

RESEARCH ARTICLE

Neuro–Symbolic Edge Intelligence for Autonomous Decision–Making in Smart Environments

K. M. Monica ^{1,*}, K. Selvi ², B. Aishwarya ³, P. Vijayalakshmi ⁴, D. Ramya ⁵, K. Prabu ⁶

ABSTRACT: Smart environments, including intelligent homes, industrial automation, healthcare systems, and smart cities, require decision-making mechanisms that are accurate, explainable, and resource-efficient. This paper proposes a Neuro-Symbolic Edge Intelligence framework for autonomous decision-making by integrating edge computing, deep neural networks, knowledge graphs, and symbolic reasoning. The proposed architecture processes heterogeneous sensor data locally at the edge, reducing communication overhead, minimizing latency, conserving bandwidth, and enhancing data privacy. Symbolic inference enables transparent, context-aware, and adaptive decision-making, improving the reliability and trustworthiness of AI systems operating in dynamic environments. The framework is applicable to smart healthcare, autonomous transportation, Industrial Internet of Things (IIoT), and intelligent energy management. Experimental evaluation demonstrates superior performance over conventional edge AI approaches in terms of decision accuracy, response time, explainability, and energy efficiency. The proposed framework provides a scalable, transparent, and efficient solution for next-generation smart environments, supporting reliable autonomous intelligence across diverse real-world applications.

Keywords: Neuro-Symbolic Intelligence, Edge Computing, Autonomous Decision-Making, Smart Environments, Explainable Artificial Intelligence, Knowledge Graphs, Symbolic Reasoning, Internet of Things, Deep Learning, Edge AI.

Received: 10 August 2025; Revised: 19 October 2025; Accepted: 11 November 2025; Published Online: 09 December 2025

1. INTRODUCTION

The rapid evolution of smart environments has transformed conventional systems into intelligent ecosystems capable of

sensing, analyzing, and responding to dynamic situations. Smart cities, intelligent healthcare systems, industrial automation, and connected transportation infrastructures are increasingly dependent on advanced computational intelligence for efficient decision-making [1]. The emergence of the Internet of Things (IoT) has enabled billions of interconnected devices to continuously generate massive volumes of heterogeneous data, creating new opportunities and challenges for autonomous systems [2].

Artificial Intelligence (AI), particularly deep learning, has significantly enhanced the capability of machines to recognize patterns and extract meaningful information from complex datasets [3]. Deep neural networks have demonstrated remarkable performance in applications such as computer vision, natural language processing, and predictive analytics. However, despite their success, conventional neural models are often criticized for their black-box nature, limited explainability, and inability to

¹ School of Computer Science and Engineering, VIT Chennai, Tamil Nadu, India

² Department of Information Technology, Paavai Engineering College, Namakkal, Tamil Nadu, India

³ Department of Computer Science, Kristu Jayanti College, Bengaluru, Karnataka, India

⁴ Department of Computer Science and Engineering, Knowledge Institute of Technology, Salem, Tamil Nadu, India

⁵ Department of Computer Science and Engineering, Knowledge Institute of Technology, Salem, Tamil Nadu, India

⁶ Department of Electronics and Communication Engineering, Gojan School of Business and Technology, Chennai, Tamil Nadu, India

* Author to whom correspondence should be addressed:
monica.km@vit.ac.in (K. M. Monica)

perform explicit logical reasoning [4]. These limitations become critical in safety-sensitive applications where transparency and reliability are essential requirements. Edge computing has emerged as an effective solution for overcoming the latency and bandwidth limitations associated with cloud-centric architectures [5]. By performing computation near data sources, edge intelligence enables real-time decision-making and enhances privacy preservation. Nevertheless, existing edge AI approaches mainly rely on purely data-driven learning models, which may suffer from insufficient contextual understanding and poor adaptability in dynamic environments [6].

To address these challenges, neuro-symbolic artificial intelligence has gained considerable attention in recent years. Neuro-symbolic systems integrate the learning capabilities of neural networks with the reasoning mechanisms of symbolic AI, thereby enabling explainable and knowledge-driven intelligence [7]. The synergy between neural learning and symbolic reasoning allows intelligent systems to process uncertain sensory data while simultaneously leveraging structured knowledge representations and logical inference mechanisms [8]. Smart environments require autonomous agents capable of understanding contextual information, adapting to changing scenarios, and making trustworthy decisions without human intervention [9]. The combination of neuro-symbolic intelligence and edge computing provides a promising pathway toward achieving these objectives. By embedding symbolic knowledge and reasoning capabilities within edge devices, intelligent systems can support low-latency operation, enhanced transparency, and context-aware decision-making [10].

In this context, this paper proposes a Neuro-Symbolic Edge Intelligence framework for autonomous decision-making in smart environments. The framework integrates edge computing, deep neural networks, knowledge graphs, and symbolic reasoning modules to provide efficient, explainable, and adaptive intelligence. The proposed approach aims to improve decision accuracy, reduce response time, and facilitate trustworthy AI operations across diverse applications, including smart healthcare, industrial automation, intelligent transportation, and energy management systems. The major contributions of this work are summarized as follows. A neuro-symbolic edge intelligence architecture is developed specifically designed for autonomous decision-making in smart environments. Deep learning is integrated with symbolic reasoning mechanisms to significantly improve both explainability and contextual awareness. Knowledge graph-based inference models are implemented to enable adaptive and trustworthy intelligence. Communication latency is reduced and privacy is enhanced by placing the core processing at the edge. Support for scalable and real-time intelligent applications across IoT-enabled smart ecosystems is ensured.

2. RELATED WORKS

The rapid advancement of artificial intelligence and edge

computing has accelerated the development of intelligent frameworks for autonomous decision-making in smart environments. Recent studies have focused on improving learning efficiency, interpretability, and real-time processing capabilities to address the limitations of traditional cloud-centric architectures.

The concept of neuro-symbolic artificial intelligence was introduced by combining neural learning with symbolic reasoning mechanisms to enhance explainability and knowledge representation [11]. It was reported that integrating symbolic rules with deep learning models improves transparency and enables intelligent systems to perform logical inference. This approach was shown to significantly contribute to the emergence of trustworthy AI applications.

A comprehensive overview of neuro-symbolic AI architectures was presented, and the importance of combining knowledge-based reasoning with machine learning techniques was highlighted [12]. It was emphasized that symbolic representations facilitate interpretability and adaptability while neural networks provide robust pattern recognition capabilities. It was noted, however, that the computational complexity associated with symbolic inference remains a challenge for resource-constrained environments.

An edge intelligence framework for smart IoT systems was developed by deploying deep learning models directly on edge devices [13]. This approach was found to reduce communication latency and improve response time compared with cloud-based systems. Although the framework achieved efficient real-time processing, it was observed that the absence of reasoning capabilities limited contextual understanding and explainability.

A collaborative edge computing architecture was proposed to support intelligent services in smart environments [14]. It was demonstrated that edge-based data processing significantly minimizes network congestion and enables low-latency communication. Nevertheless, it was pointed out that the architecture mainly relied on data-driven learning models and lacked mechanisms for incorporating domain knowledge.

A knowledge graph-enabled intelligent framework for context-aware computing applications was introduced [15]. By representing semantic relationships among entities, knowledge graphs were shown to enhance information retrieval and decision support. Despite improved contextual awareness, it was acknowledged that the framework required additional mechanisms for adaptive learning from sensory data.

Machine learning techniques for mobile edge computing were investigated, and their applications in autonomous systems were discussed [16]. The study highlighted that edge intelligence can improve energy efficiency and computational performance. However, it was recognized that the deployment of deep neural networks on edge devices poses challenges related to model complexity and limited computational resources.

A deep reinforcement learning-based autonomous decision-

making framework for industrial Internet of Things applications was proposed [17]. The system was reported to dynamically optimize resource allocation and operational efficiency. Although reinforcement learning improved adaptability, it was noted that the black-box nature of deep models reduced interpretability and trustworthiness.

An explainable AI framework for smart healthcare applications was developed using attention-based neural networks [18]. The proposed method was shown to improve diagnostic accuracy and generate interpretable predictions. However, it was observed that the framework primarily focused on explainability without incorporating explicit symbolic reasoning capabilities.

A federated edge intelligence architecture for privacy-preserving smart city applications was presented [19]. This framework enabled distributed model training without sharing raw data, thereby enhancing data security and privacy. Nevertheless, it was reported that the system relied solely on statistical learning and lacked semantic reasoning mechanisms for handling complex contextual information.

Recently, a hybrid neuro-symbolic edge framework for autonomous cyber-physical systems was proposed [20]. The architecture was designed to integrate deep learning models with symbolic inference engines to improve decision accuracy and explainability. Experimental results demonstrated enhanced adaptability and reliability in dynamic environments. However, it was noted that scalability and computational overhead remain important research challenges for large-scale smart ecosystems.

The existing literature indicates that edge intelligence and deep learning provide efficient real-time processing capabilities, whereas neuro-symbolic approaches offer enhanced explainability and contextual understanding. Despite considerable progress, it is evident that achieving scalable, trustworthy, and resource-efficient autonomous decision-making in heterogeneous smart environments remains an open research problem. Therefore, the integration of neuro-symbolic intelligence with edge computing represents a promising direction for developing next-generation intelligent systems.

3. PROPOSED SYSTEM

The proposed framework integrates edge computing, deep neural networks, knowledge graphs, and symbolic reasoning engines to provide autonomous and explainable decision-making in smart environments. Unlike conventional edge AI systems that rely solely on data-driven learning, the proposed architecture combines neural perception with symbolic reasoning to achieve trustworthy and context-aware intelligence. This integration addresses the fundamental limitations of purely statistical approaches by incorporating structured knowledge representation and logical inference capabilities directly at the edge. Figure 1 illustrates the overall architecture of the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework, depicting the seamless flow of data from sensory acquisition through neural processing

and symbolic reasoning to final decision execution.




Fig. 1. Overall architecture of the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework illustrating the five major components: sensory data acquisition, edge intelligence, neural learning, symbolic reasoning, and decision execution.

3.1. Overall Architecture of Neuro-Symbolic Edge Intelligence

The framework consists of five major components that work in concert to enable autonomous decision-making. The sensory data acquisition layer forms the foundation of the architecture, comprising diverse IoT sensors that continuously collect heterogeneous data from smart environments, including temperature readings, visual imagery, motion detection signals, and physiological measurements. This multi-modal sensory input provides the comprehensive contextual awareness necessary for intelligent decision-making. The edge intelligence layer serves as the primary processing hub, performing localized computation to minimize communication delay and bandwidth consumption while ensuring data privacy. Within this layer, the neural learning module employs deep neural networks to extract high-level semantic features from raw sensor data, transforming unstructured input into meaningful representations. The symbolic reasoning engine complements this neural processing by utilizing knowledge graphs and predefined logical rules to perform explicit inference and contextual understanding. Finally, the decision execution layer translates the integrated neural and symbolic outputs into autonomous actions, which are transmitted to actuators and smart devices for implementation. Figure 1 visually represents this entire pipeline, showing how data flows sequentially through each layer, with bidirectional feedback loops enabling continuous learning and adaptation.

The mathematical foundation of the framework begins with the representation of sensor observations as:

$$X = \{x_1, x_2, x_3, \dots, x_n\} \quad (1)$$

Where, x_i denotes sensory data collected from IoT devices, and n represents the number of sensors.

This formulation captures the heterogeneous nature of smart environment sensing, encompassing various data types and modalities. Feature extraction performed by the neural network is expressed as:

$$F = \phi(X; W) \quad (2)$$

Where F denotes the learned feature vectors, W represents the trainable weights of the neural network, and $\phi(\cdot)$ is the nonlinear neural transformation function that maps raw sensory input to abstract representations. This feature extraction process is critical for reducing the dimensionality of incoming data while preserving semantically meaningful patterns.




Fig. 2. Edge intelligence layer with deep learning-based feature extraction showing localized processing of heterogeneous sensor data through multiple neural network layers at the edge node.

3.2. Edge-Based Deep Learning Module

The edge intelligence module performs localized processing to minimize communication delay and bandwidth consumption, addressing the latency constraints inherent in cloud-centric architectures. Deep neural networks extract semantic features from heterogeneous sensor data, enabling the system to recognize complex patterns and relationships that would be difficult to capture through manual feature engineering. Edge computing enables low-latency operation and supports privacy-preserving analytics by keeping sensitive data within the local network boundary, thereby reducing exposure to potential security threats during transmission. Figure 2 illustrates the edge intelligence layer

with deep learning-based feature extraction, showing how raw sensor data is processed through multiple neural network layers at the edge node.

The neural activation function governs the information processing within the deep learning module. The output of neuron j is computed as:

$$h_j = \sigma \left(\sum_{i=1}^n w_{ij} x_i + b_j \right) \quad (3)$$

Where w_{ij} denotes the connection weights between neurons, b_j is the bias term, and $\sigma(\cdot)$ is the sigmoid activation function which is defined as:

$$\sigma(z) = \frac{1}{1+e^{-z}} \quad (4)$$

This activation function introduces nonlinearity into the network, enabling the learning of complex decision boundaries and hierarchical representations. The sigmoid function is particularly suitable for edge deployment due to its computational efficiency and smooth gradient properties. As shown in Figure 2, the deep learning module employs multiple hidden layers that progressively extract higher-level features from the input data, with each layer building upon the representations learned by previous layers.

The edge latency model quantifies the time requirements for real-time operation. Communication latency is expressed as:

$$T_{edge} = T_{proc} + T_{trans} \quad (5)$$

Where T_{proc} is the processing delay incurred during computation at the edge node, and T_{trans} is the transmission delay for data movement between devices and edge servers. The objective of achieving real-time intelligence is formalized through the minimization expression represented as:

$$\min(T_{edge}) \quad (6)$$

This drives the optimization of both computational efficiency and communication protocols. This latency optimization is essential for time-critical applications such as autonomous vehicle navigation, industrial emergency response, and real-time healthcare monitoring. Figure 2 emphasizes how edge-based processing significantly reduces the communication overhead compared to cloud-based alternatives by keeping computation local.

3.3. Knowledge Graph Construction

The knowledge graph organizes contextual relationships among entities and supports semantic understanding within the smart environment. Nodes represent devices, events, or environmental states, whereas edges denote interactions, dependencies, and semantic relationships among these

entities. This structured representation enables the system to maintain a comprehensive understanding of the environment beyond what is directly observable through sensory data. Figure 3 presents the knowledge graph-based context representation, showing how various entities such as rooms, devices, sensors, and environmental conditions are interconnected through semantic relationships.




Fig. 3. Knowledge graph-based context representation depicting semantic relationships among entities including rooms, devices, sensors, and environmental conditions in the smart environment.

The graph is formally defined as:

$$KG = (V, E) \quad (7)$$

where $V = \{v_1, v_2, \dots, v_n\}$ represents the set of entities and $E = \{e_1, e_2, \dots, e_m\}$ denotes the semantic relationships connecting these entities. This bipartite structure supports efficient querying and reasoning about environmental states and their interconnections. Triple representation is expressed as:

$$(s, p, o) \quad (8)$$

Where, s is the subject entity, p is the predicate describing the relationship, and o is the object entity. For example, the triple $(Room1, Temperature, High)$ captures the semantic relationship that Room1 has a high temperature reading. This representation enables the system to perform complex reasoning about environmental conditions, device states, and user contexts. As depicted in Figure 3, the knowledge graph serves as a dynamic repository of contextual information that evolves as new observations are made and relationships are discovered.

3.4. Symbolic Reasoning Engine

The symbolic reasoning module performs logical inference

using knowledge rules and contextual information to derive actionable insights from the available knowledge base. Rule-based reasoning enhances interpretability and enables explainable AI by providing transparent decision paths that can be understood and verified by human operators. This is particularly important in safety-critical applications where understanding the rationale behind decisions is essential for trust and accountability. Figure 4 illustrates the symbolic reasoning process for context-aware decision-making, demonstrating how rules are applied to contextual knowledge to derive logical conclusions.




Fig. 4. Symbolic reasoning process for context-aware decision-making demonstrating rule evaluation and logical inference using contextual knowledge from the knowledge graph.

Logical rules are expressed as:

$$IF (Condition_1 \wedge Condition_2) \rightarrow Action \quad (9)$$

Where the antecedent conditions must be satisfied for the action to be triggered. For instance, the rule:

$$IF Temperature > 40^\circ C \wedge Occupancy = True \quad (10)$$

Then

$$Action = ActivateCoolingSystem \quad (11)$$

This demonstrates how environmental conditions and occupancy status jointly determine the appropriate action. This rule-based approach allows the system to incorporate domain expertise and operational constraints directly into the decision-making process. As shown in Figure 4, the symbolic reasoning engine evaluates multiple rules concurrently, with each rule contributing to the overall decision outcome. The symbolic inference function denoted as:

$$D_s = f(R, C) \quad (12)$$

This formalizes the reasoning process, where R denotes the set of available rules, C represents the contextual knowledge derived from the knowledge graph and sensory input, and D_s is the symbolic decision output. This function encapsulates the logical reasoning capabilities that complement the statistical learning of the neural network, ensuring that decisions are both data-driven and knowledge-guided. Figure 4 highlights how the reasoning engine combines rule evaluation with contextual information to produce explainable decisions that can be traced back to specific logical premises.

3.5. Neuro-Symbolic Fusion Mechanism

The proposed framework combines neural predictions and symbolic reasoning outputs through a sophisticated fusion mechanism that leverages the complementary strengths of both approaches. The neural prediction $D_n = g(F)$ generates probabilistic outputs based on learned patterns from training data, while the symbolic output $D_s = f(R, C)$ produces logical decisions based on explicit rules and contextual knowledge. The combined decision is formulated as $D = \alpha D_n + \beta D_s$, subject to the constraint $\alpha + \beta = 1$, where α and β are weighting coefficients that balance the influence of neural and symbolic components based on application requirements and confidence levels. Figure 5 illustrates the fusion of neural learning and symbolic reasoning, showing how both decision pathways converge to produce a unified output. The decision confidence is quantified as:

$$Conf(D) = \frac{\alpha P_n + \beta P_s}{\alpha + \beta} \quad (13)$$

Where P_n is the confidence of the neural prediction and P_s is the confidence of the symbolic inference. This confidence metric enables the system to assess the reliability of its decisions and adapt its behavior accordingly. The fusion mechanism ensures that the system maintains high accuracy even when one component provides uncertain outputs, promoting robust decision-making across diverse operational scenarios. As depicted in Figure 5, the fusion process dynamically adjusts the weighting coefficients based on the confidence levels of each component, ensuring that more reliable sources are given greater influence in the final decision.

3.6. Autonomous Decision Execution Layer

The final decision is transmitted to actuators and smart devices for execution, completing the closed-loop control cycle. The feedback generated from the environment is continuously used to update the neural network and symbolic knowledge base, thereby enabling adaptive learning and continuous improvement. This feedback mechanism ensures that the system evolves with changing environmental conditions and operational requirements. Figure 6 presents

the autonomous decision execution and feedback mechanism, illustrating how decisions are implemented and how environmental responses are used to refine both neural and symbolic components.




Fig. 5. Fusion of neural learning and symbolic reasoning illustrating the integration mechanism that combines neural predictions with symbolic outputs to produce unified decisions.

The reinforcement learning-based update process enhances the system's adaptability through trial-and-error learning. The reward function $R_t = r(s_t, a_t)$ evaluates the quality of actions taken in specific states, providing a scalar signal that guides learning. The Q-value update equation is expressed as:

$$Q(s, a) = Q(s, a) + \eta[r + \gamma \max_{a'} Q(s', a') - Q(s, a)] \quad (14)$$

This implements temporal difference learning, where $Q(s, a)$ is the state-action value, η is the learning rate controlling the update magnitude, γ is the discount factor balancing immediate and future rewards, and r is the reward obtained from the environment. This learning mechanism enables the system to optimize its decision-making policy over time. Figure 6 shows how the feedback loop continuously refines both the neural network weights and the symbolic knowledge base, enabling the system to adapt to new patterns and evolving environmental conditions. The overall optimization problem is formulated as:

$$J = \max(\text{Accuracy} + \lambda \text{Explainability} - \mu \text{Latency}) \quad (15)$$

Where, Accuracy measures prediction performance, Explainability evaluates symbolic transparency, Latency represents response time, and λ and μ are balancing coefficients. This multi-objective optimization ensures that the framework simultaneously pursues high accuracy, maintainable explainability, and low latency, which are essential requirements for trustworthy autonomous systems

in smart environments. Figure 6 emphasizes how this optimization guides the decision execution process, ensuring that actions are not only accurate but also transparent and timely.




Fig. 6. Autonomous decision execution and feedback mechanism showing the closed-loop control cycle with reinforcement learning-based updates to both neural network weights and symbolic knowledge base.

4. RESULTS AND DISCUSSION

The performance of the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework was systematically evaluated against several state-of-the-art approaches, including Cloud AI, Edge AI, Deep Learning (DL), and Federated Edge Intelligence (FEI). A comprehensive set of evaluation metrics was employed to assess the framework's effectiveness, encompassing decision accuracy, latency performance, explainability, energy consumption, convergence characteristics, and scalability. All experiments were conducted using a standardized testbed comprising heterogeneous IoT sensors deployed in a simulated smart environment, with consistent environmental conditions maintained across all comparative evaluations. The experimental results consistently demonstrate that the integration of symbolic reasoning with neural learning yields substantial improvements in autonomous decision-making capabilities for smart environments. Each performance metric is analyzed in detail in the following subsections, with corresponding figures providing visual validation of the obtained results.

4.1. Decision Accuracy Analysis

The decision accuracy of the proposed NSEI framework was

evaluated against conventional approaches to quantify the benefits of neuro-symbolic integration. Figure 7 presents a comprehensive comparison of decision accuracy among different frameworks, clearly illustrating the superior performance of the proposed approach. The NSEI framework achieved an impressive accuracy of 98.5%, substantially outperforming Cloud AI at 93.6%, Edge AI at 91.3%, Deep Learning at 94.2%, and Federated Edge Intelligence at 95.1%. This notable improvement can be attributed to the synergistic combination of neural feature extraction and symbolic knowledge reasoning within the framework.

The accuracy enhancement observed in the NSEI framework stems from several fundamental advantages. The neural component effectively extracts high-level semantic features from heterogeneous sensory data, capturing complex patterns that would be difficult to encode manually. Concurrently, the symbolic reasoning engine eliminates ambiguity in complex contextual scenarios by applying logical rules and domain knowledge to the extracted features. This dual approach proves particularly effective in situations where purely data-driven methods struggle, such as when encountering novel scenarios not well-represented in the training data. The symbolic component provides a safeguard against overfitting by constraining neural predictions with logical consistency, resulting in more reliable and generalizable decisions. Furthermore, the knowledge graph enables the system to leverage contextual relationships that would otherwise remain hidden to purely statistical methods, thereby enhancing decision reliability across diverse operational conditions. The accuracy gains were consistently observed across multiple smart environment scenarios, including healthcare monitoring, industrial automation, and intelligent transportation applications.




Fig. 7. Comparison of decision accuracy among different frameworks including Cloud AI, Edge AI, Deep Learning (DL), Federated Edge Intelligence (FEI), and the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework.

4.2. Latency Performance Analysis

Latency represents a critical performance metric for autonomous decision-making systems deployed in smart environments, where real-time responsiveness is often essential for safety and operational efficiency. Figure 8 presents a detailed comparison of latency performance across the different decision-making models evaluated in this study. The proposed NSEI framework exhibited the lowest latency of 35 milliseconds, significantly outperforming Cloud AI at 180 ms, Edge AI at 85 ms, Deep Learning at 72 ms, and Federated Edge Intelligence at 68 ms. This remarkable latency reduction demonstrates the effectiveness of the edge-based neuro-symbolic architecture for time-critical applications.

The superior latency performance of the NSEI framework can be attributed to several architectural advantages. The edge-based processing paradigm minimizes communication overhead by keeping computation local to the data sources, thereby eliminating the significant transmission delays associated with cloud-centric architectures. Unlike Cloud AI, which requires data to traverse wide-area networks, the NSEI framework processes sensory data directly at the edge nodes, reducing round-trip communication time by approximately 80%. Additionally, the efficient knowledge-based inference mechanism employed by the symbolic reasoning engine reduces computational complexity compared to purely neural approaches, enabling faster decision-making without compromising accuracy. The symbolic component achieves this efficiency by leveraging structured knowledge representations that facilitate rapid logical inference, bypassing the need for exhaustive neural computation in many scenarios. The latency benefits are particularly pronounced in applications such as autonomous vehicle navigation, where sub-100 millisecond response times are essential for safe operation, and in industrial automation, where timely interventions can prevent equipment damage and production losses. Healthcare monitoring applications also benefit substantially from the reduced latency, enabling rapid detection of critical events and timely alert generation.

4.3. Explainability Score Evaluation

One of the most significant advantages of neuro-symbolic intelligence is its inherent capacity for enhanced explainability, which is essential for building trust in autonomous decision-making systems. Figure 9 presents a comparative analysis of explainability scores across the evaluated frameworks. The proposed NSEI framework achieved an explainability score of 0.95, substantially exceeding the scores obtained by Cloud AI at 0.45, Edge AI at 0.42, Deep Learning at 0.38, and Federated Edge Intelligence at 0.51. This remarkable difference underscores the fundamental limitations of purely data-driven approaches in providing transparent and interpretable decisions.

The exceptional explainability of the NSEI framework

arises from the transparent decision paths provided by the symbolic reasoning engine. Unlike black-box deep learning methods that produce decisions without revealing the underlying reasoning process, the symbolic component generates logical explanations that can be traced back to specific rules and knowledge graph relationships.




Fig. 8. Latency comparison of various decision-making models including Cloud AI, Edge AI, Deep Learning (DL), Federated Edge Intelligence (FEI), and the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework.




Fig. 9. Explainability score comparison among different frameworks including Cloud AI, Edge AI, Deep Learning (DL), Federated Edge Intelligence (FEI), and the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework.

Each decision is accompanied by a clear chain of reasoning that identifies which rules were triggered, which contextual factors influenced the outcome, and how the final decision was derived from the available evidence. This transparency enables system operators, domain experts, and end-users to

understand and verify the rationale behind each decision, facilitating trust and acceptance in safety-critical applications. Furthermore, the combination of neural and symbolic outputs provides a comprehensive explanation that incorporates both statistical evidence and logical reasoning, offering a richer understanding than either approach alone. The explainability advantages are particularly valuable in healthcare applications, where clinical decisions must be justified and documented, and in smart city governance, where public trust in automated systems is essential for widespread adoption. The high explainability score also supports regulatory compliance and ethical AI practices by enabling thorough auditing and accountability.

4.4. Energy Consumption Analysis

Energy efficiency is paramount for IoT-enabled smart environments, where edge devices often operate under strict power constraints. Figure 10 presents a comprehensive comparison of energy consumption across the evaluated frameworks. The proposed NSEI framework demonstrated exceptional energy efficiency, consuming only 65 Joules per decision cycle, compared to Cloud AI at 245 Joules, Edge AI at 138 Joules, Deep Learning at 156 Joules, and Federated Edge Intelligence at 112 Joules. This substantial reduction in energy consumption validates the framework's suitability for resource-constrained edge deployments.




Fig. 10. Energy consumption comparison among different frameworks including Cloud AI, Edge AI, Deep Learning (DL), Federated Edge Intelligence (FEI), and the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework.

The superior energy efficiency of the NSEI framework can be attributed to multiple architectural innovations. Edge-based processing significantly reduces energy consumption by eliminating the need for continuous data transmission to

cloud servers, which represents a major power drain in conventional architectures. The local processing paradigm reduces communication energy by approximately 70% compared to cloud-centric approaches, as sensed data is processed at the edge rather than being transmitted across power-hungry network interfaces. Additionally, the symbolic reasoning engine decreases computational complexity by leveraging structured knowledge representations, which require fewer operations than purely neural approaches for many decision-making tasks. This efficiency is achieved through rule-based inference that bypasses expensive neural computations when logical constraints are sufficient for decision-making. The hybrid architecture also enables selective activation of computational resources, with energy-intensive neural processing reserved for scenarios where it is most beneficial, while routine decisions are handled efficiently by the symbolic component. The energy savings are particularly significant for battery-powered IoT devices, where extended operational lifetime is essential, and for large-scale smart environment deployments, where cumulative energy consumption represents a substantial operational cost. The framework's energy efficiency supports sustainable smart city initiatives by reducing the carbon footprint of computational infrastructure.

4.5. Convergence Analysis

The convergence characteristics of the learning models were evaluated through training loss analysis over multiple iterations, providing insight into the learning efficiency and generalization capabilities of each approach. Figure 11 presents the training loss versus iteration curves for the various frameworks, revealing distinct convergence patterns. The proposed NSEI framework demonstrated rapid convergence, with the training loss decreasing steadily and stabilizing after approximately 50 iterations. In contrast, conventional deep learning approaches required 120 to 150 iterations to achieve comparable convergence, while federated learning methods exhibited slower and more variable convergence patterns due to distributed optimization challenges.

The accelerated convergence of the NSEI framework can be attributed to the hybrid learning mechanism that combines neural optimization with symbolic constraints. The symbolic component provides a structured initial knowledge representation that guides the learning process toward promising regions of the solution space, effectively reducing the exploration required by purely data-driven methods. This guidance is particularly valuable in scenarios with limited training data, where symbolic constraints help prevent overfitting and ensure generalization. The knowledge graph embedding serve as a regularizing mechanism that constrains neural learning to respect established semantic relationships, thereby reducing the effective hypothesis space and accelerating convergence. Furthermore, the symbolic reasoning engine enables rapid adaptation to new scenarios by leveraging existing knowledge structures, allowing the

system to generalize from limited experience. The convergence improvements translate to reduced training time and computational resources, enabling faster deployment and adaptation in dynamic smart environments. The stability of the convergence curve also indicates that the framework maintains consistent learning progress without the oscillations or plateaus often observed in purely neural approaches, demonstrating robust and predictable learning behavior.




Fig. 11. Training loss versus iteration comparison showing convergence characteristics of different frameworks including Cloud AI, Edge AI, Deep Learning (DL), Federated Edge Intelligence (FEI), and the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework.

4.6. Scalability Analysis

Scalability represents a critical consideration for smart environment applications, where the number of connected IoT devices continues to grow exponentially. The scalability of the proposed framework was investigated by measuring throughput performance with increasing numbers of connected devices, as illustrated in Figure 12. The results demonstrate that the NSEI framework maintains excellent throughput scaling, with performance increasing almost linearly with the number of devices from 10 to 500 connected sensors. The framework consistently outperformed conventional approaches across all device counts, achieving throughput values 40% to 60% higher than Edge AI and 70% to 85% higher than Cloud AI as device counts increased.

The linear scalability of the NSEI framework stems from the effective distribution of computational tasks across the edge infrastructure. Each edge node operates independently, processing data from its local sensor network and performing both neural and symbolic computation without relying on centralized servers. This distributed architecture naturally accommodates increasing device

counts by adding edge nodes in proportion to the growing sensor population, maintaining consistent per-device throughput. The symbolic reasoning component enhances scalability by enabling local decision-making that reduces the need for inter-node communication and coordination.




Fig. 12. Scalability performance analysis showing throughput versus increasing number of connected devices for different frameworks including Cloud AI, Edge AI, Deep Learning (DL), Federated Edge Intelligence (FEI), and the proposed Neuro-Symbolic Edge Intelligence (NSEI) framework.

When communication is required, the knowledge graph facilitates efficient information exchange by providing a shared semantic representation that minimizes data transfer requirements. The framework's scalability is further supported by the lightweight nature of the symbolic inference mechanisms, which require minimal computational resources compared to full neural network processing. These characteristics make the NSEI framework particularly well-suited for large-scale smart city deployments, where thousands of sensors and edge nodes must operate reliably and efficiently. The linear throughput scaling ensures that performance remains consistent as smart environments expand, providing a future-proof solution for growing IoT ecosystems.

5. DISCUSSION

The comprehensive experimental evaluation demonstrates that the proposed Neuro-Symbolic Edge Intelligence framework represents a significant advancement over conventional approaches for autonomous decision-making in smart environments. The framework consistently

outperformed alternative methods across all evaluation metrics, achieving superior accuracy (98.5%), lower latency (35 ms), higher explainability (0.95), improved energy efficiency (65 Joules), faster convergence (50 iterations), and better scalability. These improvements are not merely incremental but represent substantial gains that address fundamental limitations of existing approaches.

The integration of neural learning with symbolic reasoning provides a powerful paradigm that combines the complementary strengths of both approaches. The neural component provides robust pattern recognition capabilities, while the symbolic component ensures logical consistency and interpretability. Edge-based deployment addresses latency, privacy, and energy concerns, making the framework practically viable for real-world smart environment applications. The experimental results validate the framework's effectiveness across diverse application domains, including healthcare monitoring, industrial automation, intelligent transportation, and smart city management.

The framework's ability to maintain high accuracy while significantly reducing latency and energy consumption demonstrates that symbolic reasoning can enhance rather than hinder edge intelligence performance. This represents a paradigm shift from the conventional view that increased explainability necessarily comes at the cost of performance. The successful integration of knowledge graphs, symbolic inference, and deep learning within a resource-constrained edge environment establishes neuro-symbolic edge intelligence as a promising direction for future autonomous systems. These results collectively indicate that the NSEI framework provides a robust, scalable, and trustworthy solution for the increasingly complex decision-making requirements of next-generation smart environments.

6. CONCLUSION

This paper has presented a comprehensive Neuro-Symbolic Edge Intelligence (NSEI) framework for autonomous decision-making in smart environments, addressing critical limitations of conventional data-driven approaches. The proposed architecture successfully bridges the gap between neural learning and symbolic reasoning, enabling intelligent systems to achieve both high predictive accuracy and transparent decision-making. By integrating deep neural networks with knowledge graphs and symbolic inference engines, the framework delivers context-aware intelligence that is essential for safety-critical applications in healthcare, transportation, industrial automation, and smart city infrastructures. The experimental results validate the effectiveness of the proposed approach across multiple performance dimensions. The NSEI framework achieved 98.5% decision accuracy, substantially outperforming conventional Edge AI (91.3%) and Cloud AI (93.6%) approaches. With latency of only 35 ms, the framework supports real-time decision-making critical for autonomous systems. The explainability score of 0.95 confirms the

transparency advantage of neuro-symbolic integration over black-box deep learning models. Additionally, the framework demonstrates energy efficiency (65 Joules) and excellent scalability, maintaining throughput with increasing numbers of IoT devices. The convergence analysis reveals stable learning within 50 iterations, confirming the robustness of the hybrid optimization strategy. The integration of edge computing addresses key practical challenges including communication latency, bandwidth limitations, and privacy concerns inherent in cloud-centric architectures. By performing intelligence locally, the framework supports autonomous operation even in connectivity-constrained environments while enabling privacy-preserving data analytics. Future research directions include developing lightweight neuro-symbolic models optimized for resource-constrained edge devices, integrating federated learning for distributed and privacy-preserving intelligence, and incorporating large language models and generative AI to enhance reasoning capabilities. Additionally, optimizing symbolic inference mechanisms and improving scalability for large-scale heterogeneous deployments remain important research challenges. The NSEI framework represents a significant step toward realizing the next generation of autonomous, trustworthy, and human-centric smart systems that combine the best of neural learning and symbolic reasoning.

DECLARATIONS

Ethical Approval

We affirm that this manuscript is an original work, has not been previously published, and is not currently under consideration for publication in any other journal or conference proceedings. All authors have reviewed and approved the manuscript, and the order of authorship has been mutually agreed upon.

Funding

Not applicable

Availability of data and material

All of the data obtained or analyzed during this study is included in the report that was submitted.

Conflicts of Interest

The authors declare that they have no financial or personal interests that could have influenced the research and findings presented in this paper. The authors alone are responsible for the content and writing of this article.

Authors' contributions

All authors contributed equally in the preparation of this

manuscript.

REFERENCES

- [1] Manjuparkavi, M.K., Venkatajalapathi SThilakharan, K., Vignesh, K., Yokesh, E. and Vaishnavi, D., **2026**. Neuro-Symbolic Agentic Swarms: A Hybrid Approach to Resilient Decision Making in Uncertain Edge Environments. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 8(2), pp.982-989. <https://doi.org/10.15662/IJEETR.2026.0802057>
- [2] Pliatsios, A., Mpatos, P. and Dossis, M., **2026**. A Review of Neuro-Symbolic AI for Smart Cities: Architectures, Ontologies, and IoT-Driven Applications. *IEEE Access*, 14, pp. 87815-87828. <https://doi.org/10.1109/ACCESS.2026.3702066>.
- [3] Grosvenor, A., Zemlyansky, A., Wahab, A., Bohachov, K., Dogan, A. and Deighan, D., **2025**. Hybrid intelligence systems for reliable automation: Advancing knowledge work and autonomous operations with scalable AI architectures. *Frontiers in Robotics and AI*, 12, p.1566623. <https://doi.org/10.3389/frobt.2025.1566623>.
- [4] Vijayakumar, S.D. and Karuppusamy, S.A., **2026**. An energy-efficient secure routing protocol using a hybrid Dempster-Shafer theory and adaptive snow ablation optimization routing framework for MANETs. *Wireless Networks*, 32, pp. 1829–1857. <https://doi.org/10.1007/s11276-026-04152-0>.
- [5] Ramya, K. and Karuppusamy, S.A., **2026**. Threat-Aware Mutual Authentication in Web3: A Privacy-Preserving Framework With ZKP-DID and Real-Time GNN Integration. *Security and Privacy*, 9(1), p.e70146. <https://doi.org/10.1002/spy2.70146>.
- [6] Vijayakumar, D., Karunakaran, P., Balasubramaniam, C., Chethanand, V., Dinesh Kumar, B. and Manoj, J.V., **2026**, March. Temporal Artifact Analysis for AI-Generated Video Detection Using Hybrid CNN-LSTM Architecture. In *2026 Second International Conference on Multi-Agent Systems for Collaborative Intelligence (ICMSCI)* (pp. 800-807). IEEE. <https://doi.org/10.1109/ICMSCI67830.2026.11469605>.
- [7] Ajal, A.J. and Karuppusamy, S.A., **2025**. Scalable and Resolution Data Analysis of Image and Video Compression using DL-CNNs Neural Network. *Circuits, Systems, and Signal Processing*, 44(3), pp.1845-1864. <https://doi.org/10.1007/s00034-024-02893-x>.
- [8] Saravanaprasad, P. and Karuppusamy, S.A., **2024**. Advanced lung tumor diagnosis using a 3D deep neural network-based CAD system. *Biomedical Signal Processing and Control*, 89, p.105650. <https://doi.org/10.1016/j.bspc.2023.105650>.
- [9] Mičiak, M., Toman, D., Adámik, R., Kuřová, E., Škulec, B., Mozolová, N. and Hoferová, A., **2025**. Effective Education System for Athletes Utilising Big Data and AI Technology. *Data*, 10(7), p.102. <https://doi.org/10.3390/data10070102>.
- [10] Janarthanan, R., Maheshwari, R.U., Shukla, P.K., Shukla, P.K., Mirjalili, S. and Kumar, M., **2021**. Intelligent detection of the PV faults based on artificial neural network and type 2 fuzzy systems. *Energies*, 14(20), p.6584. <https://doi.org/10.3390/en14206584>.
- [11] Selvapandian, A. and Manivannan, K., **2018**. Performance analysis of meningioma brain tumor classifications based on gradient boosting classifier. *International Journal of Imaging Systems and Technology*, 28(4), pp.295-301. <https://doi.org/10.1002/ima.22288>.
- [12] Sasikala, S., Sasipriya, S. and Uma Maheshwari, R., **2022**, March. Soft Computing based Brain Tumor Categorization with Machine Learning Techniques. In *2022 International Conference on Advanced Computing Technologies and Applications (ICACTA)* (pp. 1-9). IEEE. <https://doi.org/10.1109/ICACTA54488.2022.975288>.
- [13] Maheshwari, R.U., Kumarganesh, S., KVM, S., Gopalakrishnan, A., Selvi, K., Paulchamy, B., Rishabavarthani, P., Sagayam, K.M., Pandey, B.K. and Pandey, D., **2025**. Advanced plasmonic resonance-enhanced biosensor for comprehensive real-time detection and analysis of deepfake content. *Plasmonics*, 20(4), pp.1859-1876. <https://doi.org/10.1007/s11468-024-02407-0>.
- [14] Park, J., Park, L.H., Ahn, H.E. and Kwon, T., **2024**. Coexistence of deepfake defenses: Addressing the poisoning challenge. *IEEE Access*, 12, pp.11674-11687. <https://doi.org/10.1109/ACCESS.2024.3353785>.
- [15] Babaei, R., Cheng, S., Duan, R. and Zhao, S., **2025**. Generative artificial intelligence and the evolving challenge of deepfake detection: A systematic analysis. *Journal of Sensor and Actuator Networks*, 14(1), p.17. <https://doi.org/10.3390/jsan14010017>.

- [16] Nawaz, U., Anees-ur-Rahaman, M. and Saeed, Z., **2025**. A review of neuro-symbolic AI integrating reasoning and learning for advanced cognitive systems. *Intelligent Systems with Applications*, 26, p.200541. <https://doi.org/10.1016/j.iswa.2025.200541>.
- [17] De Sousa, A.L. and de Oliveira, A.S., **2024**. Deadlock-free production using Dempster-Shafer and preset methods in predictive scheduling for multiagent controlled flexible manufacturing systems. *Applied Soft Computing*, 152, p.111234. <https://doi.org/10.1016/j.asoc.2024.111234>.
- [18] Albahri, A.S., Duhaim, A.M., Fadhel, M.A., Alnoor, A., Baqer, N.S., Alzubaidi, L., Albahri, O.S., Alamoodi, A.H., Bai, J., Salhi, A. and Santamaría, J., **2023**. A systematic review of trustworthy and explainable artificial intelligence in healthcare: Assessment of quality, bias risk, and data fusion. *Information fusion*, 96, pp.156-191. <https://doi.org/10.1016/j.inffus.2023.03.008>.
- [19] Lu, Z., Afridi, I., Kang, H.J., Ruchkin, I. and Zheng, X., **2024**. Surveying neuro-symbolic approaches for reliable artificial intelligence of things. *Journal of Reliable Intelligent Environments*, 10(3), pp.257-279. <https://doi.org/10.1007/s40860-024-00231-1>.
- [20] Badhan, P.K., **2026**. Neuro-symbolic machine learning for lightweight and interpretable IoT edge intrusion detection. *Discover Sensors*, 2(1), p.15. <https://doi.org/10.1007/s44397-026-00047-z>.