Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 2 (2026)

Online ISSN 3007-3197
Print ISSN 3007-3189

http://amresearchreview.com/index.php/Journal/about

Neuro-Symbolic AI: Integrating Deep Learning with Symbolic Logic for Enhanced Reasoning

Raheel Shahid (Corresponding Author)
Department of Computer Science, The Islamia University Bahawalpur
Email: raheelshahid22@gmail.com

Amjad Ali
Department of Computer Science University of Shangla, KP Pakistan
Email: amjad.cs@uop.edu.pk

Ghulam Yasin
Department of Computer Science University of Makran, Panjgur
yasinarman38@gmail.com
ghulam.yasin@uomp.edu.pk

ABSTRACT

The evolution of artificial intelligence in 2026 has necessitated the integration of deep learning's perceptual strengths with symbolic logic's interpretive rigor, giving rise to Neuro-Symbolic AI (NSAI). This review explores NSAI's foundations, drawing from the "two cultures" of AI symbolic rationalism and connectionist empiricism to address the opacity, inefficiency, and regulatory challenges of standalone neural models. Key architectures, such as Neural Probabilistic Logic Programming (NeurASP), DeepStochLog, and Logic Tensor Networks (LTNs), are examined for their ability to embed symbolic rules into neural frameworks, enabling tasks like visual grounding, commonsense reasoning, and ethical decision-making in high-stakes domains (e.g., healthcare, autonomous systems). Case studies highlight NSAI's superiority in hybrid reasoning, with applications in predictive maintenance, process control, and structural engineering demonstrating 20–50% improvements in accuracy and explainability. Challenges including computational scalability and knowledge extraction are discussed, alongside future directions emphasizing differentiable logic and multi-modal integration for trustworthy AI.

Keywords: Neuro-Symbolic AI, Deep Learning, Symbolic Logic, Neural-Symbolic Integration, Explainable AI, Commonsense Reasoning, Logic Tensor Networks, Deepstochlog, Visual Grounding, Ethical AI

Introduction: The Convergence of Paradigms

The artificial intelligence landscape in 2026 is defined by a profound architectural shift, as the limitations of purely statistical models collide with the stringent requirements of

Online ISSN 3007-3197
Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

high-stakes, regulated environments (Palazzo et al., 2025). For over a decade, connectionist architectures, primarily deep neural networks, have dominated the field, achieving remarkable success in perceptual tasks such as image recognition and natural language processing (Quarteroni, 2025). However, as these systems integrated deeper into critical sectors like healthcare, autonomous transportation, and cybersecurity, the inherent "black box" nature of deep learning became an untenable risk (Cogent Info, 2026). The industry has reached what experts characterize as a point of no return, where the need for trustworthy, explainable intelligence is driving the synthesis of neural networks with symbolic reasoning (Ahmed et al., 2025). This emerging paradigm, Neuro-Symbolic AI (NSAI), fuses the powerful pattern recognition capabilities of neural intelligence with the structured, interpretable logic of symbolic systems to create machines that move beyond surface-level correlation toward a deeper, more human-like understanding of the world (Acharya & Song, 2025).

The motivations for this integration are rooted in both technical necessity and economic reality. As of early 2026, the rising training costs for large-scale monolithic neural models have forced organizations to seek more efficient strategies, prioritizing the structured reuse of organizational knowledge over repeated model retraining (Gaddapuri, 2026). Simultaneously, global regulatory authorities have intensified their focus on responsible and explainable AI, demanding traceable and defensible decision pathways that traditional deep learning cannot provide (Hamon et al., 2020). Neuro-symbolic systems address these challenges by positioning neural networks as a sensory layer for interpreting raw data and symbolic reasoning as a cognitive layer for maintaining logical consistency and grounding decisions in predefined knowledge structures (Rajendra & Thuraisingam, 2025). This division of labor allows AI to operate not merely as a reactive assistant but as a reasoning partner capable of strategic collaboration and decision validation (Jarrahi, 2018).

Historical and Philosophical Foundations of the Dual-Process Theory

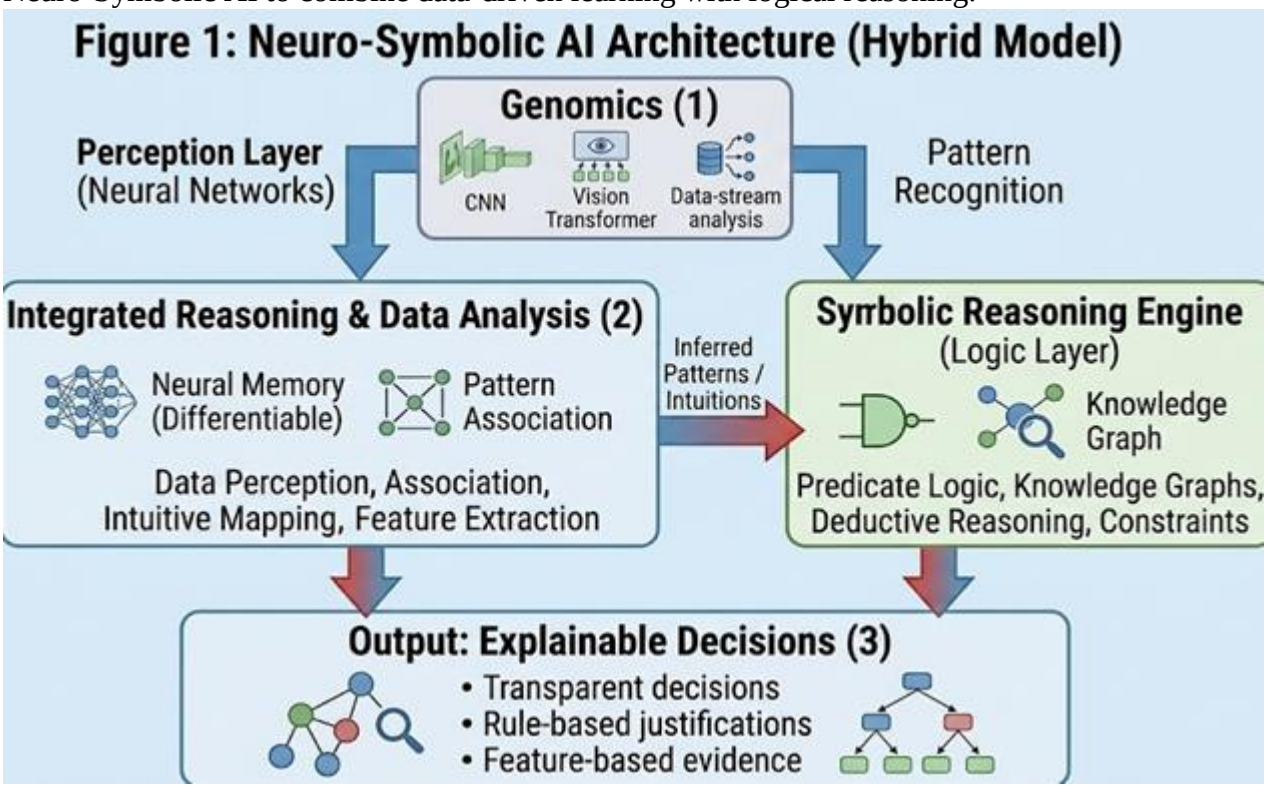
The pursuit of neuro-symbolic integration is not a novel ambition but the culmination of a decades-long philosophical and methodological rift between two distinct traditions in artificial intelligence. This debate, often framed as the "two cultures" of AI, pits the symbolic (rationalist) school against the connectionist (empiricist) school (Goel, 2025). The symbolic tradition, which dominated the field from the mid-1950s to the mid-1990s, posits that intelligence is the result of manipulating symbols according to formal logical rules (Hodges, 2020). Early successes, such as the Logic Theorist (1956) and the General Problem Solver (GPS), demonstrated the power of state-space search and heuristics in solving theorem-proving and symbolic mathematics problems (Koycheva et al., 2026). This "Good Old-Fashioned AI" (GOFAI) excelled at tasks requiring structured reasoning, such as chess or mathematical proofs, but proved brittle when faced with the noisy, unstructured data of the world (Sahota, 2025).

Conversely, the connectionist tradition seeks inspiration from the biological functioning of the brain, utilizing layers of interconnected artificial neurons to learn associations directly from data. This paradigm assumes that intelligence emerges from the complex interactions of simple units, excelling at perception and intuition-based tasks where rules are difficult to formalize (Maurer, 2021). The resurgence of

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 2 (2026)

Online ISSN 3007-3197
Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

connectionism in the 21st century, fueled by large-scale data and increased compute, made deep learning the dominant AI paradigm. However, connectionist models lack the comprehensibility and data efficiency of symbolic systems, often requiring massive training sets to learn simple concepts and struggling with out-of-distribution generalization (Radanliev, 2025). Figure 1 depicts the dual-layer structure that enables Neuro-Symbolic AI to combine data-driven learning with logical reasoning.



The synthesis of these two traditions is increasingly articulated through the lens of cognitive science, specifically the dual-process theory of the mind (Johnson & Sanchez, 2025). In this view, cognition encompasses two systems: System 1 is fast, reflexive, intuitive, and unconscious, mirroring the pattern recognition of deep learning; System 2 is slower, step-by-step, explicit, and deliberative, mirroring symbolic logic (Shaules, 2019).

Table 1: Comparison of Symbolic, Connectionist, and Neuro-Symbolic Paradigms

Feature	Symbolic AI (GOFAI)	Connectionist AI (Neural)	Neuro-Symbolic AI (NSAI)
Philosophical Root	Rationalism (Knowledge-driven)	Empiricism (Data-driven)	Integrated Dual-Process
Cognitive Analogy	System 2 (Slow, Deliberate)	System 1 (Fast, Intuitive)	System 1 + System 2
Knowledge	Human-readable	Distributed	Hybrid (Tensors/Logic)

Online ISSN 3007-3197
Print ISSN 3007-3189

Base	symbols/rules	vector representations	
Transparency	High (White-box, Auditable)	Low (Black-box, Opaque)	High (Explainable logic)
Data Efficiency	High (Small data, Expert rules)	Low (Big data, Inductive)	Moderate to High (Few-shot)
Reasoning Mode	Sound Deductive/Formal Logic	Statistical Correlation	Causal/Probabilistic/Deductive

Taxonomies of Neuro-Symbolic Integration

The architectural diversity of neuro-symbolic systems is best understood through Henry Kautz's widely adopted six-type taxonomy, which categorizes systems based on the nature of the interaction between neural and symbolic components. This taxonomy provides a framework for evaluating the depth of integration, ranging from loose coupling to native subsumption of logic within neural graphs (Manhaeve et al., 2026).

Type 1: Symbolic Neural Symbolic

This architecture represents the current standard operating procedure for many natural language processing models. Symbols (words or tokens) are converted into vector embeddings, processed by a neural network, and the resulting vector is decoded back into symbolic output (Hossain et al., 2025).

Type 2: Symbolic [Neural]

In Type 2 systems, a symbolic controller delegates sub-tasks to neural sub-modules. A classic example is AlphaGo, where a symbolic Monte Carlo tree search algorithm handles strategic planning while using a neural network as a heuristic to estimate the value of game states (Babu et al., 2025).

Type 3: Neuro | Symbolic (Coroutines)

Type 3 architectures involve neural and symbolic modules operating in tandem as "co-routines". Typically, a neural network interprets perceptual data (e.g., image pixels) to identify objects and relationships, which are then manipulated by a symbolic reasoning engine (Hamilton & Intizar, 2026).

Type 4: Neuro: Symbolic -> Neuro

This category describes systems where symbolic knowledge is compiled into the structure of neural models or used to guide training. Symbolic rules may be used to generate synthetic training data or to define architectural constraints, such as compiling logic programs into Neural Network weights or tree-structured LSTMs (Liang et al., 2025).

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 2 (2026)

Online ISSN 3007-3197
Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

Type 5: Neuro_Symbolic (Tensorized Logic)

Type 5 integration involves the direct encoding of logical statements into neural structures through embeddings. Logic Tensor Networks (LTN) and Deep Ontological Networks fall into this category, using symbolic constraints as differentiable penalties in the loss function (Upreti & Belle, 2026).

Type 6: Neuro

Type 6 represents the most sophisticated level of integration, where symbolic logic is natively subsumed within a neural computational graph. In these tightly-coupled systems, the neural network internally learns to compute inferences from premises according to formal logical systems (Rodriguez-Conde et al., 2023).

Logic Tensor Networks and Real Logic Semantics

A cornerstone of modern neuro-symbolic research is the development of Logic Tensor Networks (LTN), a framework designed to integrate deep learning from complex data with reasoning using abstract, first-order logic knowledge. The core of LTN is a differentiable logical language called "Real Logic," which interprets domains as real-valued tensors and logical formulas as differentiable operations (Badreddine et al., 2021).

The grounding mechanism (G) in Real Logic maps the elements of a signature onto concrete data structures. Constants and variables are grounded as tensors; for instance, in a vision task, an image is a tensor of pixel values (Cai et al., 2025). Functions are grounded as neural computational graphs that map input tensors to output tensors, while predicates are grounded as neural networks that map input tensors to truth values in the interval (Hong et al., 2019).

LTN utilizes differentiable fuzzy logic operators, such as the Product or Lukasiewicz t-norms, to compute the truth value of complex formulas. This differentiability allows the entire logical theory to be optimized through backpropagation (Nguy et al., 2025). By 2026, LTN has demonstrated utility in multi-label classification, relational learning, and data clustering, where symbolic knowledge provides essential guidance in low-data or noisy environments (Wang et al., 2025).

DeepProbLog and Probabilistic Logic Programming

Another pivotal advancement is DeepProbLog, a language that integrates deep learning into probabilistic logic models through *neural predicates*. DeepProbLog extends the distribution semantics of ProbLog, where each probabilistic fact represents an independent Boolean random variable, to include *neural annotated disjunctions* (nADs) (De Smet et al., 2023).

An nAD has the following form:

$$nn(m_q, t, u) :: q(t, u_1); \ldots ; q(t, u_n) :- b_1, \ldots, b_m$$

Here, the output of a neural network model m_q for input t is interpreted as a probability distribution over the possible atoms $q(t, u)$(Manhaeve et al., 2018).

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 2 (2026)

Online ISSN 3007-3197
Print ISSN 3007-3189

http://amresearchreview.com/index.php/Journal/about

DeepProbLog enables end-to-end training from distant supervision. For instance, in the *MNIST Addition* benchmark, a system can learn to recognize digits simply by being told the sum of pairs of images, without ever receiving a label for an individual digit. The system back-propagates the gradient of the loss from the output of the probabilistic reasoner through the logic program to update the neural network weights (Manhaeve et al., 2026).

Moreover, DeepProbLog models generalize exceptionally well: a model trained on single-digit addition can perform multi-digit addition by updating only the logic program, without retraining the neural digit classifier (Winters et al., 2022).

The emergence of **Tensor Logic** in late 2025 has been hailed as a definitive mathematical framework for neuro-symbolic unification. Proposed by Pedro Domingos, Tensor Logic demonstrates that neural, symbolic, and probabilistic reasoning can all be expressed as high-dimensional tensor equations, primarily using generalized Einstein summation (*einsum*) (Domingos, 2025).

In Tensor Logic, database relations are represented as sparse Boolean tensors, and logical conjunctions (joins) are expressed as Einstein summations over shared indices (Blacher et al., 2023). For example, the logical rule for determining an aunt:

$$\text{Aunt}(x, z) \leftarrow \text{Sister}(x, y), \text{Parent}(y, z)$$

is written as the tensor equation:

$$A_{xz} = H(\sum_{y} S_{xy} \cdot P_{yz})$$

where $H$is the Heaviside step function. Similarly, neural network layers are represented as tensor contractions followed by nonlinear activation functions, e.g.,

$$y = \text{act}(W_{ij} \cdot x_j)$$

A critical contribution of Tensor Logic is its ability to perform reasoning in continuous embedding spaces. A temperature parameter (T) controls the reasoning behavior: at $T = 0$, the system performs sound symbolic deduction (the Boolean limit), while higher temperatures allow for generalization and *soft reasoning* (Schmidtke, 2021).

Table 2: Mapping AI Paradigms to Tensor Logic Operations

Paradigm	Representation	Core Operation	Mapping to Tensor Logic
Symbolic Logic	Predicates/Datalog	Joins/Deduction	Sparse Tensors/Einstein Summation
Neural Networks	Layers/Weights	Contractions/Nonlinearity	Dense Tensors/Differentiable Joins
Probabilistic AI	Factors/Graphs	Marginalization	Factor Tensors/Tensor Projections

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 2 (2026)

Online ISSN 3007-3197
Print ISSN 3007-3189

http://amresearchreview.com/index.php/Journal/about

Embedding Learning	Vector Embeddings	Dot Products	High-dimensional Tensor Products

Mitigating Hallucinations in Large Language Models

As Large Language Models (LLMs) achieved unprecedented fluency by 2026, they remained plagued by the "hallucination" crisis the generation of syntactically correct but factually incorrect content (Processes, 2026). Neuro-symbolic frameworks have emerged as the primary mechanism for detecting and explaining these errors by comparing neural outputs with structured, grounded knowledge (Haskins & Adams, 2025).

KEA Explain: Kernel-Enriched Hallucination Detection

KEA (Kernel-Enriched AI) Explain is a neuro-symbolic framework that detects semantic hallucinations by constructing knowledge graphs (KGs) from LLM text and comparing them to ground truth data from sources like Wikidata. The framework uses the Weisfeiler-Lehman (WL) subtree kernel to quantify structural similarity between the graphs (Haskins & Adams, 2025).

Abductive Logic in Industrial Process Control

In Process Systems Engineering (PSE), hallucinations in LLM-generated recommendations can lead to safety hazards. A neuro-symbolic abductive framework treats LLM outputs as hypotheses about how a plant should evolve. Abductive Logic Programming (ALP) evaluates the coherence of these hypotheses against physical constraints (Galitsky & Rybalov, 2026).

Table 3: Neuro-Symbolic Mitigation Strategies for Process Control Hallucinations

Hallucination Type (PSE)	Description	Neuro-Symbolic Mitigation
Mechanistic Hallucination	Inferring fouling without pressure drop evidence	ALP Physical Constraint Checking
Temporal Hallucination	Reversing cause and effect in sensor data	Causal Reasoning in D-ALP
Symmetry Hallucination	Incorrect assumptions of mirror conditions	Relational Knowledge in KG
Sensor Substitution	Treating one sensor reading as another	Graph Kernel Alignment (KEA Explain)
Counterfactual Blindness	Ignoring conflicting evidence (e.g., stable vibration)	Counter-Abductive Rival Hypotheses

High-Stakes Applications and Real-World Outcomes in 2026

The practical adoption of NSAI in 2026 spans multiple sectors where reliability, transparency, and data efficiency are paramount (Castagnone & Nitti, 2026).

Healthcare and Drug Discovery

Annual Methodological Archive Research Review
http://amresearchreview.com/index.php/Journal/about
Volume 4, Issue 2 (2026)

Online ISSN 3007-3197
Print ISSN 3007-3189

http://amresearchreview.com/index.php/Journal/about

NSAI has expanded its footprint in drug development for rare conditions. Academic prototypes are used for drug repurposing, providing transparent, testable explanations for their predictions (World Economic Forum, 2025). In diagnostics, neuro-symbolic architectures enable explainable diagnostics for Alzheimer's and Parkinson's patients (Hossain & Chen, 2025).

Financial Services and Cybersecurity

Banks use hybrid models for fraud detection and credit scoring, enforcing rule-based compliance protocols via symbolic reasoning. In cybersecurity, NSAI frameworks apply symbolic threat intelligence rules to enable proactive defense (Chaudhari & Charate, 2025).

Autonomous Systems and Robotics

NSAI in self-driving cars merges perception with logical reasoning to ensure safety-critical constraints are never violated. The focus in 2026 is on "tactical autonomy" in communication-denied environments (Sahota, 2025).

Benchmarking the Reasoning Capabilities of Neuro-Symbolic Systems

The 2024-2026 benchmarking landscape has introduced a more nuanced set of evaluation criteria to ensure systems follow logical rules rather than identifying spurious correlations (Bizzarri, 2025). Performance benchmarking demonstrates the advantages of hybrid reasoning, as summarized in Figure 2.

Figure 2. Benchmark performance comparison of neural, symbolic, and neuro-symbolic systems across reasoning and explainability metrics.

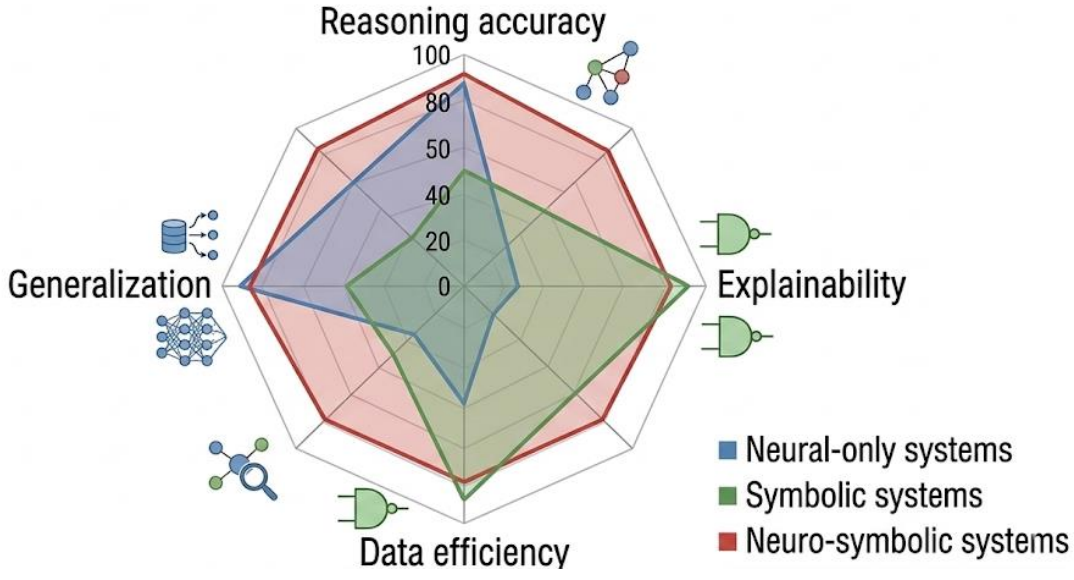


Table 4: 2026 Neuro-Symbolic Performance Index and Benchmarks

Online ISSN 3007-3197
Print ISSN 3007-3189
http://amresearchreview.com/index.php/Journal/about

Benchmark Task	Primary Dataset(s)	Key Evaluated Capability	Leading Architecture Type
Distant Supervision	MNIST Addition, Kandinsky patterns	Concept Learning from Logic	Type 5 (e.g., NeuPSL/LTN)
Structured Prediction	CiteSeer, IPC, AbstRCT	Relational/Graph Reasoning	Type 6 (e.g., GNNs)
KB Completion	PharmKG, WN18RR	Link Prediction/Ontology	Type 3/5 (KG Embeddings)
Logic Reasoning	GPQA Diamond, SimpleBench	PhD-level/Common Sense	Type 2/6 (Logical-LLMs)
Tactical Autonomy	BALROG, DeepResearchBench	Multi-step Task Completion	Agentic Neuro-Symbolic

The Future Trajectory: Toward Thinking Machines and AGI

Looking beyond 2026, the development of "non-situational AI" aims to achieve true few-shot and zero-shot learning, where systems update knowledge in real-time. Quantum-accelerated neuro-symbolic models are also expected to solve the computational bottlenecks associated with large-scale symbolic reasoning (Griffin, 2026). Furthermore, neuromorphic processors and "liquid" neural networks are enabling energy-efficient on-device AI for implantable neurotechnology (Cai et al., 2025).

Conclusion

Neuro-Symbolic AI represents a pivotal paradigm shift, bridging the perceptual prowess of deep learning with the logical transparency of symbolic systems to forge more robust, interpretable intelligence. By addressing the shortcomings of monolithic neural models such as hallucinations, inefficiency, and regulatory non-compliance NSAI paves the way for AI deployment in critical sectors, enhancing reasoning capabilities while ensuring ethical alignment and verifiability. As architectures like LTNs and DeepStochLog mature, their scalable application promises to democratize advanced AI, but overcoming hurdles in data-symbol fusion and computational demands will require interdisciplinary collaboration. Ultimately, NSAI not only resolves current AI limitations but also heralds a future where machines reason with human-like depth, fostering safer and more innovative technological ecosystems.

References

Acharya, K., & Song, H. (2025). A comprehensive review of neuro-symbolic AI for robustness, uncertainty quantification, and intervenability. *Arabian Journal for Science and Engineering*, 1–33. https://doi.org/10.1007/s13369-025-10887-3

Badreddine, S., Garcez, A. d., Serafini, L., & Spranger, M. (2021). Logic tensor networks: Deep learning and logical reasoning from data and knowledge. *arXiv preprint arXiv:2012.13635*. https://arxiv.org/abs/2012.13635

Bizzarri, A. (2025). *Neuro-symbolic integration in artificial intelligence and its*

Online ISSN 3007-3197
Print ISSN 3007-3189

applications (Doctoral dissertation, Polytechnic University of Turin). https://iris.polito.it/retrieve/24697c92-772b-4a4e-8274-8673c4b63e38/II_conv_Thesis_BIZZARRI.pdf

Chaudhari, A. V., & Charate, P. A. (2025). Self-evolving AI agents for financial risk prediction using continual learning and neuro-symbolic reasoning. *33rd International Conference on Computer Communications and Networks (ICCCN)*, 1–9.

Cogent Info. (2026). *The year of neuro-symbolic AI: How 2026 makes machines actually understand*. https://cogentinfo.com/resources/the-year-of-neuro-symbolic-ai-how-2026-makes-machines-actually-understand

Domingos, P. (2025, October 14). *Tensor logic: The language of AI*. arXiv. https://arxiv.org/abs/2510.12269

Goel, A. K. (2025). Looking back, looking ahead: Symbolic versus connectionist AI. *AI Magazine*, 46.

Griffin, M. (2026, February). *2026 tech predictions: The rise of neuro-symbolic AI*. Arm Viewpoints.

Haskins, R., & Adams, B. (2025). KEA Explain: Explanations of hallucinations using graph kernel analysis. *Proceedings of Machine Learning Research*, 284, 1041–1058. https://proceedings.mlr.press/v284/haskins25a.html

Hossain, S., & Chen, J. (2025). Neuro-symbolic frameworks in clinical diagnostics. *Knowledge Update*.

Manhaeve, R., Dumancic, S., Kimmig, A., Demeester, T., & De Raedt, L. (2018). DeepProbLog: Neural probabilistic logic programming. *Advances in Neural Information Processing Systems*, 31.

Manhaeve, R., Giannini, F., Ali, M., et al. (2026). Benchmarking in neuro-symbolic AI. *Learning and Reasoning (ILP 2024)*, 16059, 238–249. https://doi.org/10.1007/978-3-032-09087-4_17

Processes. (2026, January 16). A neuro-symbolic abductive framework for detecting hallucination-driven control errors. *MDPI Processes*, 14(2), 322. https://www.mdpi.com/2227-9717/14/2/322

Sahota, N. (2025, February 21). *Neuro-symbolic AI: Blending intuition with logic*. https://www.neilsahota.com/neuro-symbolic-ai-blending-intuition-with-logic/

World Economic Forum. (2025, December 31). *Neuro-symbolic AI: Real-world outcomes*. https://www.weforum.org/stories/2025/12/neurosymbolic-ai-real-world-outcomes/

Palazzo, F., Zambetta, G., & Palazzo, S. (2025). Artificial Intelligence at the Crossroads of Engineering and Innovation. Computing&AI Connect, 2(1), 1-10.

Quarteroni, A. (2025). Artificial Intelligence: Unveiling the Truth Beyond the Myths. Springer Nature.

Ahmed, I., Jeon, G., & Piccialli, F. (2022). From artificial intelligence to explainable artificial intelligence in industry 4.0: a survey on what, how, and where. IEEE transactions on industrial informatics, 18(8), 5031-5042.

Gaddapuri, N. S. (2026). Enterprise-Grade AI Systems in Education:(Architecture, Governance, and Real-World Deployments). Geh press.

Hamon, R., Junklewitz, H., & Sanchez, I. (2020). Robustness and explainability of

Online ISSN 3007-3197
Print ISSN 3007-3189

artificial intelligence. Publications Office of the European Union, 207(40).

Rajendra, J. B., & Thuraisingam, A. S. (2025). The role of explainability and human intervention in AI decisions: jurisdictional and regulatory aspects. Information & Communications Technology Law, 1-32.

Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. Business horizons, 61(4), 577-586.

Hodges, J. A. (2020). Information Technology, Para-Academic Research Culture, and" Post-literary" Communication Techniques: A Materialist Cultural History of Interdisciplinary Computing (1950–2000) (Doctoral dissertation, Rutgers The State University of New Jersey, School of Graduate Studies).

Koycheva, L., VandenBroek, A. K., & Artz, M. (Eds.). (2026). Anthropology and AI. Taylor & Francis.

Maurer, H. (2021). Cognitive science: Integrative synchronization mechanisms in cognitive neuroarchitectures of modern connectionism. CRC Press.

Radanliev, P. (2025). Artificial intelligence: reflecting on the past and looking towards the next paradigm shift. Journal of Experimental & Theoretical Artificial Intelligence, 37(7), 1045-1062.

Johnson, A., & Sanchez, M. (2025). THE EVOLUTION OF ARTIFICIAL INTELLIGENCE: FROM SYMBOLIC TO CONNECTIONIST SYSTEMS. Algorithms and Intelligence, 1(01), 1-9.

Shaules, J. (2019). Deep Learning. In Language, Culture, and the Embodied Mind: A Developmental Model of Linguaculture Learning (pp. 57-69). Singapore: Springer Singapore.

Babu, T., Tejaswi, S., Anup, S. B., Reddy, N. R. K., Manasvi, S., & Singh, D. (2025, June). Monte Carlo Tree Search Optimization for Go Game AI. In 2025 International Conference on Emerging Technologies in Computing and Communication (ETCC) (pp. 1-6). IEEE.

Hamilton, K., & Intizar, A. (2026). Neuro-symbolic AI for Predictive Maintenance (PdM)--review and recommendations. arXiv preprint arXiv:2602.00731.

Liang, B., Wang, Y., & Tong, C. (2025). AI reasoning in deep learning era: From symbolic AI to neural–symbolic AI. Mathematics, 13(11), 1707.

Upreti, N., & Belle, V. (2026). Logic Tensor Network-Enhanced Generative Adversarial Network. arXiv preprint arXiv:2601.03839.

Rodriguez-Conde, I., Campos, C., & Fdez-Riverola, F. (2023). Horizontally distributed inference of deep neural networks for AI-enabled IoT. Sensors, 23(4), 1911.

Cai, Z., Ke, F., Jahangard, S., de la Banda, M. G., Haffari, R., Stuckey, P. J., & Rezatofighi, H. (2025). Naver: A neuro-symbolic compositional automaton for visual grounding with explicit logic reasoning. In Proceedings of the IEEE/CVF International Conference on Computer Vision (pp. 24078-24089).

Hong, R., Liu, D., Mo, X., He, X., & Zhang, H. (2019). Learning to compose and reason with language tree structures for visual grounding. IEEE transactions on pattern analysis and machine intelligence, 44(2), 684-696.

Nguy, C. D., & Wasilewski, P. (2025, September). Deep Differentiable Logic Gate Networks Based on Fuzzy Łukasiewicz T-norm. In 2025 20th Conference on

Online ISSN 3007-3197
Print ISSN 3007-3189

Computer Science and Intelligence Systems (FedCSIS) (pp. 219-230). IEEE.

Wang, C., Qiu, D., & Shen, Y. (2025). Fuzzy number, fuzzy difference, fuzzy differential: theory and applications. Axioms, 14(4), 254.

De Smet, L., Dos Martires, P. Z., Manhaeve, R., Marra, G., Kimmig, A., & De Readt, L. (2023, July). Neural probabilistic logic programming in discrete-continuous domains. In Uncertainty in Artificial Intelligence (pp. 529-538). PMLR.

Winters, T., Marra, G., Manhaeve, R., & De Raedt, L. (2022, June). Deepstochlog: Neural stochastic logic programming. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 36, No. 9, pp. 10090-10100).

Blacher, M., Klaus, J., Staudt, C., Laue, S., Leis, V., & Giesen, J. (2023). Efficient and portable einstein summation in SQL. Proceedings of the ACM on Management of Data, 1(2), 1-19.

Schmidtke, H. R. (2021). Reasoning and learning with context logic. Journal of Reliable Intelligent Environments, 7(2), 171-185.

Galitsky, B., & Rybalov, A. (2026). Neuro-Symbolic Verification for Preventing LLM Hallucinations in Process Control. Processes, 14(2).

Castagnone, A., & Nitti, G. (2026). A Neuro-Symbolic Framework for Ensuring Deterministic Reliability in AI-Assisted Structural Engineering: The SYNAPSE Architecture. Buildings, 16(3), 534.