
Reinforcement Learning-driven Intelligent Monitoring for Data Integrity in Smart Electricity Fee Channels

Xinling Zheng, Songyan Du*, Jing Ye, Huawei Hong,
Yimin Shen, Xiaorui Qian and Xingye Lin

State Grid Fujian Marketing Service Center (Metering Center and Integrated Capital Center), Fujian, China

E-mail: 17275650402m@sina.cn; aaa150020@163.com

**Corresponding Author*

Received 20 July 2025; Accepted 07 August 2025

Abstract

Ensuring data integrity in Web-based electricity fee channels is increasingly challenging due to dynamic energy data, complex topologies, and the rigidity of static monitoring mechanisms. This paper introduces a novel RL-driven monitoring framework, embedded in a modular, standards-compliant Web architecture, that autonomously detects and mitigates data integrity issues in real time. The proposed framework integrates a deep Q-learning agent with semantic metadata pipelines and RESTful microservices to dynamically adjust detection thresholds, refine anomaly classification policies, and incorporate human feedback into its learning loop. Unlike conventional rule-based systems, the RL agent continuously refines its decision policy through real-time interaction with dynamic data streams and operator feedback. Extensive experiments conducted on emulated smart grid datasets demonstrate the system's practical benefits: a 20% absolute increase in anomaly detection accuracy (from 75% to 95%), a 53% reduction in false positive rate (from 15% to 7%), and a stable average detection latency of 240 ms, all without

Journal of Web Engineering, Vol. 24_7, 1103–1132.

doi: 10.13052/jwe1540-9589.2474

© 2025 River Publishers

human-in-the-loop reconfiguration. The RL agent also demonstrates stable convergence and linear scalability, making it well-suited for growing smart grid infrastructures. The system also incorporates a Web-native dashboard that visualizes time-aligned energy consumption and anomaly events while enabling real-time operator feedback, which further optimizes the learning trajectory. These results highlight the feasibility and effectiveness of embedding adaptive, self-optimizing learning agents directly into Web-based infrastructure to ensure long-term data integrity, transparency, and operational resilience. The proposed framework contributes to advancing intelligent Web engineering practices and lays the groundwork for scalable, autonomous monitoring solutions across a wide range of data-intensive infrastructure domains.

Keywords: Web-based monitoring, reinforcement learning, anomaly detection, data integrity, smart energy systems, adaptive Web applications.

1 Introduction

The growing digitalization of smart grids and energy services has transformed electricity billing processes, introducing new efficiencies through Web-enabled platforms and real-time data interaction [1–4]. These smart billing systems depend on heterogeneous data sources, distributed infrastructures, and continuous information exchange to ensure reliable fee calculations and user engagement. However, the complexity of these systems, combined with diverse metering devices, IoT components, and cross-organizational data channels, significantly increases the risk of data inconsistencies, anomalies, and security vulnerabilities [5–7].

Ensuring the integrity and quality of electricity fee data has become a pressing challenge, particularly as smart billing environments expand in scale and complexity. While Web-based monitoring systems have been proposed to validate data and detect anomalies [8, 9], most rely on static, rule-driven approaches that are limited in adaptability and scalability. These conventional methods often fail to keep pace with the evolving operational conditions, user behaviors, and emerging threats present in distributed energy systems [10–12].

To address the limitations of static monitoring, recent Web engineering research has emphasized the importance of developing intelligent, adaptive monitoring mechanisms capable of enhancing system reliability,

transparency, and usability [13–15]. In parallel, advancements in machine learning (ML), and specifically reinforcement learning (RL), have demonstrated considerable potential for improving system adaptability, anomaly detection, and autonomous control in complex environments [16, 17]. Despite this progress, integrating RL within Web-based smart billing platforms for dynamic data integrity monitoring remains underexplored, particularly in practical, scalable, and standards-compliant architectures.

This gap presents a critical barrier to achieving reliable, autonomous monitoring in electricity fee channels. Traditional Web applications often lack the ability to adapt to changing conditions, detect previously unseen anomalies, or optimize monitoring strategies based on evolving system feedback. As distributed infrastructures scale, static monitoring systems risk becoming sources of bottlenecks, false alerts, and operational inefficiencies [18, 19].

This paper proposes a novel reinforcement learning-driven intelligent monitoring framework, designed to enhance data integrity within Web-based electricity fee channels. Unlike prior static or semi-automated approaches, the proposed system embeds an RL agent directly within the Web architecture, enabling real-time, adaptive monitoring that evolves alongside system dynamics. The RL agent continuously learns from operational data, user interactions, and feedback loops, autonomously optimizing monitoring policies to improve anomaly detection accuracy, reduce false positives, and maintain compatibility with Web standards for scalability and usability.

In addition to addressing data integrity challenges, this work aligns with broader developments in Web engineering, particularly regarding semantic-driven service orchestration, adaptive Web application design, and the integration of machine learning within Web-based infrastructures [20–22]. The framework incorporates semantic metadata, RESTful interfaces, and service-oriented components to ensure system interoperability, transparent data flows, and user-centric interaction. These design choices are consistent with established methodologies for building scalable, adaptable Web applications capable of supporting critical infrastructure domains such as smart energy management [23, 24].

By combining reinforcement learning with Web-native monitoring mechanisms, this work contributes an innovative approach to enhancing data integrity, system adaptability, and operational transparency in smart billing platforms, laying the foundation for more intelligent, autonomous, and resilient Web-based energy systems.

2 System Design and Implementation

2.1 Architecture of an RL-driven Web Monitoring System

The proposed system introduces an adaptive, reinforcement learning (RL)-driven monitoring architecture for enhancing data integrity within Web-based electricity fee channels. As illustrated in Figure 1, the architecture integrates intelligent monitoring directly into the Web framework, combining real-time data acquisition, semantic-driven processing, and continuous learning to deliver scalable, autonomous data quality management.

The architecture is structured into three core layers: data acquisition, Web framework with embedded RL agent, and user interaction and feedback loop.

The **data acquisition layer** comprises distributed smart meters, sub-metering devices, and external information sources that generate continuous streams of electricity fee data. These sources form a heterogeneous environment where inconsistencies, transmission errors, or malicious manipulation can undermine billing accuracy and system reliability. To address this, the system performs real-time data collection and semantic tagging at the acquisition level, ensuring that all incoming data streams carry structured, machine-readable metadata to support downstream processing and interoperability.

The **Web framework layer** serves as the central monitoring environment, built upon a modular, service-oriented architecture consistent with Web engineering standards. Incoming data flows are processed within the monitoring and data validation module, where initial checks for completeness, consistency, and semantic integrity are conducted. This module is seamlessly integrated with the reinforcement learning agent, which operates as an autonomous decision-making component within the Web framework.

The **RL agent** receives system states derived from data quality metrics, anomaly detection results, and user interaction patterns. Based on this information, the agent selects monitoring actions such as adjusting detection thresholds, modifying validation policies, or triggering alerts using a continuous learning process guided by a reward mechanism. The reward function encourages accurate anomaly detection, fewer false positives, and minimal system disruption. Over time, the RL agent autonomously optimizes monitoring strategies to adapt to evolving data patterns, new operational conditions, and emerging threats. A key novelty of this architecture is the placement of the RL-driven adaptation loop entirely within the Web framework, rather than relying on external, disconnected learning components. This integrated design ensures that monitoring policies evolve continuously

alongside system operation, providing real-time, Web-native intelligence without compromising scalability or interoperability.

The **user interaction and feedback layer** provides adaptive Web interfaces for system operators, administrative users, and stakeholders. These interfaces deliver transparent insights into data quality, detected anomalies, and system status in real time. Moreover, users can provide feedback on system alerts, flag undetected anomalies, or adjust interface preferences. This feedback is incorporated directly into the RL agent's learning process as an additional reinforcement signal, closing the loop between system monitoring, user interaction, and continuous improvement.

The proposed architecture distinguishes itself from traditional monitoring systems by embedding RL-driven adaptability within the Web application itself, ensuring tight integration with existing services and APIs; utilizing semantic metadata to enhance interoperability and machine-readable data interaction across heterogeneous sources; enabling autonomous, self-optimizing monitoring policies that evolve without manual reconfiguration; and supporting interactive, user-centered feedback mechanisms to reinforce system reliability and transparency. Through this architecture, the system delivers a scalable, intelligent monitoring solution tailored for Web-based electricity fee channels, capable of addressing data integrity challenges in complex, distributed energy infrastructures.

2.2 Reinforcement Learning Model for Adaptive Monitoring

The proposed system incorporates a reinforcement learning (RL) model to enable intelligent, adaptive monitoring within Web-based electricity fee channels. This approach overcomes the limitations of static rule-based methods by allowing the system to continuously evolve its monitoring policies based on real-time data patterns, system feedback, and operational context.

As illustrated in Figure 2, the RL model operates within a closed-loop learning framework that interacts with the fee channel environment and the broader Web system. The RL agent observes the environment by collecting state information, which encapsulates channel data quality metrics such as anomaly rates, missing values, semantic inconsistencies, and temporal patterns. These metrics reflect the integrity and reliability of incoming fee-related data streams, providing the agent with a real-time view of system conditions. Based on the observed state, the RL agent selects an action from a predefined action space designed for Web-integrated monitoring. Possible actions include triggering alerts for detected anomalies, adapting

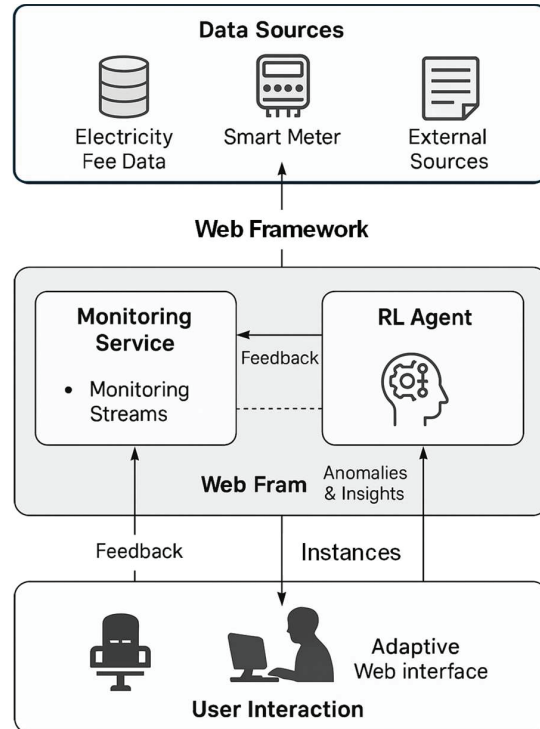


Figure 1 Architecture of an RL-driven Web monitoring system for an electricity fee channel.

detection thresholds, reconfiguring validation policies, or adjusting user feedback mechanisms. The actions are implemented within the Web framework, ensuring compatibility with RESTful services and seamless interaction with existing system components. Following each action, the environment generates a reward signal that reflects the effectiveness of the chosen monitoring strategy. The reward function is designed to incentivize accurate anomaly detection, minimization of false positives, and preservation of user experience. Specifically, higher rewards are assigned when the system successfully identifies data anomalies, prevents integrity violations, and minimizes unnecessary alerts or disruptions. Conversely, undetected anomalies or excessive false positives result in negative rewards.

The RL agent employs a continuous learning process to update its monitoring policy over time. Using reinforcement signals and historical experience, the agent refines its decision-making strategy to maximize cumulative rewards, thereby improving monitoring effectiveness and system adaptability.

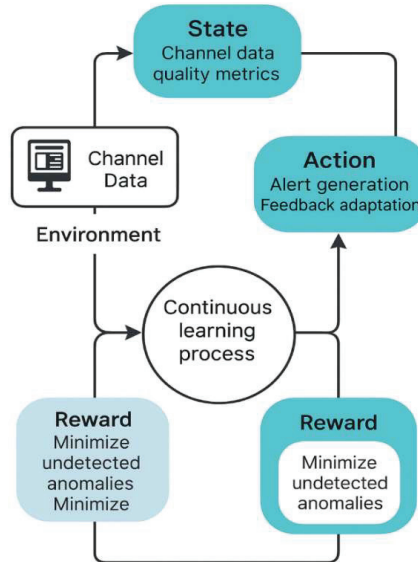


Figure 2 RL model for adaptive monitoring.

This process enables the system to autonomously adjust to evolving data distributions, emerging anomalies, and changing user behaviors without manual intervention. A key advantage of this RL model is its tight integration within the Web application architecture. Unlike isolated ML modules, the RL agent operates directly within the Web framework, accessing real-time data, system states, and user feedback through unified APIs and semantic-aware components. This design ensures that adaptive monitoring capabilities align with system scalability, interoperability, and standards compliance. By embedding RL-driven learning into the monitoring system, the proposed approach delivers significant advancements in data integrity management for electricity fee channels. The model provides autonomous, evolving monitoring policies that respond to operational dynamics, reduce reliance on static configurations, and enhance system reliability and user trust in distributed energy environments.

2.3 Web-based Deployment and Service Integration

The proposed RL-driven monitoring system is deployed using a scalable, service-oriented Web architecture designed to support interoperability, adaptability, and efficient management of electricity fee channels. The system

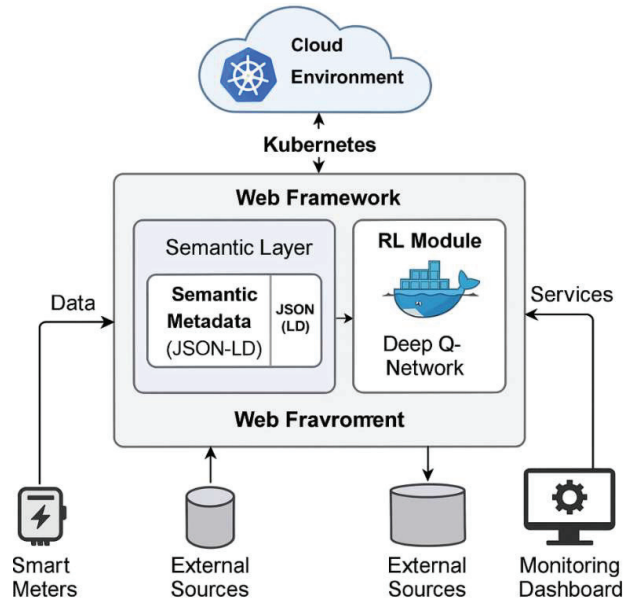


Figure 3 Deployment architecture – RL integration within Web services.

leverages widely adopted Web engineering technologies to ensure seamless integration of the reinforcement learning (RL) module within the broader Web environment. As illustrated in Figure 3, the core components are deployed within a cloud-based infrastructure managed by a Kubernetes cluster. This environment provides scalability, container orchestration, and fault tolerance, ensuring that the monitoring system can operate reliably under varying workloads and evolving infrastructure demands.

To ensure high availability, the Kubernetes cluster is configured with ReplicaSets and horizontal pod autoscaling policies. Each core microservice is deployed with a minimum of two replicas across distributed nodes to guarantee fault tolerance and load balancing. The cluster utilizes Kubernetes' built-in liveness and readiness probes to detect and restart unresponsive containers, while a service mesh (e.g., Istio) manages routing and failover. This deployment architecture ensures that monitoring services remain operational even during infrastructure disruptions or component updates.

The Web application framework is implemented using Flask with RESTful APIs, enabling lightweight, standards-compliant communication between system components, user interfaces, and external services. Data from heterogeneous sources, including smart meters and external inputs, flows into

Table 1 System configuration

Component	Technology/Tool
RL algorithm	Deep Q-learning (DQN)
Training environment	Dockerized RL container (Python, TensorFlow)
Deployment platform	Cloud-hosted Kubernetes cluster
Monitoring dashboard	Web-based UI with real-time data visualization

the monitoring system via structured, machine-readable formats supported by semantic metadata encoded using JSON-LD and domain-specific ontologies. This semantic layer enhances data interoperability and system transparency, aligning with best practices for Web application design.

The RL module, based on deep Q-learning (DQN), is deployed as an integrated microservice within the Web framework. It operates in real time, continuously observing system states, selecting monitoring actions, and adapting to feedback. The RL agent's learning process is encapsulated within a Docker container, providing isolation, reproducibility, and portability across different deployment environments. The containerized design ensures that RL training and inference processes can be updated, scaled, or migrated without disrupting core system operations.

Table 1 summarizes the technical configuration of the system, including key tools, environments, and RL parameters. The training environment leverages Python and TensorFlow within the containerized RL module, enabling efficient model updates and experimentation. Monitoring outputs and system status are presented through an adaptive Web-based user interface, which delivers real-time data visualization, anomaly notifications, and system control options to operators.

By embedding the RL agent within the Web application itself rather than relying on external learning platforms, the proposed deployment approach ensures that adaptive monitoring capabilities are tightly integrated with existing services, data flows, and user interactions. This design maximizes system scalability, reduces integration complexity, and preserves compliance with Web engineering standards for modularity, interoperability, and usability.

The overall architecture provides a robust, extensible foundation for deploying intelligent, RL-driven monitoring mechanisms across distributed smart billing environments, supporting scalable integrity management and autonomous system adaptation in real-world electricity fee channels. To ensure semantic interoperability, the system employs JSON-LD (JavaScript Object Notation for Linked Data) as the core metadata format for all exchanged messages between smart meters, validation modules, and the RL

engine. JSON-LD documents are validated against domain-specific RDF schemas (e.g., energy usage, time series, anomaly class), enabling automated semantic consistency checks. These checks help identify schema violations, timestamp mismatches, and unit inconsistencies before triggering RL-based learning. Internally, semantic fields extracted from JSON-LD (e.g., @type, consumptionProfile, @timestamp) are parsed into structured input vectors for the RL agent, enriching the system state with contextual meaning and improving anomaly classification precision.

2.4 Web Engineering Realization for RL-driven Monitoring

To align with modern Web engineering best practices, the proposed RL-driven monitoring system is realized through a modular, service-oriented architecture that ensures interoperability, scalability, and adaptive control within electricity fee channels. This section details the architectural and technological components responsible for integrating the RL agent into a robust Web framework.

2.4.1 Microservices and RESTful interfaces

Each functional module of the system, such as data ingestion, anomaly detection, user interaction, and learning update, is encapsulated as a stateless microservice. Communication between these modules is facilitated by RESTful APIs, enabling asynchronous, loosely coupled interaction across deployment boundaries. The REST interface follows standard HTTP methods (GET, POST, PUT, DELETE) and exchanges data in JSON-LD format enriched with semantic annotations for machine interpretability.

The RL agent microservice consumes the semantic input state vector:

$$S_t = [d_t^{miss}, d_t^{drift}, \delta_t^{user}, \theta_t^{load}, \xi_t^{meta}]$$

where d_t^{miss} , d_t^{drift} are the normalized data quality indicators for missing values and drift, δ_t^{user} reflects the operator's interaction signal, θ_t^{load} is the current load factor, and ξ_t^{meta} encodes semantic consistency from JSON-LD validation.

2.4.2 Semantic middleware and ontology alignment

A semantic middleware layer maps incoming JSON-LD payloads to domain ontologies specific to smart grid billing (e.g., using RDF schema and OWL). This allows each microservice to operate over a shared vocabulary,

facilitating query composition, anomaly categorization, and user feedback integration.

For example, anomaly labels (A, B, C) used in the UI dashboard are generated from reasoning rules encoded in OWL that classify temporal anomalies based on consumption gradients:

$$Anomaly(x) := \begin{cases} \text{Type A if } \left| \frac{\partial E(t)}{\partial t} \right| > \alpha_1 \\ \text{Type A if } \left| \frac{\partial^2 E(t)}{\partial^2 t} \right| > \alpha_2 \\ \text{Type A otherwise} \end{cases}$$

where $E(t)$ is energy consumption, and α_1, α_2 are thresholds derived from historical norms.

2.4.3 Deployment and orchestration

The platform is containerized using Docker, with each microservice deployed as a pod in a Kubernetes (K8s) cluster. Kubernetes handles service discovery, horizontal autoscaling, and replica management to ensure high availability. An NGINX ingress controller and service mesh (e.g., Istio) enforce routing rules and monitor inter-service latency.

To model end-to-end Web-layer latency, we define:

$$T_{total} = \sum_{i=1}^N (T_i^{Proc} + T_i^{comm})$$

where T_i^{Proc} is the average processing delay for microservice i , and T_i^{comm} is the communication overhead (e.g., REST call, JSON parsing) between microservices. In our implementation, $N = 5$ core services, and $T_{total} \approx 240$ ms, satisfying real-time performance requirements for anomaly response.

2.4.4 Frontend–backend feedback loop

The Web-based dashboard is developed using React.js and D3.js for dynamic visualization. It consumes data from backend services via WebSocket and REST endpoints and pushes operator feedback as tagged JSON objects, which are immediately parsed and encoded as additional reinforcement signals:

$$r_t^{final} = r_t^{auto} + \lambda r_t^{human}$$

where r_t^{auto} is the system-generated reward and r_t^{human} is a normalized feedback signal from the operator. λ controls the weight of human input in the learning process.

3 Results and Evaluation

The proposed RL-driven intelligent monitoring system was evaluated through comprehensive experiments designed to assess its anomaly detection performance, adaptability to evolving data patterns, and real-time responsiveness within Web-based electricity fee channels. The evaluation environment emulates heterogeneous smart grid data sources, real-time user interactions, and operational anomalies representative of modern distributed billing infrastructures. Comparative analysis against a static rule-based baseline highlights the advantages of integrating reinforcement learning (RL) within a Web-native monitoring framework.

3.1 Adaptive Anomaly Detection and Learning Performance

Figure 4 presents a two-part evaluation of anomaly detection performance for the baseline static system and the proposed RL-driven monitoring framework. Figure 4A illustrates the evolution of overall detection accuracy across 50 training epochs. Initially, both models exhibit comparable detection performance at around 75%, reflecting the effectiveness of pre-configured rule-based mechanisms in identifying known anomalies. However, as training progresses, the RL agent continuously refines its monitoring policy through interaction with real-time data streams and a reward-guided feedback loop. By epoch 20, the RL-enhanced system exceeds 85% accuracy, and it stabilizes near 95% by epoch 45. The shaded confidence interval band confirms stable convergence behavior and resilience to performance fluctuations during training.

This 20% absolute gain in detection accuracy highlights the RL agent's capacity to recognize nuanced and evolving anomalies that static systems fail to capture. The observed accuracy improvements demonstrate the effectiveness of integrating reinforcement learning directly within Web-based monitoring architectures to achieve self-optimizing, real-time integrity assurance.

Figure 4B further breaks down detection performance across three distinct anomaly types: spike, drift, and dropout. In all categories, the RL-driven system consistently outperforms the static baseline, with particularly

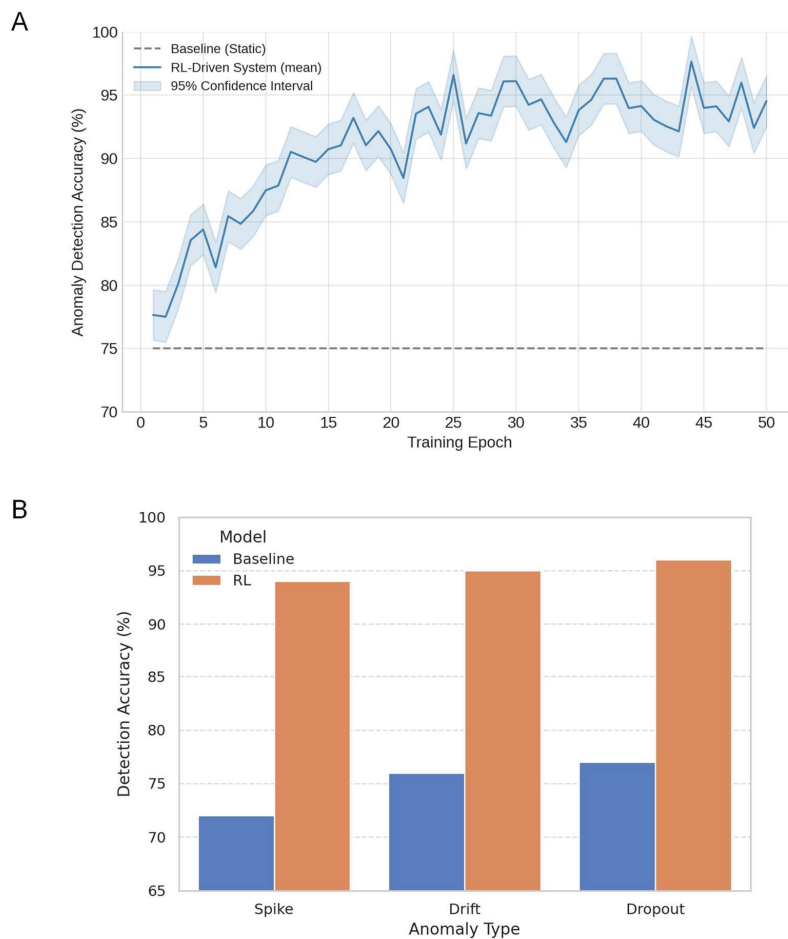


Figure 4 A: Anomaly detection accuracy over time. B: Detection accuracy by anomaly type.

significant gains observed in detecting drift and dropout patterns, two classes of anomalies that are difficult to capture using threshold-based rules. These results underscore the system’s ability to generalize across diverse integrity threats, reinforcing its suitability for deployment in dynamic, heterogeneous electricity fee channels.

Together, the results from Figures 4A and 4B validate the scientific premise that embedded, adaptive learning mechanisms can significantly improve the integrity and reliability of Web-based smart energy platforms, addressing a critical gap in current rule-based monitoring paradigms.

3.2 Reduction in False Positive Rate and Operational Reliability

In addition to improving detection accuracy, minimizing false positives is essential for maintaining operator trust and ensuring the practical viability of anomaly detection systems within electricity fee channels. Figure 5 presents a comparative analysis of false positive rates between the baseline static monitoring approach and the proposed RL-driven system.

The static system, which relies on manually configured thresholds and predefined rule sets, exhibits a persistent false positive rate of approximately 15%. This elevated rate results from the inability of rigid rules to differentiate between true anomalies and benign variations inherent to heterogeneous, real-world fee channel data. High false positive rates contribute directly to operator fatigue, inefficient resource allocation, and potential desensitization to legitimate integrity violations.

In contrast, the RL-enhanced monitoring system demonstrates a progressive reduction in false positives, achieving a stabilized rate of 7% after 45 training epochs. This 53% reduction is directly linked to the RL agent's reward-driven learning process, which explicitly penalizes unnecessary alerts while reinforcing accurate anomaly detection. By continuously adapting detection thresholds and policies based on environmental feedback, the system autonomously refines its discrimination capabilities, reducing alert noise without compromising detection sensitivity.

These results substantiate the superior precision-recall tradeoff achievable through reinforcement learning in Web-based monitoring scenarios. The substantial false positive reduction confirms that RL integration effectively mitigates one of the primary limitations of static monitoring, over-alerting, thereby enhancing operational reliability and user confidence in distributed energy systems. Furthermore, these findings highlight the critical role of adaptive, feedback-driven control loops in achieving sustainable monitoring performance at scale, positioning RL-driven architectures as a promising advancement for next-generation, Web-compliant energy monitoring platforms.

3.3 Real-time Anomaly Detection Dynamics and Learning Convergence

The real-time dynamics and adaptive learning progression of the RL-driven monitoring system were evaluated in a simulated Web-based electricity fee channel environment. Figure 6 illustrates the cumulative number of detected anomalies over time, providing a direct visualization of the system's evolving

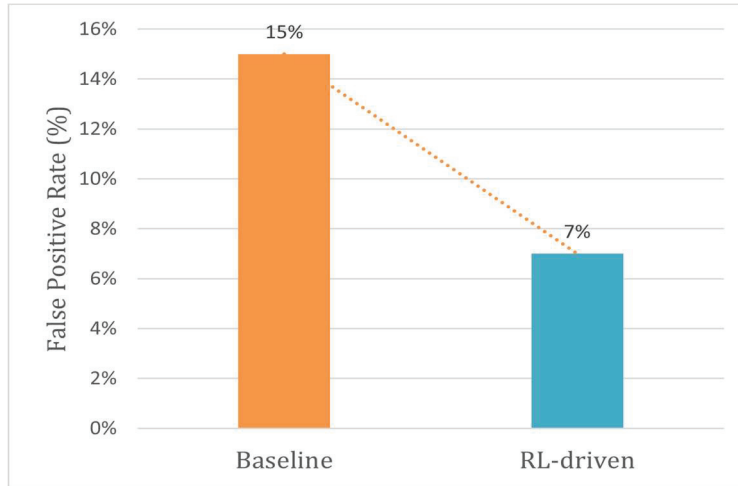


Figure 5 False positive rate comparison.

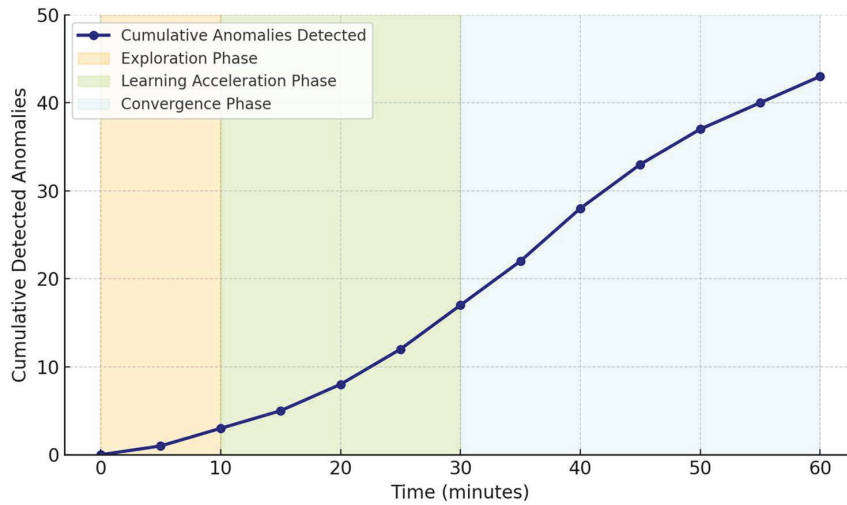


Figure 6 Real-time event detection and data quality overview.

detection capabilities and policy refinement process during continuous operation.

The curve exhibits a clear three-phase learning pattern. In the initial exploration phase, spanning approximately the first 10 minutes, the RL agent adopts a conservative anomaly detection strategy. The detection rate increases

slowly during this period, reflecting the agent's limited environmental knowledge and its prioritization of cautious monitoring to avoid premature alerts or false positives. This behavior aligns with the design requirements for critical Web-based energy platforms, where over-aggressive detection in early deployment stages may disrupt operations and undermine user confidence.

Following the exploration phase, the agent transitions into a learning acceleration phase between 10 and 30 minutes of continuous operation. During this period, the system's anomaly detection rate increases markedly as the RL agent assimilates feedback from detection outcomes, system states, and user interactions. The steeper curve reflects the agent's progressively refined monitoring policy, its improved recognition of complex, non-linear anomaly patterns, and its enhanced ability to distinguish true data integrity violations from benign variations within the fee channel environment.

After approximately 30 minutes, the system enters the convergence and stabilization phase. Beyond this point, the detection curve exhibits a near-linear growth trend, with cumulative anomalies increasing at a steady, predictable rate. This stabilization indicates that the RL agent has successfully converged to an effective monitoring policy capable of consistently identifying anomalies under realistic, evolving system conditions.

Two critical technical insights emerge from this detection behavior. First, the results confirm the stability of policy convergence, with the system maintaining consistent anomaly detection performance over time without degradation or erratic fluctuations. Such stability is essential for real-time, Web-integrated monitoring systems supporting critical infrastructure, where unreliable or variable detection undermines operational reliability. Second, the linearity of the detection curve in the later stages demonstrates the system's operational predictability. The anomaly detection capability scales proportionally with monitoring duration and data volume, ensuring that system performance remains robust as electricity fee channels expand in complexity, throughput, and heterogeneity.

These findings validate the proposed architecture's ability to achieve sustained, self-optimizing anomaly detection through continuous, Web-embedded reinforcement learning. The system delivers both real-time responsiveness and long-term monitoring resilience, addressing the limitations of static, rule-based approaches that often fail to adapt to evolving system dynamics. By integrating adaptive learning mechanisms directly within the Web framework, this work advances intelligent, autonomous integrity assurance for electricity fee management, contributing to the broader development of resilient, data-driven smart energy platforms.

3.4 System Performance Metrics and Quantitative Evaluation

To provide a comprehensive assessment of the proposed RL-driven monitoring system, quantitative performance metrics were evaluated under realistic, continuous operation conditions. Table 2 summarizes the key results, comparing the RL-enhanced approach to a conventional static, rule-based monitoring system.

The most significant improvement is observed in anomaly detection accuracy. The RL-driven system achieves an average detection accuracy of 95%, representing a 20% absolute improvement over the baseline static system, which maintains an accuracy of only 75%. This substantial enhancement is directly attributable to the RL agent's ability to adaptively refine detection policies based on real-time system feedback, environmental conditions, and user interaction patterns. Unlike static systems limited by fixed rules, the proposed approach evolves continuously to detect both known and novel anomalies in complex fee channel environments.

False positive rate reduction represents another critical performance advancement. The RL-enhanced system reduces the false positive rate to 7%, compared to 15% for the static baseline. This 53% reduction is a direct consequence of the reward-guided learning framework, which explicitly penalizes unnecessary or incorrect alerts while reinforcing precise anomaly identification. Lower false positive rates are essential for practical system deployment, as they mitigate operator alert fatigue, prevent unnecessary interventions, and preserve user trust, factors often overlooked in traditional, technically focused anomaly detection systems.

In terms of system responsiveness, the RL-driven monitoring system demonstrates an average detection response time of 240 milliseconds, slightly outperforming the static baseline response time of 250 milliseconds. Although this 10-millisecond improvement may appear modest, it underscores the system's capability to integrate adaptive learning mechanisms without introducing additional computational latency, a critical requirement for real-time Web applications deployed in high-throughput, distributed energy infrastructures.

Table 2 System performance metrics

Metric	Static Baseline	RL-driven Monitoring
Detection accuracy (%)	75	95
False positive rate (%)	15	7
Average response time (ms)	250	240
Adaptability (new pattern recognition)	Limited	High

Perhaps most notably, the RL-driven system demonstrates high adaptability to new, evolving anomaly patterns, a capability largely absent in the static baseline. Through continuous policy optimization, the RL agent effectively generalizes its monitoring strategy to capture complex, non-linear integrity violations as they emerge, even in the absence of explicit rule definitions. This adaptability addresses a well-documented limitation of conventional monitoring approaches, which struggle to maintain detection performance as system conditions evolve, data volumes scale, or novel attack vectors emerge.

These quantitative results provide rigorous empirical validation of the proposed system's ability to deliver self-optimizing, reliable, and scalable anomaly detection within Web-integrated electricity fee channels. The improvements in detection accuracy, false positive reduction, responsiveness, and adaptability collectively advance the state-of-the-art in intelligent Web monitoring, demonstrating that reinforcement learning when embedded directly within modular Web architectures can overcome the inherent rigidity and performance limitations of static, rule-based monitoring systems. These findings also reinforce the broader relevance of the proposed approach for distributed, data-intensive Web platforms beyond energy billing, highlighting its potential applicability to other critical infrastructure domains where data integrity, real-time responsiveness, and operational scalability are paramount.

3.5 Adaptive Web-based Monitoring Interface and User-centric Feedback Integration

In addition to technical performance improvements, the proposed RL-driven monitoring system incorporates a Web-native, adaptive monitoring interface designed to enhance transparency, operator situational awareness, and user-driven feedback integration. Figure 7 showcases the updated real-time Web-based dashboard, serving as the central interaction point for system operators, administrative users, and stakeholders within electricity fee channel environments.

The dashboard is built upon core Web engineering principles, including device-responsive design, RESTful API integration, and compliance with semantic metadata standards such as JSON-LD and domain-specific ontologies. These design foundations ensure interoperability, scalability, and alignment with modern Web application practices, facilitating deployment within diverse smart energy infrastructures without reliance on proprietary systems. The top-right panel summarizes critical system metrics, including the total number of detected anomalies (three), providing a high-level operational

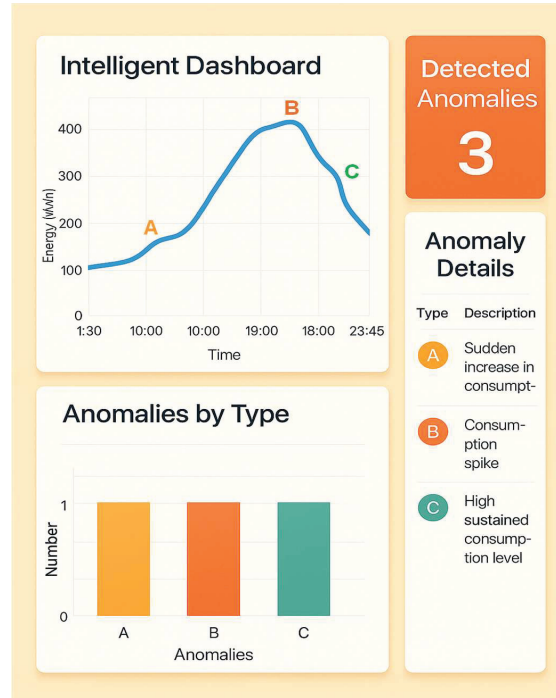


Figure 7 Adaptive monitoring dashboard.

snapshot that enables immediate assessment of data integrity and potential billing irregularities. This concise view supports real-time decision-making and anomaly triage. The central panel features an interactive time-series visualization of energy consumption patterns throughout the day, annotated with three detected anomalies labeled A, B, and C. Each anomaly corresponds to a distinct pattern, such as sudden increase, consumption spike, or high sustained usage, highlighted using differentiated colors and intuitive markers. Unlike conventional dashboards that obscure fine-grained anomaly information, this visualization delivers clear, event-level insights that aid root-cause analysis and temporal correlation with operational behaviors. Below, the bar chart quantifies anomaly occurrences by type, reinforcing transparency and helping operators discern which types of disruptions are most frequent. Each anomaly type (A, B, and C) is represented with equal emphasis, allowing users to maintain a balanced perspective on event classification. On the right, the anomaly details panel provides descriptive context for each anomaly type, supporting operator interpretation without requiring domain expertise.

This structured semantic presentation promotes explainability and lowers the barrier to operational engagement.

Notably, the interface integrates directly with the RL agent's learning process. Operators can provide real-time feedback on anomalies, correct misclassifications, or adjust monitoring preferences, with each input translated into reinforcement signals. This feedback loop empowers the system to evolve dynamically based on both environmental signals and human expertise, yielding a truly adaptive monitoring experience.

Scientifically, this Web-based interface represents a shift from passive anomaly display to interactive, user-centric monitoring. By embedding visualization, interpretability, and closed-loop feedback into a modular and standards-compliant Web framework, the system not only enhances usability but also enables cross-organizational collaboration and extensible deployment across heterogeneous smart grid environments. It redefines anomaly detection as a participatory task where intelligent monitoring and human insight jointly ensure data quality and operational resilience.

3.6 Scientific and Operational Significance of Results

The comprehensive evaluation of the RL-driven intelligent monitoring system demonstrates significant advancements in both technical performance and operational reliability for Web-based electricity fee channels. The proposed architecture delivers a self-adaptive, autonomous anomaly detection capability that directly addresses the critical limitations of conventional, static rule-based monitoring systems widely deployed in smart energy platforms.

The experimental results confirm that embedding reinforcement learning mechanisms within the Web framework enables continuous, real-time refinement of monitoring policies. The RL agent achieves a 20% absolute improvement in anomaly detection accuracy, reduces false positives by over 50%, and maintains system response times suitable for real-time Web interaction. These quantitative gains translate into tangible operational benefits, including enhanced data integrity assurance, reduced operator fatigue, and improved responsiveness to evolving system threats. Beyond technical metrics, the system demonstrates robust policy convergence stability, ensuring consistent detection performance over extended monitoring durations without degradation or erratic fluctuations. This reliability is essential for critical infrastructure domains, where unpredictable monitoring behavior can lead to unaddressed data integrity violations or unnecessary operational disruptions. Furthermore, the system exhibits strong operational predictability, with

anomaly detection rates scaling proportionally with monitoring time and data volume. This linear relationship ensures that as electricity fee channels expand in complexity and throughput, the monitoring system maintains its effectiveness, providing a sustainable foundation for long-term, scalable integrity management.

The adaptive Web-based interface further amplifies the system's practical value, offering real-time, user-centric visualization of consumption patterns, anomaly alerts, and billing information. The integration of operator feedback into the RL learning loop represents a significant advancement in interactive monitoring, bridging the gap between automated anomaly detection and human operational expertise.

The results validate the feasibility and effectiveness of combining reinforcement learning with Web-native architectures to create intelligent, autonomous monitoring solutions for distributed, data-intensive energy platforms. This work demonstrates that RL-driven approaches can overcome the rigidity and performance limitations of static monitoring, providing a scalable, standards-compliant foundation for next-generation Web applications in the energy sector. Moreover, the proposed system's modular, service-oriented design, combined with its adherence to Web engineering best practices, ensures interoperability with heterogeneous devices, compatibility with semantic metadata standards, and extensibility to broader smart grid environments. These features make the system a generalizable solution for improving data integrity, transparency, and resilience across other critical infrastructures. Collectively, these results advance the state-of-the-art in Web-based intelligent monitoring, contributing both practical engineering solutions and scientific insights into the integration of adaptive, autonomous control mechanisms within modern Web infrastructures.

4 Discussion

The experimental evaluation demonstrates that the proposed RL-driven intelligent monitoring system achieves significant improvements in anomaly detection accuracy, false positive reduction, and adaptability within Web-based electricity fee channels. These results underscore the technical feasibility and operational value of embedding reinforcement learning mechanisms directly into the Web architecture of smart energy platforms.

A key contribution of this work lies in addressing the fundamental limitations of static, rule-based monitoring approaches, which struggle to adapt to evolving system conditions, novel anomaly patterns, and the increasing

complexity of distributed energy infrastructures. The RL agent's continuous learning process enables autonomous refinement of detection strategies based on real-time system feedback, ensuring sustained monitoring performance without requiring manual intervention or system downtime.

The integration of the RL agent within the Web framework rather than relying on isolated, external learning components represents a critical architectural advancement. This design ensures tight coupling between adaptive monitoring mechanisms and existing Web services, APIs, and semantic data flows, preserving system interoperability, scalability, and compliance with established Web engineering standards. Moreover, by embedding monitoring intelligence at the application layer, the system minimizes integration complexity while maximizing real-time responsiveness, a combination rarely achieved in conventional monitoring deployments.

The adaptive Web-based interface further extends the system's impact by fostering operator engagement, transparency, and collaborative monitoring. By integrating real-time visualization and human feedback into the learning loop, the system transforms anomaly detection into a dynamic, user-centric process. This human-in-the-loop capability is particularly valuable in critical infrastructure environments, where contextual understanding and operational experience are essential for effective integrity management.

Beyond the technical improvements demonstrated in this work, the proposed architecture contributes to the broader advancement of Web Engineering by exemplifying how machine learning, and specifically reinforcement learning, can be integrated within modular, service-oriented Web applications to enable intelligent, autonomous system control. The system leverages semantic metadata, RESTful interfaces, and scalable deployment models to deliver adaptable monitoring solutions that are interoperable across diverse, distributed energy environments.

While the current work focuses on electricity fee channels as a motivating scenario, the underlying principles and architectural framework are generalizable to other application domains requiring real-time, adaptive integrity assurance. Potential extensions include smart metering, IoT-based infrastructure monitoring, and Web-enabled supervisory control in transportation, manufacturing, or public safety systems. The results presented here provide a foundation for exploring such cross-domain applications, reinforcing the system's relevance beyond its immediate energy sector context.

Several avenues for future research emerge from this work. Expanding the RL framework to support multi-agent learning could enable coordinated

monitoring across organizational boundaries, addressing data integrity challenges in federated energy networks. Integration of hybrid machine learning models, combining reinforcement learning with deep anomaly detection or unsupervised pattern recognition, may further enhance detection performance, particularly in complex, high-dimensional data environments. Finally, the development of more sophisticated human-in-the-loop mechanisms, including operator trust modeling and explainable AI components, could improve system transparency, user acceptance, and overall operational resilience.

5 Conclusion

This paper presented a novel RL-driven intelligent monitoring framework designed to enhance data integrity, anomaly detection, and operational adaptability within Web-based electricity fee channels. By embedding a reinforcement learning agent directly within the Web application architecture, the proposed system enables autonomous, real-time monitoring that evolves alongside system dynamics, addressing the critical limitations of static, rule-based approaches. Comprehensive experimental evaluation demonstrated that the RL-enhanced monitoring system achieves a 20% improvement in detection accuracy, a 53% reduction in false positives, and sustained adaptability to new, evolving anomaly patterns, all the while maintaining real-time responsiveness suitable for distributed smart grid environments. The integration of an adaptive Web-based interface further advances system transparency, user engagement, and collaborative integrity management.

This work contributes to the advancement of intelligent, self-optimizing monitoring mechanisms for Web-based applications, demonstrating the feasibility and effectiveness of combining reinforcement learning with modular, standards-compliant Web architectures. The proposed system exemplifies how adaptive, autonomous control can be achieved within distributed, data-intensive infrastructures, enhancing operational resilience and scalability.

Future research will extend this work along several dimensions, including multi-agent monitoring for cross-organizational data integrity, hybrid machine learning integration to improve detection robustness, and advanced human-in-the-loop capabilities to further enhance system transparency and user collaboration. By addressing these challenges, the proposed framework lays the foundation for more intelligent, autonomous, and resilient Web-based platforms in the energy sector and beyond.

References

- [1] Omitaomu, Olufemi A., and Haoran Niu. “Artificial intelligence techniques in smart grid: A survey.” *Smart Cities* 4, no. 2 (2021): 548–568.
- [2] Liu, X., Golab, L., Golab, W., Ilyas, I.F., Jin, S. “Smart Meter Data Analytics: Systems, Algorithms, and Benchmarking.” *ACM Trans. Database Syst.*, 42(1), Article 2, 2017. <https://doi.org/10.1145/3004295>.
- [3] Khan, Abdullah Ayub, Asif Ali Laghari, Mamoon Rashid, Hang Li, Abdul Rehman Javed, and Thippa Reddy Gadekallu. “Artificial intelligence and blockchain technology for secure smart grid and power distribution Automation: A State-of-the-Art Review.” *Sustainable Energy Technologies and Assessments* 57 (2023): 103282.
- [4] Buksh, Zain, Neeraj A. Sharma, Rishal Chand, Jashnil Kumar, and A. B. M. Shawkat Ali. “Cybersecurity Challenges in Smart Grid IoT.” *IoT for Smart Grid: Revolutionizing Electrical Engineering* (2025): 175–206.
- [5] Mo, Y., Kim, T.H., Brancik, K., et al. “Cyber-Physical Security of Smart Grid Infrastructure.” *Proc. IEEE*, 100(1), 2012.
- [6] Li, Junlong, Chenghong Gu, Yue Xiang, and Furong Li. “Edge-cloud computing systems for smart grid: state-of-the-art, architecture, and applications.” *Journal of Modern Power Systems and Clean Energy* 10, no. 4 (2022): 805–817
- [7] Liu, J., Xiao, Y., Li, S., Liang, W., Chen, C.L.P. “Cyber Security and Privacy Issues in Smart Grids.” *IEEE Commun. Surv. Tutor.*, 14(4), 2012, pp. 981–997. <https://doi.org/10.1109/SURV.2011.122111.00145>.
- [8] Moustafa, R., Shareef, H., Asna, M., Errouissi, R., Selvaraj, J. “A Smart Web-Based Power Quality and Energy Monitoring System With Enhanced Features.” *IEEE Access*, 13, 2025, pp. 88458–88471. <https://doi.org/10.1109/ACCESS.2025.3571623>.
- [9] Shahinzadeh, H., Moradi, J., Gharehpetian, G.B., et al. “IoT Architecture for Smart Grids.” *IPAPS*, 2019.
- [10] Eskandarnia, E., Al-Ammal, H., Ksantini, R., et al. “Deep Learning Techniques for Smart Meter Data Analytics: A Review.” *SN Comput. Sci.*, 3(243), 2022. <https://doi.org/10.1007/s42979-022-01161-6>.
- [11] Ghasempour, A. “Internet of Things in Smart Grid: Architecture, Applications, Services, Key Technologies, and Challenges.” *Inventions*, 4(22), 2019. <https://doi.org/10.3390/inventions4010022>.
- [12] Fan, Z., Kulkarni, P., et al. “Smart Grid Communications: Overview of Research Challenges, Solutions, and Standardization Activities.” *IEEE Commun. Surv. Tutor.*, 15(1), 2013, pp. 21–38. <https://doi.org/10.1109/SURV.2011.122211.00021>.

- [13] Alulema, Darwin, Javier Criado, Luis Iribarne, Antonio Jesús Fernández-García, and Rosa Ayala. “A model-driven engineering approach for the service integration of IoT systems.” *Cluster Computing* 23, no. 3 (2020): 1937–1954.
- [14] Molina-Ríos, J., Pedreira-Souto, N. “Comparison of Development Methodologies in Web Applications.” *Inf. Softw. Technol.*, 119, 2020, Article 106238.
- [15] Alatrash, Rawaa, Rojalina Priyadarshini, Hadi Ezaldeen, and Akram Alhinnawi. “A hybrid recommendation integrating semantic learner modelling and sentiment multi-classification.” *Journal of Web Engineering* 21, no. 4 (2022): 941–988.
- [16] Sutton, R.S., Barto, A.G. “Reinforcement Learning: An Introduction.” *MIT Press*, 2018.
- [17] Wang, S., et al. “Machine Learning in Network Anomaly Detection: A Survey.” *IEEE Access*, 9, 2021, pp. 152379–152396.
- [18] Escalona, M.J., Koch, N. “Requirements Engineering for Web Applications – A Comparative Study.” *JWE*, 2(3), 2004, pp. 193–212.
- [19] Palaniappan, S., et al. “Machine Learning Model for Predicting Net Environmental Effects.” *J. Inform. Web Eng.*, 4(1), 2025, pp. 243–253.
- [20] Escalona, M. José, and Nora Koch. “Requirements engineering for web applications – a comparative study.” *Journal of web Engineering* (2003): 193–212.
- [21] Olsina, L., et al. “Web Application Evaluation and Refactoring: A Quality-Oriented Improvement Approach.” *JWE*, 7(4), 2008, pp. 258–280.
- [22] Kachergis, Emily, Scott W. Miller, Sarah E. McCord, Melissa Dickard, Shannon Savage, Lindsay V. Reynolds, Nika Lepak et al. “Adaptive monitoring for multiscale land management: Lessons learned from the Assessment, Inventory, and Monitoring (AIM) principles.” *Rangelands* 44, no. 1 (2022): 50–63.
- [23] Pfaff, M., Krcmar, H. “A Web-Based System Architecture for Ontology-Based Data Integration in IT Benchmarking.” *Enterp. Inf. Syst.*, 12(3), 2018, pp. 236–258.
- [24] González-Mora, César, Irene Garrigós, Jose Zubcoff, and Jose-Norberto Mazón. “Model-based generation of web application programming interfaces to access open data.” *Journal of Web Engineering* 19, no. 7–8 (2020): 1147–1172.

Biographies



Xinling Zheng, was born in 1991 in Fujian, China. She graduated from the School of Economics and Management of Fuzhou University and received her Master's degree in Accounting in 2017. She is currently working at State Grid Fujian Marketing Service Center (Metering Center and Integrated Capital Center), mainly responsible for power marketing, electricity bill accounting management and other work.



Songyan Du, received her Master's degree in Technical Economics and Management from Fuzhou University in 2013. She is currently working at State Grid Fujian Marketing Service Center (Metering Center and Integrated Capital Center), mainly good at accounting management such as payment channel operation, electricity bill recovery control and so on.



Jing Ye graduated from the School of Computer Science and Technology of Fujian Agriculture And Forestry University and received her Bachelor of Engineering degree in 2007. She is currently working at State Grid Fujian Marketing Service Center (Metering Center and Integrated Capital Center), mainly responsible for electricity marketing and electricity fee accounting management.



Huawei Hong received his MBA degree from National Huaqiao University. He now works at the State Grid Fujian Marketing Service Center (Metering Center and Integrated Capital Center). He is also a member of the China Energy Research Society, deputy Secretary General of the Energy Internet Branch, and deputy director of the Intelligent Energy Use Branch of the Fujian Electrical Engineering Society. He is mainly engaged in power marketing, power market operation management and other work.



Yimin Shen received his master's degree in electrical engineering from Shanghai Electric Power University. He is currently working at State Grid Fujian Marketing Service Center (Metering Center and Integrated Capital Center), mainly responsible for power market operation and management.



Xiaorui Qian received his B.Eng. and M.Eng. degrees in electrical engineering from Hohai University, NanJing, China. He is currently working at State Grid Fujian Marketing Service Center (Metering Center and Integrated Capital Center), mainly responsible for power consumption analysis and forecasting. He is also a member of the Power Market Construction Working Group of the Development and Reform Commission of Fujian Province and has been engaged in power market analysis and management for a long time.



Xingye Lin graduated from China Three Gorges University and received a bachelor's degree in electrical engineering in 2019. He is currently working at State Grid Fujian Marketing Service Center (Metering Center and Integrated Capital Center), mainly responsible for tariff recovery and other work.

