applied sciences

MDPI

Editorial

Applications in Neural and Symbolic Artificial Intelligence

Bikram Pratim Bhuyan [1,2,*], Manolo Dulva Hina [2] and Amar Ramdane-Cherif [1]

[1] LISV Laboratory, UVSQ, University of Paris-Saclay, 10–12 Avenue de l'Europe, 78140 Vélizy, France; amar.ramdane-cherif@uvsq.fr
[2] LyRIDS Laboratory, École Centrale d'Électronique (ECE), 10 Rue Sextius Michel, 75015 Paris, France; manolo-dulva.hina@ece.fr
* Correspondence: bpbhuyan@ece.fr

1. Introduction

The past decade has witnessed the remarkable ascent of neural network-based artificial intelligence, and deep learning in particular, as a transformative force across science, engineering, and society (with Generative AI becoming a household name). Architectures such as convolutional neural networks, recurrent networks, transformer models, and graph neural networks have achieved unprecedented performance in image recognition, natural language processing, autonomous systems, and scientific discovery [1,2]. Their ability to learn complex, high-dimensional representations directly from raw data has made them the dominant paradigm in modern AI. Yet, despite these remarkable capabilities, purely data-driven neural approaches carry limitations that increasingly constrain their deployment in high-stakes, safety-critical, trust-worthy, and resource-scarce environments. Neural networks are notoriously opaque, as their internal representations resist interpretation, making it difficult for practitioners to understand, validate, or trust their decisions. They are data-hungry, requiring large annotated datasets that are costly or impossible to assemble in specialised domains such as rare disease management or structural mechanics. They lack explicit mechanisms for logical reasoning, causal inference, or the direct encoding of structured domain knowledge, leading to fragile generalisation when encountering out-of-distribution inputs [3]. These limitations reflect a fundamental architectural tension between the statistical nature of neural learning and the structured, knowledge-rich character of many real-world problems.

Symbolic artificial intelligence, whose roots trace back to the foundational decades of the discipline, offers a complementary set of strengths [4]. By representing knowledge through explicit symbols, formal rules, ontologies, and logical inference mechanisms, symbolic systems excel precisely where neural approaches struggle. They are transparent by design, their reasoning can be inspected and audited step by step, they can encode expert knowledge compactly without requiring large training datasets, and they support formal verification of correctness and consistency [5]. Expert systems, fuzzy logic controllers, knowledge graphs, constraint satisfaction solvers, and rule-based decision engines have delivered decades of value in medical diagnosis, industrial control, legal reasoning, and scientific simulation. However, symbolic AI in turn faces its own fundamental limitations. It is brittle in the presence of noise and ambiguity, requires substantial manual effort to construct and maintain knowledge bases, struggles to process raw perceptual data such as images or sensor streams, and lacks the adaptability to learn from experience and update representations autonomously. The brittleness and scalability constraints of purely symbolic systems became increasingly apparent as applications moved from narrow,

Check for updates

Academic Editor: Andrea Prati

Received: 19 March 2026
Accepted: 25 March 2026
Published: 27 March 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

well-defined expert domains into the open, dynamic, and data-rich environments that characterise contemporary technology.

The recognition that neither paradigm alone is sufficient has motivated a growing and increasingly vigorous research programme aimed at their principled integration, variously described as neuro-symbolic AI, hybrid AI, or knowledge-guided machine learning [6,7]. The central insight of this paradigm is that neural and symbolic approaches are not competitors but complementary tools with respective strengths and weaknesses. Their integration in a common architecture also promises systems that are simultaneously capable, interpretable, and robust. The relevance of this integration extends far beyond computer science. As AI is deployed in domains where decisions carry direct consequences for human welfare, like healthcare, public safety, environmental management, critical infrastructure, and education, the ability to combine the learning power of neural networks with the transparency and knowledge-encoding capacity of symbolic reasoning becomes not a technical preference but a practical necessity. These requirements collectively define the scope and motivation of the present Special Issue.

This Special Issue of *Applied Sciences*, entitled "Applications in Neural and Symbolic Artificial Intelligence", was launched to gather original research at precisely this frontier: the design, development, and evaluation of AI systems that leverage the integration of, or productive interplay between, neural and symbolic approaches across diverse scientific and technological domains. The response to the call demonstrated the breadth and vitality of the field. Articles were accepted for publication following rigorous peer review, spanning applications in healthcare, environmental hazard assessment, structural mechanics, energy management, cybersecurity, organisational decision-making, and computer science education, as well as contributions to the mathematical foundations of higher-order network learning. Despite their disciplinary diversity, the papers together demonstrate that the principled combination of data-driven neural learning and structured symbolic reasoning yields AI systems that are more accurate, more interpretable, and more robust than either paradigm achieves in isolation. The following section provides a structured overview of each contribution, organised by thematic cluster, before the concluding section draws cross-cutting insights from the collection as a whole.

2. Overview of Contributions

This Special Issue compiles original research articles that collectively advance the theory and practice of neural and symbolic artificial intelligence across a wide range of scientific and technological domains. Rather than presenting these contributions as an undifferentiated sequence, we organise them into six thematic clusters that reflect the natural affinities among the works and allow the cross-cutting themes of the collection to emerge clearly. The clusters progress from theoretical foundations to applied systems, reflecting the spectrum from mathematical abstraction to real-world deployment that characterises the current state of neural and symbolic research.

2.1. Theoretical Foundations in Neural and Symbolic Learning

Any principled integration of neural and symbolic AI must rest on sound mathematical foundations. Moh Ousellam et al. [8] proposed the Directed Higher-Ordered Neural Network (HONN) framework for learning on directed hypergraphs, which formalises within a unified spectral framework the integration of directional symbolic knowledge structures with neural feature propagation. The authors introduce three spectral Laplacian formulations with a tunable parameter that continuously interpolates between local identity preservation and global structural diffusion. Extensive experiments on five benchmark datasets demonstrate that HONN consistently matches or surpasses the state of the art.

2.2. Physics-Informed and Engineering Applications

Two contributions exemplify complementary strategies in engineering and computational science, where physical laws (which can be expressed as differential equations, constitutive relations, or conservation principles) constitute a rich source of symbolic knowledge that can constrain, regularise, and guide neural learning.

Mouratidou and Stavroulakis [9] address one of the most computationally demanding challenges in multiscale structural analysis, namely the need to evaluate nonlinear constitutive relations at every integration point of a macroscale finite element model. The authors propose replacing this computationally prohibitive procedure with a neural network constitutive metamodel combined with a physics-informed variant. Applied to masonry composite structures and hyperelastic materials including polymers and rubbers, the proposed framework achieved better results.

Dimitrova-Angelova et al. [10] conducted a rigorous experimental comparison between a physics-derived symbolic model and three neural network architectures for predicting photovoltaic power generation in the context of digital twin development for smart buildings. The results show an important practical insight that symbolic models embed genuine and useful domain knowledge, but neural models are superior when the system's behaviour becomes nonlinear and variable. The study provides empirically grounded guidance for practitioners designing digital twins for renewable energy systems.

2.3. Healthcare Decision Support and Clinical Intelligence

The healthcare domain presents some of the most demanding requirements for AI, as decisions must be simultaneously accurate, interpretable, auditable, and aligned with established clinical guidelines. Two contributions in this Special Issue develop sophisticated systems for paediatric and remote patient management that achieve this balance through complementary hybrid architectures.

Benabderrahmane et al. [11] present MedGuard-FL, a context-aware federated learning framework for secure remote patient monitoring that addresses the need to protect sensitive patient data through privacy-preserving learning while simultaneously ensuring model accuracy and resilience against adversarial attacks. Comparative analysis establishes MedGuard-FL as capable of unifying patient-aware symbolic adaptation with comprehensive, multi-layer adversarial defences to enhance both the security and the clinical responsiveness of neural federated systems.

Zendaoui et al. [12] present NS-Assist, a hybrid web–mobile decision support system for the management of a rare chronic kidney disorder in children. NS-Assist realises neuro-symbolic interaction through two complementary reasoning modules. The first is a rule-based expert system encoding clinically validated 'IF–THEN' protocols, and the second is a fuzzy inference engine that handles the inherent clinical uncertainty. Together, these symbolic and fuzzy components provide real-time, interpretable recommendations to both clinicians through a web interface and caregivers through a mobile application.

2.4. Environmental Intelligence and Geospatial Risk Assessment

The management of environmental hazards provides a domain where neural and symbolic AI must reconcile large-scale spatial data, domain-expert knowledge about causal mechanisms, and the demands of operational decision-makers for both predictive accuracy and actionable interpretability. Bouzeraa et al. [13] address this challenge through a GIS-integrated machine learning framework for wildfire susceptibility mapping. Symbolic domain knowledge is embedded at multiple stages: in the expert-guided selection and justification of conditioning factors, in the GIS-based encoding of spatial relationships, in a k-means clustering step that removes label noise from the low-susceptibility class by detecting

instances with feature profiles resembling high-risk zones, and in the post hoc application of SHAP (SHapley Additive exPlanations) to interpret the predictions of the best-performing model. Four machine learning classifiers are trained and compared. The result is a GIS-ready, operationally usable risk map that local authorities and firefighting services can directly employ for resource pre-positioning, patrol routing, and land-use planning.

2.5. Security, Trust, and Governance

As AI systems are deployed in adversarial, high-stakes, and socially consequential settings, the demands for trustworthiness, accountability, and transparency intensify. Two contributions in this Special Issue address these demands in the complementary contexts of network security and organisational decision-making.

Almohaimeed and Albalwy [14] address the challenge of securing resource-constrained Internet of Things (IoT) networks against increasingly sophisticated intrusion attempts. The work demonstrates that the principled combination of symbolic feature selection and neural classification delivers measurable gains in both detection accuracy and computational efficiency, two properties that are jointly essential for operational IoT security.

Manolache and Popescu [15] approach the interface from the domain of organisational governance, presenting the design and real-world evaluation of a collaborative decision-making platform that integrates the generative reasoning capabilities of GPT-4o with the accountability mechanisms of blockchain technology.

2.6. Artificial Intelligence Education

The long-term vitality of neural and symbolic AI depends not only on the development of new systems but also on the cultivation of practitioners who understand the paradigms. Jeon et al. [16] propose a four-stage pedagogical framework for computer vision education that deliberately integrates neural network-based AI and symbolic AI at each learning level, grounded in Bloom's Taxonomy of Educational Objectives. At each stage, symbolic AI is incorporated alongside neural learning in rule-based criteria for classification, logical sequencing and annotation for captioning, predefined algorithmic criteria for segmentation, and explicit decision logic for action recognition. The framework was validated through two rounds of expert validity assessments. A pilot study applied the framework, achieving statistically significant improvements across all five survey domains, including scientific knowledge, understanding of AI in inquiry-based activities, and interest in computer vision education.

3. Conclusions

This Special Issue set out to showcase the breadth and vitality of research in neural and symbolic artificial intelligence. The contributions collected here fulfil that ambition, spanning domains as diverse as structural mechanics, paediatric telemedicine, wildfire risk assessment, IoT cybersecurity, organisational governance, energy engineering, and computer science education. In summary, this Special Issue demonstrates that neural and symbolic artificial intelligence is a broad and productive scientific agenda with immediate relevance across multiple domains of applied science and technology. The Guest Editors thank all authors for their rigorous and creative contributions, all reviewers for their careful and constructive evaluations, and the editorial team of *Applied Sciences* for their support throughout the publication process.

Author Contributions: Conceptualization, B.P.B., M.D.H. and A.R.-C.; methodology, B.P.B., M.D.H. and A.R.-C.; software, B.P.B., M.D.H. and A.R.-C.; validation, B.P.B., M.D.H. and A.R.-C.; formal analysis, B.P.B., M.D.H. and A.R.-C.; investigation, B.P.B., M.D.H. and A.R.-C.; resources, B.P.B., M.D.H. and A.R.-C.; data curation, B.P.B., M.D.H. and A.R.-C.; writing—original draft preparation,

B.P.B., M.D.H. and A.R.-C.; writing—review and editing, B.P.B., M.D.H. and A.R.-C.; visualization, B.P.B., M.D.H. and A.R.-C.; supervision, B.P.B., M.D.H. and A.R.-C.; project administration, B.P.B., M.D.H. and A.R.-C.; funding acquisition, B.P.B., M.D.H. and A.R.-C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [CrossRef] [PubMed]
2. Goodfellow, I.; Bengio, Y.; Courville, A.; Bengio, Y. *Deep Learning*; MIT Press: Cambridge, UK, 2016; Volume 1.
3. Marcus, G.; Davis, E. *Rebooting AI: Building Artificial Intelligence We Can Trust*; Vintage: New York, NY, USA, 2019.
4. Russell, S. Rationality and intelligence: A brief update. In *Fundamental Issues of Artificial Intelligence*; Springer: Berlin/Heidelberg, Germany, 2016; pp. 7–28.
5. Brachman, R.; Levesque, H. *Knowledge Representation and Reasoning*; Elsevier: Amsterdam, The Netherlands, 2004.
6. Bhuyan, B.P.; Ramdane-Cherif, A.; Singh, T.P.; Tomar, R. *Neuro-Symbolic Artificial Intelligence*; Studies in Computational Intelligence; Springer: Singapore, 2024.
7. Bhuyan, B.P.; Ramdane-Cherif, A.; Tomar, R.; Singh, T. Neuro-symbolic artificial intelligence: A survey. *Neural Comput. Appl.* **2024**, *36*, 12809–12844. [CrossRef]
8. Moh Ousellam, Y.; Bhuyan, B.P.; Fissoune, R.; Ivanova, G.; Ramdane-Cherif, A. Learning Directed Knowledge Using Higher-Ordered Neural Networks: Building a Predictive Framework. *Appl. Sci.* **2025**, *15*, 11085. [CrossRef]
9. Mouratidou, A.D.; Stavroulakis, G.E. A neural network constitutive model and automatic stiffness evaluation for multiscale finite elements. *Appl. Sci.* **2025**, *15*, 3697. [CrossRef]
10. Dimitrova-Angelova, D.; González-González, J.F.; Carmona-Fernández, D.; Jaramillo-Morán, M.A. Photovoltaic Digital Twins: Mathematical Modeling vs. Neural Networks for Energy Management in Smart Buildings. *Appl. Sci.* **2025**, *15*, 8883. [CrossRef]
11. Benabderrahmane, F.; Kerkouche, E.; Bouchemal, N. Risk-Aware Privacy-Preserving Federated Learning for Remote Patient Monitoring: A Multi-Layer Adaptive Security Framework. *Appl. Sci.* **2025**, *16*, 29. [CrossRef]
12. Zendaoui, N.; Bouchemal, N.; Bouchemal, N.; Boussebough, I.; Ivanova, G. NS-Assist: Nephrotic Syndrome Assistance System for Pediatric Decision-Making in Pandemic Situations. *Appl. Sci.* **2025**, *15*, 11433. [CrossRef]
13. Bouzeraa, Y.; Bouchemal, N.; Djaaboub, S.; Hristov, G.; Zahariev, P. Machine Learning-Based Wildfire Susceptibility Mapping: A GIS-Integrated Predictive Framework. *Appl. Sci.* **2025**, *15*, 12188. [CrossRef]
14. Almohaimeed, M.; Albalwy, F. Enhancing IoT network security using feature selection for intrusion detection systems. *Appl. Sci.* **2024**, *14*, 11966. [CrossRef]
15. Manolache, S.; Popescu, N. The development of an OpenAI-based solution for decision-making. *Appl. Sci.* **2025**, *15*, 3408. [CrossRef]
16. Jeon, I.S.; Kang, S.J.; Kang, S.J. A staged framework for computer vision education: Integrating AI, data science, and computational thinking. *Appl. Sci.* **2024**, *14*, 9792. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.