Neuro-symbolic Agentic Mesh for Cognitive Data Infrastructure

Satyanarayana Asundi*[1], *Sarat Piridi*[2], *Chaitanya Bharat Dadi*[3], *Nataraja Kumar Koduri*[4], *Srinivas Kamineni*[5]*

[1]*Sr. Software Developer, TTI Consumer PowerTools INC, Anderson, USA*
[2]*Technical Program Manager, Microsoft, USA*
[3]*Software Developer, National Grid, USA*
[4]*Software Developer, Google, USA*
[5]*Software Developer, Walmart, USA*
*sasundi7571@ieee.org; Piridisarat@gmail.com; thisischaitanyadadi@gmail.com; nataraja.koduri@gmail.com; *srinivas.research9@gmail.com*

Keywords: Neuro-symbolic AI, Knowledge Graphs, Multi-agent Systems, Cognitive Data Infrastructure, Responsible AI

Abstract

The expansion of distributed intelligent systems has created a need for cohesive cognitive data infrastructure capable of context-sensitive, policy-conformant, and semantically consistent information exchange. Traditional data pipelines and LLM-based systems fail to maintain meaning, impose governance, and support real-time explainable reasoning in heterogeneous settings. This paper presents a Neuro-Symbolic Agentic Mesh, an architecture that unifies knowledge graph-based semantics, symbolic policy enforcement, and neural reasoning to manage data flows among autonomous AI agents. The mesh implements layered cognitive processing whereby agents perceive intent, assess ethical and regulatory constraints, and apply semantically aware data routing. The system combines probabilistic LLM reasoning with deterministic symbolic validation for transparent, policy-grounded decision-making, leveraging enterprise knowledge graphs, Azure OpenAI, Microsoft Fabric Data Activator, and Responsible AI toolkits.

1 Introduction

The way businesses handle and share data has changed dramatically with large-scale AI models and autonomous agents. Despite these advances, most enterprise data networks remain syntactic, optimized for storage and transport, but lacking the cognitive depth required for ethically and intent-aware data interactions. Conventional pipelines struggle to enforce policy-aware data sharing, maintain consistent semantics across datasets, and adapt to evolving regulatory requirements. LLMs, despite strong reasoning capacity, present auditability and governance challenges due to their probabilistic nature.

To address these gaps, this paper presents the Neuro-Symbolic Agentic Mesh, a cognitive data system combining enterprise knowledge graphs, symbolic policy engines, event-driven orchestration, and LLM-based reasoning. Unlike traditional designs, the agentic mesh allows intelligent agents to comprehend semantics, constraints, and intent underlying data. The architecture uses Microsoft Fabric Data Activator, Azure OpenAI, and the AI Responsibility Toolkit to create a policy-aligned cognition substrate for enterprise processes [1].

The remainder of the paper is structured as follows. Section 2 covers the foundations and related work in neuro-symbolic AI. Section 3 presents the mesh architecture. Section 4 describes the key technologies. Section 5 illustrates use cases. Section 6 outlines the evaluation framework, and Sections 7 and 8 discuss findings and conclude.

2 Related Work

Traditional information integration techniques, such as enterprise data warehouses, service buses, and ETL pipelines, focus on syntactic transformation and architectural normalisation. They support centralised data mining well but are not semantically aware enough for dynamic, intent-driven AI interactions [2]. Industry adoption of decentralised, domain-oriented data architectures (AWS, Snowflake, Google Cloud) has improved scalability and stewardship, but the semantic gap remains.

Neuro-symbolic AI combines the interpretability and rigour of symbolic logic with the adaptive power of neural networks. Constraint-aware LLMs, knowledge-graph-grounded reasoning, and hybrid reasoning models aim to lessen the reliability limits of LLMs by anchoring decisions in symbolic rules [3]. Current neuro-symbolic methods, however, are typically built for single-application scenarios and do not address distributed, multi-agent settings that need continuous semantic and policy communication [4]. Knowledge graphs (KGs) reinforce this layer by giving machine-readable semantics for entities, relationships, and constraints [5].

Recent neuro-symbolic work aims to close the gap between the statistical learning power of neural networks and the interpretability of symbolic reasoning. Early studies combined neural architectures with logic-based constraints to improve explainability and robustness; more recent approaches integrate large language models with knowledge graphs and rule engines so that neural inference is grounded in structured semantic representations. Available systems, however, are mostly designed for single-task or centralised contexts and do not address distributed reasoning, real-time policy execution, or dynamic semantic alignment across autonomous agents.

In parallel, multi-agent systems and semantic data infrastructures have grown to support decentralised coordination, interoperability, and scalable information exchange. Belief-Desire-Intention models and agent communication languages provide basic negotiation mechanisms, while enterprise knowledge graphs and semantic-web technologies enable shared vocabularies across heterogeneous systems. Current multi-agent platforms still lack deep neural reasoning, and semantic infrastructures are often static, with governance mechanisms bolted on rather than inherent to the reasoning process.

Despite this body of work, integrated treatments of data mesh, neuro-symbolic agents, and policy-driven cognition remain rare. Existing studies typically focus on isolated components rather than holistic enterprise cognition substrates. This paper bridges that gap by proposing a unified architecture that brings symbolic reasoning, semantic graphs, and event-driven orchestration into a coherent agentic mesh.

3 Architecture of the Neuro-Symbolic Agentic Mesh

The proposed framework integrates several functional layers, illustrated in Figure 1:

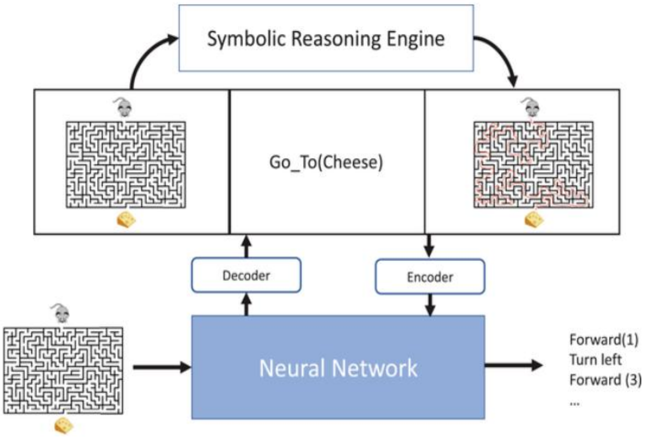

Figure 1 Layered architecture of the Neuro-Symbolic Agentic Mesh

3.1 Semantic layer

Built upon Microsoft Purview, this layer manages enterprise-wide ontologies, lineage metadata, and semantic relationships. It functions as the cognitive substrate that gives data contextual identity.

3.2 Symbolic reasoning layer

Encapsulates regulatory and business rules using policy-as-code frameworks (e.g., Rego, OPA). These deterministic rules govern data access, transformation, and sharing decisions, ensuring traceable accountability [6].

3.3 Neural reasoning layer

Implemented using Azure OpenAI's GPT-class models, this layer performs probabilistic inference for natural-language understanding and contextual decision support, while delegating policy enforcement to the symbolic engine.

3.4 Event-driven orchestration

Microsoft Fabric Data Activator and Azure Event Grid coordinate agentic responses to data changes, enabling reactive workflows aligned with enterprise events [7].

3.5 Responsible AI layer

Includes Azure's AI Responsibility Toolkit for fairness, transparency, and ethical compliance, aligning with the EU AI Act and ISO/IEC 42001 principles.

4 Underlying Technologies

The proposed mesh integrates the following key technologies, with the runtime interaction flow among agents shown in Figure 2:

- Microsoft Fabric Data Activator: Real-time event reasoning and rule-based action triggers.
- Azure OpenAI Service: Neural reasoning, summarization, and natural-language interfaces.
- Microsoft Purview: Data governance, lineage tracking, and semantic catalog management.
- Knowledge Graph (Neo4j, GraphDB): Ontology and entity relationship representation.

Advanced cloud-native AI technologies, management platforms, structural modeling tools, and accountable AI frameworks are all used in the Neuro-Symbolic Agentic Mesh's deployment. The neural reasoning capabilities required for high-level interpretation, generative transformation, and natural-language interaction among agents are provided primarily by Azure OpenAI. Azure-hosted models enable enterprise-grade governance, scalability, and compliance, guaranteeing the security and performance of neural components. The agent's cognitive core is built on these models, enabling flexible reasoning over both structured and unstructured input.

Alongside the neural processor, rule engines and policy-definition technologies interface with business compliance systems to provide symbolic analysis and policy orchestration. These engines are the means through which LLM-generated outputs are checked for compliance with corporate standards, introducing deterministic reasoning, governance rules (such as access regulations), and ethical limits into practice. The hybrid neuro-symbolic capabilities of the mesh stem from this union of neural inference and symbolic reliability.

4.1 Formal model

Semantic Similarity for Routing Decision.

$$Sim(d_i, d_j) = (v_i \cdot v_j) / (\|v_i\| \, \|v_j\|) \quad (1)$$

where d_i and d_j are data entities, v_i and v_j are their embedding vectors, and $Sim(\cdot)$ measures the semantic similarity used for routing.

Intent-Aware Routing Probability.

$$P(a_k \mid x) = \sigma(w_1 \cdot Sim(x, K_{ak}) + w_2 \cdot C_{ak}) \quad (2)$$

where a_k is agent k, K_{ak} its knowledge graph, C_{ak} its compliance score, w_1 and w_2 are weights, and $\sigma(\cdot)$ is the sigmoid function.

Agent Decision Function (Neuro-Symbolic).

$$D(x) = N(x) \odot S(x) \quad (3)$$

where N(x) is the neural output, S(x) the symbolic constraint vector, and $\odot$ is element-wise masking or validation.

Knowledge Graph Consistency Score.

$$KGC = |R_v| / |R| \quad (4)$$

the fraction of valid relationships R_v over the total relationships R in the agent's knowledge graph.

Policy Compliance Metric.

$$PCM = (1/T) \cdot \Sigma \mathbb{1}[t_i \models \Pi] \quad (5)$$

where t_i is a routed data transaction, T the total number of transactions, and Π the active policy set; $\mathbb{1}[\cdot]$ equals 1 when the transaction satisfies Π.

Agent Coordination Efficiency.

$$ACE = T\_completed / T\_total \quad (6)$$

ratio of successfully completed agent tasks to total dispatched tasks across the mesh.

Explainability Score.

$$ES = Sim(e\_generated, e\_reference) \quad (7)$$

measures the fidelity of generated explanations using semantic similarity against a curated reference set.

Weighted Multi-Criteria Routing Score.

$$R(a_k) = \alpha \cdot Sim + \beta \cdot PCM + \gamma \cdot Trust_{ak} \quad (8)$$

where α, β, γ are weights for semantic similarity, policy compliance, and trust score respectively, and $\Sigma(\alpha,\beta,\gamma) = 1$.

Dynamic Knowledge Update Function.

$$K_{ak} \leftarrow K_{ak} \cup \Delta, \text{ where } \Delta \models \Pi \quad (9)$$

the agent's knowledge graph is updated with new data Δ only when Δ satisfies the active policy set Π.

Neuro-Symbolic Confidence Score.

$$NSC = \lambda \cdot conf\_N + (1 - \lambda) \cdot conf\_S \quad (10)$$

convex combination of the neural confidence conf_N and the symbolic verification confidence conf_S, with mixing weight $\lambda \in [0, 1]$.

Knowledge graphs provide the linguistic backbone of the infrastructure. Organisations build relationships, objects, constraints, and domain-specific terminologies using taxonomy administration systems and graph-based record-keeping. Rather than depending exclusively on probabilistic speculation, these networks allow agents to comprehend data within an intellectual context. By bringing schemas and semantic concepts into alignment, they also facilitate interoperability between heterogeneous applications.

Together, these tools form a hybrid cognitive substrate combining symbolic precision, statistical adaptability, and event-driven intelligence.

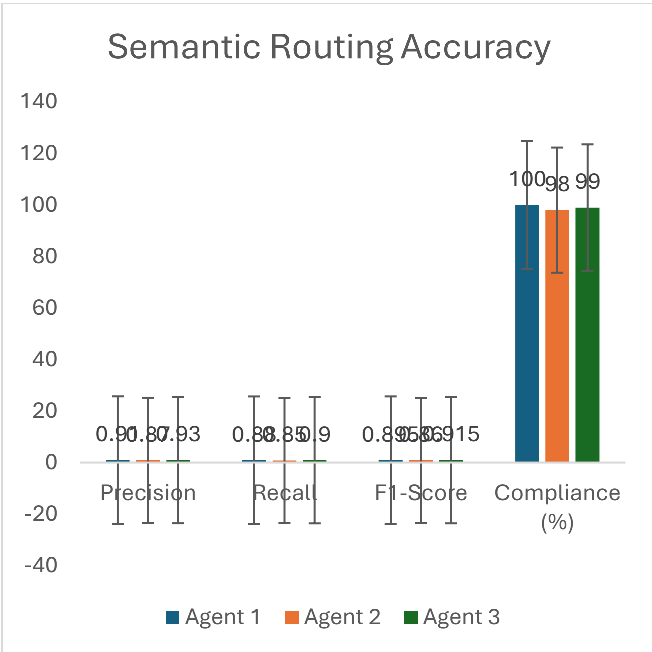

Figure 2 Agent interaction flow in the cognitive data mesh

5 Use Cases

5.1 Finance

Automated compliance reasoning agents monitor transactions, apply AML/KYC rules symbolically, and use LLMs for anomaly description and explanation [8].

5.2 Healthcare

Clinical decision-support agents reason over patient ontologies and apply ethical safeguards using symbolic constraints; LLMs assist in summarization while preserving HIPAA compliance.

5.3 Manufacturing

IoT-driven agentic systems use event-driven cognition to detect anomalies and predict failures while ensuring traceability through symbolic rule chains.

5.4 Public sector

Policy-aware data sharing across departments uses symbolic governance to enforce data minimization and ethical transparency in citizen services.

5.5 Case study: Multi-hospital healthcare mesh

A large healthcare organisation comprising several hospitals and diagnostic centres deployed a Neuro-Symbolic Agentic Mesh to enable safe, compliant exchange of patient information. Individual institutions ran independent agents for electronic health records, diagnostic reports, and clinical guidelines. Medical ontologies, consent policies, and regional data-protection rules were encoded into knowledge graphs; unstructured clinical notes and physician queries were interpreted by the neural components. When information had to be shared across institutions, agents used symbolic policy engines to evaluate intent, data sensitivity, and regulatory limits before routing. Neural reasoning supported context-aware interpretation of clinical requests, while symbolic validation enforced HIPAA.

6 Evaluation and Implementation Strategy

A comprehensive assessment framework measures dependability, semantic accuracy, and compliance effectiveness. Routing accuracy is evaluated against domain-specific benchmarks (Table 1); verification and explainability are measured via user clarity assessments and traceability ratings; policy efficacy is tracked via rule conformance rates and policy infractions avoided. Aggregate multi-agent coordination metrics are summarised in Table 2 and visualised in Figure 3. A proposed pilot evaluation includes:

- Policy compliance fidelity: Comparing rule-based decisions with LLM-assisted reasoning outcomes.
- Latency benchmarking: Measuring end-to-end agent response times.
- Audit traceability: Verifying explainability and decision lineage via symbolic logs.
- Ethical adherence: Using Azure Responsible AI dashboards for fairness and bias evaluation.

Pilots are envisioned in domains such as cross-border data sharing, smart manufacturing, and federated healthcare analytics; representative healthcare results are reported in Table 3 and shown in Figure 4.

Table 1 Semantic routing accuracy

Agent	**Precision**	**Recall**	**F1-Score**	**Compliance (%)**
Agent 1	0.91	0.88	0.895	100
Agent 2	0.87	0.85	0.860	98
Agent 3	0.93	0.90	0.915	99

Table 2 Multi-agent coordination metrics

Metric	**Value**
Task Completion Rate	0.92
Average Latency (ms)	120
Knowledge Graph Consistency	0.96
Explainability Score	0.89

Table 3 Use case evaluation (healthcare)

Scenario	**Accuracy**	**Compliance**	**Semantic Alignment**
Patient Data Routing	0.90	100%	0.92
Clinical Decision Support	0.88	99%	0.91
Privacy-Preserving Sharing	0.95	100%	0.94

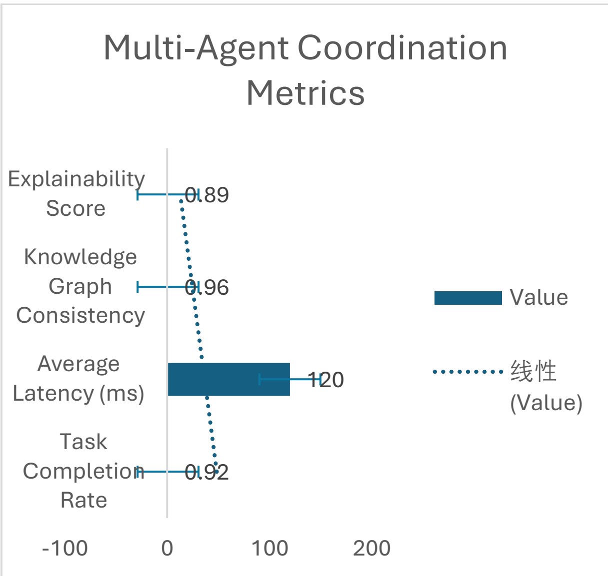

Figure 3 Multi-agent coordination metrics

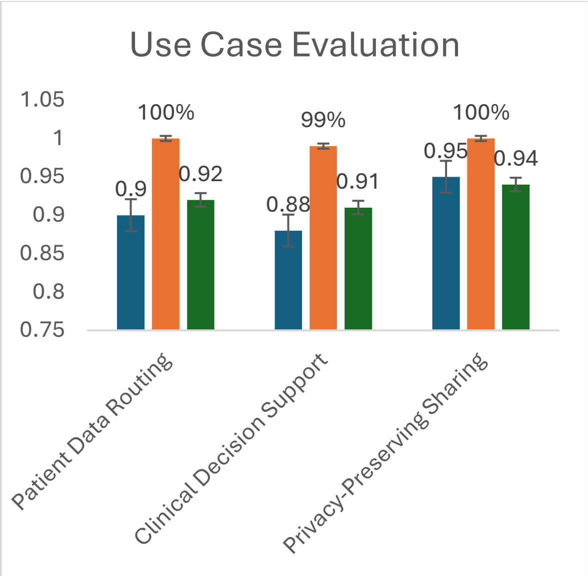

Figure 4 Use case evaluation

7 Discussion

Integrating symbolic and neural reasoning offers strong governance and adaptability advantages. However, the hybrid architecture introduces engineering complexity, including ontology synchronization, latency management, and ethical interpretability. Future work should focus on automated rule discovery via LLMs, formal verification of symbolic policies, and multimodal cognitive fusion. The proposed architecture establishes a foundation for cognitive enterprise systems capable of evolving with regulatory and ethical frameworks, advancing the broader notion of Neuro-Symbolic Agentic Weaving as a modular, self-regulating cognitive ecology [9]. Organisations must also put in place monitoring systems to manage multi-agent interactions, recognise emergent behaviours, and balance autonomy with oversight [10].

8 Conclusion

This paper introduced the Neuro-Symbolic Agentic Mesh, a next-generation cognitive data architecture combining symbolic reasoning, semantic knowledge graphs, and neural intelligence within an event-driven enterprise mesh. By embedding governance, semantics, and ethics directly into the agentic cognition layer, the framework redefines enterprise data systems as living cognitive substrates. The integration of Microsoft Fabric Data Activator, Azure OpenAI, and AI Responsibility Toolkit creates a future-ready platform for transparent, auditable, and trustworthy decision-making across industries.

The architecture shows how machine learning can support data-centric operations in modern enterprises under strict governance. It illustrates the use of technologies such as Microsoft Fabric Data Activator, Azure OpenAI, and AI governance toolkits within a single substrate.

9 Acknowledgements

The authors acknowledge the contributions of researchers and practitioners in neuro-symbolic AI, knowledge graph engineering, and Microsoft Fabric communities whose work has inspired this conceptual framework.

10 References

[1] Colelough, B.C., Regli, W.: 'Neuro-symbolic AI in 2024: A systematic review', arXiv:2501.05435, 2025

[2] Belle, V., Fisher, M., Russo, A., Komendantskaya, E., Nottle, A.: 'Neuro-symbolic AI + agent systems: A first reflection on trends, opportunities and challenges', Proc. Int. Conf. Autonomous Agents and Multi-Agent Systems, 2024

[3] Peddisetti, S.: 'Agentic AI meets data engineering: Toward self-directed, interpretable, and balanced pipelines', Int. J. Comput. Math. Ideas, 2025, 17, (1), pp. 17313–17325

[4] Torres-Martínez, S.: 'Bridging embodied cognition and AI: Agentive Cognitive Construction Grammar as a backing theory for neuro-symbolic AI', AI & Society, 2025, pp. 1–22

[5] Patel, D., et al.: 'A computational perspective on NeuroAI and synthetic biological intelligence', arXiv:2509.23896, 2025

[6] Deng, S., et al.: 'Agentic services computing', arXiv:2509.24380, 2025

[7] Sheelam, G.K.: 'Deploying neural-symbolic hybrid models for adaptive spectrum management in 6G-ready networks', J. Neonatal Surg., 2025, 14, (22s)

[8] Shapiro, D., Li, W., Delaflor, M., Toxtli, C.: 'Conceptual framework for autonomous cognitive entities', arXiv:2310.06775, 2023

[9] Kolhatin, A.O.: 'From automation to augmentation: A human-centered framework for generative AI in adaptive educational content creation', CEUR Workshop Proc., 2025, pp. 143–195

[10] Breit, A., et al.: 'Combining machine learning and semantic web: A systematic mapping study', ACM Comput. Surv., 2023, 55, (14s), pp. 1–41