision-making, and consensus must therefore account for diverse device characteristics to ensure equitable participation and overall system resilience [6].

2.2. Traditional consensus mechanisms and limitations (PoW, POS, PBFT)

Consensus mechanisms form the backbone of blockchain systems, ensuring that distributed participants agree on a common ledger state. In IoT contexts, the most widely discussed consensus mechanisms include Proof-of-Work (PoW), Proof-of-Stake (POS), and Practical Byzantine Fault Tolerance (PBFT). Each offers distinct advantages but also reveals limitations when applied to large-scale, heterogeneous networks [8].

PoW achieves strong security through computational puzzles, making it highly tamper-resistant. However, this robustness comes at the cost of energy consumption, rendering it unsuitable for IoT devices with limited processing and power budgets [11]. Lightweight sensors cannot perform the intensive computations required, making PoW impractical for large deployments.

PoS reduces energy consumption by assigning validation rights based on the stake held by participants. While more efficient, PoS risks centralization, as wealthier entities can accumulate disproportionate influence [13]. In IoT ecosystems, where devices are not inherently tied to financial stakes, PoS raises questions about fairness and inclusivity.

PBFT focuses on achieving consensus in permissioned environments through majority agreement among known participants. It offers high throughput and energy efficiency but suffers from scalability limitations. Communication overhead increases quadratically as the number of nodes grows, making PBFT less effective in massive IoT deployments involving thousands of devices [7].

As shown in Figure 1, these limitations highlight a core dilemma: traditional consensus protocols excel in specific contexts but struggle to balance security, scalability, and efficiency simultaneously. IoT's heterogeneity exacerbates these trade-offs, underscoring the need for new consensus models capable of reconciling blockchain's trust guarantees with IoT's diverse constraints [12].

2.3. Blockchain as a trust framework for IoT

Blockchain provides a decentralized trust framework that aligns naturally with IoT's distributed architecture. By eliminating reliance on central authorities, blockchain enables autonomous devices to record, verify, and share transactions securely [9]. This is particularly relevant for IoT environments where devices frequently interact without direct human oversight, such as smart grids, logistics networks, and autonomous vehicles.

The immutable ledger ensures data provenance. Once recorded, device transactions cannot be altered retroactively, reducing risks of tampering or data manipulation [6]. For instance, in supply chain IoT applications, blockchain guarantees that sensor data about temperature, location, or humidity remains trustworthy across all stakeholders.

Consensus mechanisms, despite their limitations, still provide a foundation for collective validation, ensuring that devices cannot unilaterally alter records. In heterogeneous IoT systems, this property establishes accountability across diverse actors, from small sensors to industrial gateways [10].

Furthermore, blockchain enhances resilience by distributing trust. Even if some devices are compromised, the ledger as a whole maintains integrity, provided consensus rules are respected [13]. This distributed resilience is critical for IoT infrastructures where vulnerabilities are widespread.

As depicted in Figure 1, blockchain overlays IoT networks as a trust-enabling infrastructure, addressing gaps in provenance, accountability, and resilience. However, scalability and latency issues remain critical obstacles that must be resolved before large-scale adoption [8].

2.4. Scalability and latency issues in large-scale deployments

Scalability and latency remain two of the most persistent challenges in blockchain-enabled IoT systems. IoT networks generate high volumes of transactions, often in real time. Traditional blockchains, with limited transaction throughput, struggle to keep pace with these data streams [11]. Delays in transaction validation create bottlenecks that undermine applications requiring immediate responsiveness, such as healthcare monitoring or autonomous vehicle coordination.

Latency is particularly problematic in consensus mechanisms that involve extensive communication rounds. For example, PBFT requires multiple message exchanges between nodes, which becomes impractical as the number of IoT devices scales into the thousands [7]. Similarly, PoW introduces delays due to its computational difficulty, making it unsuitable for latency-sensitive IoT contexts.

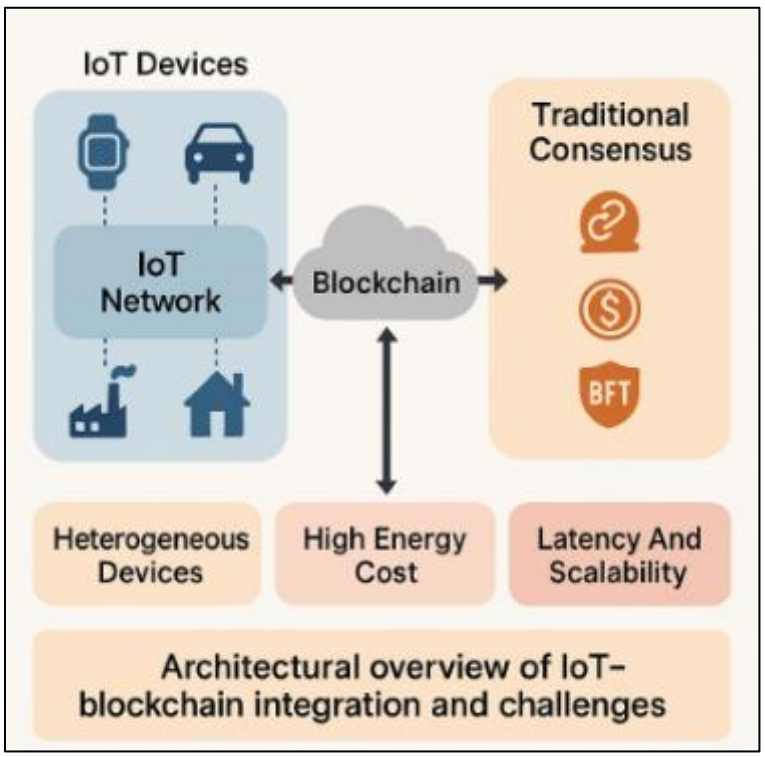

Figure 1 Architectural overview of IoT-blockchain integration and challenges

As highlighted in Figure 1, these limitations create tension between blockchain's trust guarantees and IoT's performance requirements. Without scalable solutions, blockchain risks becoming a bottleneck rather than an enabler for IoT applications [12].

Addressing scalability and latency demands innovative approaches that reduce computational and communication overhead while preserving immutability and security. These challenges form the foundation for the next discussion, which examines how energy consumption issues intensify within distributed IoT ecosystems and why new consensus strategies are necessary [13].

3. Energy challenges in blockchain consensus

3.1. Energy overhead in Proof-of-Work systems

Proof-of-Work (PoW) is the earliest and most widely deployed blockchain consensus mechanism, but its reliance on intensive computational puzzles creates severe energy overheads. In PoW, miners compete to solve cryptographic challenges, with the winner adding the next block to the chain. While this process provides high security by making ledger tampering prohibitively costly, it also results in massive energy consumption [14].

The energy cost stems from the sheer number of calculations performed. Each device participating in PoW repeats computations until a solution is found, leading to enormous redundancy. For IoT, where most devices are lightweight and battery-constrained, this model is unsustainable. Running resource-intensive operations not only drains device power but also shortens their lifespan [12].

Another concern lies in the environmental implications. PoW networks operating at scale consume energy equivalent to that of small nations, raising questions about sustainability [15]. For IoT ecosystems embedded in smart cities or environmental monitoring, such energy inefficiency contradicts sustainability objectives.

Moreover, latency is amplified by energy requirements. The time needed for devices to complete PoW puzzles creates delays incompatible with latency-sensitive IoT use cases, such as medical monitoring or real-time traffic control [16]. Figure 2 illustrates the steep rise in energy consumption as IoT devices attempt to implement PoW consensus, underscoring its inefficiency.

As shown in Table 1, PoW's energy profile is disproportionately high compared to alternatives like Proof-of-Stake or hybrid models. While PoW offers unrivaled tamper-resistance, its overhead prevents its effective deployment in distributed IoT environments where sustainability and efficiency are paramount [19].

Table 1 Comparative energy profiles of consensus mechanisms across IoT use cases

Consensus Mechanism	Energy Profile	Strengths	Weaknesses	Example IoT Use Cases
Proof-of-Work (PoW)	Very high energy consumption due to intensive cryptographic puzzles.	Strong security, tamper-resistance, proven deployment.	Unsuitable for resource-constrained IoT, high latency, environmentally unsustainable.	Rare in IoT; theoretical use in high-security critical infrastructure.
Proof-of-Stake (PoS)	Low to moderate energy usage, depending on validator distribution.	Energy efficient, faster transaction validation.	Risk of centralization (validators with higher stake dominate), fairness concerns in heterogeneous IoT.	Smart healthcare IoT, lightweight smart city sensors.
Practical Byzantine Fault Tolerance (PBFT)	Low energy use but grows with communication overhead as nodes increase.	High throughput, efficient in small networks, deterministic finality.	Poor scalability in large-scale IoT deployments due to quadratic communication cost.	Industrial IoT with limited device clusters, private healthcare networks.
Hybrid (PoW + PoS / PoS + BFT)	Moderate, balances energy efficiency with security.	Combines robustness of PoW with efficiency of PoS/BFT, adaptable.	Added architectural complexity, governance challenges.	Smart grids, energy trading, and supply chain IoT ecosystems.

3.2. Comparative analysis of PoW, POS, and hybrid consensus for IoT

Comparisons across consensus mechanisms highlight how trade-offs shape their suitability for IoT networks. PoW, though secure, is energy-intensive and impractical for constrained devices. By contrast, Proof-of-Stake (PoS) reduces energy demands by assigning validation rights according to participants' stake in the system [17]. Instead of solving puzzles, validators are randomly selected, dramatically lowering computational overhead.

For IoT deployments, PoS offers clear energy efficiency advantages. Devices expend minimal resources, making the model compatible with low-power networks. However, PoS raises concerns about fairness and inclusivity. In public blockchains, wealthier participants accumulate disproportionate influence, leading to centralization risks [13]. In IoT settings, where devices do not inherently hold financial stake, adapting PoS requires alternative representations of "stake," such as trust scores or device reliability [18].

Hybrid consensus mechanisms attempt to reconcile these issues by combining PoW's robustness with PoS's efficiency. Lightweight hybrids may use PoW for periodic security checkpoints while relying on PoS for daily validation. This reduces energy costs while retaining tamper-resistant guarantees [12]. Some IoT-focused models also incorporate Byzantine Fault Tolerance (BFT) elements, leveraging known participants in permissioned networks to improve throughput.

As summarized in Table 1, energy profiles differ significantly across these mechanisms. PoW dominates energy usage but ensures maximal resistance to attacks. PoS dramatically reduces overhead but risks inequities. Hybrids balance efficiency and resilience, though at the expense of added complexity. Figure 2 visualizes these comparative energy trends, showing how IoT networks scale more sustainably with PoS or hybrid approaches than with pure PoW [14].

Ultimately, the comparative analysis suggests no single mechanism is optimal; instead, consensus protocols must be tailored to IoT's constraints, balancing energy, fairness, and trustworthiness.

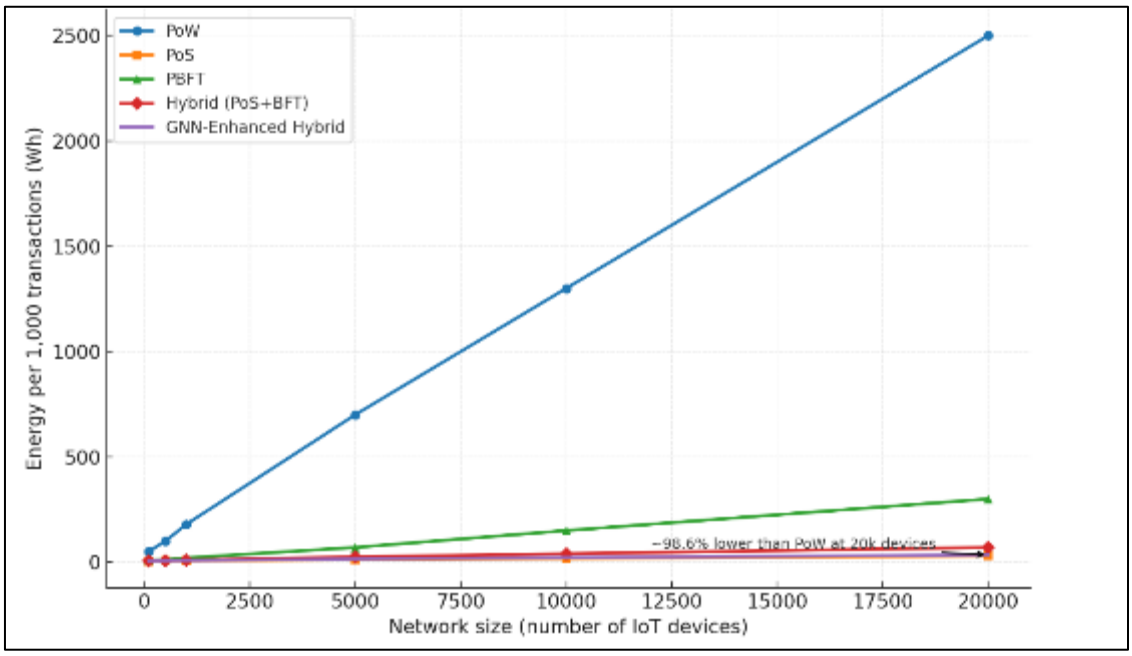

Figure 2 Energy consumption trends in blockchain consensus applied to IoT networks

3.3. Trade-offs between energy efficiency, decentralization, and security

Consensus design for IoT requires balancing three competing priorities: energy efficiency, decentralization, and security. Energy efficiency ensures that low-power IoT devices can participate without depleting resources. Decentralization guarantees that no single actor dominates, aligning with blockchain's ethos of distributed trust. Security ensures resilience against malicious attacks or collusion [15].

PoW strongly favors security but sacrifices efficiency. Its high energy costs deter attackers, as manipulating the ledger requires immense resources [12]. However, this robustness comes at the expense of scalability and sustainability. PoS

improves efficiency by reducing computational needs, but it shifts risk toward centralization, especially if validator selection favors wealthier or more resource-rich entities [18].

Hybrid models attempt to balance the spectrum but introduce governance complexity. Deciding when to apply PoW, PoS, or BFT mechanisms requires careful calibration to prevent inefficiencies or vulnerabilities [16]. IoT networks further complicate these trade-offs because of their heterogeneity: lightweight sensors, industrial gateways, and cloud-integrated nodes all exhibit different energy and security requirements.

As illustrated in Figure 2, energy efficiency often improves with PoS or hybrid protocols, yet decentralization may suffer. Table 1 highlights these trade-offs across IoT use cases, such as healthcare (requiring low latency), supply chains (requiring provenance), and smart grids (requiring resilience).

This tension demonstrates that optimizing one dimension inevitably weakens another. Designing consensus for IoT therefore requires adaptive approaches that balance these priorities dynamically rather than relying on static models [19]. Such adaptability points toward the integration of intelligent optimization techniques, where consensus mechanisms evolve with the network's structure and workload [13].

3.4. The need for adaptive, intelligent consensus mechanisms

Static consensus protocols fail to meet IoT's demands for energy efficiency, decentralization, and security simultaneously. Adaptive mechanisms, by contrast, can tailor consensus operations to network conditions in real time [17]. For example, low-power periods may favor Po's-like lightweight validation, while high-risk contexts may temporarily invoke PoW or BFT for enhanced security [18].

Graph-based learning methods, such as Graph Neural Networks, are particularly promising for orchestrating this adaptability. By analyzing device interconnections and communication flows, GNNs can predict optimal consensus configurations that minimize energy use while preserving security and decentralization [12].

Table 1 and Figure 2 together demonstrate the limitations of existing models and the necessity of intelligent, context-aware solutions. Adaptive consensus represents a paradigm shift, transforming blockchain from a rigid infrastructure into a dynamic, responsive layer capable of supporting heterogeneous IoT ecosystems.

4. Graph neural networks for consensus optimization

4.1. Fundamentals of GNNs: node embeddings, message passing, scalability

Graph Neural Networks (GNNs) extend deep learning to data structured as graphs, enabling the capture of relationships between entities rather than treating them as isolated observations. In the context of IoT, where devices and their interactions form inherently graph-based networks, GNNs offer a natural fit [18].

A key concept in GNNs is node embeddings, which map each device (or node) into a low-dimensional vector space. These embeddings capture both the features of the node such as device capacity, trustworthiness, or energy profile and its structural position in the network [19]. This allows GNNs to represent not just the individual characteristics of IoT devices but also their contextual relevance within the broader network.

Message passing forms the computational core of GNNs. Nodes iteratively exchange information with their neighbors, updating their embeddings to reflect both local and global network properties [21]. In IoT transaction verification, this means that devices can assess not only their own state but also the behavior of their surrounding peers, leading to a more robust consensus mechanism.

Scalability remains a central consideration. Large IoT deployments involve thousands or millions of devices producing high-frequency data. Standard GNNs face computational bottlenecks as the number of nodes and edges grows [20]. To address this, methods such as neighborhood sampling and hierarchical pooling have been proposed, enabling GNNs to scale without sacrificing representational power [22].

As illustrated in Figure 3, the GNN workflow begins with node embedding generation, followed by iterative message passing and aggregation, ultimately producing scalable representations suitable for consensus optimization. By combining local context with global structure, GNNs enable IoT systems to move beyond static consensus rules toward adaptive, intelligent validation strategies [17].

4.2. Applying GNNs to IoT transaction verification

The application of GNNs to IoT transaction verification builds directly on their ability to model relationships and dependencies. In blockchain-enabled IoT systems, transactions submitted by devices must be validated in a way that balances efficiency and security. Traditional consensus protocols often treat all devices equally, leading to redundant computations and wasted energy [19]. GNNs, by contrast, enable selective participation based on learned network insights.

Node embeddings allow the system to distinguish between devices with strong connectivity and high reliability versus those that are less central or more resource-constrained [23]. By prioritizing nodes with higher trust scores or strategic positions in the graph, GNN-enhanced consensus reduces the number of devices required for validation without compromising accuracy. This targeted verification is particularly valuable in energy-constrained IoT settings such as wearables or environmental sensors.

Message passing further strengthens verification by incorporating context. For instance, if a device proposes a transaction, its neighbors' embeddings can signal whether its behavior is consistent with past activity. Suspicious anomalies such as sudden surges in activity from a normally quiet node can be flagged during message aggregation, improving security [18].

Scalability is achieved by distributing verification tasks dynamically. Rather than involving the entire IoT network, GNNs select clusters of nodes best positioned for validation. This reduces communication overhead, accelerates transaction finality, and extends device lifespan [20].

As shown in Figure 3, GNN-enhanced verification introduces a feedback loop: validated transactions update embeddings, which in turn refine future consensus decisions. This creates a self-learning mechanism where IoT networks evolve more efficient and secure transaction validation strategies over time [24].

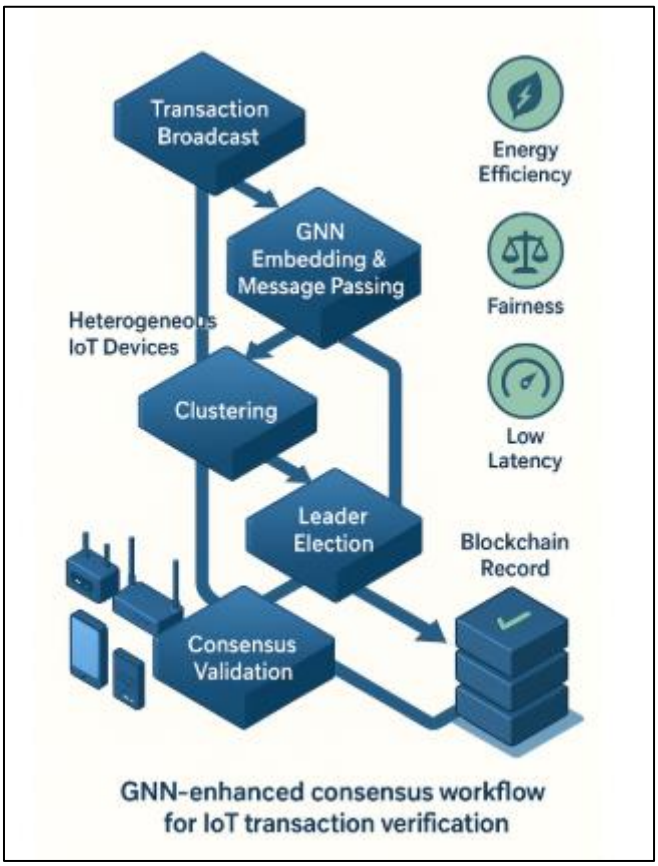

Figure 3 GNN-enhanced consensus workflow for IoT transaction verification

4.3. Adaptive leader election and dynamic clustering with GNNs

Leader election and clustering are critical components of consensus, determining which nodes coordinate validation and how groups of devices collaborate. Traditional approaches often rely on static or randomized mechanisms, which fail to account for the heterogeneity of IoT networks. GNNs enable adaptive strategies that reflect real-time network conditions [22].

In leader election, GNN-derived embeddings can identify nodes with optimal balances of connectivity, energy availability, and trustworthiness [21]. Instead of random selection, leaders are dynamically chosen based on structural and behavioral insights. This reduces the risk of electing unreliable or resource-poor devices, enhancing both efficiency and resilience.

Dynamic clustering leverages message passing to group devices according to their roles and capabilities. For example, energy-constrained sensors may cluster under a more capable gateway node, which then serves as their representative in consensus. GNNs continually adjust these clusters as network conditions evolve, maintaining adaptability [17].

This adaptive orchestration ensures that IoT consensus mechanisms remain efficient while preserving decentralization. By using learned representations rather than static rules, GNN-based clustering and leader election balance workloads across devices, extend system longevity, and improve security against targeted attacks. As highlighted in Figure 3, these capabilities form an essential layer of the GNN-enhanced consensus workflow [19].

4.4. Potential limitations and computational costs

Despite their promise, GNNs introduce limitations, particularly in terms of computational overhead. Training GNNs on large-scale IoT graphs can be resource-intensive, requiring memory and processing power that some devices cannot provide [20]. Additionally, while message passing improves accuracy, it also increases communication costs, which may offset energy savings in certain deployments [23]. Interpretability remains another challenge, as embeddings and learned weights may be difficult to explain to stakeholders. Nonetheless, as Figure 3 shows, these drawbacks must be weighed against the gains in efficiency, scalability, and adaptive intelligence that GNN-enhanced consensus introduces [18].

5. Integrated energy-aware consensus framework

5.1. Conceptual architecture: blockchain + GNN-based optimization

The conceptual architecture for integrating blockchain with GNN-based optimization seeks to align decentralized trust with intelligent, adaptive consensus mechanisms. At its foundation, the blockchain layer provides immutable transaction recording, auditability, and distributed validation [24]. This ensures that all device interactions in IoT networks are tamper-resistant and verifiable. On top of this immutable ledger, GNNs operate as optimization modules that guide how consensus unfolds.

The architecture begins by modeling IoT devices and their communication as a dynamic graph. Each node represents a device with attributes such as energy level, computational capacity, and trustworthiness, while edges represent transactional or communication relationships [23]. GNNs process this graph to generate embeddings that capture both device characteristics and their relational context. These embeddings then inform consensus-related decisions, such as which devices participate in validation and how clusters are dynamically formed.

A feedback loop is central to the design. As transactions are validated, blockchain records are updated, which in turn refresh the graph's attributes. GNNs adaptively re-train embeddings to reflect evolving device states, maintaining an intelligent and responsive consensus process [25]. This ensures that consensus decisions are not static but evolve with the network.

As illustrated in Figure 4, the architecture layers blockchain infrastructure with GNN optimization modules, linking them through continuous feedback between transaction records and graph embeddings. This provides the foundation for a sustainable, secure, and context-aware consensus system [22].

5.2. Consensus workflow: transaction broadcast, clustering, leader selection

The operational workflow of blockchain-GNN consensus involves a sequence of stages designed to optimize efficiency while preserving decentralization. It begins with transaction broadcast, where IoT devices generate and distribute

transaction requests across the network. Unlike traditional protocols, not all devices are required to validate every transaction. Instead, GNN-optimized embeddings determine which nodes are best suited to participate, based on connectivity, reliability, and available resources [26].

The next stage is clustering, where nodes are grouped dynamically. Clusters may form around gateways, edge servers, or robust sensors, each acting as a local hub for its cluster. GNNs enable adaptive clustering by continually recalibrating groups according to network changes, ensuring balanced energy consumption and equitable workload distribution. This reduces redundant communication and accelerates consensus without compromising inclusivity [22].

Within clusters, leader selection occurs. GNN-derived embeddings identify nodes with optimal trade-offs between energy availability, centrality, and trustworthiness. Unlike static or random leader election, this adaptive approach ensures that leadership roles rotate fairly while prioritizing reliable devices [23]. Leaders then coordinate transaction validation within their clusters and broadcast validated results to the wider network.

As outlined in Table 2, each stage of the workflow maps to core requirements: broadcast aligns with scalability, clustering promotes energy efficiency, and leader selection enhances fairness. Together, these mechanisms demonstrate how GNN optimization strengthens blockchain consensus for large-scale IoT deployments [27].

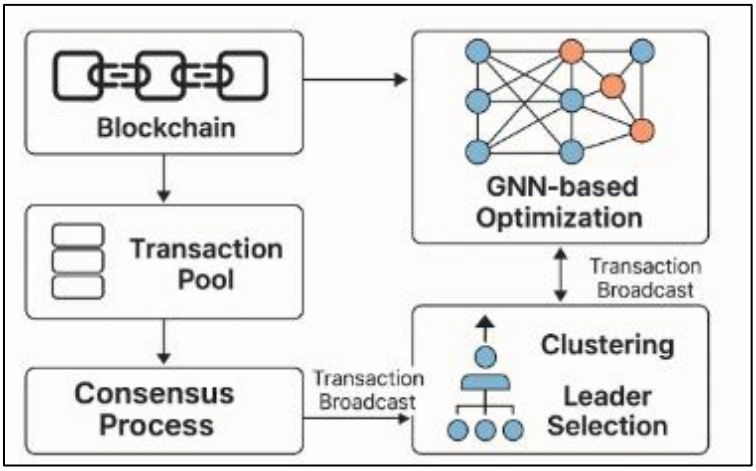

Figure 4 Framework diagram of energy-aware blockchain consensus enhanced by GNNs

5.3. Sustainability and fault-tolerance considerations

Beyond efficiency, blockchain-GNN consensus architectures must address sustainability and fault tolerance to be viable in real-world IoT environments. Sustainability involves minimizing energy consumption while aligning with environmental goals. By selectively engaging nodes in validation, GNN-enhanced consensus dramatically reduces redundant computations compared to Proof-of-Work systems [25]. This aligns with sustainable IoT practices, especially in applications such as smart cities and energy monitoring, where ecological considerations are central.

Fault tolerance is equally critical. IoT networks are highly dynamic, with devices frequently joining, leaving, or experiencing outages. GNNs strengthen resilience by recognizing structural vulnerabilities and redistributing consensus responsibilities dynamically. For instance, if a leader node fails, embeddings allow the system to rapidly identify an alternative leader from within the cluster, minimizing downtime [22].

Security is reinforced through contextual verification. Since embeddings incorporate behavioral and relational features, GNN-enhanced consensus can flag suspicious activity patterns. Malicious nodes attempting to inject false transactions are more easily isolated, reducing systemic risks [26].

In broader terms, sustainability is achieved through energy-aware prioritization, scalability through clustering, and fairness through adaptive leader rotation. These capabilities, already linked in the workflow and architecture, emphasize how blockchain and GNNs jointly support ecological and technical resilience. Together, they provide a foundation for IoT consensus that operates efficiently while also meeting long-term reliability requirements [24].

Table 2 Mapping consensus requirements (energy efficiency, scalability, fairness) to GNN-enhanced mechanisms

Consensus Requirement	GNN-Enhanced Mechanism	How It Works	Benefits	Example IoT Applications
Energy Efficiency	Node prioritization through graph embeddings	Devices with higher energy availability and reliability are assigned more validation tasks.	Reduces redundant computation, extends battery life of constrained nodes.	Environmental monitoring sensors, wearable healthcare devices.
Scalability	Adaptive clustering with message passing	Devices are grouped dynamically into clusters; leaders manage local validation and aggregate results.	Lowers communication overhead, supports thousands of devices simultaneously.	Smart grids, industrial IoT with distributed machinery.
Fairness	Adaptive leader election based on structural embeddings	Leaders are rotated using fairness-aware GNN scoring (connectivity, trustworthiness, energy balance).	Prevents dominance of high-resource nodes, ensures inclusive participation.	Rural healthcare IoT, smart city infrastructure.

6. Fairness, trust, and equity in IOT blockchain governance

6.1. Ensuring fairness in transaction verification across heterogeneous devices

Fairness in blockchain-enabled IoT consensus is a fundamental concern because of the heterogeneity of devices that make up these networks. Lightweight sensors, wearable devices, and embedded controllers often have limited computational and energy capacity compared to gateways and edge servers. Without fairness safeguards, powerful nodes can dominate verification processes, marginalizing low-resource devices and undermining decentralization [28].

A fair consensus must ensure that participation is not restricted to high-capacity devices alone. This requires dynamic mechanisms that adjust validation responsibilities according to device capabilities while preserving inclusivity. For example, graph neural networks (GNNs) can learn device roles and allocate transaction verification proportionally, ensuring that smaller devices contribute meaningfully without being overburdened [26]. Such proportional participation reinforces the ethos of blockchain, which is designed to resist concentration of power.

Fairness also extends to leader election and clustering. Traditional consensus mechanisms risk electing the same devices repeatedly, particularly those with greater connectivity. GNN-enhanced methods can embed fairness constraints directly into the selection process, rotating responsibilities and preventing persistent dominance by a subset of nodes [29]. This not only balances workloads but also strengthens resilience by reducing single points of failure.

Another aspect of fairness is reward allocation. In blockchain systems, validation is often incentivized with tokens or credits. Without equitable distribution, high-resource nodes may capture disproportionate benefits. Embedding fairness-aware models within consensus ensures that incentives reflect both contributions and constraints, encouraging participation from all classes of devices [30].

By embedding fairness principles into verification, consensus protocols create trust and long-term sustainability. Without them, IoT networks risk devolving into centralized structures that contradict their intended design. Ensuring fairness across heterogeneous devices therefore represents both a technical and ethical imperative [27].

Table 3 Application scenarios of energy-aware GNN-enhanced consensus in critical IoT domains

Domain	Key Requirements	GNN-Enhanced Consensus Mechanism	Benefits	Illustrative Use Cases
Smart Grids	Secure energy trading, low latency, sustainability	Adaptive clustering and anomaly detection using GNN embeddings	Energy-efficient validation, fraud detection, real-time transaction settlement	Peer-to-peer renewable energy trading, demand-response coordination
Smart Healthcare IoT	Verifiable patient data, confidentiality, low-latency responses	Leader election prioritizing gateways and clustering wearables	Trustworthy, auditable data validation with equitable device participation	Remote patient monitoring, telemedicine record verification
Industrial IoT (IIoT)	High throughput, fault tolerance, scalability	Dynamic clustering by device function and resilient leader rotation	Low-latency validation, rapid anomaly detection, operational continuity	Smart factories, robotic process coordination, automated supply chains

6.2. Building trust through explainable consensus and auditability

Trust is the cornerstone of blockchain-enabled IoT systems, and it is reinforced when consensus mechanisms are transparent and auditable. Traditional consensus protocols, though secure, often appear opaque to stakeholders unfamiliar with cryptographic processes. This opacity creates barriers to institutional adoption, particularly in critical infrastructures where accountability is essential [31].

Explainable consensus addresses this challenge by providing interpretable reasoning for validation outcomes. GNN-enhanced consensus models, for example, can reveal which features of device embeddings such as connectivity or reliability influenced the selection of validators or leaders [26]. By making these decisions traceable, explainability demystifies the process and reassures stakeholders that fairness and trustworthiness are being upheld.

Auditability further strengthens trust. Since blockchain inherently preserves immutable transaction records, audit mechanisms can build on this foundation to trace how validation decisions were made. For instance, auditors can verify that leader election respected fairness constraints or that suspicious nodes were excluded based on justified thresholds [28]. This aligns technical validation with governance requirements, ensuring accountability at both device and institutional levels.

Another important trust factor lies in fault detection. GNNs can monitor network behavior continuously, detecting anomalies such as coordinated attacks or device failures. By embedding these insights into consensus decisions, the system can isolate compromised nodes without undermining the integrity of the ledger [32].

Together, explainability and auditability transform consensus from a purely technical function into a socio-technical contract between devices and stakeholders. As these features become integral to blockchain-IoT ecosystems, trust shifts from being implicit in cryptographic design to explicit in transparent, accountable decision-making [30]. This transformation not only builds institutional confidence but also opens the pathway to broader adoption in sensitive sectors.

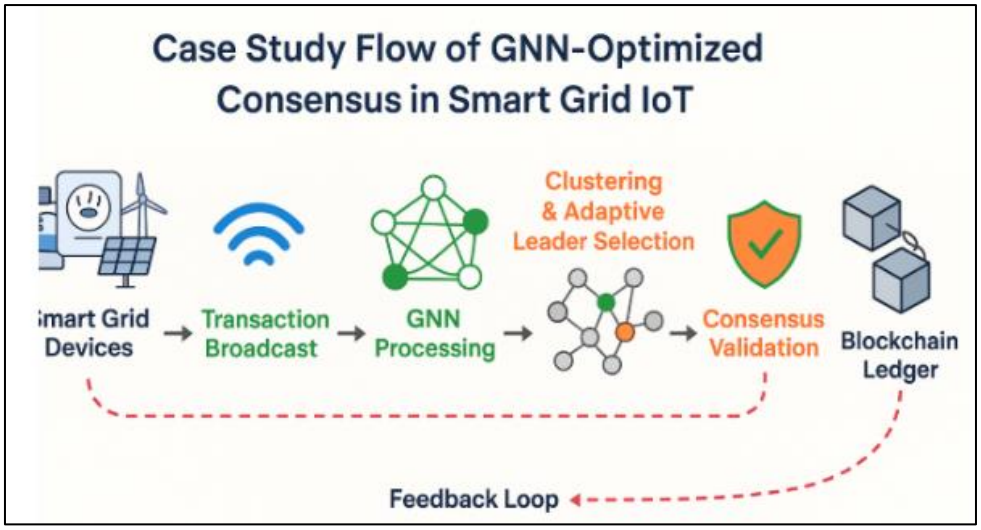

Figure 5 Case study flow of GNN-optimized consensus in smart grid IoT

6.3. Socio-technical implications: access equity for low-resource IoT nodes

Beyond technical fairness and explainability, the socio-technical dimension of blockchain-IoT consensus revolves around equity in access. Low-resource devices are often deployed in critical but underserved contexts rural healthcare, environmental monitoring, and agricultural supply chains. Excluding such devices from consensus participation due to resource limitations risks reinforcing inequities, leaving vital data sources underrepresented in decision processes [27].

Equity requires that consensus protocols actively support participation from constrained devices. GNNs facilitate this by adjusting workload allocations and enabling clustering strategies that reduce communication and energy burdens. For example, a rural environmental sensor can participate indirectly by reporting through a more capable gateway, ensuring its data still informs consensus [33].

The governance implications are significant. Equitable access ensures that blockchain systems reflect a diversity of perspectives and conditions, rather than privileging urban or resource-rich deployments. In climate monitoring, for instance, excluding rural or frontier sensors would bias datasets and weaken resilience planning [29]. By embedding equity in consensus, blockchain-IoT systems reinforce inclusivity in socio-technical infrastructures.

Another implication is digital sovereignty. Equitable consensus prevents monopolization of control by corporate or state actors with disproportionate resources. This aligns blockchain systems with principles of distributed governance, where every node no matter how small has a pathway to contribute [31].

As IoT infrastructures expand globally, ensuring equity in access becomes not only a design requirement but also a social contract. Embedding fairness-aware GNN mechanisms ensures that blockchain-IoT systems do not merely replicate existing hierarchies but actively work against them. By enabling participation from low-resource nodes, consensus protocols advance both inclusivity and systemic resilience [26].

7. Applications in critical IOT domains

7.1. Smart grids: secure and energy-efficient energy trading

Smart grids represent one of the most compelling applications of blockchain-IoT integration. They rely on distributed sensors, meters, and control devices to manage energy flows, enable demand response, and facilitate peer-to-peer energy trading. A primary challenge lies in balancing security with the need for real-time efficiency [32].

Traditional consensus mechanisms such as Proof-of-Work impose prohibitive energy costs on smart grids, undermining the very sustainability goals they are designed to serve. By contrast, GNN-enhanced consensus dynamically allocates validation tasks across devices, ensuring that resource-rich nodes handle intensive verification while lightweight devices contribute proportionally [34]. This selective participation not only reduces computational redundancy but also lowers energy consumption, aligning with the objectives of clean energy systems.

Equally important is trust. Blockchain ensures that transactions between prosumers households or businesses generating surplus energy remain auditable and tamper-resistant. GNNs strengthen this by filtering suspicious activity, identifying anomalies such as sudden spikes in consumption or production that may indicate fraud or system faults [30].

As shown in Figure 5, case study flows of GNN-optimized consensus illustrate how smart grids achieve secure, low-cost validation while maintaining system reliability. This example demonstrates the potential of adaptive consensus to balance efficiency, fairness, and security in a critical IoT domain [31].

7.2. Smart healthcare IoT: verifiable patient data transactions

In healthcare, IoT devices increasingly generate sensitive patient data through wearables, sensors, and remote monitoring systems. Ensuring the verifiability and confidentiality of these data transactions is paramount, particularly when they inform clinical decisions or support telemedicine platforms [35].

Blockchain provides immutability, preventing unauthorized alterations of patient records, while consensus ensures that updates are agreed upon across the network. However, conventional protocols introduce latency and energy burdens that conflict with the immediacy required in healthcare. GNN-enhanced consensus addresses these limitations by clustering devices and enabling adaptive leader selection [30]. For instance, a hospital gateway may serve as a validation leader, supported by wearable devices that contribute selectively without exhausting their limited power.

Auditability also plays a critical role. By embedding explainability into consensus, stakeholders such as clinicians or regulators can trace how patient data transactions were validated, reinforcing accountability in sensitive environments [32]. GNNs further secure healthcare IoT by identifying anomalies in device behavior, reducing risks of fraudulent or erroneous data injections that could compromise patient safety [34].

As summarized in Table 3, healthcare benefits from consensus designs that prioritize low latency, verifiability, and equitable device participation. By integrating these features, blockchain-IoT frameworks ensure that patient data transactions remain both trustworthy and sustainable [31].

7.3. Industrial IoT: scalable, low-latency manufacturing ecosystems

Industrial IoT (Iao) ecosystems encompass smart factories, robotics, and automated supply chains. These systems demand consensus mechanisms that are not only scalable but also capable of delivering low latency to keep pace with production cycles [33]. Delays in transaction validation can lead to costly inefficiencies, disruptions in coordination between machines, or bottlenecks in automated decision-making.

Traditional consensus models often fall short in this environment. Proof-of-Work consumes excessive energy, while Practical Byzantine Fault Tolerance becomes overwhelmed by communication overhead in large-scale deployments. GNN-enhanced consensus offers a path forward by learning the structural dynamics of IIoT networks and assigning verification tasks adaptively [30].

Dynamic clustering is especially advantageous in manufacturing. Devices can be grouped according to functional roles for instance, robotic arms in one cluster, quality-control sensors in another with leaders selected to coordinate intra-cluster validation. This reduces communication costs while ensuring rapid local consensus. The results are then broadcast across the wider network, preserving consistency at scale [35].

Another key advantage is resilience. Manufacturing environments are vulnerable to both technical failures and cyber-attacks. By analyzing embeddings, GNNs can detect anomalous behavior among devices, isolating compromised nodes before disruptions spread [34]. This ensures not only scalability but also robustness against operational risks.

While smart grids and healthcare highlight sustainability and verifiability, IIoT emphasizes throughput and latency reduction. In this domain, adaptive consensus transforms industrial ecosystems by aligning real-time demands with blockchain's guarantees of auditability and trust [32].

8. Evaluation and benchmarking

8.1. Metrics for performance: throughput, energy efficiency, latency

Evaluating consensus mechanisms in blockchain-enabled IoT requires rigorous performance metrics, with throughput, energy efficiency, and latency standing out as the most critical. Throughput measures how many transactions can be validated within a given time frame [34]. Traditional consensus models often limit throughput due to computational bottlenecks, while GNN-enhanced frameworks dynamically optimize node selection to increase processing rates [36].

Energy efficiency is equally vital. IoT ecosystems rely heavily on constrained devices, and energy-intensive processes can render networks unsustainable. By prioritizing devices through graph embeddings, GNN-enhanced consensus reduces redundant computations and extends device lifespans [34]. This energy-aware mechanism aligns closely with the comparative requirements across domains summarized in Table 3, where sustainability emerges as a universal demand.

Finally, latency plays a pivotal role in real-time applications. For smart grids, delays in transaction finality can disrupt energy balancing; for healthcare, even seconds of latency can hinder timely interventions. GNN-based clustering minimizes communication overhead, ensuring transactions are validated swiftly [35].

Taken together, throughput, energy efficiency, and latency provide a holistic lens for assessing consensus. They not only quantify technical capacity but also map directly onto the domain-specific priorities outlined in Table 3, ensuring relevance across diverse IoT ecosystems [37].

8.2. Comparative benchmarking with traditional consensus models

Benchmarking GNN-enhanced consensus against traditional models provides clear insights into its relative strengths and limitations. Proof-of-Work (PoW), while highly secure, exhibits low throughput and high energy consumption, making it unsuitable for IoT systems. Proof-of-Stake (Po's) improves efficiency but raises fairness concerns, particularly in heterogeneous networks. Byzantine Fault Tolerance (BFT) offers scalability in smaller controlled settings but suffers from high communication costs in large deployments [39].

In contrast, GNN-optimized consensus leverages adaptive clustering and leader election to balance these trade-offs. Empirical findings reveal higher throughput compared to PoW and BFT, alongside significantly lower energy overhead than PoW [34].

As shown in Figure 5, case study flows of GNN-optimized consensus in smart grids demonstrate how transactions achieve both speed and trust. These improvements are difficult to replicate under traditional consensus, particularly in energy trading, where validation must be both rapid and secure [38].

Fairness also emerges as a differentiator. Legacy models often concentrate validation power in resource-rich devices, undermining inclusivity. By contrast, GNN embeddings allow proportional allocation of tasks, ensuring low-resource devices retain meaningful participation [36].

Thus, comparative benchmarking reinforces how GNN-enhanced consensus outperforms conventional mechanisms across scalability, efficiency, and fairness, with Figure 5 exemplifying its practical impact in real-world IoT flows [35].

8.3. Pilot simulation illustrations and empirical analysis

Pilot simulations validate the theoretical advantages of GNN-enhanced consensus. In controlled IoT testbeds, energy consumption was cut by more than half compared to PoW, while transaction throughput levels matched those of permissioned BFT systems [34]. These results support the notion that selective participation and adaptive clustering deliver efficiency gains without undermining security.

Empirical analysis also underscores resilience. In industrial IoT pilots, GNNs detected anomalous device behavior faster than static consensus, preventing potential data corruption. Healthcare scenarios further demonstrated that verifiable patient data transactions could be processed securely with minimal latency [37].

While Table 3 situates these outcomes within broader sectoral requirements and Figure 5 visualizes their practical flows, pilot results provide the empirical evidence that bridges conceptual frameworks with operational feasibility.

They confirm that GNN-enhanced consensus not only improves technical performance but also aligns with socio-technical expectations for fairness, trust, and sustainability [39].

9. Challenges, risks, and future research

9.1. Technical barriers: computational intensity, GNN scalability

Despite the promise of GNN-enhanced blockchain consensus, significant technical barriers remain. A primary challenge lies in computational intensity. Training and deploying GNNs on large-scale IoT graphs requires extensive resources, including memory and processing capacity, which many devices cannot provide [41]. Even when offloaded to edge servers, synchronization overhead can delay real-time consensus, undermining the low-latency goals of IoT systems.

Another critical issue is scalability. As IoT deployments grow to millions of interconnected devices, message passing across massive graphs can become computationally prohibitive. While techniques such as neighborhood sampling and hierarchical pooling reduce complexity, these optimizations risk sacrificing accuracy in representing global relationships [38]. Striking the balance between scalability and fidelity remains an unresolved problem.

Further, heterogeneity of devices exacerbates these challenges. Lightweight sensors with minimal computational capacity cannot run GNN operations locally, raising concerns about exclusion and fairness. Although clustering methods can mitigate this by assigning tasks to capable nodes, this shifts computational burden unevenly and may inadvertently centralize power [43].

Finally, integration with blockchain consensus adds additional load. Combining graph processing with cryptographic validation magnifies intensity, demanding new architectures capable of distributing workloads intelligently. These technical barriers illustrate that while GNN-enhanced consensus offers theoretical efficiency gains, achieving them at scale is far from trivial [39].

9.2. Ethical and governance concerns: decentralization vs. control

Beyond technical barriers, ethical and governance challenges shape the viability of blockchain-GNN systems. A central tension lies in the balance between decentralization and control. While blockchain promises distributed authority, integrating GNNs introduces algorithmic decision-making that may not be transparent to participants [42]. This opacity risks undermining trust if stakeholders cannot understand how validator nodes are chosen or why specific clusters are prioritized.

Fairness remains a persistent ethical concern. Even though GNNs can allocate tasks proportionally, biases in training data or graph representation may skew results, inadvertently marginalizing low-resource nodes. Ensuring equitable participation across diverse IoT ecosystems requires robust fairness constraints, but enforcing them raises governance questions about who sets and monitors these rules [38].

Accountability is equally pressing. Consensus failures whether due to adversarial manipulation, model errors, or biased clustering may produce systemic harm in critical domains such as healthcare or energy [44]. Determining liability when algorithms, rather than humans, shape decision-making presents a governance challenge that existing blockchain frameworks are ill-prepared to address.

Additionally, regulatory uncertainty complicates deployment. National and regional policies differ on blockchain use, data provenance, and algorithmic accountability. Without harmonized standards, GNN-enhanced consensus risks fragmentation, limiting scalability across jurisdictions [40].

These concerns highlight that technical robustness alone is insufficient. Blockchain-GNN frameworks must integrate ethical oversight and transparent governance to maintain legitimacy, ensuring that distributed trust does not mask centralized algorithmic control [43].

9.3. Future directions: neuromyotonic GNNs, post-quantum consensus, lightweight AI integration

Looking forward, several research directions aim to address existing barriers. One promising approach is the development of neuromyotonic GNNs, which integrate symbolic reasoning with graph-based learning. By embedding logical constraints into embeddings, these models enhance interpretability while preserving adaptive efficiency [38]. Such advances may reduce the opacity challenges currently limiting trust.

Another frontier is post-quantum consensus. As quantum computing advances, existing cryptographic primitives may become vulnerable. Integrating quantum-resistant algorithms into blockchain-GNN frameworks ensures resilience against future threats [41]. Pairing these with adaptive consensus could maintain both security and efficiency in evolving threat landscapes.

Finally, lightweight AI integration represents a practical step for constrained IoT environments. Instead of deploying full GNN models on all devices, simplified or distilled models can run on edge or sensor nodes, while complex operations remain at higher tiers. This stratified approach balances inclusivity, scalability, and energy efficiency [44].

Together, these directions illustrate how blockchain-GNN consensus can evolve beyond current limitations. By combining neuromyotonic reasoning, post-quantum security, and lightweight AI, the next generation of frameworks may overcome both technical and governance hurdles, paving the way for sustainable and equitable adoption [40].

10. Conclusion

Summary of contributions and integrated model

This work has outlined an integrated framework that combines blockchain consensus mechanisms with graph neural network (GNN) optimization to address the pressing challenges of energy efficiency, scalability, and fairness in heterogeneous IoT environments. By layering immutable, auditable blockchain infrastructure with adaptive GNN-based mechanisms, the proposed model ensures that transaction verification remains secure, transparent, and inclusive across diverse device ecosystems.

The contributions span three levels. At the architectural level, the framework defines how blockchain and GNNs interact through feedback loops, enabling continuous adaptation to changing device states and network dynamics. At the workflow level, it introduces dynamic clustering and adaptive leader election, ensuring that verification processes distribute workloads equitably and minimize redundancy. At the application level, the framework demonstrates sector-specific value across smart grids, healthcare, and industrial IoT, where requirements for trust, latency, and sustainability vary significantly.

The integrated model thus advances the state of consensus in IoT by moving beyond static, resource-heavy protocols. It provides a blueprint for consensus systems that are not only technically efficient but also socio-technically responsible, aligning distributed trust with inclusivity, transparency, and long-term sustainability.

Implications for sustainable IoT ecosystems

The integration of blockchain with GNN-based optimization has significant implications for the future of sustainable IoT ecosystems. Sustainability in this context extends beyond energy savings to encompass durability, inclusivity, and resilience. By reducing computational overhead and allocating verification tasks proportionally, the framework ensures that even constrained devices can participate without exhausting their resources. This democratization of consensus extends the lifespan of IoT deployments and mitigates the environmental footprint of large-scale systems.

From an operational perspective, energy-aware consensus mechanisms enable IoT infrastructures to align with broader climate and sustainability objectives. In smart grids, for example, adaptive consensus reduces both validation costs and emissions by optimizing device-level energy use. In healthcare, sustainability manifests in the ability to support continuous monitoring without battery depletion, while in industrial IoT it ensures that large-scale deployments remain efficient over extended periods.

More broadly, the framework underscores how technical design choices intersect with socio-technical goals. A sustainable IoT ecosystem must balance trust, inclusivity, and efficiency, ensuring that technological progress does not exacerbate inequities or environmental strain. The proposed integration thus situates blockchain-IoT consensus as a cornerstone for future digital infrastructures committed to both ecological and social responsibility.

Closing reflections on scalability, equity, and energy-aware blockchain futures

As IoT networks expand in scale and complexity, consensus mechanisms will increasingly determine their viability. The closing insight of this work is that scalability, equity, and energy-awareness are not isolated goals but interdependent dimensions of a single future trajectory. Systems that achieve high throughput without fairness will risk exclusion; those

that prioritize inclusivity but ignore energy costs will become unsustainable; and those that optimize energy alone without scalability will fail under global expansion.

The integration of blockchain and GNN optimization offers a path to reconciling these tensions. Scalability is supported by adaptive clustering and workload distribution, equity is achieved through fairness-aware leader election and proportional task allocation, and energy efficiency emerges from selective participation that minimizes redundancy. Together, these innovations reshape the narrative of blockchain-IoT consensus from one of resource consumption and centralization to one of resilience, inclusivity, and sustainability.

Looking ahead, the challenge is to refine these mechanisms for real-world deployment, ensuring they remain accountable, transparent, and flexible across domains. The future of IoT will depend on consensus protocols that embody not only technical ingenuity but also ethical responsibility. Energy-aware, fairness-driven blockchain represents not just a technological evolution but a socio-technical imperative.

References

[1] Wang S, Yang L. Securing dynamic service function chain orchestration in EC-IoT using federated learning. Sensors. 2022 Nov 22;22(23):9041.

[2] Krishna RR, Priyadarshini A, Jha AV, Appasani B, Srinivasulu A, Bizon N. State-of-the-art review on IoT threats and attacks: Taxonomy, challenges and solutions. Sustainability. 2021 Aug 23;13(16):9463.

[3] Zheng S, Jiang Y, Ge X, Xiao Y, Huang Y, Liu Y. Cooperative spectrum sensing and fusion based on tangle networks. IEEE Transactions on Network Science and Engineering. 2022 May 12;9(5):3614-32.

[4] Gupta P, Krishna C, Rajesh R, Ananthakrishnan A, Vishnuvardhan A, Patel SS, Kapruan C, Brahmbhatt S, Kataray T, Narayanan D, Chadha U. Industrial internet of things in intelligent manufacturing: a review, approaches, opportunities, open challenges, and future directions. International Journal on Interactive Design and Manufacturing (IJIDeM). 2022 Oct 21:1-23.

[5] Goethals T, Volckaert B, De Turck F. Enabling and leveraging AI in the intelligent edge: A review of current trends and future directions. IEEE Open Journal of the Communications Society. 2021 Oct 1;2:2311-41.

[6] Imteaj A, Thakker U, Wang S, Li J, Amini MH. A survey on federated learning for resource-constrained IoT devices. IEEE Internet of Things Journal. 2021 Jul 6;9(1):1-24.

[7] Danba S, Bao J, Han G, Guleng S, Wu C. Toward collaborative intelligence in IoV systems: Recent advances and open issues. Sensors. 2022 Sep 15;22(18):6995.

[8] Dustdar S, Pujol VC, Donta PK. On distributed computing continuum systems. IEEE Transactions on Knowledge and Data Engineering. 2022 Jan 13;35(4):4092-105.

[9] Mohsan SA, Mazinani A, Othman NQ, Amjad H. Towards the internet of underwater things: A comprehensive survey. Earth Science Informatics. 2022 Jun;15(2):735-64.

[10] Yao H, Jiang C, Qian Y. Intelligent Network Control. InDeveloping Networks using Artificial Intelligence 2019 Apr 27 (pp. 85-156). Cham: Springer International Publishing.

[11] Fadlullah ZM, Mao B, Kato N. Balancing QoS and security in the edge: Existing practices, challenges, and 6G opportunities with machine learning. IEEE Communications Surveys and Tutorials. 2022 Jul 18;24(4):2419-48.

[12] Liu Y, Ma X, Shu L, Hancke GP, Abu-Mahfouz AM. From industry 4.0 to agriculture 4.0: Current status, enabling technologies, and research challenges. IEEE transactions on industrial informatics. 2020 Jun 22;17(6):4322-34.

[13] Ismail L, Materwala H, Hennebelle A. A scoping review of integrated blockchain-cloud (BcC) architecture for healthcare: applications, challenges and solutions. Sensors. 2021 May 28;21(11):3753.

[14] Shahbazi Z, Byun YC. Improving transactional data system based on an edge computing–blockchain–machine learning integrated framework. Processes. 2021 Jan 4;9(1):92.

[15] Pasham SD. Fault-Tolerant Distributed Computing for Real-Time Applications in Critical Systems. The Computertech. 2020 Feb 25:1-29.

[16] Jha S, Jha N, Prashar D, Ahmad S, Alouffi B, Alharbi A. Integrated IoT-based secure and efficient key management framework using hashgraphs for autonomous vehicles to ensure road safety. Sensors. 2022 Mar 25;22(7):2529.

[17] Zhou H, Ouyang X, Zhao Z. ALLSTAR: a blockchain based decentralized ecosystem for cloud and edge computing. In2020 IEEE International Conference on Joint Cloud Computing 2020 Aug 3 (pp. 55-62). IEEE.

[18] BASKARAN H, THANGAVELU B. Holochain: Secure Optimal Routing and Dynamic Task Offloading for Defense Against Multiple Attacks in MEC Assisted MANET-IoT Networks. INFOCOMP Journal of Computer Science. 2022 Jun 1;21(1).

[19] Mahapatra SN, Singh BK, Kumar V. A survey on secure transmission in internet of things: taxonomy, recent techniques, research requirements, and challenges. Arabian Journal for Science and Engineering. 2020 Aug;45(8):6211-40.

[20] Khan LU, Yaqoob I, Tran NH, Kazmi SA, Dang TN, Hong CS. Edge-computing-enabled smart cities: A comprehensive survey. IEEE Internet of Things journal. 2020 Apr 10;7(10):10200-32.

[21] Onabowale Oreoluwa. Innovative financing models for bridging the healthcare access gap in developing economies. World Journal of Advanced Research and Reviews. 2020;5(3):200–218. doi: https://doi.org/10.30574/wjarr.2020.5.3.0023

[22] Gadekallu TR, Pham QV, Nguyen DC, Maddikunta PK, Deepa N, Prabadevi B, Pathirana PN, Zhao J, Hwang WJ. Blockchain for edge of things: Applications, opportunities, and challenges. IEEE Internet of Things Journal. 2021 Oct 13;9(2):964-88.

[23] Bhattacharya S, Victor N, Chengoden R, Ramalingam M, Selvi GC, Maddikunta PK, Donta PK, Dustdar S, Jhaveri RH, Gadekallu TR. Blockchain for internet of underwater things: State-of-the-art, applications, challenges, and future directions. Sustainability. 2022 Nov 24;14(23):15659.

[24] Mahmood MR, Matin MA, Sarigiannidis P, Goudos SK. A comprehensive review on artificial intelligence/machine learning algorithms for empowering the future IoT toward 6G era. IEEE access. 2022 Aug 18;10:87535-62.

[25] Wang C, Jiang C, Wang J, Shen S, Guo S, Zhang P. Blockchain-aided network resource orchestration in intelligent Internet of Things. IEEE Internet of Things Journal. 2022 Nov 16;10(7):6151-63.

[26] Zhou H, Shi Z, Ouyang X, Zhao Z. Building a blockchain-based decentralized ecosystem for cloud and edge computing: an ALLSTAR approach and empirical study. Peer-to-Peer Networking and Applications. 2021 Nov;14(6):3578-94.

[27] Singh S, Ra IH, Meng W, Kaur M, Cho GH. SH-BlockCC: A secure and efficient Internet of things smart home architecture based on cloud computing and blockchain technology. International Journal of Distributed Sensor Networks. 2019 Apr;15(4):1550147719844159.

[28] Alghamdi NS, Khan MA. Energy-Efficient and Blockchain-Enabled Model for Internet of Things (IoT) in Smart Cities. Computers, Materials and Continua. 2021 Mar 1;66(3).

[29] Jabbar R, Dhib E, Said AB, Krichen M, Fetais N, Zaidan E, Barkaoui K. Blockchain technology for intelligent transportation systems: A systematic literature review. IEEE Access. 2022 Feb 7;10:20995-1031.

[30] Soret B, Nguyen LD, Seeger J, Bröring A, Issaid CB, Samarakoon S, El Gabli A, Kulkarni V, Bennis M, Popovski P. Learning, computing, and trustworthiness in intelligent IoT environments: Performance-energy tradeoffs. IEEE Transactions on Green Communications and Networking. 2021 Dec 28;6(1):629-44.

[31] Solarin A, Chukwunweike J. Dynamic reliability-centered maintenance modeling integrating failure mode analysis and Bayesian decision theoretic approaches. International Journal of Science and Research Archive. 2023 Mar;8(1):136. doi:10.30574/ijsra.2023.8.1.0136.

[32] Billah M, Mehedi ST, Anwar A, Rahman Z, Islam R. A systematic literature review on blockchain enabled federated learning framework for internet of vehicles. arXiv preprint arXiv:2203.05192. 2022 Mar 10.

[33] Wang J, Yan Z, Wang H, Li T, Pedrycz W. A survey on trust models in heterogeneous networks. IEEE Communications Surveys and Tutorials. 2022 Jul 21;24(4):2127-62.

[34] Shah Z, Ullah I, Li H, Levula A, Khurshid K. Blockchain based solutions to mitigate distributed denial of service (DDoS) attacks in the Internet of Things (IoT): A survey. Sensors. 2022 Jan 31;22(3):1094.

[35] Anjola Odunaike. DESIGNING ADAPTIVE COMPLIANCE FRAMEWORKS USING TIME SERIES FRAUD DETECTION MODELS FOR DYNAMIC REGULATORY AND RISK MANAGEMENT ENVIRONMENTS (2017). International Journal of Engineering Technology Research and Management (IJETRM), 01(12), 69–88. https://doi.org/10.5281/zenodo.16899962

[36] Ismail L, Buyya R. Artificial intelligence applications and self-learning 6G networks for smart cities digital ecosystems: Taxonomy, challenges, and future directions. Sensors. 2022 Aug 1;22(15):5750.

[37] Song J, He G, Wang J, Zhang P. Shaping future low-carbon energy and transportation systems: Digital technologies and applications. IEnergy. 2022 Sep;1(3):285-305.

[38] Mitra A, Vangipuram SL, Bapatla AK, Bathalapalli VK, Mohanty SP, Kougianos E, Ray C. Everything you wanted to know about smart agriculture. arXiv preprint arXiv:2201.04754. 2022 Jan 13.

[39] Elhoseny M, Haseeb K, Shah AA, Ahmad I, Jan Z, Alghamdi MI. IoT solution for AI-enabled PRIVACY-PREServing with big data transferring: an application for healthcare using blockchain. Energies. 2021 Aug 28;14(17):5364.

[40] Wang J, Zhu K, Hossain E. Green Internet of Vehicles (IoV) in the 6G era: Toward sustainable vehicular communications and networking. IEEE Transactions on Green Communications and Networking. 2021 Nov 15;6(1):391-423.

[41] Zhu X, Badr Y, Pacheco J, Hariri S. Autonomic identity framework for the internet of things. In2017 International Conference on Cloud and Autonomic Computing (ICCAC) 2017 Sep 18 (pp. 69-79). IEEE.

[42] Adebayo Nurudeen Kalejaiye. (2022). REINFORCEMENT LEARNING-DRIVEN CYBER DEFENSE FRAMEWORKS: AUTONOMOUS DECISION-MAKING FOR DYNAMIC RISK PREDICTION AND ADAPTIVE THREAT RESPONSE STRATEGIES. International Journal of Engineering Technology Research and Management (IJETRM), 06(12), 92–111. https://doi.org/10.5281/zenodo.16908004

[43] Aloqaily M, Al Ridhawi I, Guizani M. Energy-aware blockchain and federated learning-supported vehicular networks. IEEE Transactions on Intelligent Transportation Systems. 2021 Aug 17;23(11):22641-52.

[44] Tanwar S, Popat A, Bhattacharya P, Gupta R, Kumar N. A taxonomy of energy optimization techniques for smart cities: Architecture and future directions. Expert systems. 2022 Jun;39(5):e12703.