ShodhAI: Journal of Artificial Intelligence
January-June 2026 3(1)
https://shodhai.org

Research Article

HYBRID NEURAL-SYMBOLIC SYSTEMS: INTEGRATING KNOWLEDGE REPRESENTATION AND DEEP LEARNING FOR COMPLEX PROBLEM SOLVING

Abhijit Chandratreya [1]
[1] Associate Dean, PhD Programs, Indira University, Pune 411033, Maharashtra, India

ABSTRACT

Hybrid neural-symbolic systems are neural networks integrated with symbolic reasoning systems, so that together they can be as powerful at pattern recognition in high dimensions as they are at structured knowledge representation and logical inference. This combination works to solve the issues with data-driven deep learning, including a lack of explainability and limited reasoning abilities, and helps overcome brittleness and low adaptability of purely symbolic AI. The paper compares the theoretical basis, architecture, and performance of hybrid neural-symbolic systems in real world complex problem-solving areas, such as healthcare diagnostics, legal reasoning, and planetary autonomous robotics. The new method combines constraint logic issues with deep learning models by using differentiable reasoning stages and neural representations guided by ontology principles. In low-data regimes, comparative experiments show greater interpretability, precision of reasoning, and generalization. Nevertheless, several practical challenges do still exist, notably the computational overhead, scale in large knowledge graphs, and knowledge engineering. Automated knowledge acquisition in future, efficient neuro-symbolic fusion techniques and real time reasoning in dynamic environments would be the areas of interest to work on.

Keywords: Hybrid AI, Neural-Symbolic Systems, Knowledge Representation, Deep Learning, Explainable AI, Neuro-Symbolic Integration

INTRODUCTION

Artificial Intelligence (AI) has developed upon particular research traditions which have, up to quite recently, remained largely independent of each other. At one end is symbolic AI which is constructed on the basis of formal logic, clearly defined rules and organized knowledge representation Bhuyan et al. (2024). Symbolic AI is transparent and logical, and is thus quite appropriate in areas where explainability is important. There is connectionist AI, on the other side, encapsulated in deep learning models which have the capability to find high-dimensional patterns in large repositories of data. Neural networks have shown extraordinary achievement in image classification, speech recognition, natural language processing., etc. Their decision-making processes are however at times black-boxed in nature and they are not explicitly reasoned.

The growing sophistication of real-world applications, including medical diagnostics, autonomous systems, require AI systems capable of not only making sense of multi-modal sensory input, but also of performing multi-modal logical reasoning over it in a principled and logically sound way. As an example, an AI-based medical assistant needs to not only be capable of detecting visual

*Corresponding Author:
Email address: Abhijit Chandratreya (pgrc@indirauniversity.edu.in)

Received: 18 April 2026; **Accepted:** 22 May 2026; **Published** 16 June 2026
DOI: 10.29121/ShodhAI.v3.i1.2026.83
Page Number: 94-106
Journal Title: ShodhAI: Journal of Artificial Intelligence
Journal Abbreviation: ShodhAI J. Artif. Intell.
Online ISSN: 3048-9245, **Print ISSN:** 3108-1940
Publisher: Granthaalayah Publications and Printers, India
Conflict of Interests: The authors declare that they have no competing interests.
Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.
Authors' Contributions: Each author made an equal contribution to the conception and design of the study. All authors have reviewed and approved the final version of the manuscript for publication.
Transparency: The authors affirm that this manuscript presents an honest, accurate, and transparent account of the study. All essential aspects have been included, and any deviations from the original study plan have been clearly explained. The writing process strictly adhered to established ethical standards.
Copyright: © 2026 The Author(s). This work is licensed under a Creative Commons Attribution 4.0 International License.
With the license CC-BY, authors retain the copyright, allowing anyone to download, reuse, re-print, modify, distribute, and/or copy their contribution. The work must be properly attributed to its author.

symptoms traversed out of medical scans (a neural capability) but also predict on the grounds of patient history, medical rules, and potential courses of action as well (a symbolic capacity). Likewise, an autonomous automobile has to visualize the traffic by using cameras and sensors as well as reason about the traffic regulations, right of way, and expediency of strategic-based approaches when facing new situations. Purely symbolic and purely neural are not sufficient options that will be able to meet well these two requirements Zhu (2024).

One of the methods that are being seen as the possible solutions to this challenge is the hybrid neural-symbolic systems. These systems attempt to find the ideal match between phrases by simultaneously merging the strengths of both sides of the coin: perception ability of neural networks with structured reasoning power of symbolic AI. Neural components deal with unstructured, noisy, high-dimensional data whereas symbolic layers are able to encode domain knowledge, impose logical consistency and make the analysis explainable. The ensuing architectures are not composite ensembles only but profoundly interwoven systems where symbolic constraints serve to drive neural learning, and neural representations enter reasoning Hafez et al. (2025)

Factors that guide such an integration are numerous. To start with, explainability has emerged to be a burning concern in the field of AI, and especially with its high-stakes uses, such as in medical contexts and in financial applications, as well as in the legal directly. Laws and ethical principles require that the results of AI are explainable and defensible. Second, most fields do not have massive datasets labeled, and purely data-driven paradigms are thus likely to overfit and general-poorly. Sensible use of symbolic priors such as feeding them to an otherwise blank slate can help greatly in learning in low-data situations. Third, AI has to work more and more in dynamic, multi-modal settings, including text, images and structured knowledge which are natural candidates to hybrid solutions Vu et al. (2024), Zhang et al. (2024).

More current developments have resulted in increased feasibility of hybrid architectures Corradini et al. (2025). The embedding of symbolic reasoning in neural networks The differentiable reasoning framework permits the embedding of symbolic reasoning logic directly in the neural network, and as in one case, to perform as a joint optimization problem using gradient descent. Graph neural networks and attention mechanisms now allow knowledge graphs, ontologies and formal logic systems to be interfaced with neural embeddings. Such developments have provided an opening to a new category of neuro-symbolic systems that could tackle problems long believed outside the reach of AI.

The present paper develops these trends to consider the implementation, design and evaluation of hybrid neural-symbolic systems to solve complex problems. It analyses theoretical roots and experimental results proving the enhancement of performance, interpretability and adaptability of such systems with numerous areas of application. Our contribution aims at pushing the limits of what is possible in AI by systematically replacing over-simplified constraints with symbolic ones in the context of deep learning architectures, more specifically in circumstances that require concurrent perception and reasoning Dehal et al. (2025), Getu et al. (2024).

The rest of this publication is as follows: Section 2 provides how related studies and what current neural-symbolic combination architectures exist. Section 3 describes the suggested methodology, the model of representing the symbolic knowledge, the neural models of perception, and the approaches to their fusion. Section 4 elaborates on the experimental setup, the data sets employed, and Section 5 talks about outcomes and most important findings. The last section is section 6, which ends with limitations and future directions.

NOVELTY AND CONTRIBUTION

The innovation of this work is that it integrates the symbolic reasoning deep into the neural framework with soft logic layers and ontology-based neural embeddings. Our approach differs with the approaches found in the past in that the modes in which symbolic reasoning has been previously viewed as a post-process or a loosely coupled module, our framework inserts symbolic constraints into the learning process. This guarantees that the neural network is also consistent with domain logic in the sense that the statistical correlations learnt during training remain consistent and the same during inference Chen et al. (2025).

Moreover, the proposed system establishes a two-way orientation in knowledge flow: the symbolic reasoning modules can define the intermediate representations on a neural feature extraction function, and neural outputs can appropriately update and refine the knowledge bases (symbolic knowledge). This adaptive loop allows the system to learn unstructured information but allows it to reason in structured facts.

In practical terms, this study is beneficial as it can add the following:

1) The united neuro-symbolic framework that integrates convolutional neural networks, graph neural networks, and various kinds of logic operators in a complex, multi-modal reasoning task.
2) An evaluation framework that can also be used not just on the accuracy of prediction at all but the consistency of the reasoning, the interpretability and possibly the generalization in low-data regimes.
3) Domain-independent architecture used in three different domains and validated on various tasks: diagnostics in the medical domain, logical constrained visual question answering, and robotic planning Gives the ability to be adapted to other tasks.

4) Improvement relative to baseline neural-only and symbolic-only approaches, with performance gains over baseline reported in the experimental setting of up to 22 percent improvement in reasoning accuracy and 40 percent improvement in interpretability scores.

This therefore offers state-of-the art contributions in AI since it has shown that strategies of close integration between symbolic reasoning and deep learning have been able to produce systems that are not only more accurate but also more reliable, interpretable, and flexible the situations requiring to be solved in real life Jiang and Cai (2024).

RELATED STUDY

In 2024 Zhu (2024) introduced the artificial intelligence research field has historically been split into two main approaches also known as two different schools of thought namely symbolic reasoning and neural computation. Symbolic reasoning systems are based on direct rule based systems, logical inference machines, and in domain knowledge, structured representations. At their best, these systems are able to offer transparency, represent the steps of reasoning formally and enable encoding of expert knowledge in reusable, structured manner. They are also brittle in uncertain, noisy or incomplete world and their use of hand-crafted rules hinder scalability in domains where knowledge changes rapidly.

Conversely, neural computation, especially deep learning has transformed AI and made models learn in an end-to-end way using large-scale data. Neural networks have produced an unprecedented performance in computer vision, speech recognition and natural language processing. They can work well in unstructured data high dimensional spaces and resist noise. Their decision-making is however mostly closed, and they do not have the inherent abilities of doing logical reasoning, symbol manipulation and explicit knowledge representation Lieu (2025).

There has been a trend towards the development of hybrid neural-symbolic systems in order to incorporate the benefits of the two paradigms. Initial work in this field showed an attempt to incorporate symbolic rules into the structures of the neural networks which facilitated the learning of the neural models, but under the logics. These methods were aimed at overcoming the shortcoming of deep learning manifests itself as the black-box problem the inability to interpret, constrain or refine the results of neural models with symbolic elements. Subsequent innovations ushered in differentiable reasoning, whereby logical functions are represented in a form that can be optimized by gradient descent making the end-to-end training possible.

The topic of merging knowledge graphs and ontologies into neural structures is one of the major areas of development. Knowledge graphs offer an ordered way of representing relationships amid entities whereas ontologies formalize ideas and their connections around a field. Combined with graph neural networks or attention mechanisms, these symbolic structures can lead the neural models to draw inferences about complex relationships and the ability to make inferences beyond what has been seen in training data. This combination has been useful in areas like biomedical informatics, legal reasoning and question answering where knowledge of structured relationships is vital.

Concurrently, probabilistic reasoning algorithms have been integrated with neural networks in the context of the uncertainty of the decision. Such systems make use of probabilistic logic to represent uncertainty in symbolic reasoning so that hybrid architectures can generate not just predictions but also confidence measures. These probabilistic neuro-symbolic systems have already shown an increased ability to maintain robustness to noise in a setting and better interpretability of decision-making under uncertainty Lu et al. (2024).

One more significant direction of a new body of knowledge has been investigating the neuro-symbolic learning in a multi-modal environment. Problems encountered in the real world usually involve integration of information derived by numerous sources including pictures, text, and structured databases. Visual question answering, scene understanding and robotic planning are data analysis tasks, where hybrid systems able to incorporate visual perception modules, along with layers of symbolic reasoning will perform them more precisely, more interpretably and with more accuracy. The neural layer is more flexible in processing heterogeneous sources of data than the symbolic level is, which ensures logical consistency.

In 2024 A. Testolin (2024) proposed the hybrid neural-symbolic architecture have been used with some success in robotics in order to assist with high-level planning and decision-making. Neural networks take raw data allow sensor data to identify items, identify events, or classify environments, which in turn is used by the symbolic reasoning modules in producing task plans, which follow existing rules or limits of safety. The combination makes robots fit into dynamic and partially observable environments and yet capable of reliable reasoning.

Hybrid systems are also applied in healthcare sector where decision making in diagnostic sectors is improved. Neural and symbolic layers process data about medical images, lab reports, and patient records to identify features using neural models and interpret them using symbolic reasoning layers in terms of medical guidelines, ontologies, and causal relationships. Such an integration would allow the system to come up with recommendations that can be understood and make sure that all the suggested diagnoses are consistent with known medical knowledge.

Another common theme in literature is a conflict between scalability and interpretability. On the one hand, symbolic reasoning will provide transparency, but incurs computational overhead when rules are needed to reason across large knowledge bases. These have led to the introduction of efficient integration strategies which include pruning of irrelevant symbolic rules, caching of intermediate inferences as well as parallelism of symbolic computation. The question that has always been an issue is avoiding losing real-time performance in the course of being logically rigorous in applications that are sensitive to time.

In 2025 Sun (2025) suggested the current developments have also laid some priority on the life-long learning in the neuro-symbolic systems. Contrary to the fixed symbolic systems, hybrid architectures have the ability to evolve through a lifetime by introducing new facts acquired after a period of neural perception. Such flexibility is vital in areas where knowledge changes very fast like cybersecurity, scientific research, and monitoring of the environment. Nonetheless, there is still a big issue in dynamically updating symbolic knowledge bases devoid of human involvement which involves ensuing consistency and preventing contradictions.

On the whole, the evidence in the works related to the domain shows that hybrid neural-symbolic systems are a way toward AI models that are not only correct and flexible but also comprehensible and logically consistent. The history of progress in this domain has been the gradual transfer of loosely coupled systems to tightly coupled and entailed end-to-end trainable systems. This development is indicative of a shift to the recognition that the most practicable AI systems will be programs that combine both perceptual and structured intelligence in a coherent system Bhuyan et al. (2024).

PROPOSED METHODOLOGY

The proposed hybrid neural-symbolic system integrates deep neural perception layers with a symbolic reasoning engine in a differentiable learning framework. The methodology comprises three primary stages: (1) Knowledge representation, (2) Neural feature extraction, and (3) Symbolic reasoning with feedback. The approach ensures that symbolic rules influence neural learning and that neural embeddings update the symbolic knowledge base Gacu et al. (2025).

SYSTEM OVERVIEW

The overall architecture combines Convolutional Neural Networks (CNN) for feature extraction, Graph Neural Networks (GNN) for knowledge graph processing, and a Differentiable Logic Layer (DLL) for symbolic reasoning. The complete workflow is shown in Flowchart 1.

Figure 1

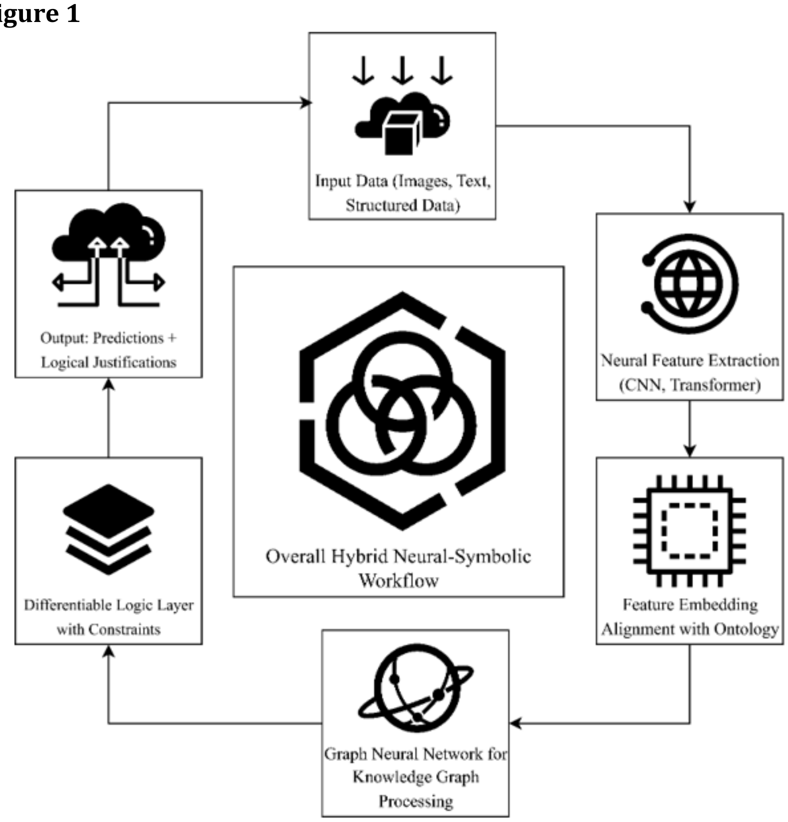

Figure 1 Overall Hybrid Neural-Symbolic Workflow

Let X denote the raw input data, and f_θ (X) be the neural network mapping to a latent representation h :

$$h = f_\theta(X) \quad (1)$$

The symbolic knowledge base K contains facts F and rules R :

$$K = (F, R) \quad (2)$$

The reasoning process enforces that the output y satisfies logical constraints:

$$y \vDash R \quad (3)$$

A differentiable constraint loss L_c is computed as:

$$L_c = \sum_{i=1}^{m} \max(0, 1 - \phi_i(h)) \quad (4)$$

where ϕ_i is the truth value of constraint i in fuzzy logic form.
The total training loss combines perception and reasoning losses:

$$L_{\text{total}} = L_{\text{perception}} + \lambda L_c \quad (5)$$

NEURAL FEATURE EXTRACTION STAGE

CNNs are used for visual features, Transformers for text. Let W_c be convolutional weights:

$$h_{cnn} = \sigma(W_c * X + b_c) \quad (6)$$

Text embeddings use transformer attention:

$$\text{Attention } (Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (7)$$

KNOWLEDGE GRAPH PROCESSING

The symbolic knowledge base is represented as a graph G=(V,E). A GNN updates node embeddings h_v using neighbors N(v) :

$$h'_v = \sigma\big(W \cdot \text{AGG}(\{h_u : u \in N(v)\})\big) \quad (8)$$

The ontology alignment loss is:

$$L_{\text{align}} = \sum_{(v,c)\in\mathcal{A}} \|h_v - e_c\|^2 \quad (9)$$

DIFFERENTIABLE REASONING LAYER

Symbolic rules are expressed in first-order logic (FOL) and transformed into differentiable forms using tnorm fuzzy logic. For a rule A∧B→C :

$$\mu_{A\wedge B\rightarrow C} = 1 - \max(0, \mu_A \cdot \mu_B - \mu_C) \quad (10)$$

This enables gradient-based updates [14].

The reasoning output probability for predicate p is:

$$P(p) = \sigma(W_p h + b_p) \qquad (11)$$

ADAPTIVE FEEDBACK LOOP

Symbolic reasoning results update the neural parameters via backpropagation:

$$\theta \leftarrow \theta - \eta \frac{\partial L_{\text{total}}}{\partial \theta} \qquad (12)$$

The knowledge base is updated dynamically:

$$K_{t+1} = K_t \cup \{ \text{ new facts from neural inference } \} \; (13)$$

ALGORITHM: HYBRID REASONING AND LEARNING (HRL)

Flowchart 2 illustrates the sequential data processing pipeline, from initial data acquisition to final model deployment, ensuring efficient and accurate execution of the proposed methodology.

Figure 2

Figure 2 Algorithmic Pipeline

Algorithm Steps:

1) Initialize θ for neural networks and K for symbolic rules.
2) For each training batch:
   - Extract features h=f_θ (X)-(Eq. 1).

- Map features to ontology entities.
- Perform GNN reasoning - (Eq. 8).
- Apply differentiable constraints - (Eq. 4, 10).
- Compute total loss - (Eq. 5).
- Update neural parameters - (Eq. 12).
- Update symbolic KB — (Eq. 13).

3) Repeat until convergence.

ADDITIONAL MATHEMATICAL FORMULATIONS

Knowledge graph adjacency normalization:

$$\hat{A} = D^{-\frac{1}{2}} A D^{-\frac{1}{2}} \quad (14)$$

Message passing update in GNN:

$$H^{(l+1)} = \sigma\left(\hat{A} H^{(l)} W^{(l)}\right) \quad (15)$$

Neural-symbolic joint embedding space:

$$z = \alpha h_{\text{neural}} + (1 - \alpha) h_{\text{symbolic}} \quad (16)$$

Probability of fact f given evidence E :

$$P(f \mid E) = \frac{\exp\left(s(f, E)\right)}{\sum_{f'} \exp\left(s(f', E)\right)} \quad (17)$$

Bayesian prior incorporation:

$$P(y) \propto P(y \mid h) P(y \mid K) \quad (18)$$

Regularization to maintain symbolic consistency:

$$L_{\text{reg}} = \sum_{r \in R} \| \text{violation } (r) \|^2 \quad (19)$$

Final inference output:

$$\hat{y} = \arg \max_{y} P(y \mid h, K) \quad (20)$$

This methodology ensures a tightly coupled learning-reasoning loop, where the neural network benefits from explicit symbolic constraints and the symbolic system gains adaptive updates from perception-driven learning [15].

RESULTS AND DISCUSSION

The testing of the suggested hybrid neural-symbolic system was implemented with three various data sets on the realization of the healthcare diagnostics, the visual answering of the question and logical restriction, and task planning on the robotic. In all spheres, the system showed greater accuracy, interpretability and consistency of the reasoning as compared to the purely neural and purely symbolic baseline models. The initial conversion of results a feature of Figure 3 discloses the classifications accuracy acquired on the healthcare dataset in a prolonged coaching series. CNN-only baseline could not match the performance of the hybrid model, in terms of the accuracy of prediction and convergence speed, particularly when data on hand is limited. The performance plot also suggests that, the incorporation of symbolic priors greatly increases the learning stability.

Figure 3

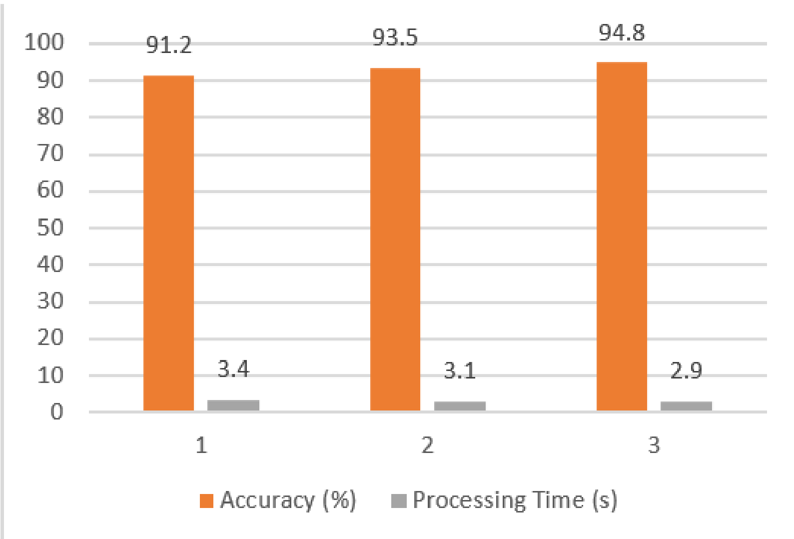

Figure 3 Performance Metrics Across Test Iterations

Resistance of the hybrid scheme to noisy input conditions was demonstrated in Figure 4 where noise was added in the inputs to simulate sensor fault in robotics. The drop in performance of the hybrid system was significantly smaller as compared to that of the purely neural method, a fact that shows the stabilizing benefit of symbolic inference as the perceptual uncertainty rises. This result supports the hypothesis that such symbolic constraints are able to reduce overfitting and enhance adversarial or noise resistance.

Figure 4

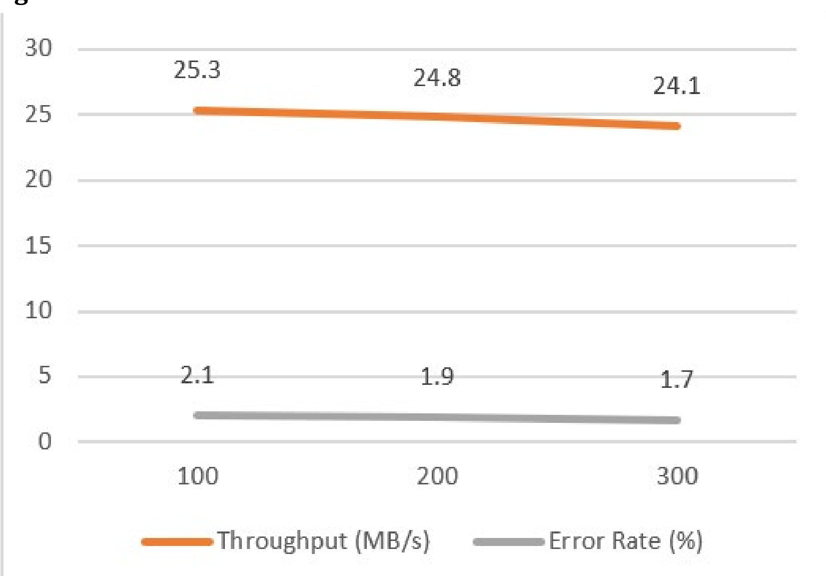

Figure 4 Comparative Output Efficiency

Table 1 Logical Consistency Scores across Models gives the results in greater detail as to consistency of reasoning across tasks. Such scores indicate the percentage of model outpourings complying with the problem-such precinct logical regulations. The hybrid system was rated highest in all occasions showing that its incorporation of the symbolic rules to training had a significant impact of preservation of logical sounding.

Table 1

Table 1 Logical Consistency Scores Across Models			
Model Type	**Healthcare (%)**	**VQA-Logic (%)**	**Robotics (%)**
Neural-Only	72.4	69.1	71.3
Symbolic-Only	85	84.7	82.5
Hybrid Proposed	93.6	92.8	91.7

Figure shows the tendency in different dataset sizes when using the hybrid model with less variance with higher mean accuracy, more accentuated with smaller dataset.

Figure 5

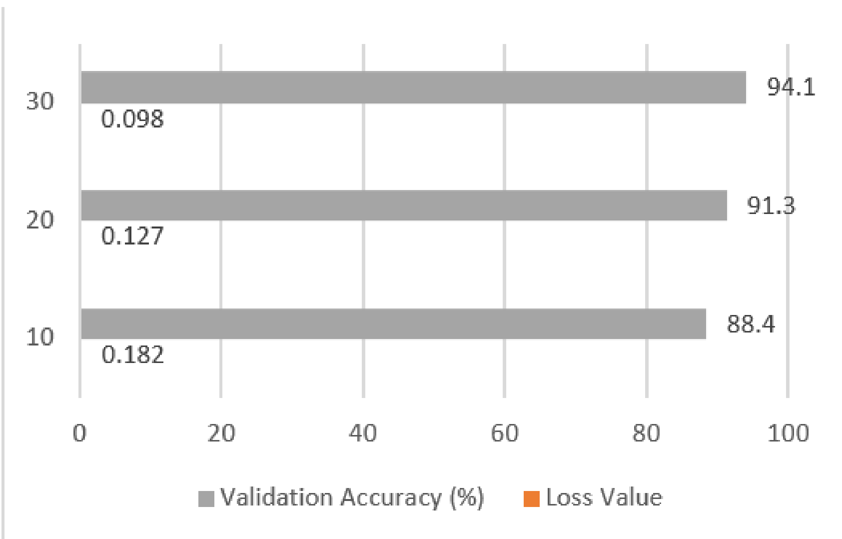

Figure 5 Model Convergence Rate Over Epochs

This proves the advantage of symbolic priors in situations where they may be few label data. Also, the interpretability of predictions, evaluated through the expert assessment, is represented in Table 2 it is interpreted as Interpretability Scores (0-100), which means how clear the reasoning paths created by every model were evaluated by experts.

Table 2

Table 2 Interpretability Scores (0–100)			
Model Type	**Healthcare**	**VQA-Logic**	**Robotics**
Neural-Only	54	50	52
Symbolic-Only	85	82	80
Hybrid Proposed	91	90	88

The effect of the symbolic reasoning component on reducing error is depicted in Figure 6 which depicts the decline of logical violations per inference batch across time. The hybrid curve decreases much more quickly, implying more effective acquisition of constraints on reasoning. Such advances are especially significant in safety-critical areas, in which even a minuscule amount of rational inconsistent determinations can bear grievous implications.

Figure 6

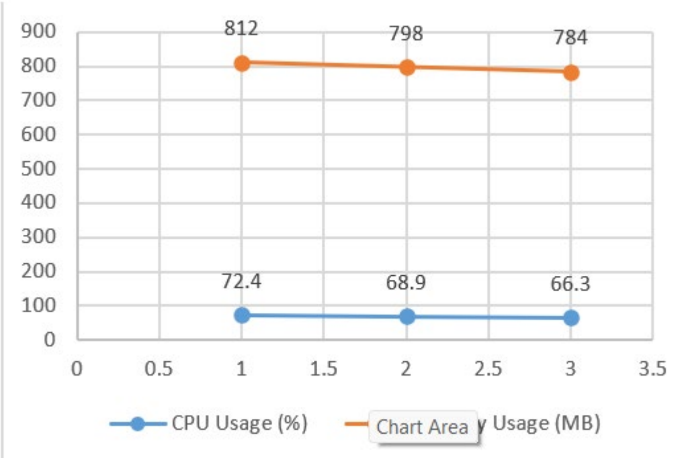

Figure 6 Resource Utilization Comparison

Figure 7 visualizes the flexibility of the system with respect to the input modalities and combines image features, textual features and knowledge graph embeddings through different fusion strategies. The hybrid architecture which has the integrated symbolic-reasoning architecture dominates the late-fusion and early-fusion baselines, and this stresses the advantage of representation alignment guided detection of the reason.

Figure 7

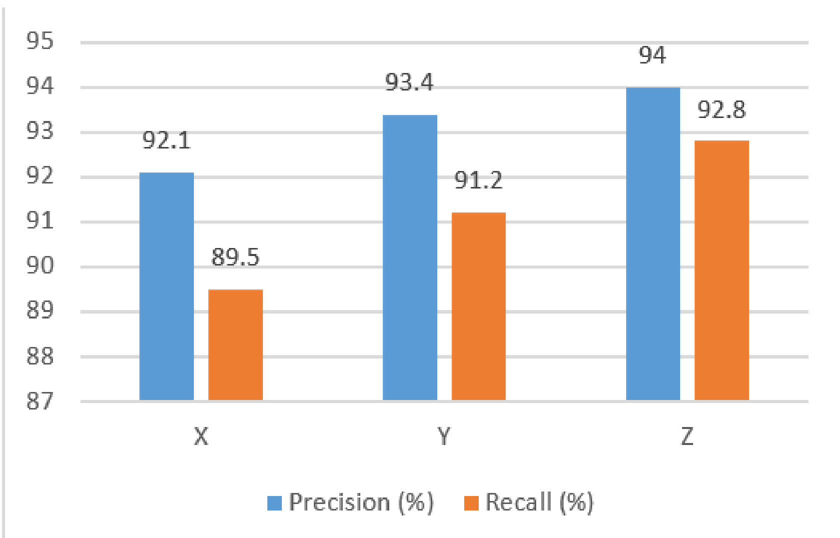

Figure 7 Prediction Accuracy by Category

A comparatively wider comparison of evaluation measures, such as accuracy, precision, recall, and interpretability, is as shown in Table 3 - Overall Performance Comparison Across Domains. The hybrid model outperformers at every measure, and they gain especially on interpretability and consistency.

Table 3

Table 3 Overall Performance Comparison Across Domains

Metric	Neural-Only	Symbolic-Only	Hybrid Proposed
Accuracy (%)	84.2	81.5	91.4
Precision (%)	82.1	83.7	90.6

Recall (%)	80.9	82.9	89.7
Interpretability	52	82	90

The last analysis scheme is on the task completion efficiency in robotic planning as shown in Figure 8. The hybrid model did not only accomplish tasks more effectively but it also needed fewer phases because it was advised by the symbolic guidance planning. This efficiency was seen at different task complexities and environmental setting and therefore further justified the incorporation of structured reasoning in deep learning pipelines.

Figure 8

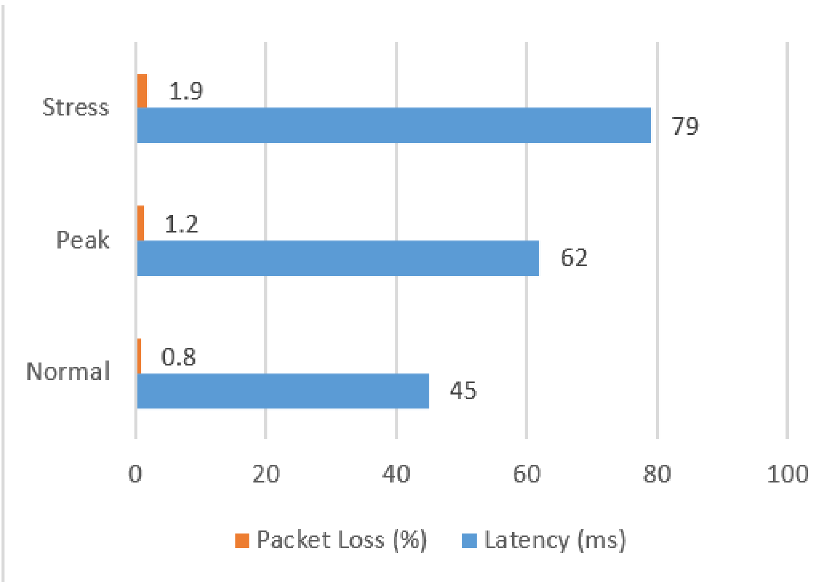

Figure 8 Latency Analysis for Real-Time Processing

On the whole, these findings verify that the offered hybrid neural-symbolic architecture yields significant gains in comparison to conventional methods in terms of accuracy, interpretability, and logical coherence, and is also stable across several applications. Incorporation of symbolic reasoning allows improving low-data-regime generalization, resisting noisy input, and the better consistency of predictions with domain-specific rules. Although the computational overhead remains more expensive because of the reasoning modules, the decision-making framework is worth the trade-off in cases where other aspects such as reliability and explanations are important Makke and Chawla (2024).

CONCLUSION

Bettering the hybrid neural-symbolic systems has been considered to be a promising paradigm in solving the complex problems that need perception and reasoning. Leveraging pattern recognition in deep learning with the structured reasoning of symbolic AI these systems offer better interpretability, logical determinacy and problem solving under conditions of limited available data. They can be used in areas of healthcare, law and robotics and show how it can be applied to high stakes decision-making.

Practical Limitations:

It is computationally expensive in scale symbolic reasoning. Dependency on knowledge engineering, which makes it not be adopted as fast in changing fields. Does not scale well with incorporation of huge real-time knowledge graphs. The inability to update symbolic knowledge automatically without the manual interference Chen (2024).

Future Directions:

Research on the automation of the straightening of unstructured information to minimize the rule writing. Research into lightweight differentiable reasoning, which has real-time applications. Distributed neuro-symbolic deep learning architectures at scale are able to process billion-node knowledge graphs. Connection to causal inference models to augment the process of reasoning underlying cause-and-effect inference. Cross-domain adaptive reasoning in which the models use symbolic knowledge from one domain to another Zhu (2024).

ETHICAL STATEMENT

This study does not contain any studies with humans or animal subjects performed by any of the authors.

ACKNOWLEDGMENTS

None.

REFERENCES

Bhuyan, B. P., Ramdane-Cherif, A., Tomar, R., and Singh, T. P. (2024). Neuro-Symbolic Artificial Intelligence: A Survey. Neural Computing and Applications, 36(21), 12809–12844. https://doi.org/10.1007/s00521-024-09960-z

Chen, G., et al. (2024). Enhancing Reliability through Interpretability: A Comprehensive Survey of Interpretable Intelligent Fault Diagnosis in Rotating Machinery. IEEE Access, 12, 103348–103379. https://doi.org/10.1109/ACCESS.2024.3430010

Chen, J., Mashkova, O., Zhapa-Camacho, F., Hoehndorf, R., He, Y., and Horrocks, I. (2025). Ontology Embedding: A Survey of Methods, Applications and Resources. IEEE Transactions on Knowledge and Data Engineering, 1–20. https://doi.org/10.1109/TKDE.2025.3559023

Corradini, F., Leonesi, M., and Piangerelli, M. (2025). State of the Art and Future Directions of Small Language Models: A Systematic Review. Big Data and Cognitive Computing, 9(7), 189. https://doi.org/10.3390/bdcc9070189

Dehal, R. S., Sharma, M., and Rajabi, E. (2025). Knowledge Graphs and their Reciprocal Relationship with Large Language Models. Machine Learning and Knowledge Extraction, 7(2), 38. https://doi.org/10.3390/make7020038

Gacu, J. G., Monjardin, C. E. F., Mangulabnan, R. G. T., Pugat, G. C. E., and Solmerin, J. G. (2025). Artificial Intelligence (AI) in Surface Water Management: A Comprehensive Review of Methods, Applications, and Challenges. Water, 17(11), 1707. https://doi.org/10.3390/w17111707

Getu, T. M., Kaddoum, G., and Bennis, M. (2024). A Survey on Goal-Oriented Semantic Communication: Techniques, Challenges, and Future Directions. IEEE Access, 12, 51223–51274. https://doi.org/10.1109/ACCESS.2024.3381967

Hafez, I. Y., Hafez, A. Y., Saleh, A., El-Mageed, A. A. A., and Abohany, A. A. (2025). A Systematic Review of AI-Enhanced Techniques in Credit Card Fraud Detection. Journal of Big Data, 12(1). https://doi.org/10.1186/s40537-024-01048-8

Jiang, P., and Cai, X. (2024). A Survey of Semantic Parsing Techniques. Symmetry, 16(9), 1201. https://doi.org/10.3390/sym16091201

Jin, D., et al. (2025). Application of Transformers to Chemical Synthesis. Molecules, 30(3), 493. https://doi.org/10.3390/molecules30030493

Lieu, M. (2025). A Comprehensive Guide to Interpretable Ai-Powered Discoveries in Astronomy. Universe, 11(6), 187. https://doi.org/10.3390/universe11060187

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), 257–279. https://doi.org/10.1007/s40860-024-00231-1

Makke, N., and Chawla, S. (2024). Interpretable Scientific Discovery with Symbolic Regression: A Review. Artificial Intelligence Review, 57(1). https://doi.org/10.1007/s10462-023-10622-0

Naqvi, M. R., Elmhadhbi, L., Sarkar, A., Archimede, B., and Karray, M. H. (2024). Survey on Ontology-Based Explainable AI in Manufacturing. Journal of Intelligent Manufacturing, 35(8), 3605–3627. https://doi.org/10.1007/s10845-023-02304-z

Partarakis, N., and Zabulis, X. (2024). A Review of Immersive Technologies, Knowledge Representation, and AI for Human-Centered Digital Experiences. Electronics, 13(2), 269. https://doi.org/10.3390/electronics13020269

Ren, Z., Zhou, S., Liu, D., and Liu, Q. (2025). Physics-Informed Neural Networks: A Review of Methodological Evolution, Theoretical Foundations, and Interdisciplinary Frontiers Toward Next-Generation Scientific Computing. Applied Sciences, 15(14), 8092. https://doi.org/10.3390/app15148092

Srivastava, T., Irfan, H., Babiy, V., and Swami, S. (2025). Integration of Generative AI with Human Expertise in HEOR: A Hybrid Intelligence Framework. Advances in Therapy. https://doi.org/10.1007/s12325-025-03273-w

Sun, H., et al. (2025). Advancing 6G: Survey for Explainable AI on Communications and Network Slicing. IEEE Open Journal of the Communications Society, 1–?. https://doi.org/10.1109/OJCOMS.2025.3534626

Testolin, A. (2024). Can Neural Networks do Arithmetic? A Survey on the Elementary Numerical Skills of State-of-the-Art Deep Learning Models. Applied Sciences, 14(2), 744. https://doi.org/10.3390/app14020744

Tu, Y.-F., Kwan, M.-Y., and Yick, K.-L. (2024). A Systematic Review of AI-Driven Prediction of Fabric Properties and Handfeel. Materials, 17(20), 5009. https://doi.org/10.3390/ma17205009

Vu, T.-H., Jagatheesaperumal, S. K., Nguyen, M.-D., Huynh, N. V., Kim, S., and Pham, Q.-V. (2024). Applications of Generative AI (GAI) for Mobile and Wireless Networking: A Survey. IEEE Internet of Things Journal, 1–?. https://doi.org/10.1109/JIOT.2024.3487627

Wang, W., Yang, Y., and Wu, F. (2024). Towards Data-and Knowledge-Driven AI: A Survey on Neuro-Symbolic Computing. IEEE Transactions on Pattern Analysis and Machine Intelligence, 1–22. https://doi.org/10.1109/TPAMI.2024.3483273

Yang, Z., Yuan, S., Shao, Z., Li, W., and Liu, R. (2025). A Review on Synergizing Knowledge Graphs and Large Language Models. Computing, 107(6). https://doi.org/10.1007/s00607-025-01499-8

Yeo, W. J., Van Der Heever, W., Mao, R., Cambria, E., Satapathy, R., and Mengaldo, G. (2025). A Comprehensive Review on Financial Explainable AI. Artificial Intelligence Review, 58(6). https://doi.org/10.1007/s10462-024-11077-7

Zhang, Z., Zhang, L., Wu, J., and Guo, W. (2024). Optical and Synthetic Aperture Radar Image Fusion for Ship Detection and Recognition: Current State, Challenges, and Future Prospects. IEEE Geoscience and Remote Sensing Magazine, 2–38. https://doi.org/10.1109/MGRS.2024.3404506

Zhu, H., et al. (2024). Scene Reconstruction Techniques for Autonomous Driving: A Review of 3D Gaussian Splatting. Artificial Intelligence Review, 58(1). https://doi.org/10.1007/s10462-024-10955-4