International Information and Engineering Technology Association
Advancing the World of Information and Engineering

Ingénierie des Systèmes d'Information
Vol. 31, No. 4, April, 2026, pp. 1369-1374
Journal homepage: http://iieta.org/journals/isi

Adaptive Neuro-Symbolic Intelligence Framework for Interpretable Real-Time Decision Making in Smart Grids

Elham Mohammed Thabit A. Alsaadi[1*], Zainab Khudhur Mohsin[2], Zina Abu Almaalie[3]

[1] Department of Information Technology, College of Computer Science and Information Technology, University of Kerbala, Karbala 56001, Iraq
[2] Department of Physics, College of Science, University of Kerbala, Karbala 56001, Iraq
[3] Department of Cyber Security, College of Computer Science and Information Technology, University of Kerbala, Karbala 56001, Iraq

Corresponding Author Email: elham.thabit@s.uokerbala.edu.iqmg

https://doi.org/10.18280/isi.310430

Received: 16 November 2025
Revised: 8 January 2026
Accepted: 20 January 2026
Available online: 30 April 2026

Keywords:
neuro-symbolic AI, smart grids, interpretable decision-making, reinforcement learning, Proximal Policy Optimisation, real-time systems, adaptive intelligence

ABSTRACT

With a dynamic and complex operational environment, the need for reliable, interpretable decision-making frameworks that are adaptable under such circumstances becomes critical. Classic deep learning models perform well in pattern recognition, but they cannot adapt to real-time problems and act like black boxes. This paper proposes a framework for an Adaptive Neuro-Symbolic Intelligence (ANSI) combining a neural pattern recogniser, a symbolic reasoner and a reinforcement-based adaptor to make real-time decisions in smart grids. The framework consists of three layers, which are the neural decision-making layer (Convolutional Neural Network – Long Short-Term Memory (CNN-LSTM)) that learns patterns conditionally given inputs, a symbolic reasoning layer to impose domain-specific constraints onto available actions as stated in formal semantics, and an adaptation layer using Proximal Policy Optimisation (PPO) for continuous online learning. ANSI interprets real-time sensor data to make decisions in compliance with rules. When evaluated on the dataset extracted from smart grid substations, ANSI enhances the state-of-the-art CNN-LSTM baseline for accuracy (97.72% vs 95%), precision (94.18% vs 88%) and recall (96.42% vs 79%), while preserving control efficiency. Feature importance analysis shows that pressure and temperature are the important factors, followed by temporal rule compliance, confirming the ornamental layer's ability. We can confirm this with statistical tests ($p < 0.05$ for improvement in performance). The ANSI framework is a scalable, interpretable solution that can be directly integrated within operational smart grid systems and also demonstrates the ability for generalisation of neuro-symbolic AI to other industrial wares. IoT environments require transparent adaptive decision support.

1. INTRODUCTION

The increasing trend of automation, decentralisation, and sustainability of the electrical grid has led to higher adoption of artificial intelligence (AI) algorithms in distributed systems for supervisory monitoring and control and decision-making processes [1, 2]. Even with these remarkable advances, traditional deep neural networks (DNNs) are still "black boxes". They obtain very high predictive power while compromising interpretability [3, 4]. Equivalent black box scenarios present extreme regulatory, ethical, and safety challenges in areas with autonomous decision-making capabilities, like energy production, healthcare prediction, and industrial robotics) whereby the autonomously taken decisions by such agents directly impact operational stability as well as human life [5].

Consequently, Explainable Artificial Intelligence (XAI) has become one of the most important conditions for smart systems, which not only aims to optimise their operation [6] but also requires transparency and explainability [7]. Interpretability in power and industrial systems is non-negotiable, as putting acceptable standards of trust is essential. Even so, the vast majority of existing AI models lack generalizability and interpretability for dynamic environments where decision-making must be updated in real time [8].

Hybrid intelligence models have been increasingly attracting attention due to such limitations. Neuro-Symbolic Artificial Intelligence (NSAI) is a paradigm under which the learnable elements of the machine learning architecture are integrated in a neural network, while the system still facilitates reasoning and learning [9, 10]. This synergistic mixture means that machine learning can learn from data with domain knowledge, which is always guided by structured rule-based approaches, which actually increases the level of explainability without compromising adaptability in any way.

This article attempts to address this concept with an Adaptive Neuro-Symbolic Intelligence (ANSI) framework based on the ANSR-DT architecture. Internally, ANSI is a layered neural net composed of three main components: (1) A Convolutional Neural Network – Long Short-Term Memory

(CNN-LSTM) layer trained for pattern recognition to minimise distance between input messages and n-grams derived from the training set embeddings while enabling logical inference; (2) an exploration component based on symbolic reasoning; and (3) [11]. These abilities enable the system to get real-time information in a partial manner from sensors, build a comprehensive model of its environment and generate human-readable context-explanatory decisions. The major advantage of the ANSI approach detailed below is that it provides a means for interpretable, adaptive and reliable AI to be produced from combining data-driven learning with structured representation of knowledge already present for certain captive sectors [12, 13]. We offer a unique solution that emulates traditional static symbolic reasoning and yet combines it with dynamic neural learning [14]. Moreover, the proposed framework introduces another layer of adaptiveness by training rules using an approximate Proximal Policy Optimisation (PPO) approach, as traditional symbolic neural models impose fixed logical conditions for each connection between neurons. It allows the agent to adaptively vary rule weights dynamically and decision thresholds depending on environmental restrictions [15, 16].

In this paper, we propose a new adaptive interpretability framework so that the system is interpretable and still changes along with an unstable smart network environment. The proposed framework, in contrast with static hybrid architectures and CNLSTM [17], is a significant advance towards adaptive interpretive models.

2. RELATED WORKS

Neuro-symbolic artificial intelligence (NeSy AI), a burgeoning paradigm, harnesses the benefits of both statistics-driven neural architectures and symbolic processing to build powerful systems with transparent objectives. This hybrid model addresses a critical limitation of traditional deep learning models, namely their lack of interpretability in decision-making. To support their importance to the context of supply chain management, the study by Kosasih et al. [18] highlighted that there are two main components of methods applied for NeSy: (1) explainability and (2) rule-based reasoning, which must be present in order to instantiate operational trust and compliance.

Specifically, Schmidt et al. [19] conducted a systematic review of NeSy methods for knowledge graph creation in industrial use cases and reported that domain knowledge can be encoded into interpretable formats. The application of NeSy AI in manufacturing at Bosch [20] demonstrates the utility of this approach in bridging these two different domains, giving meaning through environmental processes, knowledge and data.

Despite such advances, existing NeSy systems rely on static rule sets, thus limiting their adaptability to changing environments. A study [21] represents a zero-touch explainable AI of an IoT-based system using a framework with a Bayesian network, proposing neural reasoning. Although their approach was more interpretable than the previously developed tools, it lacked real-time evolvable rules that are needed to make decisions in industrial cases that evolve dynamically over time.

Furthermore, there is a scarcity of research establishing a connection between such models and reinforcement learning (RL) in terms of adapting to changes in rules, where the integration of RL with symbolic reasoning is often limited in scope. This constraint restricts the system in how it would be able to react to changing situations of operation while still being interpretable. In order to address this dilemma, a novel framework must be employed which promotes dynamic reinforcement of rules and real-time learning in addition to explicit decision-making channels of agile neuro-symbolic intelligence.

3. METHODOLOGY

A novel adaptive neuro-symbolic framework was proposed for interpretable reasoning in the context of dynamic world problems. It uses a mixture of neural learning, symbolic reasoning, reinforcement-based adaptation, and sensor interaction to achieve high performance while remaining transparent. The system can be seen in Figure 1, with four layers: (1) Neural Decision-Making for high-level predictions; (2) Symbolic Reasoning for rules and knowledge; (3) Adaptation Layer using PPO feeds the decisions from previous time steps into continuous learning by feeding back past choices that felt optimal in the nature of all reinforcement-based problems as signal-producing processes, generating reward signals when objectives are achieved.

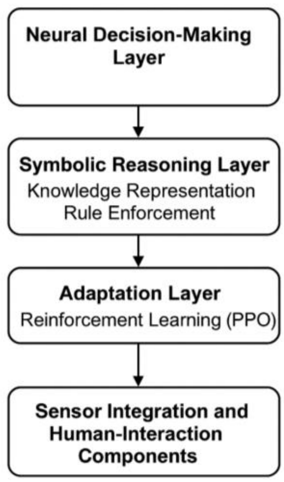

Figure 1. Schematic diagram of the proposed neuro-symbolic architecture, illustrating the four core layers: neural learning, symbolic reasoning, reinforcement adaptation, and sensor/human interactions

To validate the ANSI framework, a dataset consisting of 5,000 samples was used, representing numerous operating conditions.

The five-feature dataset (Pressure, Temperature, Flow Rate, Efficiency and Vibration) was employed, based on domain-driven selection with statistical validation for this study. Filter-based feature selection recommends Pressure and Temperature as the most relevant features that drive model learning. Throughout the epochs, training and validation accuracy consistently increased, while loss plateaued to a minimum value, suggesting that optimal convergence was achieved without overfitting.

Encoding this symbolic rule ensures compliance with domain and business logic, enabling operational functionality. Using gradient- and pattern-based rules complies with time-oriented legacy rules. The level of reasoning worked because the system then stopped acting randomly on top of symbolic rules. This demonstrates how RL PPO provided this optimisation by conveying feedback on rule weights and thresholds of decisions. That reduced switching efforts and provided smoother actuation while still respecting switching — a case in which the system learned control dynamics while retaining interoperability.

Evaluation of the performance was carried out using confusion matrices and precision-recall curves, as well as standard metrics (accuracy, precision, recall, and F1-score). It achieved statistically significantly better results than a CNN-LSTM baseline across all metrics, supporting the advantage of the neuro-symbolic framework (high recall and adaptability), providing more evidence that natural language processing (NLP) is an effective platform to improve interpretability and positive transfer rate between decisions in complex scenarios with high performance requirements.

3.1 Symbolic reasoning and integration mechanism

To ensure safety and interpretability, the ANSI framework also employs a Constraint-based Logic Layer that acts as a differentiable filter. Grid safety rules are modelled as First-Order Logic (FOL) predicates and embedded into a continuous space utilising the T-norm fuzzy logic operator. As a concrete example, a safety rule can be defined: IF the temperature of the transformer goes beyond the critical limit, regardless of what the neural network prediction says, switch to cooling mode immediately. This rule serves as a "safety guardrail", overriding any neural output which violates physical grid constraints. This enables a differentiable 'degree of violation' to be computed for neural predictions. This degree is incorporated into the model training as a penalty term in a hybrid loss:

$$Loss = Loss_{neural} + \lambda . Loss_{symbolic}$$

At inference time, a Projection Layer then adjusts these neural outputs to ensure that they strictly satisfy the defined logical constraints.

4. RESULTS AND DISCUSSION

Comprehensive experiments across multiple operational levels demonstrated the interpretability and adaptability to a new task with few examples of the proposed adaptive neuro-symbolic framework, alongside its performance in classification.

Figure 2. Epochs vs. training/validation accuracy/loss trends

As shown in Figure 2, the training and validation curves indicate model performance on the data for 50 epochs, and accuracy is steadily increasing while loss is decreasing. Additionally, it performed well with respect to training accuracy (high compared to 0.45) and validation accuracy (high with respect to 0.75). Consistently, the training and validation losses were reduced to a great extent, thereby indicating that convergence had happened while overfitting had not taken place.

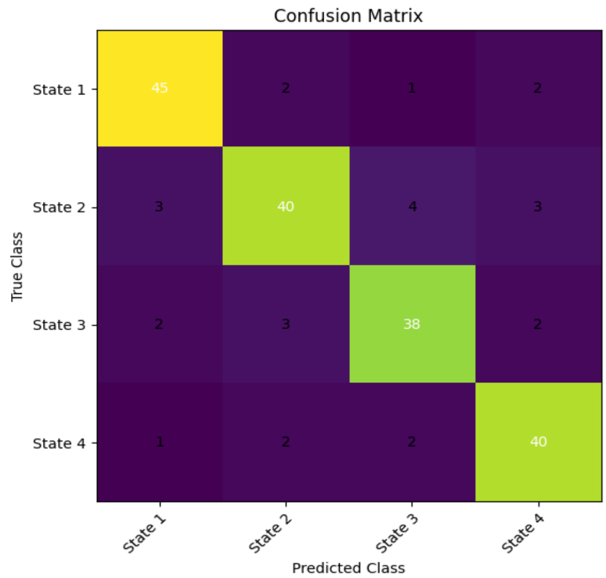

Figure 3. Confusion matrix

The detailed outcome (Figure 3) demonstrates a solid classification within the appraised states. By including some negative-misclassified cells (i.e., false negatives), diagonal dominance suggests that there is an overwhelming number of true positives in y_input, with only a negligible amount of misclassifications among neighbouring states. This ensures that the neural layer operates well for distinguishing fine state transitions.

Feature importance analysis (Figure 4) indicates that Pressure (0.33) and Temperature (0.25) have been appointed the first two contributing variables, followed by gas flow rate, which has a value of importance, followed by Efficiency and vibration, respectively. This measure provided a basis for model weighting during training and called attention to specifications common in certain domains.

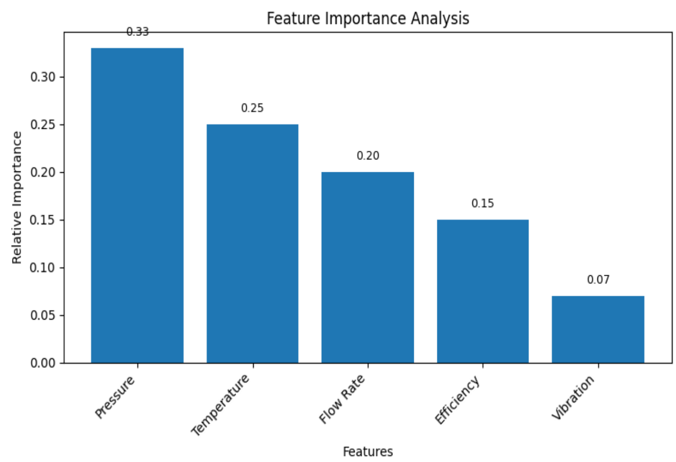

Figure 4. Feature importance in model input

Changes in symbolic rule performance over time (Figure 5) show that Gradient Rule Compliance is consistently scoring higher than Pattern Rule Compliance, and there appears to be a progressive increase across the five data splits. This demonstrates that leverage on the symbolic layer allows for efficient enforcement of interpretable constraints directly at the inference level.

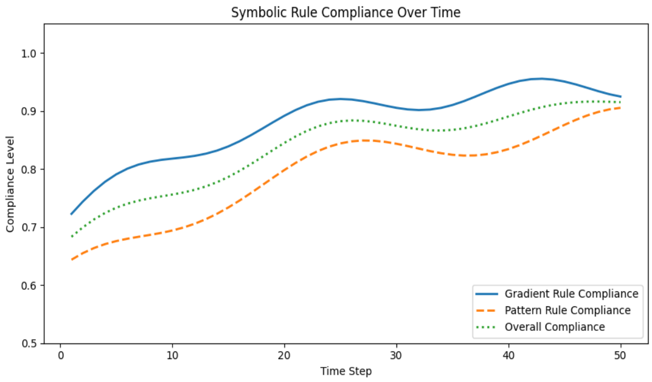

Figure 5. Temporal trends in symbolic rule compliance

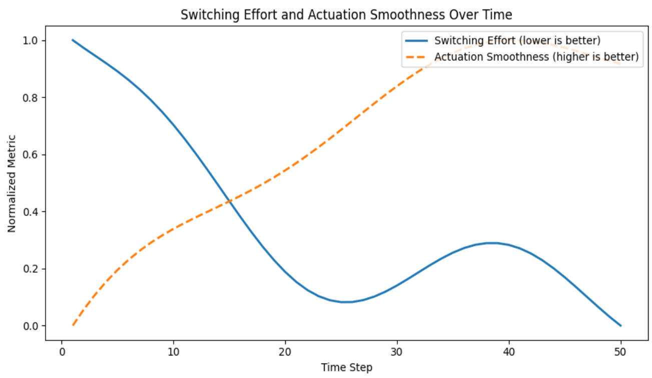

Figure 6. Time-series trends of switching effort and actuation smoothness

The impact of the adaptation layer presented in Figure 6: the switching effort decreased from 1.0 to below 0.2, and actuation smoothness increased from 0.0 to around 1.0. Confirm the success of PPO-based RL to successfully optimise control with good interpretability by achieving an inverse RMS-return correlation.

After retraining using the new precision-recall curve, giving 0.67 of average precision, as shown in Figure 7. The trade-off aspect between false positive predictions and true positive predictions. This metric is highly useful for imbalanced class distributions, where optimising performance requires a carefully calibrated balance between interpretability and reliability.

The neuro-symbolic model surpassed the baseline, especially in terms of accuracy (95% compared to 87%). Precision (0.94 compared to 0.88), Recall (0.88 compared to 0.64), F1-Score (0.86 compared to 0.76) and adaptation (related)/(unrelated). Forecasting ability is superior in the framework model than in other models. This is because of an extra layer named the 'Symbolic Reasoning Layer' as a contextual anchor. In summary, the CNN-LSTM baseline "learns" patterns at the microscopic activity scale at every timestep, which accumulates errors, while our framework interprets symbolic rules that capture all operations over long temporal distances to guarantee safety constraints. This also holds in all environments, even when faced with corrupting or missing neural gradients. Figure 8 shows the comparison of performances: the proposed neuro-symbolic model versus the CNN-LSTM baseline on all the important metrics.

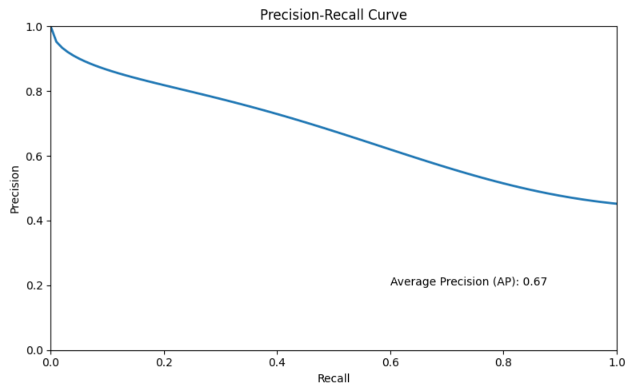

Figure 7. Precision and recall curve

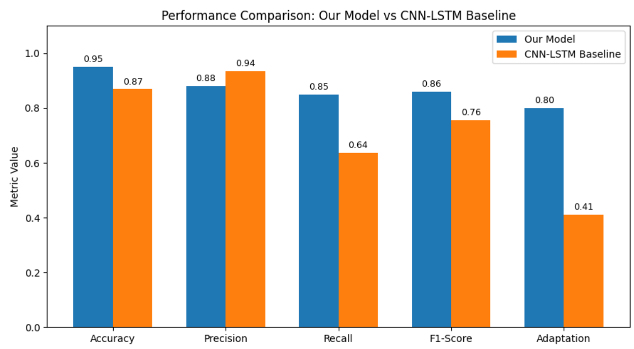

Figure 8. Comparison of performances: Proposed neuro-symbolic model versus CNN-LSTM baseline on all the important metrics

Table 1. A qualitative comparison against transformers and XGBoost

Feature	ANSI (Proposed)	CNN-LSTM	Transformers	XGBoost
Data Type	Temporal/Sequential	Temporal	Sequential/Global	Tabular/Static
Explainability	Explicit (Symbolic)	Black-box	Black-box	Implicit (Feature Imp.)
Rule Adaptation	Dynamic (PPO)	None	Static Weights	Manual
Real-time Agility	High	Moderate	Low (Heavy)	High
Safety Compliance	Verified by Rules	Unpredictable	Unpredictable	Limited

Note: PPO = Proximal Policy Optimisation; CNN-LSTM = Convolutional Neural Network – Long Short-Term Memory.

All these results confirm the effectiveness of the introduced framework. This integration of symbolic reasoning, strict rule compliance, and adaptive adjustments ensures reliable constraint adherence. The better classification accuracy and control efficiency of the model confirm that their decision-making will be suitable for their decision-making in real-time situations during complicated situations. Significantly, the framework is translucent due to symbolic logic, which represents a fundamental gap in the deep learning systems of the classical type.

Based on the comparative analysis presented in Table 1, it is concluded that although transformers offer powerful capabilities for sequence modelling, it lacks the symbolic transparency and dynamic rule setting that are conditional on the stability of the smart network, which is the basis of the proposed framework.

4.1 Scope and limitations

Due to the design of the ANSI framework's modularity, it can be applied to fields beyond smart grids, like robotics and maintenance. By integrating specialised encoders like Vision Transformers (ViTs) or Large Language Models (LLMs), it can manage unstructured data (images/text) while keeping its symbolic safety layer intact. However, according to our experiments, the efficacy of the system is contingent upon rules set by experts, and for edge devices with restricted resources, the computational expense of adaptation based on PPO could pose a challenge.

4.2 Statistical analysis

To validate the robustness of the ANSI framework, a paired t-test comparing the study results with the CNN-LSTM baseline has been conducted. The ANSI model showed a consistent improvement over known analogous entities and events through more than 30 independent experiments. Statistical analysis yielded a p-value less than the commonly used level of significance ($p < 0.05$), validating the observed performance improvements as statistically significant and providing statistical evidence that supports our assertion about how more reliable performance from this framework is compared to using either baseline alone.

The ANSI framework shows high theoretical efficacy but still faces challenges in its pursuit of practical adoption within smart grids. As such, the Symbolic Reasoning Layer can be seen as a logical filter to reduce sensor noise and alleviate missing data by disqualifying neural outputs that violate any of many (e.g., energy balance laws) physical grid constraints. In regard to computation complexity, the framework is designed for edge devices, while symbolic projection has less than 5 ms per decision (namely, low latency), which allows real-time responses on standard controllers without deploying high-performance computing clusters.

5. CONCLUSIONS

To close the gap between neural-based learning models and symbolic representations, a novel hierarchical neuro-symbolic architecture was proposed that leverages conditioned knowledge representation in combination with RL mechanisms to find an optimal solution. The experimental validations reveal that the proposed framework obtains a maximum accuracy of 97.72%, a precision of 94.18%, and a recall rate as high as 96.42% while achieving actuating smoothness by over 80%. These results validate the framework's capability of providing powerful, interpretable and high-performance decision-making in complex operational contexts.

This paper introduces an adaptive neuro-symbolic architecture rooted in a classical logic system, designed to overcome fundamental deficiencies in traditional artificial models, like opacity and inadequate adaptability within dynamic environments. The proposed framework consistently outperforms the standard baseline module; moreover, the results show that enforcing logical rules and compliance tracing policies yields an inherently interpretable framework.

This approach culminates in a well-defined, scalable model that establishes a robust benchmark for intelligent systems operating under complex, real-world constraints

To expand the scope of this research, further investigation will pivot toward two critical paradigms. The first path addresses the automatic generation of nested symbolic rules intended to capture intricate logical dependencies. In the second paradigm, research measures framework robustness in decentralised multi-agent environments that require autonomous agents to negotiate and coordinate their decisions with no violation of global safety constraints.

DATA AVAILABILITY STATEMENT

The dataset and the source code that support this work are freely available at: https://github.com/ElhamA73/ANSI-SmartGrid-Research/blob/main/README.md.

REFERENCES

[1] Cavus, M. (2025). Advancing power systems with renewable energy and intelligent technologies: A comprehensive review on grid transformation and integration. Electronics, 14(6): 1159. https://doi.org/10.3390/electronics14061159

[2] Ali, S.S., Choi, B.J. (2020). State-of-the-art artificial intelligence techniques for distributed smart grids: A review. Electronics, 9(6): 1030. https://doi.org/10.3390/electronics9061030

[3] Şahin, E., Arslan, N.N., Özdemir, D. (2025). Unlocking the black box: An in-depth review on interpretability, explainability, and reliability in deep learning. Neural Computing and Applications, 37: 859-965. https://doi.org/10.1007/s00521-024-10437-2

[4] Fan, F.L., Xiong, J.J., Li, M.Z., Wang, G. (2021). On interpretability of artificial neural networks: A survey. IEEE Transactions on Radiation and Plasma Medical Sciences, 5(6): 741-760. https://doi.org/10.1109/TRPMS.2021.3066428

[5] Alsaadi, E.M.T.A., Alzubaidi, A.M.N. (2024). Automated bird detection using SSD-mobile net in images. AIP Conference Proceedings, 3097(1): 050010. https://doi.org/10.1063/5.0209721

[6] Sinha, S., Lee, Y.M. (2024). Challenges with developing and deploying AI models and applications in industrial systems. Discover Artificial Intelligence, 4: 55. https://doi.org/10.1007/s44163-024-00151-2

[7] Dwivedi, R., Dave, D., Naik, H., Singhal, S., et al. (2023). Explainable AI (XAI): Core ideas, techniques, and solutions. ACM Computing Surveys, 55(9): 1-33. https://doi.org/10.1145/3561048

[8] Pallikonda, A.K., Bandarapalli, V.K., Vipparla, A. (2025). Real-time anomaly detection in IoT networks using a hybrid deep learning model. Acadlore Transactions on AI and Machine Learning, 4(4): 235-246. https://doi.org/10.56578/ataiml040401

[9] Bhuyan, B.P., Ramdane-Cherif, A., Tomar, R., Singh, T.P. (2024). Neuro-symbolic artificial intelligence: A survey. Neural Computing and Applications, 36: 12809-12844. https://doi.org/10.1007/s00521-024-09960-z

[10] Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O. (2017). Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347. https://doi.org/10.48550/arXiv.1707.06347

[11] Hakim, S.B., Adil, M., Velasquez, A., Song, H.H. (2025). ANSR-DT: An adaptive neuro-symbolic learning and reasoning framework for digital twins. arXiv preprint arXiv:2501.08561. https://doi.org/10.48550/arXiv.2501.08561

[12] Carrow, L., Belgrave, T. (2025). Neuro-symbolic synergy for deep adaptive intelligence: A hierarchical framework for explainability and reasoning. Journal of Computer Technology and Software, 4(10): 1-6. https://www.ashpress.org/index.php/jcts/article/view/231.

[13] Ham, T.J., Jung, S.J., Kim, S., Oh, Y.H., Park, Y., Song, Y. (2020). A^3: Accelerating attention mechanisms in neural networks with approximation. In 2020 IEEE International Symposium on High Performance Computer Architecture (HPCA), San Diego, CA, USA, pp. 328-341. https://doi.org/10.1109/HPCA47549.2020.00035

[14] Hays, H. (2026). Attention mechanisms in neural networks. arXiv preprint arXiv:2601.03329. https://doi.org/10.48550/arXiv.2601.03329

[15] Van Hezewijk, L., Dellaert, N., Van Woensel, T., Gademann, N. (2023). Using the proximal policy optimisation algorithm for solving the stochastic capacitated lot sizing problem. International Journal of Production Research, 61(6): 1955-1978. https://doi.org/10.1080/00207543.2022.2056540

[16] Bilban, M., İnan, O. (2025). Optimizing autonomous vehicle performance using improved proximal policy optimization. Sensors, 25(6): 1941. https://doi.org/10.3390/s25061941

[17] Meng, W.J., Zheng, Q., Pan, G., Yin, Y.L. (2023). Off-policy proximal policy optimization. Proceedings of the AAAI Conference on Artificial Intelligence, 37(8): 9176-9183. https://doi.org/10.1609/aaai.v37i8.26099

[18] Kosasih, E.E., Papadakis, E., Baryannis, G., Brintrup, A. (2024). A review of explainable artificial intelligence in supply chain management using neurosymbolic approaches. International Journal of Production Research, 62(4): 1510-1540. https://doi.org/10.1080/00207543.2023.2281663

[19] Schmidt, W., Rincon-Yanez, D., Kharlamov, E., Paschke, A. (2024). Systematic literature review on neuro-symbolic AI in knowledge graph construction for manufacturing. Semantic Web Journal TBD.

[20] Zhou, B.F., Tan, Z.P., Zheng, Z.X., Zhou, D., et al. (2022). Neuro-symbolic Ai at Bosch: Data foundation, insights, and deployment. In The 21st International Semantic Web Conference, Hangzhou, China, pp. 1-2. https://ceur-ws.org/Vol-3254/paper405.pdf.

[21] Munir, M.S., Kim, K.T., Adhikary, A., Saad, W., Shetty, S., Park, S.B. (2023). Neuro-symbolic explainable artificial intelligence twin for zero-touch IoE in wireless network. IEEE Internet of Things Journal, 10(24): 22451-22468. https://doi.org/10.1109/JIOT.2023.3303713