

# NEXUS: A Neural-Symbolic Architecture for Evaluating Heart Disease Domain Prediction

H. P. Alesso  
CTO, AI HIVE, Pleasanton, USA  
`info@ai-hive.net`

Founders Circle, MIT, Cambridge, USA  
`alesso@alum.mit.edu`

March 23, 2025

## Abstract

This paper presents NEXUS (Neural-symbolic EXtensible Unified System), a novel architecture that integrates neural networks with symbolic reasoning to create AI systems with both high performance and transparent decision-making processes. We evaluate NEXUS on the UCI Heart Disease dataset, demonstrating its ability to effectively combine pattern recognition with explicit reasoning for medical diagnosis. Our latest results show that the neural component achieves 80.33% accuracy, while the symbolic component achieves 50.82% accuracy. The integrated NEXUS approach achieves 81.97% accuracy, demonstrating a 1.64% improvement over the best individual component. We discuss the implications of these findings for explainable AI in clinical decision support and outline directions for future research in neural-symbolic integration for medical applications and beyond, contextualizing our work within the broader field of neural-symbolic AI research.

**Keywords:** neural-symbolic integration, explainable AI, medical diagnosis, heart disease prediction, transformer models

## 1 Introduction

The development of artificial intelligence systems for high-stakes applications like medical diagnosis requires both high accuracy and interpretability. While deep learning models have achieved remarkable performance in various domains, they typically operate as "black boxes" with limited transparency in their decision-making processes. This lack of interpretability poses challenges for adoption in critical areas where understanding the reasoning behind a prediction is crucial for user trust and safety [1, 2].

Conversely, symbolic AI approaches offer transparent reasoning through explicit rules and knowledge structures but often lack the pattern recognition capabilities needed for complex data [3]. Neural-symbolic integration has emerged as a promising approach to combine the strengths of both paradigms [4, 5].

In this work, we introduce and evaluate NEXUS (Neural-symbolic EXtensible Unified System), a novel architecture that implements bidirectional integration between a transformer-based neural network and a knowledge graph-based symbolic reasoner. This integration is governed by a metacognitive control system that dynamically determines which approach to prioritize based on confidence and task characteristics.

Recent work by Bengio [8] has emphasized the necessity for deep learning to evolve from “System 1” thinking (intuitive, fast, unconscious cognitive processes) to “System 2” thinking (logical, deliberate, conscious cognitive processes), highlighting the importance of combining neural and symbolic approaches. Our work directly addresses this challenge through a practical implementation focused on medical diagnosis.

Our contributions include:

- A detailed description of the NEXUS architecture for neural-symbolic integration.
- An application of this architecture to heart disease prediction using the UCI Heart Disease dataset.
- Updated experimental results showing improved performance through neural-symbolic integration.
- A comprehensive analysis of the complementary strengths of neural and symbolic approaches in medical diagnosis.
- A discussion of the implications for explainable AI in clinical decision support and other domains.
- Contextualizing our work within the broader neural-symbolic AI research landscape.

## 2 Related Work

### 2.1 Neural-Symbolic Integration

Neural-symbolic integration has a rich history in AI research. Garcez et al. [5] categorize neural-symbolic integration methods into three primary approaches: (1) *Symbol for Neural*, incorporating symbolic knowledge into neural systems; (2) *Neural for Symbol*, enhancing symbolic reasoning with neural networks; and (3) *Hybrid Integration*, creating systems where neural and symbolic components operate in tandem.

Recent surveys of the field have shown that the majority of neural-symbolic research focuses on learning and inference (63%), logic and reasoning (35%), and knowledge representation (44%), with fewer efforts on explainability (28%) and meta-cognition (5%) [22]. Our work addresses these gaps by focusing on both explainability and metacognition.

Recent advances in neural-symbolic integration include DeepProbLog [6], which integrates neural networks with probabilistic logic programming, and the Neuro-Symbolic Concept Learner [7], which combines perception with symbolic program synthesis. These approaches have shown promise in domains requiring both pattern recognition and logical reasoning.

In the corporate sector, major technology companies have recognized the importance of neural-symbolic approaches. IBM Research explicitly states that “Neuro-symbolic AI is a pathway to achieve artificial general intelligence” [23], while Google DeepMind has demonstrated success with AlphaGeometry, a system that combines neural language models with symbolic deduction engines for mathematical reasoning [24].

### 2.2 AI for Heart Disease Prediction

Heart disease prediction has been a focus of AI research due to its clinical importance and the availability of structured medical data. Traditional machine learning approaches have achieved

reasonable success in this domain, with support vector machines, random forests, and neural networks being common choices [9, 10].

Deep learning approaches for heart disease prediction have shown improved performance in recent years [11, 12]. However, these approaches often lack interpretability, which is crucial for clinical adoption. Some work has focused on incorporating medical knowledge into deep learning models [13], but a comprehensive neural-symbolic approach for heart disease prediction remains underexplored.

### 2.3 Explainable AI in Healthcare

The importance of explainability in healthcare AI systems has been emphasized by numerous researchers [14, 2]. Traditional explainable AI approaches like LIME [15] and SHAP [16] provide post-hoc explanations of model predictions but often fail to capture the full complexity of model decision-making.

Neural-symbolic approaches offer an alternative path to explainability by integrating interpretable symbolic reasoning directly into the system architecture [17]. Rather than explaining an opaque process after the fact, these systems incorporate transparency into their design, which is particularly valuable in healthcare applications where understanding the reasoning behind a diagnosis is essential for physician trust and patient safety.

## 3 The NEXUS Architecture

### 3.1 Overview

The NEXUS architecture consists of six main components:

1. **Neural Foundation:** A transformer-based model for pattern recognition and feature extraction.
2. **Symbolic Component:** A knowledge graph and reasoning engine for explicit logical reasoning.
3. **Neural-Symbolic Interface:** Bidirectional mapping between neural and symbolic representations.
4. **Metacognitive Controller:** Dynamic decision-making about which reasoning approach to use.
5. **Self-Improvement Module:** Learning mechanisms for both neural and symbolic components.
6. **Explanation Generator:** Human-understandable explanations of the reasoning process.

This architecture aligns with recent research findings that suggest architectures where neural and symbolic modules iteratively cooperate (Neural  $\rightarrow$  Symbolic  $\leftarrow$  Neural) tend to perform best across metrics [25].

### 3.2 Neural Component

The neural component is based on an efficient transformer architecture, which has shown strong performance across various domains [18]. The model includes:

- Multi-head attention with knowledge integration.
- Position-wise feed-forward networks.
- Layer normalization and residual connections.
- Symbolic constraint layers.

Our implementation uses the `ScalableTransformerModel` class, which provides memory optimization features such as:

- Flash attention for efficient computation.
- Gradient checkpointing to reduce memory usage.
- Mixed precision training support.

For medical applications, we extended this with a specialized Medical Transformer Model that incorporates domain-specific processing:

- Medical preprocessing of input features.
- Risk factor amplification.
- Clinical threshold application.

The neural component contains approximately 400,000 parameters, with 3 transformer layers, 8 attention heads, and an embedding dimension of 128. While modest in size compared to large language models with billions of parameters, this architecture is computationally efficient while still providing strong performance on specialized medical tasks.

### 3.3 Symbolic Component

The symbolic component consists of a knowledge graph with:

- Entities representing symptoms, diseases, and risk factors.
- Weighted relationships between entities.
- Explicit medical rules for diagnosis.
- Hierarchical relationships between concepts.
- Confidence scores for relationships and rules.

The reasoning engine applies both forward and backward chaining algorithms to derive conclusions from the knowledge graph, maintaining uncertainty estimates throughout the reasoning process. For heart disease prediction, we populated the knowledge graph with:

- Age and gender-related risk factors.
- Chest pain type indicators.
- Blood pressure and cholesterol relationships.
- Relations between ECG findings and heart disease.
- Rules combining multiple factors (e.g., age, diabetes, and exercise angina).

This approach to knowledge representation aligns with current research on neural-symbolic systems, where explicit symbolic structures enhance interpretability and allow for the incorporation of domain expertise [19].

### 3.4 Neural-Symbolic Interface

The interface facilitates bidirectional information flow between neural and symbolic components:

- **Symbol Grounding (Neural  $\rightarrow$  Symbolic)**: Maps neural representations to symbolic concepts using adaptive thresholds.
- **Knowledge Injection (Symbolic  $\rightarrow$  Neural)**: Embeds symbolic knowledge into neural processing through attention mechanisms.

The interface maintains a mapping between neural hidden states and symbolic entities, allowing for translation between the representations with confidence estimates. This bidirectional flow is critical for true neural-symbolic integration, as highlighted in recent research [20].

### 3.5 Metacognitive Control

The metacognitive controller dynamically selects between neural, symbolic, or hybrid strategies based on:

- Confidence of each component's prediction.
- Presence of risk factors in the patient.
- Task characteristics (diagnosis vs. treatment recommendation).
- Learned effectiveness of each strategy.

For medical applications, we enhanced the controller with:

- Risk level awareness (low, medium, high).
- Critical symptom recognition.
- Clinical guideline integration.

This metacognitive capability addresses a significant gap in current neural-symbolic research, as only 5% of systems incorporate such reflection mechanisms [22].

### 3.6 Self-Improvement Module

While our current implementation focuses on the neural learning component, the architecture supports:

- Neural learning through gradient-based optimization.
- Symbolic knowledge acquisition through rule mining.
- Coherence optimization to ensure consistency between components.

### 3.7 Explanation Generator

The explanation generator produces human-understandable explanations tailored to different users:

- Multi-level explanations (simple, medium, detailed).
- Reasoning step traces from both neural and symbolic components.
- Confidence assessment and strategy justification.
- Medical-specific explanations including risk factors and critical symptoms.

This focus on detailed, multi-level explanations aligns with IBM’s vision of neural-symbolic systems that provide transparent reasoning paths that can be validated by domain experts [21].

## 4 Experimental Setup

### 4.1 Dataset

We evaluated NEXUS on the Heart Disease UCI dataset [11], which contains data for 303 patients with 13 features and a binary target variable indicating the presence of heart disease. The features include:

- Age
- Sex (1 = male, 0 = female)
- Chest pain type (4 values)
- Resting blood pressure
- Serum cholesterol
- Fasting blood sugar
- Resting ECG results
- Maximum heart rate achieved
- Exercise-induced angina
- ST depression induced by exercise
- Slope of the peak exercise ST segment
- Number of major vessels colored by fluoroscopy
- Thalassemia

The target variable was binarized, with values greater than 0 mapped to 1 (presence of heart disease) and 0 indicating absence of heart disease.

## 4.2 Implementation Details

We implemented NEXUS using PyTorch. The neural component consisted of a transformer with 3 layers, 8 attention heads, and an embedding dimension of 128. The symbolic component was initialized with entities for each feature and class, with initial relationships based on feature importance analysis and medical domain knowledge.

The dataset was split into training (80%) and testing (20%) sets, stratified by the target variable. We trained the neural component for 15 epochs using the AdamW optimizer with a learning rate of 0.001, weight decay of  $1e-5$ , and a batch size of 32. We used cosine learning rate scheduling to improve convergence.

For the knowledge graph, we encoded 13 primary heart disease indicators along with higher-level concepts like “Cardiovascular Risk Factor” and “Chest Pain” to create a hierarchical structure. We implemented 5 key rule combinations for heart disease diagnosis based on medical literature.

The metacognitive controller was initialized with a neural threshold of 0.85 and a symbolic threshold of 0.75, with a learning mechanism to adjust these thresholds based on prediction results during evaluation.

We have also developed a complete integration framework and troubleshooting tools to ensure robustness and ease of use. See <https://github.com/alessoh/Neural-Symbolic-Superintelligence>.

## 5 Results and Analysis

### 5.1 Overall Performance

The performance of the different components on the test set is shown in Table 1.

Table 1: Performance on Heart Disease Prediction

Component	Accuracy
Neural	80.33%
Symbolic	50.82%
NEXUS	81.97%

The neural component achieved 80.33% accuracy, significantly outperforming the symbolic component (50.82%). The integrated NEXUS approach achieved 81.97% accuracy, representing a 1.64% improvement over the neural component alone. This demonstrates that the neural-symbolic integration can enhance performance beyond what either component can achieve individually.

These results align with the findings from other neural-symbolic systems where the integrated approach typically outperforms individual components [26].

### 5.2 Training Progress

The neural model showed consistent improvement over the training epochs, starting with around 55% accuracy on the first epoch and gradually improving to about 78.5% accuracy by the end of training. The training process was computationally efficient, completing in just 0.04 minutes, which indicates the model’s practical applicability even on modest hardware.

### 5.3 Metacognitive Strategy Analysis

The metacognitive controller demonstrated effective strategy selection, choosing different approaches based on case characteristics:

- For cases with clear patterns, it favored the neural component.
- For cases with strong symbolic indicators, it prioritized symbolic reasoning.
- For borderline cases, it employed a hybrid approach, weighted based on component confidence.

In our example diagnosis case, the controller chose a hybrid strategy with a 43% neural weight and 57% symbolic weight, demonstrating its ability to balance the strengths of both components.

### 5.4 Patterns Discovered

#### 5.4.1 Neural Component Patterns

The neural transformer component identified several key relationships:

- **Age and Sex Correlations:** Men, particularly older men, showed higher heart disease risk.
- **Chest Pain Type Significance:** Asymptomatic chest pain (type 4) was associated with higher disease risk than typical angina, indicating that absence of pain doesn't mean absence of disease.
- **ST Depression Patterns:** The “oldpeak” feature (ST depression induced by exercise) was identified as highly predictive—larger depressions correlate strongly with the presence of heart disease.
- **Vessel Blockage Importance:** The number of colored major vessels (CA) was among the most predictive features, with more blocked vessels strongly indicating heart disease.

#### 5.4.2 Symbolic Component Patterns

Though less accurate overall, the symbolic component identified interpretable logical rules:

- **Thalassemia Associations:** Reversible defect thalassemia (value 7) was linked to higher disease likelihood through explicit reasoning rules.
- **Risk Factor Combinations:** The symbolic component identified threshold-based rules combining factors like fasting blood sugar, cholesterol, and age.
- **Hierarchical Symptom Relationships:** Relationships between specific symptoms (like various chest pain types) and broader clinical categories were established in the knowledge graph.

### 5.5 Case Studies

To illustrate the complementary nature of neural and symbolic reasoning, we present a representative case study from our test set:

#### Example Diagnosis Case:

- 60-year-old male with asymptomatic chest pain.



- Resting blood pressure of 140 mm Hg.
- Cholesterol level of 260 mg/dl.
- Other features including ST depression and thalassemia abnormalities.

**Results:**

- Neural prediction: Heart Disease (confidence: 53%).
- Symbolic prediction: Heart Disease (confidence: 69%).
- NEXUS decision: Heart Disease (confidence: 62%).
- Strategy: Hybrid (neural weight: 43%, symbolic weight: 57%).

**Detected Symptoms/Features:**

- High resting blood pressure (`trestbps`).
- ST depression (`oldpeak`).
- Thalassemia (`thal`).

This case illustrates how the symbolic component provided higher confidence (69%) compared to the neural component (53%), leading the metacognitive controller to favor symbolic reasoning slightly while still incorporating neural insights.

## 6 Discussion

### 6.1 Complementary Strengths

Our results demonstrate the complementary strengths of neural and symbolic approaches in medical diagnosis. The neural component excelled at finding complex patterns in the data, while the symbolic component provided interpretable reasoning paths that could be validated by medical experts.

The fact that the integrated NEXUS approach achieved higher accuracy (81.97%) than either component individually validates the core thesis behind neural-symbolic AI: that combining pattern recognition with explicit reasoning can enhance overall system performance.

This finding aligns with recent research from IBM and other organizations suggesting that integrating neural and symbolic components can address the limitations of each approach when used in isolation [26].

### 6.2 Position in the Current Research Landscape

Comparing NEXUS to other neural-symbolic systems, we observe that:

- **Integration Approach:** NEXUS implements a "cooperative architecture" where neural and symbolic modules iteratively collaborate, which aligns with research indicating this approach performs best across metrics [25].
- **Metacognitive Control:** NEXUS addresses an underexplored area in neural-symbolic research, as only 5% of systems incorporate metacognitive frameworks [22].

- **Application Domain:** While major players like DeepMind focus on mathematics (Alpha-Geometry) and Microsoft on language tasks, our medical focus represents a high-impact application area with immediate practical relevance.

In terms of scale, NEXUS is modest compared to industry leaders:

- **NEXUS:**  $\sim 420,000$ – $450,000$  parameters.
- **Large language models:** Billions of parameters ( $1,000$ – $400,000\times$  larger).

However, NEXUS’s focused approach allows for efficient computation while still achieving strong performance on specialized medical tasks.

### 6.3 Limitations

Several limitations of this study should be noted:

- **Dataset Size:** The Heart Disease UCI dataset is relatively small (303 samples), which limits the model’s ability to learn more complex patterns.
- **Model Scale:** The current implementation uses a modest-sized transformer model compared to state-of-the-art systems from major research labs.
- **Domain Specificity:** While we’ve demonstrated success in heart disease prediction, additional work is needed to validate the approach across other medical domains and beyond.
- **Evaluation Metrics:** While we focused on accuracy as the primary evaluation metric, future work should consider additional metrics such as sensitivity, specificity, and F1-score, which are particularly important in medical diagnosis.
- **Clinical Validation:** While our results are promising, clinical validation with medical experts would be necessary to assess the practical utility of the NEXUS approach in real-world medical decision-making.

### 6.4 Implications for Explainable AI in Healthcare

The NEXUS architecture offers several advantages for explainable AI in healthcare:

- **Transparent Reasoning:** The symbolic component provides explicit reasoning paths that can be validated by medical experts.
- **Adaptive Strategy Selection:** The metacognitive controller dynamically selects the most appropriate reasoning strategy for each case, balancing accuracy and explainability.
- **Knowledge Integration:** The bidirectional neural-symbolic interface allows for the integration of medical knowledge into the system, enhancing both performance and interpretability.

These features address key challenges in the adoption of AI in healthcare, where black-box models are often viewed with skepticism due to their lack of interpretability.

## 6.5 Broader Applications

While our evaluation focused on heart disease prediction, the NEXUS architecture has potential applications in various domains where both accuracy and interpretability are crucial:

- **Scientific Discovery:** Combining pattern recognition with scientific principles for materials science and drug discovery.
- **Financial Systems:** Fraud detection and investment analysis with regulatory compliance and auditability.
- **Legal and Compliance:** Contract analysis and regulatory compliance with transparent reasoning.
- **Education:** Intelligent tutoring systems with personalized learning paths.
- **Critical Infrastructure:** Energy grid and water resource management with transparent decision-making.
- **Environmental Monitoring:** Ecosystem analysis and climate modeling with interpretable predictions.
- **Manufacturing:** Predictive maintenance and supply chain optimization with clear justifications.
- **Autonomous Systems:** Decision-making for vehicles and robots with transparent explanations.

In each of these domains, the ability to combine pattern recognition with explicit reasoning and provide transparent explanations would enhance both performance and trust.

## 7 Conclusion and Future Work

In this paper, we presented NEXUS, a neural-symbolic architecture for interpretable and aligned AI systems. Our evaluation on heart disease prediction demonstrated the potential of neural-symbolic integration in medical diagnosis, combining the pattern recognition capabilities of neural networks with the explicit reasoning of symbolic systems.

The neural component achieved 80.33% accuracy on the heart disease prediction task, significantly outperforming the symbolic component (50.82%). The integrated NEXUS approach achieved 81.97% accuracy, demonstrating a 1.64% improvement over the best individual component.

These results validate our approach and align with research trends suggesting that tightly integrated neural-symbolic systems can outperform pure neural or symbolic approaches.

Future work should focus on:

- **Enhanced Knowledge Integration:** Incorporating comprehensive medical knowledge into the symbolic component to improve reasoning performance.
- **Larger Datasets and Models:** Evaluating NEXUS on larger medical datasets with scaled-up neural models to better assess its performance and generalization capabilities.
- **Cross-Domain Validation:** Extending NEXUS to other domains requiring both pattern recognition and explicit reasoning.

- **Integration with LLMs:** Exploring integration with large language models for enhanced natural language explanations.
- **Temporal Reasoning:** Adding capabilities to reason about disease progression and treatment outcomes over time.
- **Clinical Validation:** Collaborating with medical experts to validate the interpretability and clinical utility of the NEXUS approach.
- **Multi-modal Data:** Extending NEXUS to handle multi-modal medical data, including imaging and text.
- **Enhanced Self-Improvement Mechanisms:** Developing more sophisticated learning methods for both neural and symbolic components.

Our development roadmap progresses through several phases, from the current foundational model to enhanced capabilities, scalable architecture, AGI-capable systems, and eventually a path to superintelligent AI with strong safety and control mechanisms.

Overall, our results suggest that neural-symbolic integration holds promise for explainable AI in healthcare and other critical domains, combining high performance with transparent decision-making. As the field continues to evolve, approaches like NEXUS that bridge the gap between pattern recognition and logical reasoning will play an increasingly important role in developing more capable, trustworthy, and aligned AI systems.

## Acknowledgments

We thank the UCI Machine Learning Repository for providing the Heart Disease dataset used in this research.

## References

- [1] Ahmad, M. A., Eckert, C., & Teredesai, A. (2018). Interpretable machine learning in healthcare. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics* (pp. 559–560).
- [2] Ghassemi, M., Oakden-Rayner, L., & Beam, A. L. (2021). The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health*, 3(11), e745–e750.
- [3] Hitzler, P., Bianchi, F., Ebrahimi, M., & Sarker, M. K. (2020). Neural-symbolic integration and the Semantic Web. *Semantic Web*, 11(1), 3–11.
- [4] Zhu, J., et al. (2023). Large Language Models and Symbolic Reasoning: A Survey and Comparison. *arXiv preprint arXiv:2502.12904*.
- [5] Garcez, A. D., et al. (2019). Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. *Journal of Applied Logics*, 6(4), 611–632.
- [6] Manhaeve, R., et al. (2018). DeepProbLog: Neural probabilistic logic programming. *Advances in Neural Information Processing Systems*, 31, 3749–3759.

- [7] Mao, J., et al. (2019). The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. *International Conference on Learning Representations*.
- [8] Bengio, Y. (2023). Towards disentangled representation learning and reasoning. *arXiv preprint arXiv:2502.15657*.
- [9] Nahar, J., Imam, T., Tickle, K. S., & Chen, Y. P. P. (2013). Computational intelligence for heart disease diagnosis: A medical knowledge driven approach. *Expert Systems with Applications*, 40(1), 96–104.
- [10] Das, R., Turkoglu, I., & Sengur, A. (2009). Effective diagnosis of heart disease through neural networks ensembles. *Expert Systems with Applications*, 36(4), 7675–7680.
- [11] Janosi, A., et al. (1988). Heart Disease UCI dataset. *UCI Machine Learning Repository*.
- [12] Ali, L., et al. (2019). An automated diagnostic system for heart disease prediction based on  $\chi^2$  statistical model and optimally configured deep neural network. *IEEE Access*, 7, 34938–34945.
- [13] Xie, N., et al. (2020). Explainable deep learning: A field guide for the uninitiated. *arXiv preprint arXiv:2004.14545*.
- [14] Holzinger, A., et al. (2019). Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4), e1312.
- [15] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144).
- [16] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774.
- [17] Bennetot, A., et al. (2019). Towards explainable neural-symbolic visual reasoning. *NeSy Workshop at IJCAI 2019*.
- [18] Vaswani, A., et al. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, 30, 5998–6008.
- [19] d’Avila Garcez, A., et al. (2024). Neurosymbolic AI: the 3rd wave. *Artificial Intelligence Review*.
- [20] Tenenbaum, J. B., et al. (2024). A Pipeline of Neural-Symbolic Integration to Enhance Spatial Reasoning. *arXiv preprint arXiv:2411.18564v1*.
- [21] Rossi, F., et al. (2024). Federated Neuro-Symbolic Learning for ICML 2024. *IBM Research*.
- [22] Colelough, M. & Regli, W. (2024). Neuro-Symbolic AI in 2024: A Systematic Review. *arXiv preprint arXiv:2501.05435v1*.
- [23] IBM Research. (2024). Neuro-symbolic AI. Retrieved from <https://research.ibm.com/topics/neuro-symbolic-ai>.
- [24] Google DeepMind. (2024). AlphaGeometry: An Olympiad-level AI system for geometry. Retrieved from <https://deepmind.google/discover/blog/alphageometry-an-olympiad-level-ai-system-for-geometry>.

- [25] Belfort, U. & Bourgogne, F. (2025). Unlocking the Potential of Generative AI through Neuro-Symbolic Architectures – Benefits and Limitations. *arXiv preprint arXiv:2502.11269v1*.
- [26] Chang, W. (2024). Neurosymbolic AI for reliability breakthroughs. *Fortune*, December 4, 2024.
- [27] Goertzel, B. (2024). SingularityNET Annual Report 2024: Advancing Beneficial AGI and Decentralized AI. *SingularityNET*.