Information Fusion 127 (2026) 103752

Contents lists available at ScienceDirect

Information Fusion

journal homepage: www.elsevier.com/locate/inffus

Verifying energy generation via edge LLM for web3-based decentralized clean energy networks

Shan Jiang [a,1], Wenchang Chai [b,1], Mingjin Zhang [b,*], Jiannong Cao [b], Shichang Xuan [c], Jiaxing Shen [d]

[a] School of Software Engineering, Sun Yat-sen University, Zhuhai, China
[b] Department of Computing, The Hong Kong Polytechnic University, Hong Kong SAR, China
[c] Information Security Research Center, Harbin Engineering University, Harbin, China
[d] School of Data Science, Lingnan University, Hong Kong SAR, China

ARTICLE INFO

Keywords:
Decentralized energy networks
Web3
Large language models
Edge LLM

ABSTRACT

The global transition to clean energy is critical to achieving climate goals, yet traditional centralized systems face challenges in flexibility, grid resilience, and equitable access. While decentralized web3-based energy networks offer promising alternatives, existing solutions lack robust architectures to integrate distributed generation with real-time demand and fail to provide trustworthy energy verification mechanisms. This work introduces DeCEN, a decentralized clean energy network that synergizes collaborative edge computing and web3 technologies to address these gaps. DeCEN leverages autonomous edge devices to collect and process sensory data from renewable generators, enabling localized decision-making and verification of energy production. A layer-2 blockchain solution establishes a transparent web3 ecosystem, connecting clean energy generators and consumers through tokenized incentives for green energy activities. To combat fraud, DeCEN incorporates a novel large language model (LLM)-based energy verification protocol that analyzes sensory data to validate renewable claims, ensuring accountability and stabilizing token value. Additionally, a distributed LLM inference algorithm partitions LLMs into shards deployable on resource-constrained edge devices, enabling decentralized, low-latency processing while preserving data privacy and minimizing communication overhead. By integrating edge computing, blockchain, and AI-driven verification, DeCEN improves the reliability, trust, and efficiency of decentralized clean energy networks, offering a scalable pathway toward global renewable energy targets.

1. Introduction

The global transition to clean energy has emerged as a critical imperative in combating climate change, fostering sustainability, and ensuring energy security [1]. For instance, the International Renewable Energy Agency (IRENA) estimates that renewable energy must account for 90% of global electricity generation by 2050 to meet the Paris Agreement's target [2]. However, traditional clean energy systems remain predominantly centralized, relying on large-scale generation facilities such as offshore wind farms (e.g., Hornsea Project in the UK) and solar megaplants (e.g., China's Qinghai Solar Park). While effective for bulk generation, these centralized models struggle with inefficiencies in dynamic supply-demand balancing, grid resilience, and equitable access [3]. For example, during California's 2020 heatwaves, centralized grids faced rolling blackouts despite surplus solar generation, underscoring the risks of inflexible infrastructure.

In recent years, decentralized energy networks enabled by web3 technologies have gained traction as viable alternatives [4]. Projects like Brooklyn Microgrid and Power Ledger exemplify early attempts to democratize energy markets by enabling peer-to-peer solar energy trading [5]. These systems allow households with rooftop solar panels to sell excess energy directly to neighbors, bypassing traditional utilities. Similarly,the EU's NRGcoin initiative uses blockchain to incentivize renewable energy production through tokenized rewards [6]. Despite their promise, these systems often prioritize incentive mechanisms while neglecting foundational technical and operational challenges [7,8]. Specifically, most blockchain-based energy projects focus on data transparency but lack frameworks to holistically

* Corresponding author.
E-mail addresses: jiangsh73@mail.sysu.edu.cn (S. Jiang), wenchang.chai@connect.polyu.hk (W. Chai), cs-mingjin.zhang@polyu.edu.hk (M. Zhang), jiannong.cao@polyu.edu.hk (J. Cao), xuanshichang@hrbeu.edu.cn (S. Xuan), jiaxingshen@LN.edu.hk (J. Shen).
[1] Equal contributions

https://doi.org/10.1016/j.inffus.2025.103752
Received 2 May 2025; Received in revised form 9 September 2025; Accepted 13 September 2025
Available online 16 September 2025

integrate distributed energy resources with real-time consumer demand [9].

Two critical gaps persist in existing web3-based decentralized clean energy networks. On one hand, they lack a systematic architecture to bridge decentralized generation and consumption. For example, Australia's Decentralized Energy Exchange, while successful in small-scale trials, struggles to scale due to its reliance on cloud-based coordination, which introduces latency and centralization bottlenecks [10]. On the other hand, current systems lack robust energy verification mechanisms [11]. Without proof of renewable origin, consumers can hardly trust claims of green energy [12]. Besides, the energy generators tend to overclaim the generated clean energy to fake the contributions to clean energy networks and earn more web3 tokens [13]. Existing systems fail to provide a protocol to verify the amount of generated clean energy in an accurate and trustworthy manner.

In this work, we propose DeCEN, a collaborative edge computing and web3-based decentralized clean energy network with key functions as follows. First, edge devices autonomously collaborate to collect and process data from internet of things (IoT) devices attached to clean energy generators, e.g., solar radiation and temperature sensors. These data are critical to verify the amount of clean energy generated. Second, we propose a layer-2 blockchain solution to connect clean energy generators, electricity consumers, and other sustainability participants to form a web3 ecosystem, using native tokens to incentivize green activities.

On top of DeCEN, we propose a large language model (LLM)-based energy generation verification protocol with generator-attached sensory data as input. The protocol is essential to resist fraudulent energy attacks and stabilize the token value. Traditionally, LLMs are deployed on the cloud, which brings centralization concerns and is not in line with the decentralized architecture of DeCEN. To this end, we propose a distributed LLM inference algorithm to split complex LLMs into shards deployable on resource-constrained edge devices. In this way, the edge devices collaboratively run LLM inference tasks with sensory data from attached IoT devices as input. Such an architecture achieves nearly on-device data processing and reduces communication overhead significantly.

The main contributions of this work are as follows:

- We propose DeCEN, a decentralized clean energy network integrating collaborative edge computing and web3 technologies to bridge distributed energy generation and consumption, addressing inefficiencies in centralized systems. DeCEN employs the collaborative edge computing architecture to enable localized decision-making and establishes a web3 ecosystem connecting generators, consumers, and sustainability participants, using native tokens to incentivize green activities.
- We propose an LLM-based energy generation verification protocol using sensory data and LLMs to validate renewable energy claims, preventing overclaiming and stabilizing web3 token value. We conduct extensive experiments to examine the effectiveness of the protocol on public datasets.
- We propose a distributed LLM inference algorithm to split complex LLMs into shards deployable on resource-constrained edge devices, enabling decentralized, low-latency processing while minimizing communication overhead and preserving data privacy.

The rest of this paper is organized as follows. Sec. 2 surveys the related literature and articulates the uniqueness of this work. Sec. 3 presents the system architecture of DeCEN and introduces the roles of key entities. Sec. 4.1 elaborates on the proposed distributed LLM inference algorithm over collaborative edge networks. Sec. 5 presents the extensive experimental results on public datasets. Finally, Sec. 6 concludes this work and identifies the future directions.

2. Related work

This section summarizes the related work from the perspectives of blockchain-based decentralized energy networks and renewable energy for edge computing.

2.1. Blockchain-based decentralized energy networks

Blockchain-based energy trading has been extensively studied in the literature [14–16] with a particular focus on the privacy preservation approaches, consensus mechanisms, and demand-supply games.

Privacy preservation in blockchain-based P2P energy trading has been addressed through diverse cryptographic and architectural strategies. PrGChain [17] uses a decentralized ZKOracle with ZKPs to validate off-chain energy data without exposing sensitive details, while [18] employs stealthy on-chain transmission and non-interactive ZKPs (NIZK) to anonymize trading relationships and securely aggregate demand reports. Both methods ensure regulatory compliance but face scalability trade-offs due to computational overhead. MuLPP [19] introduces a three-tiered approach: pseudonym-based anonymity, RSA accumulators for bid validation, and BLS multi-signatures for participant confidentiality. Its PADPeC protocol enhances off-chain negotiation privacy, though reliance on permissioned communication limits decentralization. [20] leverages Trusted Execution Environments (TEEs) to isolate critical computations, ensuring confidentiality for energy aggregators. In contrast, PP-BCETS [21] integrates Ciphertext-Policy Attribute-Based Encryption (CP-ABE) for fine-grained access control, enabling encrypted arbitration but introducing cryptographic trust assumptions.

While ZKPs balance transparency and privacy, TEEs and CP-ABE introduce hardware or trust dependencies. Hybrid approaches combining lightweight ZKPs, decentralized protocols, and adaptive access control may address gaps in scalability and heterogeneous deployment.

Recent consensus innovations for blockchain-based energy trading focus on scalability and efficiency. MuLCOff [22] employs multi-layer sharded blockchains [23] with off-chain computation, enabling concurrent transaction processing. It achieves 87% higher throughput and reduces latency by 75-124× compared to on-chain methods, though it relies on permissioned shards. RL-optimized consensus [24] integrates reinforcement learning (DDPG) to dynamically optimize consensus parameters and energy trading decisions, balancing costs and rewards while ensuring supplier reliability via a reputation system. HiCoOB [25] is a hierarchical protocol with separate local/global consensus layers, allowing parallel transaction execution. It reduces leader overhead by 10× and local transaction wait time by 5×, preventing double-spending.

While MuLCOff [22] and HiCoOB [25] enhance scalability through sharding and hierarchy, they sacrifice decentralization. The RL approach [24] introduces adaptability but lacks real-world validation. Hybrid models merging hierarchical architectures with adaptive optimization remain unexplored.

Some studies explore game-theoretic and blockchain-integrated approaches to optimize demand-supply dynamics in energy trading. ETradeChain [26] proposes a blockchain-based Local Energy Market (LEM) using a modified double auction protocol with pseudo-coins (Pcoins) tied to energy generation. The protocol ensures decentralized fairness, reducing consensus delay by 90% via Raspberry Pi-based IoT testbeds. Liu et al. [27] design a blockchain framework for electricity data trading in low-carbon systems. The method combines multi-objective pricing games (with market power constraints) and digital watermarking to balance data valuation, profit maximization, and copyright protection. Bao et al. [28] integrate energy trading and consensus via a two-stage mechanism. The "proof-of-energy" consensus rewards prosumers based on historical generation, while a noncooperative game optimizes block rewards and trading strategies, improving social welfare and energy efficiency. Zhang et al. [29] develop two noncooperative games for dynamic pricing and peak load reduction in P2P markets.

Off-chain processing reduces latency while maintaining integrity, achieving lower overhead than on-chain modes.

ETradeChain [26] and [29] prioritize decentralized market fairness and efficiency through auction mechanisms and dynamic pricing, respectively. [27] and [28] uniquely align incentives with sustainability goals, i.e., data monetization in low-carbon systems [27] and energy-efficient consensus [28]. While these frameworks enhance demand-supply equilibrium, challenges persist in balancing computational complexity (e.g., game-theoretic optimizations [28]) with real-time scalability.

2.2. Edge computing and energy networks

The convergence of edge computing and renewable energy networks has emerged as a pivotal area of research, focusing on sustainability, efficiency, and adaptive resource management.

Several studies leverage advanced machine learning techniques to harmonize energy consumption and computational demands. Multi-agent deep reinforcement learning (MADRL) frameworks are proposed to optimize battery scheduling and renewable energy utilization in distributed edge servers, prioritizing privacy-preserving collaboration among edge service providers [30]. These frameworks demonstrate marked reductions in carbon emissions while maintaining operational cost efficiency. Complementary work integrates hierarchical cloud-edge-end architectures with deep learning for data repair and deep reinforcement learning (DRL) for dynamic policy optimization, addressing the intermittency of renewables and enhancing fault tolerance in smart grids [31].

To address the volatility of renewable energy and computational workloads, online distributed optimization strategies are introduced [32,33]. One such approach employs stochastic dual-subgradient methods to decouple energy trading decisions from task offloading, enabling edge servers to dynamically balance renewable energy usage with grid interactions [32]. Another proposes predictive, distributed scheduling algorithms that align job deadlines with renewable availability, minimizing grid dependency while optimizing server load balancing and transmission efficiency [33].

Economic frameworks are explored to align stakeholder incentives with energy sustainability. A Stackelberg game model optimizes pricing and energy allocation between renewable providers and IoT edge operators, ensuring equilibrium in profit maximization and resource distribution [34]. Such models highlight the role of cooperative competition in fostering efficient energy markets.

Practical implementations validate theoretical frameworks. For instance, a microgrid-powered edge computing prototype combines solar-wind hybrid systems with edge devices, demonstrating high reliability in sustaining operations under renewable energy constraints [35]. Similarly, QoS-aware offloading models prioritize edge or core cloud resources based on renewable availability, validated through IoT applications like real-time video analytics, which benefit from reduced latency and energy-aware computation [36].

While existing works underscore the viability of renewable-integrated edge systems, challenges persist. These include scalability in heterogeneous environments, interoperability between decentralized energy and computing infrastructures, and long-term sustainability under fluctuating demands. Promising avenues include hybrid models that combine adaptive machine learning with robust economic frameworks, as well as standardized protocols for cross-domain energy-compute coordination. The interplay between edge computing and energy networks is increasingly governed by intelligent distributed systems that prioritize sustainability without compromising performance. Innovations in machine learning, real-time optimization, and economic modeling lay the groundwork for resilient, low-carbon edge infrastructures. Future research must bridge gaps in scalability and interoperability to realize fully autonomous, renewable-powered edge ecosystems.

In a nutshell, existing works of renewable energy networks focus on blockchain-based energy trading and renewable energy-empowered edge computing. They rarely consider connecting the renewable energy generators and consumers into a web3 ecosystem to incentivize more sustainability activities. Second, they regard renewable energy as essential for edge computing while neglecting the technological role of edge computing for renewable energy generation. Third, they overlook the critical problems of energy generation verification, bringing risks of renewable energy fraud.

3. System architecture

This section presents the system architecture of web3-based decentralized clean energy networks (DeCEN), consisting of four layers: clean energy generators and consumers, clean energy-related IoT devices, edge devices and network, and blockchain and tokens, as depicted in Fig. 1.

The system architecture of the web3-based decentralized clean energy network is designed around the core idea of incentivizing and verifying clean energy generation and green behaviors through a distributed, trust-minimized infrastructure. At the heart of this network are clean energy generators and consumers. Generators produce renewable energy from sources such as solar panels and wind turbines. At the same time, consumers actively use this clean energy and engage in environmentally responsible activities like recycling, bringing reusable containers to canteens, or even deploying their own clean energy devices. These participants form the foundational layer of the ecosystem, driving both energy production and sustainable behavior.

Supporting this layer is a network of clean energy-related IoT devices that serve as the sensory and interaction interface of the system. These include reverse vending machines that allow users to recycle cans and bottles, environmental sensors such as solar radiation and wind speed monitors that track the performance and conditions of energy generators, and smartphones that capture and digitize user behaviors. These devices continuously collect real-time data, which is essential for verifying green activities and energy output.

This data is processed and validated at the edge through a distributed network of edge devices, including Nvidia Jetson AGX Orin, Xavier NX, and Orin Nano units. These devices are strategically deployed to form a resilient edge network that handles critical tasks such as authenticating IoT devices, collecting and storing data locally, and running LLM-based algorithms to verify the authenticity and efficiency of energy generation. By processing data at the edge, the system reduces latency, enhances privacy, and ensures scalability without overburdening centralized infrastructure.

The trust and incentive layer of the architecture is anchored in a custom Ethereum layer-2 blockchain called PolyChain [37], maintained collectively by the edge network. PolyChain is responsible for recording and tracing all verified data related to energy generation and green activities. To ensure long-term immutability and transparency, the hash values of this data are periodically committed to the Ethereum mainnet. In parallel, PolyChain manages the generation and distribution of GreenCoin, a native utility token. GreenCoins are awarded to generators, consumers, IoT devices, and edge infrastructure operators based on their verified contributions to the network, i.e., whether through producing clean energy, promoting sustainability, or maintaining system integrity. This tokenized reward mechanism aligns individual incentives with the collective goal of building a decentralized, transparent, and sustainable clean energy economy.

Recently, LLMs have been more and more popular due to their capacity to effectively work on problems of complex scales, such as in chatting, virtual assistants, and content generation, equaling human performance [38,39]. This work proposes using LLMs deployed on edge devices to verify the amount of generated clean energy based on the numerous sensory data surrounding the clean energy generators. Edge

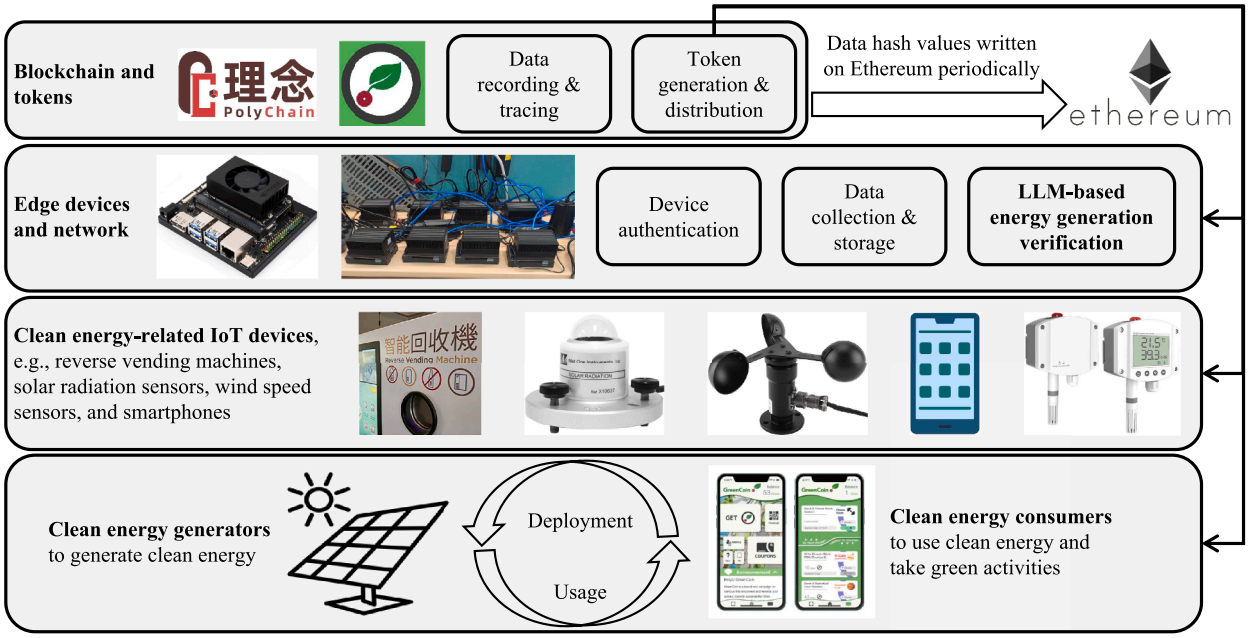

Fig. 1. System architecture of web3-based decentralized clean energy networks. PolyChain, an Ethereum layer-2 blockchain, is used to issue GreenCoin tokens and connect clean energy generators and consumers in a web3-based sustainability ecosystem. The native tokens incentivize the generation of more clean energy and participation in more green activities.

LLM for energy generation verification has the following advantages. On the one hand, the numerous IoT data of clean energy generators can be processed on-site at the network edge, contributing to reduced network congestion, improved processing efficiency, and enhanced data privacy. On the other hand, LLMs fully discover the data patterns and accurately verify the amount of generated clean energy, preventing energy fraud. However, deploying LLMs on edge devices is non-trivial [40] because LLMs are complex and large in scale, while edge devices are generally equipped with constrained resources [41].

4. Edge LLM-based energy generation verification

4.1. LLM-based energy generation verification

As shown in Fig. 2, the LLM-based verification framework consists of three main components, i.e., input processing, prompt formatting, and classification output extraction.

LLM Input Processing. The framework takes various features as input to enhance the performance.

Historical Context. We examine the influence of historical sequence length through three configurations: (DataPoint) isolated test points without context, (Previous x) multi-day historical sequences, and (CurrPrev x) hybrid representations combining historical data with current-day (previous hours) observations. This approach enables the capture of both short-term fluctuations and longer cyclical patterns while preventing temporal data leakage through timestamp-based filtering.

Conditional Variables. We systematically evaluated the contribution of domain-specific auxiliary information through three experimental conditions: (None) power data in isolation, (Weather) power data with meteorological variables, and (Irradiation) power data with irradiance measurements. This stratification allowed us to quantify the relative importance of environmental factors in anomaly detection.

Contextual Metadata. To provide deeper contextual understanding, we augmented the feature space with contextual metadata derived from historical patterns. The augmented feature spaces included month-level and hour-level power generation profiles (Month & Hour), correlation coefficients between environmental variables and power output, and structured semantic descriptions of measurement units and variable relationships (Background). These enrichment features were encoded as

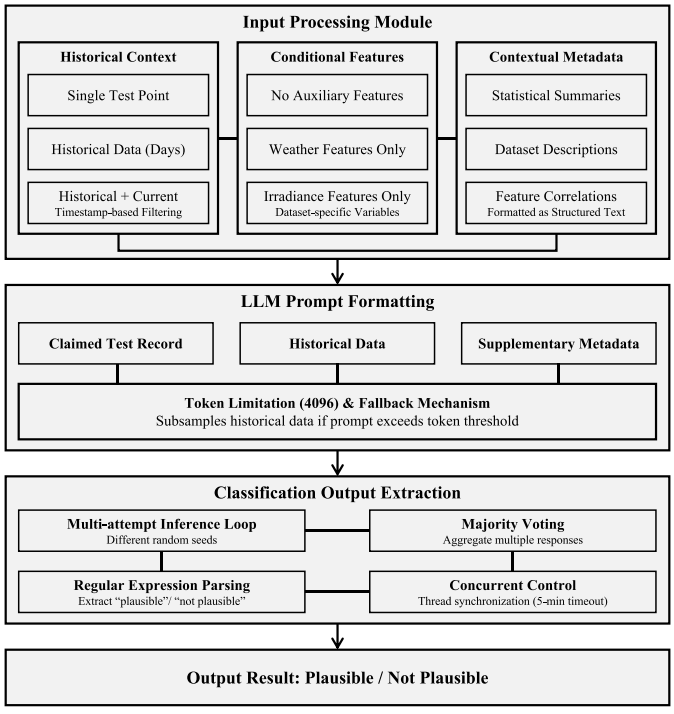

Fig. 2. Workflow of the proposed LLM-based energy generation verification framework. The framework consists of three main components: (1) Input Processing Module with historical context, conditional features, and contextual metadata; (2) LLM Prompt Formatting with token limitation handling; and (3) Classification Output Extraction with multi-attempt inference and majority voting.

structured text to leverage the LLM's natural language understanding capabilities while providing quantitative context for anomaly evaluation.

LLM Prompt Formatting. Due to input length constraints of LLMs, we developed a token-aware prompt construction methodology. Each prompt comprises three elements: the test record, corresponding historical data, and contextual metadata. When prompts exceed the 4096-token threshold, we implement adaptive historical data

subsampling proportional to the excess length, preserving critical temporal context while maintaining token limitations.

LLM Classification Output Extraction. For classification extraction, we developed a probabilistic inference framework that aggregates multiple stochastic passes. The system employs regular expression-based parsing to identify explicit classifications and implements majority voting across multiple response samples. This approach mitigates the inherent stochasticity of language model outputs and enhances classification robustness.

4.2. LLM inference over collaborative edge networks

We propose a three-phase algorithm to achieve efficient LLM inference over resource-constrained collaborative edge networks: profiling, task scheduling optimization, and collaborative inference, as depicted in Fig. 3. The detailed algorithm is presented in [42].

Profiling is an offline process that involves gathering essential runtime data needed for optimization, and it only needs to be conducted once. The profiling step collects various traces, including: 1) the execution time of each layer on different devices; 2) the size of activations and memory usage for each layer of the LLM; and 3) the available memory on each device, along with the bandwidth between devices. For execution time, we measure the duration required to generate a token during both the prefill and autoregressive stages, averaging the results. In cases where devices lack sufficient memory to hold the entire model for profiling, we employ dynamic model loading, sequentially loading model layers to accommodate memory constraints. This profiling data is crucial for developing intelligent task scheduling strategies.

During the task scheduling optimization stage, the scheduler devises a deployment strategy by determining which devices will participate, how to partition the LLM model on a layer-by-layer basis, and where to allocate each model shard. The strategy takes into account heterogeneous resources, device memory limitations, and privacy constraints, ensuring efficient LLM inference on selected devices.

Once the LLM model partitioning and allocation strategy is established, the chosen devices engage in collaborative model inference. Memory space for the key-value (KV) cache is pre-allocated on each participating device. We consider two scenarios for collaborative inference: sequential inference and pipeline parallel inference.

In sequential inference, devices process computations in turn using their allocated model shards. For example, if the LLM model is divided into three shards assigned to devices 1, 2, and 3, device 1 processes the input data first, then passes the activations or outputs to device 2, which processes the data before transmitting it to device 3. This approach is ideal for serving a single user, such as in a smart home setting where personal devices like tablets, phones, and smart speakers collaborate to perform LLM inference. Here, the user inputs a prompt, receives a response, and then inputs another prompt. The goal is to minimize latency in sequential inference. However, this method is not resource-efficient, as devices 2 and 3 remain idle while device 1 is computing. To enhance resource utilization, we implement pipeline parallelism, as seen in previous works like Gpipe and PipeDream for cloud servers. In pipeline parallel inference, input data is divided into micro-batches and fed into the system sequentially. Device 1 processes data batch B1 and then transmits the intermediate data to device 2. After processing B1, device 1 immediately begins handling data batch B2. This pipeline approach ensures that every device remains active, maximizing system resource utilization.

5. Performance evaluation

5.1. Experimental settings

Datasets. This study employs two publicly available Kaggle datasets to comprehensively evaluate the proposed LLM-based verification method.

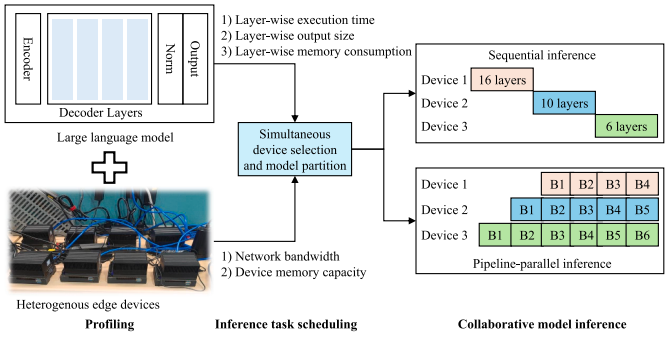

Fig. 3. Distributed LLM inference algorithm workflow over collaborative edge networks consisting of three phases: profiling, task scheduling optimization, and collaborative model inference.

- Renewable Energy and Weather Conditions (REWC) dataset[1] provides hourly meteorological measurements and associated renewable energy metrics collected over a 5-year period (2016-2021). It comprises 17 distinct features across 43,824 temporal instances, including temperature, pressure, humidity, wind speed, precipitation, and solar radiation. Additional parameters include binary sunlight indicators, daylight duration metrics, and a 26-category weather classification system.
- Solar Power Generation (SPG) dataset[2] records power generation from two Indian solar power plants monitored over a 34-day period in 2020. It comprises two components: generation data and sensor measurements. The generation data contains 527,040 entries from 39 inverters across both facilities, recording AC power output, DC power input, and daily yield at 15 min intervals. Complementary sensor data includes 2,159,592 records from monitoring stations, capturing irradiance, ambient temperature, and module temperature at one-minute resolution.

We employ chronological partitioning with an 80:20 train-test split to preserve temporal integrity. All numerical features underwent standardization according to:

$$z = \frac{x - \mu}{\sigma} \tag{1}$$

where parameters μ and σ are derived exclusively from the training partition to prevent information leakage.

Baselines. To assess the effectiveness of the proposed LLM-based verification framework, we compare its performance against several classical machine learning models. These models are implemented with fixed hyperparameters and represent a diverse range of learning paradigms.

- *Support Vector Machine (SVM)* with radial basis function kernel, regularization parameter of $C = 10$ and `gamma='scale'` [43].
- *Decision Tree (DT)* with a maximum depth of 8 and a minimum of 10 samples required to split a node [44].
- *Random Forest (RF)* composed of 100 decision trees, with each constrained to a maximum depth of 8 [45].
- *Logistic Regression (LR)* with $C = 1.0$ and a maximum of 1000 iterations [46].

Fake Data Generation. To comprehensively evaluate the proposed method for power generation verification, we generate adversarial fake power generation samples through five perturbation strategies.

[1] https://www.kaggle.com/datasets/samanemami/renewable-energy-and-weather-conditions/

[2] https://www.kaggle.com/datasets/anikannal/solar-power-generation-data/

(1) Noise-based perturbations introduce minor fluctuations by adding or subtracting a multiple of the standard deviation from the original value.

$$\tilde{x}_t = x_t \pm k\sigma_t, \quad k \in \{1, 3\} \tag{2}$$

where σ_t denotes the standard deviation of power values at temporal position t across training data.

(2) Mean-shift deviations replace the original value with the historical mean for the corresponding time, optionally adjusted upward or downward.

$$\tilde{x}_t = \mu_t \pm k\sigma_t, \quad k \in \{1, 3\} \tag{3}$$

(3) Boundary violations assign values that exceed historical maxima or fall below minima, with non-negative clipping applied when necessary.

$$\tilde{x}_t = \begin{cases} \max + \sigma_t \\ \max(x_t^{\min} - \sigma_t, 0) \end{cases} \tag{4}$$

(4) Repeat anomalies substitute the current value with a previously observed value from the same time on an earlier day.

$$\tilde{x}_t = x_{t-\tau}, \quad \tau \in \mathbb{Z}^+ \text{ (daily cycles)} \tag{5}$$

(5) Non-zero anomalies alter zero-valued entries by assigning a fixed non-zero value or enforcing zero output.

$$\tilde{x}_t = \begin{cases} \epsilon & \text{if } x_t = 0 \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

where ϵ represents physically implausible nocturnal generation values sampled from historical non-zero baselines.

Evaluation Metrics. Our evaluation framework addresses the class imbalance inherent in anomaly detection scenarios. Apart from the standard metrics (accuracy, precision, recall, F1-score), we also emphasize balanced performance indicators by introducing balanced accuracy and the Matthews correlation coefficient.

Balanced accuracy provides proportional class weighting:

$$\text{Balanced Accuracy} = \frac{1}{2}\left(\frac{\text{TP}}{\text{TP} + \text{FN}} + \frac{\text{TN}}{\text{TN} + \text{FP}}\right) \tag{7}$$

The Matthews correlation coefficient serves as our primary metric due to its robustness to class imbalance:

$$\text{MCC} = \frac{\text{TP} \times \text{TN} - \text{FP} \times \text{FN}}{\sqrt{(\text{TP} + \text{FP})(\text{TP} + \text{FN})(\text{TN} + \text{FP})(\text{TN} + \text{FN})}} \tag{8}$$

MCC produces high scores only when predictions perform well on both positive and negative classes, offering a comprehensive assessment of anomaly detection capability.

5.2. Main results

We use one-day historical context plus current-day records as input. LLM prompts dynamically handled token limits through strategic subsampling while preserving critical temporal patterns. We evaluated four distilled DeepSeek-R1 variants (8B to 70B parameters) against classical baselines under five adversarial scenarios (*Noise*, *Mean*, *Boundary*, *Repeat*, *Non-zero*).

The evaluation reveals distinct performance patterns tied to model architectures and dataset characteristics (Table 1). Qwen 2.5-based variants (R1-14B/R1-32B) demonstrate stronger anomaly detection on weather-influenced energy data (REWC), achieving superior balanced accuracy and MCC compared to larger Llama 3-based models. Conversely, Llama 3 architectures (70B) excel in solar generation scenarios (SPG), attaining perfect recall despite lower precision, suggesting enhanced sensitivity to temporal patterns in photovoltaic data.

Machine learning methods exhibit critical limitations, with LR producing negative MCC scores in REWC (an indicator of worse classification than random guessing) and RF showing systematic bias toward majority classes. These failures highlight the inadequacy of statistical correlation-based approaches for adversarial scenarios requiring joint analysis of temporal sequences and environmental covariates.

Moreover, we find that model scale alone does not explain performance variations. The Qwen-based 14B model outperforms both the larger Llama 3-70B and the smaller Llama 3-8B variants across different datasets. We analyze this architectural divergence across different adversarial simulation strategies in detail in the following subsections.

5.3. Comparison under different data faking strategies

Our analysis reveals distinct performance patterns across data faking strategies and model architectures. As shown in Table 2, LLMs demonstrate superior overall detection accuracy compared to traditional machine learning approaches, yet this advantage varies significantly across different adversarial scenarios.

LLMs exhibit remarkable effectiveness in detecting large-magnitude statistical manipulations. The R1-32B model achieves peak performance in extreme noise injection (Noise + 3σ) and mean shift (Mean + 3σ) scenarios, with R1-14B following closely. This suggests that larger models develop enhanced sensitivity to significant statistical deviations during pre-training. We observe a consistent performance degradation when the manipulation magnitude decreases, with all models showing reduced accuracy on smaller deviations compared to their larger counterparts. This degradation appears more pronounced in traditional ML models than in LLMs, indicating greater robustness of neural approaches to subtle manipulations.

Strategy-specific vulnerabilities reveal fundamental architectural limitations across model types. Traditional methods substantially outperform LLMs in repeat pattern detection, with performance on these patterns declining significantly as LLM scale increases. This counterintuitive result suggests that larger models may optimize for statistical pattern recognition at the expense of sequential anomaly detection capabilities. Similarly, boundary violations expose complementary strengths between architectures, with tree-based methods outperforming even the strongest LLM on boundary detection tasks. This advantage likely stems from the explicit threshold-based decision mechanisms in tree models that align directly with boundary violation characteristics.

The directionality of statistical manipulation significantly impacts detection performance across all models. For noise-based perturbations, negative shifts are generally more challenging to detect than positive ones for LLMs, potentially relating to differential sensitivities to upper and lower bound violations developed during training. The R1-14B model demonstrates the most balanced directional robustness among LLMs, with smaller performance differences between positive and negative manipulations compared to other variants.

Semantic consistency recognition presents perhaps the most striking capability gap between model classes. LLMs achieve near-perfect accuracy in non-zero detection tasks, significantly outperforming most traditional methods. This exceptional performance likely results from LLMs' pre-training on diverse datasets that implicitly encode physically plausible value ranges, enabling them to recognize semantically implausible patterns. Notably, the Decision Tree achieves perfect accuracy on this task as well, possibly due to its explicit decision boundary methodology.

5.4. Feature importance analysis

To evaluate the influence of different features on classification performance, we examined three principal feature categories: historical context, conditional variables, and contextual metadata (Fig. 4).

Historical context Configuration.

Temporal representation significantly influences anomaly detection performance across various configurations. Models leveraging both

Table 1
Performance comparison of classical ML and LLM-based models on SPG & REWC datasets.

Data	Model	Acc.	Bal. Acc.	Prec.	Rec.	Spec.	F1	MCC
SPG	DT	0.755	0.629	0.153	0.480	0.778	0.232	0.161
	LR	0.466	0.417	0.054	0.360	0.475	0.094	-0.088
	RF	0.660	0.587	0.113	0.500	0.673	0.185	0.098
	SVM	0.531	0.517	0.082	0.500	0.533	0.141	0.018
	R1-14B	0.815	0.885	0.290	**0.967**	0.803	0.446	0.470
	R1-32B	**0.844**	**0.900**	**0.326**	**0.967**	**0.833**	**0.487**	**0.508**
	R1-70B	0.774	0.871	0.269	**0.967**	0.776	0.420	0.443
	R1-8B	0.797	0.814	0.253	0.833	0.794	0.388	0.384
REWC	DT	0.623	0.585	0.108	0.540	0.630	0.181	0.093
	LR	0.580	0.397	0.037	0.180	0.613	0.062	-0.114
	RF	0.702	0.527	0.091	0.320	0.733	0.142	0.032
	SVM	0.562	0.469	0.066	0.360	0.578	0.112	-0.033
	R1-14B	**0.823**	**0.843**	**0.286**	0.867	**0.819**	**0.430**	**0.432**
	R1-32B	0.738	0.812	0.214	0.900	0.725	0.346	0.356
	R1-70B	0.679	0.781	0.181	0.900	0.661	0.302	0.308
	R1-8B	0.631	0.800	0.172	**1.000**	0.600	0.294	0.322

Table 2
Detection performance across process_method strategies. Values represent balanced accuracy.

Model	Original	Noise + 3σ	Noise + σ	Noise-1σ	Noise-3σ	Mean + 3σ	Mean-σ	Mean-1σ	Mean-3σ	Boundary + σ	Boundary-σ	Repeat	Non-Zero
DT	0.510	0.790	0.520	0.630	0.840	0.710	0.510	0.610	0.840	0.820	0.870	0.310	**1.000**
LR	0.270	0.750	0.710	0.400	0.470	0.770	**0.720**	0.460	0.380	0.810	0.490	**0.390**	0.180
RF	0.410	0.820	0.650	0.600	**0.860**	0.770	0.650	0.620	0.780	0.870	**0.930**	0.330	0.560
SVM	0.430	0.740	0.510	0.500	0.590	0.730	0.540	0.560	0.630	0.750	0.600	0.300	0.220
R1-14B	0.917	0.967	**0.900**	**0.733**	0.850	**0.950**	0.717	**0.850**	**0.867**	**0.950**	0.867	0.083	1.000
R1-32B	**0.933**	**0.983**	0.883	0.667	0.817	0.933	0.667	0.717	0.850	0.933	0.817	0.083	1.000
R1-70B	0.933	0.932	0.672	0.550	0.817	0.850	0.614	0.627	0.783	0.898	0.833	0.033	1.000
R1-8B	0.917	0.900	0.817	0.550	0.717	0.900	0.667	0.517	0.567	0.917	0.683	0.133	1.000
ML Avg.	0.405	0.775	0.597	0.532	0.690	0.745	0.605	0.562	0.657	0.812	0.723	0.333	0.490
LLM Avg.	0.925	0.946	0.818	0.625	0.800	0.908	0.666	0.678	0.767	0.925	0.800	0.083	1.000

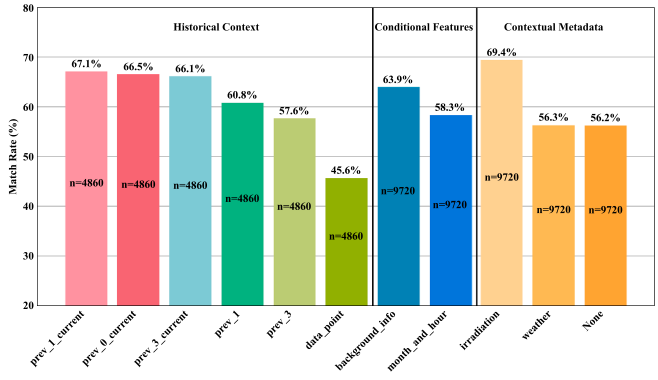

Fig. 4. Detection performance comparison across feature modalities with 95% confidence intervals (n = 29160 total samples). The overall match rate is 60.6%. **Left:** Historical context methods, where approaches combining current and historical data consistently outperform historical-only methods and single-point detection. **Center:** Conditional features showing performance differential between system background information and temporal metadata. **Right:** Contextual metadata showing irradiation data significantly outperforming both weather information and no contextual features.

current and historical data demonstrate superior performance, with approaches incorporating the immediate preceding timestep with current measurements (CurrPrev 1) achieving the highest overall accuracy. This configuration particularly excels at detecting boundary violations and negative shift perturbations, suggesting its enhanced capability to identify contextual inconsistencies.

In contrast, configurations relying solely on historical data without current measurements show reduced effectiveness, while single-point detection exhibits substantially diminished performance. This pattern underscores the critical importance of temporal context for accurate anomaly detection, with optimal results emerging from combined current and historical representations rather than from either in isolation.

Temporal context efficacy varies notably across manipulation strategies. All methods incorporating temporal context perform exceptionally well on semantic anomalies (Non-zero) and large-magnitude perturbations (Noise + 3σ), likely due to these anomalies' distinctive temporal signatures. However, for pattern-based anomalies (Repeat), simpler configurations without extensive historical context paradoxically outperform more complex temporal representations, suggesting that excessive historical information may obscure certain repetition patterns. This counterintuitive finding highlights the importance of targeted temporal sampling strategies aligned with specific anomaly detection objectives.

5.5. Study on key features

To systematically evaluate the efficacy of different information modalities on detection performance, we examined three principal categories of features: historical context, conditional features, and contextual metadata (Fig. 4). These features were utilized differently across model architectures—traditional ML models incorporated them as direct training inputs, while LLM-based detection leveraged them as contextual elements within prompts. The following analysis examines how various feature configurations influence detection accuracy across manipulation strategies and datasets.

Historical Context Sampling Methods. Historical data incorporation significantly enhances anomaly detection performance across various sampling configurations (Table 3 and Table 4). In Fig. 4 (Blue), methods leveraging both current and historical data demonstrate superior performance, with approaches incorporating the previous time step and current data achieving the highest accuracy (67.1%), closely followed by methods using only the immediate preceding data point (None, 66.5%) and three previous time steps (CurrPrev 3, 66.1%). In contrast,

Table 3
Accuracy by historical context and adversarial simulation strategy (Model: R1-32B; Conditional Var: Irradiation; Contextual Metadata: Background).

Hist	Noise + 3σ	Noise + σ	Noise-σ	Noise-3σ	Mean + 3σ	Mean + σ	Mean-σ	Mean-3σ	Boundary + σ	Boundary-σ	Non-Zero	Original	Repeat
None	0.800	**0.917**	0.583	0.400	0.917	**0.883**	0.667	0.417	0.867	0.500	**1.000**	0.667	**0.383**
Current	0.967	0.900	0.633	0.833	0.917	0.783	**0.800**	**0.867**	0.900	**0.817**	**1.000**	0.900	0.067
Previous 1	**0.983**	0.850	0.517	0.767	0.917	0.717	0.600	0.850	0.883	0.783	**1.000**	0.883	0.100
CurrPrev 1	**0.983**	0.883	**0.667**	0.817	0.933	0.667	0.717	0.850	0.933	**0.817**	**1.000**	**0.933**	0.083
Previous 3	0.933	0.733	0.600	0.750	**0.950**	0.617	0.633	0.833	0.900	0.767	**1.000**	0.900	0.117
CurrPrev 3	**0.983**	0.833	0.600	**0.850**	0.933	0.717	0.617	0.850	**0.950**	0.800	**1.000**	0.900	0.083

Table 4
Accuracy by historical context and adversarial simulation strategy (Conditional Var: Irradiation; Contextual Metadata: Background).

Hist	Noise + 3σ	Noise + σ	Noise-σ	Noise-3σ	Mean + 3σ	Mean + σ	Mean-σ	Mean-3σ	Boundary + σ	Boundary-σ	Non-Zero	Original	Repeat
None	0.783	0.792	0.496	0.392	0.800	**0.771**	0.600	0.412	0.800	0.446	0.996	0.725	**0.300**
Current	0.929	**0.825**	0.571	0.808	0.904	0.662	**0.704**	**0.808**	0.875	0.779	0.996	0.908	0.083
Previous 1	0.925	0.696	0.433	0.742	0.863	0.637	0.579	0.750	0.854	0.717	**1.000**	0.850	0.117
CurrPrev 1	0.942	0.812	**0.625**	0.800	**0.908**	0.658	0.675	0.767	**0.921**	**0.800**	**1.000**	**0.925**	0.083
Previous 3	0.912	0.675	0.496	0.742	0.887	0.621	0.579	0.746	0.846	0.708	0.992	0.871	0.150
CurrPrev 3	**0.958**	0.783	0.613	**0.812**	0.879	0.675	0.662	0.787	0.904	0.779	**1.000**	0.867	0.138

Table 5
Accuracy by conditional variables and adversarial simulation strategy.

Condition	Noise + 3σ	Noise + σ	Noise-σ	Noise-3σ	Mean + 3σ	Mean + σ	Mean-σ	Mean-3σ	Boundary + σ	Boundary-σ	Non-Zero	Original	Repeat
Irradiation	**0.983**	**0.883**	**0.667**	**0.817**	**0.933**	**0.667**	**0.717**	**0.850**	**0.933**	**0.817**	**1.000**	**0.933**	0.083
Weather	0.850	0.550	0.400	0.583	0.833	0.550	0.567	0.650	0.833	0.650	0.983	0.800	**0.200**
None	0.867	0.600	0.483	0.767	0.833	0.567	0.533	0.767	0.800	0.750	**1.000**	0.900	0.050

Table 6
Accuracy by conditional variables and adversarial simulation strategy (Historical Context: CurrPrev 1; Contextual Metadata: Background).

Condition	Noise + 3σ	Noise + σ	Noise-σ	Noise-3σ	Mean + 3σ	Mean + σ	Mean-σ	Mean-3σ	Boundary + σ	Boundary-σ	Non-Zero	Original	Repeat
Irradiation	**0.942**	**0.812**	**0.625**	**0.800**	**0.908**	**0.658**	**0.675**	**0.767**	**0.921**	**0.800**	**1.000**	**0.925**	0.083
Weather	0.792	0.542	0.358	0.558	0.792	0.529	0.508	0.562	0.792	0.588	0.904	0.825	**0.196**
None	0.808	0.558	0.429	0.700	0.808	0.529	0.529	0.679	0.775	0.721	0.967	0.883	0.079

methods relying solely on historical data without current measurements show reduced effectiveness (Previous 1: 60.8%, Previous 3: 57.6%), while single-point detection performs substantially worse (45.6%). This consistent pattern suggests that temporal context provides crucial information for accurate anomaly detection, with the combination of current and historical data offering optimal discriminative capacity.

Sampling method efficacy varies notably across different adversarial types. All methods incorporating temporal context perform exceptionally well on non-zero anomalies (92.5%-97.2%) and maximum boundary violations (84.2%-94.4%), likely due to these anomalies' distinct temporal signatures. However, for subtle manipulations like minimal noise insertion (23.1%-36.1%) and data repetition (11.1%-29.4%), even the best sampling approaches struggle to achieve high detection rates. The substantial performance differential between complex temporal sampling methods and single-point detection across most manipulation categories (up to 21.5%) highlights the critical importance of temporal context in anomaly detection systems, particularly for sophisticated attack vectors designed to mimic realistic operational patterns.

Conditional Variables. Conditional variables significantly affect detection capabilities (Table 5 and Table 6). Configurations incorporating irradiation data substantially enhance overall detection performance, outperforming both weather information and configurations without conditional features. This pattern aligns with photovoltaic system physics, where irradiance demonstrates a more direct relationship with power output than general weather conditions.

The substantial performance differential between irradiation-informed and other approaches suggests that direct solar resource measurements provide critical contextual information for anomaly detection algorithms, while the minimal difference between weather-based preprocessing and approaches without conditional features indicates limited utility of general meteorological data alone.

Table 7
Accuracy by contextual metadata and dataset.

Meatadata	REWC	SPG
Backgroud info	0.739	0.844
Month and hour	0.639	0.685
None	0.739	0.821

Table 8
Accuracy by contextual metadata and dataset (Historical Context: CurrPrev 1; Conditional Var: irradiation).

Meatadata	REWC	SPG
Backgroud info	71.79%	80.77%
Month and hour	60.77%	66.54%
None	71.92%	78.08%

Analysis across manipulation types reveals distinct effects of conditional variables. Irradiation data demonstrates superior performance for both positive and negative perturbations, with particularly enhanced detection capability for boundary violations. For non-zero anomalies, irradiation preprocessing achieves perfect detection compared to significantly lower performance with other feature configurations. However, for repetition patterns, weather features unexpectedly outperform irradiation data, suggesting that broader meteorological conditions may provide better differentiation for certain temporal anomalies—an important finding for comprehensive adversarial detection systems.

Contextual Metadata. Supplementary contextual metadata significantly influences detection performance, with distinct dataset-dependent patterns (Table 7 and Tab. Table 8). Background information about system specifications demonstrates superior overall performance

across both datasets, outperforming both temporal metadata and configurations without additional context. This pattern suggests that structural and operational characteristics provide more discriminative features for anomaly detection than temporal patterns alone.

The impact of contextual metadata exhibits dataset-specific variation, with background information providing a greater advantage in the SPG dataset compared to the REWC dataset. This dataset-dependent efficacy underscores that supplementary metadata value depends on its direct relevance to the specific generation context. Interestingly, temporal metadata (month and hour statistics) consistently underperforms both background information and configurations lacking metadata across both datasets, suggesting that certain types of contextual information may introduce noise that impedes effective anomaly detection.

6. Conclusion and future directions

This work presents DeCEN, a decentralized clean energy network that integrates edge computing and web3 technologies to address the limitations of traditional centralized systems. DeCEN enables localized energy verification through edge devices, connects stakeholders via a transparent blockchain ecosystem, and incorporates an LLM-based protocol to prevent fraud and stabilize token value. A novel distributed LLM inference algorithm ensures decentralized, low-latency processing while preserving data privacy. DeCEN's contributions include a robust architecture for decentralized energy networks, a verified LLM-based verification protocol, and an enhanced distributed LLM framework. This work demonstrates a scalable and trustworthy pathway toward global renewable energy targets.

Looking ahead, we will enhance DeCEN in two directions. First, GreenCoin will be upgraded from a static reward to a settlement-grade currency powering a full peer-to-peer marketplace where prosumers post bids and offers in GreenCoin, prices clear every minute against live feeder congestion data, and quarterly futures let households pre-sell surplus while factories hedge cloudy days, in which no utility or bank intermediaries required. Second, coin minting will be hardened through privacy-preserving cryptography: smart-meter readings are secret-shared among neighbor-edge nodes that jointly compute an aggregate kilowatt-hour attestation, then a succinct zero-knowledge proof convinces layer-2 validators that every new GreenCoin is fully backed by verifiable clean generation, yet household-level outputs and identities remain cryptographically hidden. Together, these upgrades turn DeCEN into a trustworthy, scalable and private energy-trading overlay that rewards every kilowatt of green power without compromising user confidentiality.

CRediT authorship contribution statement

Shan Jiang: Writing – original draft, Conceptualization, Methodology; **Wenchang Chai:** Writing – original draft, Data curation, Validation; **Mingjin Zhang:** Writing – review & editing, Methodology, Validation; **Jiannong Cao:** Writing – review & editing, Supervision, Funding acquisition; **Shichang Xuan:** Writing – review & editing, Funding acquisition; **Jiaxing Shen:** Funding acquisition, Data curation.

Data availability

No data was used for the research described in the article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the "Research on Oversea Web 3.0 Policies, Regulations, and Development Strategies" project under China Mobile Innovation and Research Institute, HK RGC Theme-based Research Scheme (No.T43-513/23-N), the Natural Science Foundation of Ningxia(No.2025AAC030535), the Fundamental Research Funds for the Central Universities (No.3072025ZH0604), Lingnan University (SDS24A17), and the Research Institute for Artificial Intelligence of Things, The Hong Kong Polytechnic University.

References

[1] A.G. Olabi, M.A. Abdelkareem, Renewable energy and climate change, Renew. Sustain. Energy Rev. 158 (2022) 112111.
[2] L. Chen, G. Msigwa, M. Yang, A.I. Osman, S. Fawzy, D.W. Rooney, P.-S. Yap, Strategies to achieve a carbon neutral society: a review, Environ. Chem. Lett. 20 (4) (2022) 2277–2310.
[3] A. Shinde, V. Mane, N. Ambulgekar, Z. Ansari, Renewable Energy Trading Platform using Blockchain, Procedia Comput. Sci. 233 (2024) 243–253.
[4] Y. Wang, Z. Su, N. Zhang, BSIS: Blockchain-based secure incentive scheme for energy delivery in vehicular energy network, IEEE Trans. Ind. Inf. 15 (6) (2019) 3620–3631.
[5] L. Orsini, S. Kessler, J. Wei, H. Field, How the brooklyn microgrid and transActive grid are paving the way to next-gen energy markets, in: The Energy Internet, Elsevier, 2019, pp. 223–239.
[6] M. Mihaylov, S. Jurado, N. Avellana, K. Van Moffaert, I.M. de Abril, A. Nowé, NRGcoin: virtual currency for trading of renewable energy in smart grids, in: 11th International Conference on the European Energy Market (EEM14), IEEE, 2014, pp. 1–6.
[7] S. Wang, A.F. Taha, J. Wang, K. Kvaternik, A. Hahn, Energy crowdsourcing and peer-to-peer energy trading in blockchain-enabled smart grids, IEEE Transactions on Systems, Man, and Cybernetics: Systems 49 (8) (2019) 1612–1623.
[8] Z. Guo, B. Qin, Z. Guan, Y. Wang, H. Zheng, Q. Wu, A high-efficiency and incentive-compatible peer-to-peer energy trading mechanism, IEEE Trans. Smart Grid 15 (1) (2023) 1075–1088.
[9] M. Andoni, V. Robu, D. Flynn, S. Abram, D. Geach, D. Jenkins, P. McCallum, A. Peacock, Blockchain technology in the energy sector: A systematic review of challenges and opportunities, Renew. Sustain. Energy Rev. 100 (2019) 143–174.
[10] J. Guerrero, A. Chapman, G. Verbic, A study of energy trading in a low-voltage network: Centralised and distributed approaches, in: 2017 Australasian Universities Power Engineering Conference (AUPEC), IEEE, 2017, pp. 1–6.
[11] M. Jené-Vinuesa, A. Abdullah, M. Aragüés-Peñalba, A. Sumper, Enhancing fraud detection in renewable energy grids through behind-the-meter PV disaggregation, in: 2024 IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE), IEEE, 2024, pp. 1–5.
[12] G.S. Erbuğa, Everything but Deterrence: Curbing Energy Fraud, in: Ethics and Sustainability in Accounting and Finance, Volume IV, Springer, 2024, pp. 149–164.
[13] M.J. Ashley, M.S. Johnson, Establishing a secure, transparent, and autonomous blockchain of custody for renewable energy credits and carbon credits, IEEE Eng. Manage. Rev. 46 (4) (2018) 100–102.
[14] Z. Li, J. Kang, R. Yu, D. Ye, Q. Deng, Y. Zhang, Consortium blockchain for secure energy trading in industrial internet of things, IEEE Trans. Ind. Inf. 14 (8) (2017) 3690–3700.
[15] Y. Zuo, Tokenizing renewable energy certificates (recs)—a blockchain approach for rec issuance and trading, IEEE Access 10 (2022) 134477–134490.
[16] T. AlSkaif, J.L. Crespo-Vazquez, M. Sekuloski, G. Van Leeuwen, J.P.S. Catalão, Blockchain-based fully peer-to-peer energy trading strategies for residential energy systems, IEEE Transactions on Industrial Informatics 18 (1) (2021) 231–241.
[17] A.-S. Berkani, H. Moumen, S. Benharzallah, M.T. Kechadi, A. Bounceur, PrGChain: a privacy-preserving blockchain-enabled energy trading system, Compute. Electr. Eng. 123 (2025) 110289.
[18] S. Jiang, J. Li, X. Zhang, H. Yue, H. Wu, Y. Zhou, Secure and privacy-preserving energy trading with demand response assistance based on blockchain, IEEE Trans. Netw. Sci. Eng. 11 (1) (2023) 1238–1250.
[19] J. Abdella, Z. Tari, R. Mahmud, MuLPP: A multi-level privacy preserving for blockchain-based bilateral P2P energy trading, J. Netw. Comput. Appl. 237 (2025) 104141.
[20] D. Hou, J. Zhang, S. Cui, J.K. Liu, Peer-to-peer energy trading with privacy and fair exchange, in: 2024 IEEE International Conference on Blockchain (Blockchain), IEEE, 2024, pp. 174–182.
[21] Z. Guan, X. Lu, W. Yang, L. Wu, N. Wang, Z. Zhang, Achieving efficient and Privacy-preserving energy trading based on blockchain and ABE in smart grid, J. Parallel Distrib. Comput. 147 (2021) 34–45.
[22] J.A. Abdella, Z. Tari, N. Sohrabi, R. Mahmud, MuLCOff: a multi-layer consensus and off-chain computation for efficient and privacy-aware blockchain-based peer-to-peer energy trading, IEEE Netw. 38 (5) (2024) 264–272.
[23] S. Jiang, J. Cao, C.L. Tung, Y. Wang, S. Wang, Sharon: Secure and efficient cross-shard transaction processing via shard rotation, in: IEEE INFOCOM 2024-IEEE Conference on Computer Communications, IEEE, 2024, pp. 2418–2427.
[24] R. Wang, Y. Chen, E. Li, L. Che, H. Xin, J. Li, X. Zhang, Joint optimization of energy trading and consensus mechanism in blockchain-empowered smart grids: a reinforcement learning approach, J. Cloud Comput. 12 (1) (2023) 121.

[25] J. Abdella, Z. Tari, R. Mahmud, N. Sohrabi, A. Anwar, A. Mahmood, Hicoob: Hierarchical concurrent optimistic blockchain consensus protocol for peer-to-peer energy trading systems, IEEE Trans. Smart Grid 14 (5) (2022) 3927–3943.
[26] U.R. Barbhaya, L. Vishwakarma, D. Das, Etradechain: blockchain-based energy trading in local energy market (lem) using modified double auction protocol, IEEE Trans. Green Commun. Network. 8 (1) (2023) 559–571.
[27] Z. Liu, B. Huang, Y. Li, Q. Sun, T.B. Pedersen, D.W. Gao, Pricing game and blockchain for electricity data trading in low-carbon smart energy systems, IEEE Trans. Ind. Inf. 20 (4) (2024) 6446–6456.
[28] Z. Bao, C. Tang, X. Xie, G. Chen, F. Lin, Z. Zheng, Toward green and efficient blockchain for energy trading: A noncooperative game approach, IEEE Internet Things J. 10 (22) (2023) 20021–20032.
[29] M. Zhang, F. Eliassen, A. Taherkordi, H.-A. Jacobsen, H.-M. Chung, Y. Zhang, Demand–response games for peer-to-peer energy trading with the hyperledger blockchain, IEEE Trans. Syst. Man Cybern. Syst. 52 (1) (2021) 19–31.
[30] H. Liao, G. Tang, D. Guo, Y. Wang, R. Cao, Rethinking low-carbon edge computing system design with renewable energy sharing, in: Proceedings of the 53rd International Conference on Parallel Processing, 2024, pp. 950–960.
[31] X. Yang, X. Guan, N. Wang, Y. Liu, H. Wu, Y. Zhang, Cloud-edge-end intelligence for fault-tolerant renewable energy accommodation in smart grid, IEEE Trans. Cloud Comput. 11 (2) (2021) 1144–1156.
[32] X. Chen, H. Wen, W. Ni, S. Zhang, X. Wang, S. Xu, Q. Pei, Distributed online optimization of edge computing with mixed power supply of renewable energy and smart grid, IEEE Trans. Commun. 70 (1) (2021) 389–403.
[33] G. Perin, M. Berno, T. Erseghe, M. Rossi, Towards sustainable edge computing through renewable energy resources and online, distributed and predictive scheduling, IEEE Trans. Netw. Service Manage. 19 (1) (2021) 306–321.
[34] C. Peng, D. Li, F. Tian, Y. Guo, Renewable energy powered IoT data traffic aggregation for edge computing, in: Communications, Signal Processing, and Systems: Proceedings of the 2018 CSPS Volume III: Systems 7th, Springer, 2020, pp. 861–869.
[35] W. Li, T. Yang, F.C. Delicato, P.F. Pires, Z. Tari, S.U. Khan, A.Y. Zomaya, On enabling sustainable edge computing with renewable energy resources, IEEE Commun. Mag. 56 (5) (2018) 94–101.
[36] Y. Li, A.-C. Orgerie, I. Rodero, M. Parashar, J.-M. Menaud, Leveraging renewable energy in edge clouds for data stream analysis in iot, in: 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID), IEEE, 2017, pp. 186–195.
[37] S. Jiang, J. Cao, J. Zhu, Y. Cao, PolyChain: a generic blockchain as a service platform, in: International Conference on Blockchain and Trustworthy Systems, Springer, 2021, pp. 459–472.
[38] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, Adv. Neural Inf. Process. Syst. 33 (2020) 1877–1901.
[39] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, et al., A survey on evaluation of large language models, ACM Trans. Intell. Syst. Technol. 15 (3) (2024) 1–45.
[40] X. Chen, J. Cao, Y. Sahni, S. Jiang, Z. Liang, Dynamic task offloading in edge computing based on dependency-aware reinforcement learning, IEEE Trans. Cloud Comput. 12 (2) (2024) 594–608.
[41] Y. Zheng, Y. Chen, B. Qian, X. Shi, Y. Shu, J. Chen, A review on edge large language models: Design, execution, and applications, ACM Comput. Surveys 57 (8) (2025) 1–35.
[42] M. Zhang, X. Shen, J. Cao, Z. Cui, S. Jiang, Edgeshard: Efficient llm inference via collaborative edge computing, IEEE Internet Things J. (2024).
[43] V. Jakkula, Tutorial on support vector machine (svm), School EECS, Washington State Univ. 37 (2.5) (2006) 3.
[44] L. Breiman, J. Friedman, C.J. Stone, R.A. Olshen, Classification and Regression Trees, Taylor & Francis, 1984.
[45] L. Breiman, Random Forests, Mach. Learn. 45 (1) (2001) 5–32.
[46] D.W. Hosmer, S. Lemeshow, R.X. Sturdivant, Applied Logistic Regression, Wiley Series in Probability and Statistics, John Wiley & Sons, 2013.