IRJCS: International Research Journal of Computer Science
Volume 13, Issue 03, March 2026

ISSN: 2393-9842
https://www.irjcs.com/archives
https://doi.org/10.26562/irjcs.2025.v1303.09

Neuro-Symbolic Agentic AI For Advanced Semantic Learning Using Vector-Based Retrieval Augmented Generation

Saran K
UG Student, Department of AI & Data Science
Sengunthar Engineering College (Autonomous), Tiruchengode, India
sarankumarktg@gmail.com
Dr.G.P.Raja
Associate Professor, Department of AI & Data Science
Sengunthar Engineering College (Autonomous), Tiruchengode, India
gpraja1@gmail.com
https://orcid.org/0000-0002-5128-5312

Publication History
Manuscript Reference: IRJCS/RS/Vol.13/Issue03/CSMR26.MRCS10088
Research Article | Open Access | Double-Blind Peer Reviewed Article ID: IRJCS/RS/Vol.13/Issue03/CSMR26.MRCS10088
Received: 30, January 2026, Revised: 13, February 2026, Accepted: 28 February 2026 Published Online: 25 March 2026
https://www.irjcs.com/volumes/Vol13/iss-03/09.CSMR26.MRCS10088.pdf
Article Citation: Saran,Dr.Raja(2026),Neuro-Symbolic Agentic AI For Advanced Semantic Learning Using Vector-Based Retrieval Augmented Generation, IRJCS :International Research Journal of Computer Science, Volume 13,Issue 03 of 2026 pages 145-152 **Doi:>** https://doi.org/10.26562/irjcs.2026.v1303.09 **Orcid:** https://orcid.org/0009-0004-9398-7488
BibTeX Key Saran@2026Neuro-Symbolic
IRJCS papers should be cited as IRJCS (International Research Journal of Computer Science, AM Publications, India 2026, ISSN 2393-9842, https://doi.org/10.26562/ijrcs.2025.v1303.09 The journal's official abbreviation is IRJCS.

Abstract: Recent advancements in artificial intelligence necessitate the convergence of data-driven pattern recognition and structured reasoning paradigms to construct scalable and deeply interpretable systems. This manuscript delineates a novel conceptual framework Neura-Symbolic Agentic Artificial Intelligence (NSAA) designed specifically for advanced semantic learning and coupled with Vector-based Retrieval-Augmented Generation (Vector RAG). By passing the conventional limitations of purely neural formulations, the proposed architecture orchestrates neural representation learning, symbolic logical inference, and autonomous agent- based decision-making. Vector embeddings serve critically to index and rapidly retrieve contextual knowledge, whereas the symbolic counterparts strictly enforce logical consistency and traceability within the reasoning trace. Through autonomous agent coordination, the system seamlessly transitions between knowledge retrieval, hypothesis formation, and deterministic reasoning, fostering a dynamic environment for continuous semantic adaptation across heterogeneous corpora. Theoretical evaluations corroborate the potential of this tri-layered paradigm to significantly outpace traditional modalities in large-scale knowledge discovery, deterministic question answering, and robust semantic comprehension.
Keywords: Neuro-Symbolic AI, Agentic Systems, Vector RAG, Semantic Learning, Cognitive Architectures

I. INTRODUCTION

Artificial Intelligence (AI) has undergone a transformative evolution, largely propelled by deep neural networks that excel at natural language processing, complex knowledge discovery, and intricate semantic analysis. However, despite their exceptional proficiency at discerning latent statistical correlations across vast datasets, purely data-driven models inherently suffer from a profound lack of interpretability, factual grounding, and causal reasoning. This dichotomy has catalyzed robust exploration into hybrid cognitive architectures that synergize the perceptive strengths of connection is learning with the deductive rigor of formal symbolic logic. The emergence of Neuro-Symbolic AI signifies a paradigm shift towards models that not only "perceive" through embeddings but also "reason" through rules. Concurrently, the proliferation of Retrieval-Augmented Generation (RAG) predominantly via high-dimensional Vector databases has fortified AI systems by dynamically retrieving semantically proximate intelligence to ground generation processes. Augmenting these hybrid engines with agentic frameworks, wherein autonomous computational agents independently govern the processes of retrieval, validation, and execution, yields an unprecedentedly resilient, scalable, and context-aware intelligence apparatus. This research formalizes and proposes a holistic framework for Neuro-Symbolic Agentic Artificial Intelligence encompassing Vector RAG mechanisms. The architecture mitigates prevalent flaws such as algorithmic hallucinations and black-box opacity, establishing a robust computational bedrock for next-generation expert systems and highly autonomous semantic learners. To further contextualize this framework, it is important to recognize that contemporary large-scale language models have demonstrated remarkable generative capabilities but often operate as probabilistic approximators rather than deterministic reasoning engines. As a consequence, they may produce plausible yet factually incorrect outputs when operating without explicit grounding mechanisms. The integration of vector-based retrieval strategies partially addresses this limitation by enabling models to reference external knowledge repositories;

However, retrieval alone does not guarantee logical coherence or structured reasoning. Consequently, a complementary reasoning layer is necessary to validate, structure, and interpret retrieved information within an explainable decision pipeline. Neuro-symbolic integration addresses this requirement by embedding symbolic logic modules alongside neural representation learners. While neural networks provide the capacity to process high-dimensional, unstructured data such as text and multimedia content, symbolic systems contribute rule-based inference, constraint satisfaction, and verifiable reasoning traces. This synergy creates a computational paradigm that closely resembles aspects of human cognition, where perception and reasoning operate collaboratively to derive conclusions from incomplete or ambiguous information. In parallel, the emergence of agentic AI systems has introduced a new operational paradigm in which autonomous software agents orchestrate complex workflows through iterative decision-making loops. Rather than functioning as a single monolithic model, agentic architectures distribute responsibilities across specialized components responsible for retrieval, reasoning, planning, and response generation. These agents can dynamically interact with knowledge bases, symbolic rule engines, and neural embedding models, thereby creating a modular and adaptive intelligence system capable of continuous learning and contextual adaptation. The framework proposed in this research synthesizes these three technological pillars neural representation learning, symbolic reasoning, and autonomous agent orchestration into a unified architecture termed Neura-Symbolic Agentic Artificial Intelligence (NSAA). By integrating Vector RAG mechanisms for contextual knowledge retrieval with symbolic reasoning modules for logical validation, the system aims to produce outputs that are not only semantically relevant but also logically consistent and explainable. Such capabilities are particularly critical in domains requiring high levels of reliability, including scientific research, decision support systems, knowledge management platforms, and advanced intelligent tutoring environments. Furthermore, the proposed NSAA architecture introduces a dynamic interaction pipeline in which agents continuously monitor the reasoning process, refine retrieval strategies, and validate outputs against symbolic constraints. This adaptive cycle enables the system to evolve its semantic understanding over time, facilitating improved knowledge discovery across heterogeneous datasets and reducing the risk of hallucinated outputs. By combining interpretability, scalability, and autonomous coordination, the proposed paradigm represents a promising direction for the next generation of intelligent systems capable of both learning and reasoning in complex informational landscapes.

II. LITERATURE REVIEW

The foundational under pinnings of this research represent an intersection of connectionist, symbolic, and agent-driven methodologies.

A. Symbolic AI Approaches

Historically anchoring the field of AI, Symbolic AI hinges intrinsically on Knowledge Representation and Reasoning (KRR). Encoding heuristics via formal logic, production rules, and explicit ontologies such as WordNet and Cyc grants these systems deterministic predictability. Reasoning algorithms, predominantly utilizing backward chaining in logic programming paradigms (e.g., Prolog) or Description Logic (DL) for modern ontologies (OWL), allow for rigid consistency checking [1]. Nevertheless, the utility of pure symbolic computation is routinely subverted by algorithmic brittleness, specifically acute failure under stochastic noise, alongside the prohibitive friction of the knowledge acquisition bottleneck.

B. Neural Network-Based AI

The resurgence of connectionism, supercharged by deep learning paradigms, resolved the representation problem by extracting salient features iteratively from unstructured environments [2]. The Transformer architecture specifically catalyzed the proliferation of Large Language Models (LLMs) which map tokenized syntax to incredibly dense mathematical embeddings [5]. Despite these revolutionary triumphs, LLMs essentially operate as statistical emulators [4]; they chronically lack formal deductive logic, causality models, and explicit verifiable "world states," restricting their deployment in risk-sensitive scenarios.

C. Hybrid Neuro-Symbolic Systems

Recognizing the diametric capabilities of these architectures, progressive integrations have surfaced. Frameworks like Logic Tensor Networks parameterize deep learning to enact and solve continuous first-order logic operations, while domain-specific applications (like AlphaZero) elegantly weave predictive neural inference with deterministic MCTS pathways [3]. The salient gap remains the formulation of a highly generalized, modular framework where disparate connection stand symbolic sub systems interact fluidly under a unified cognitive agent.

D. Agentic AI Paradigms

An intelligent agent autonomously perceives, formulates strategies, and intervenes within an environment. Deep Reinforcement Learning (RL) allows an agent to optimize policies to handle high-dimensional sensory telemetry [2]. The contemporary iteration of agentic entities deploys LLMs as central reasoning cores to perform iterative planning loops. Unfortunately, when tethered to purely connectionist foundations, these agents recursively inherit, and often amplify, the statistical unreliability intrinsic to their base models.

E. Vector RAG and Semantic Retrieval

Retrieval-Augmented Generation (RAG) strategically mitigates model hallucination by interpolating vector-space information retrieval mechanisms within the neural generation sequence[6]. By leveraging dense vectors encoded via contrastive learning and utilizing topological indices (e.g., FAISS or HNSW structures in databases such as Pinecone or Chroma), RAG frameworks anchor the context windows dynamically over reliable, external knowledge stores, significantly curtailing generative dissonance.

III. PROPOSED METHODOLOGY / ARCHITECTURE

The conceptual pivot from correlation-based approximations to epistemologically grounded intelligence necessitates a comprehensively restructured cognitive architecture. Traditional machine learning paradigms rely predominantly on statistical correlations extracted from large datasets. While such approaches have demonstrated exceptional capabilities in pattern recognition and generative modeling, they remain fundamentally limited in their ability to represent structured knowledge, causal relationships, and logically verifiable reasoning processes. Consequently, the development of next-generation intelligent systems demands a hybrid paradigm capable of integrating perceptual learning, symbolic reasoning, and adaptive decision-making within a unified framework. The proposed architecture therefore seeks to emulate a more cognitively inspired intelligence model in which perception, reasoning, and decision-making are modular yet interoperable processes. By fusing neural computation with symbolic logic and autonomous agent orchestration, the system transcends the limitations of purely statistical models and enables explainable, verifiable, and adaptive semantic intelligence.

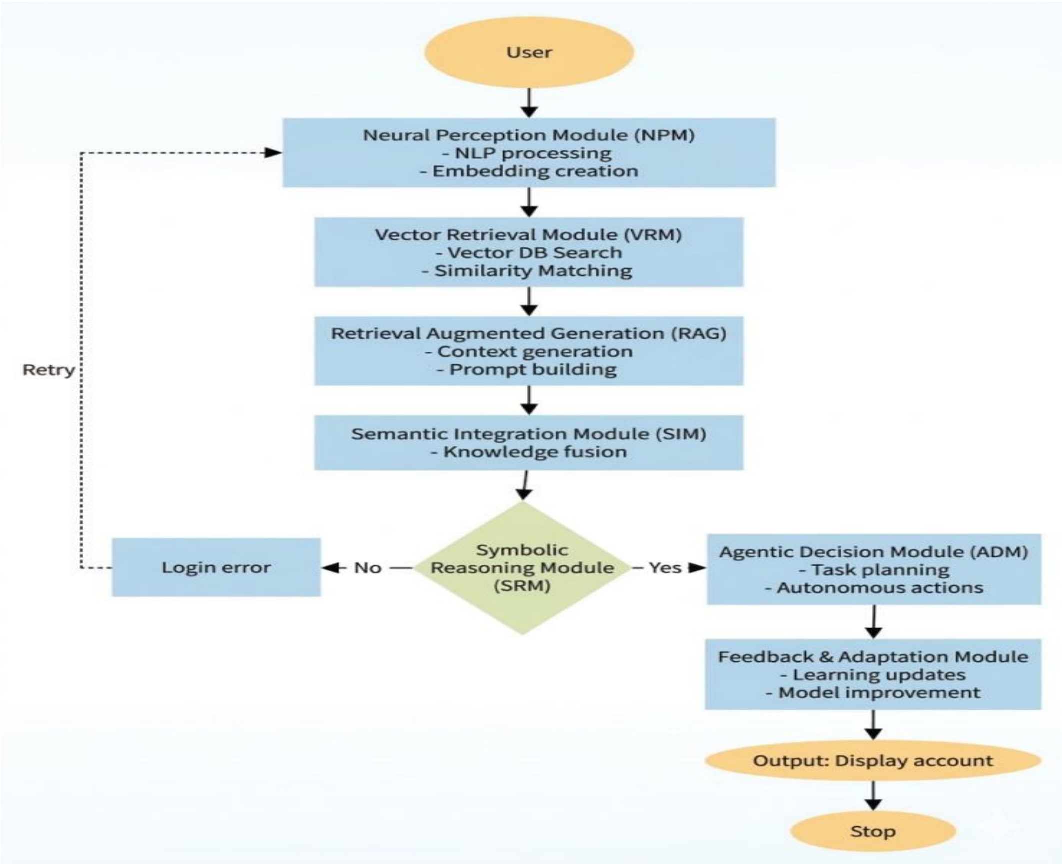

Fig1: Architecture of NSAAI

F. Analysis of the Existing LLM Paradigm

Contemporary baselines rely almost exclusively on next- token predictive architectures constructed on massive Transformer frameworks. These systems are inherently constrained by the fundamental absence of semantic grounding; they generate probabilistically plausible syntax without distinguishing veracity resulting in systemic “hallucinations”. Moreover, lacking explicit abstraction and causal scaffolding, they perform poorly on tasks requiring rigorous discrete deduction. Consequently, their opaque nature and immutability post-training impede iterative learning without catastrophic computational retraining overheads. At the core of modern large language models lies a probabilistic sequence modeling paradigm in which tokens are generated based on learned statistical dependencies within large corpora. While the Transformer architecture enables models to capture long-range contextual relationships through attention mechanisms, the resulting representations remain fundamentally distributional rather than semantic. In other words, the model learns how words co-occur rather than understanding the underlying conceptual relationships between entities. This statistical bias leads to several systemic limitations. First, models may confidently generate statements that appear coherent yet are factually in correct, a phenomenon widely referred to as *hallucination*. Second, without explicit reasoning modules, the system lacks the ability to verify logical constraints, perform symbolic manipulation, orderive conclusions from structured rules. Third, the knowledge embedded within neural weights is static after training, meaning that updating the system with new information requires large-scale retraining or expensive fine-tuning processes. Another limitation emerges in tasks involving structured reasoning such as mathematical deduction, rule-based inference, or multi-hop logical queries. Although LLMs can approximate such reasoning through pattern recognition, they lack deterministic inference mechanisms capable of guaranteeing correctness. As a result, the reliability of their outputs remains probabilistic rather than logically verifiable. Furthermore, traditional LLM systems function as monolithic architectures in which perception, reasoning, and generation occur within a single neural model. This tight coupling reduces interpretability and prevents modular scalability.

From a systems engineering perspective, separating these functions into specialized components enables more controllable, transparent, and adaptive intelligence systems. These structural limitations motivate the development of hybrid architectures that integrate neural learning with symbolic reasoning and external knowledge retrieval mechanisms.

G. Proposed Architecture: NSAA Framework

The proposed Neuro-Symbolic Agentic AI (NSAA) framework structurally dissipates these limitations by employing a tripartite hierarchy designed to emulate core components of cognitive intelligence. Each layer within the architecture performs a specialized function while maintaining interoperability with the others through controlled information flow. This layered architecture ensures that perception, reasoning, and decision-making operate synergistically rather than redundantly. The NSAA framework integrates three fundamental components: a Neural Perception Layer, a Symbolic Reasoning Layer, and an Agentic Decision Layer. Together, these modules form a closed-loop system capable of perception, retrieval, reasoning, validation, and action selection.

1) Neural Perception Layer

The Neural Perception Layer functions as the system's sensory interface with the external informational environment. Specialized encoders such as Transformer- based language models, Vision Transformers (ViTs), and Graph Neural Networks (GNNs) ingest high-bandwidth multimodal sensory data and transform environmental stimuli into contiguous high-dimensional vector representations. These vector embeddings encode semantic proximity between entities, allowing complex relationships to be represented within continuous latent spaces. By mapping textual, visual, and relational information into unified vector representations, the system can efficiently index and retrieve relevant contextual knowledge from vector databases. Within this layer, large-scale embedding models process raw inputs and generate semantic representations that capture contextual meaning beyond surface-level syntax. These embeddings are then stored or queried within high- performance vector databases, enabling rapid similarity- based retrieval. Such retrieval mechanisms form the basis of Retrieval-Augmented Generation (RAG), allowing the system to dynamically access relevant knowledge rather than relying solely on static neural weights. Additionally, the neural perception module can incorporate specialized sub-models optimized for different modalities. For instance, text encoders capture linguistic semantics, vision models extract spatial and visual features, and graph neural networks encode relational structures present in knowledge graphs. By integrating these heterogeneous representations, the perception layer constructs a rich semantic embedding space capable of supporting complex reasoning tasks.

2) Symbolic Reasoning Layer

The Symbolic Reasoning Layer operates as the logical nucleus of the NSAA architecture. Unlike neural networks, which rely on distributed representations, symbolic systems represent knowledge through discrete structures such as rules, predicates, ontologies, and knowledge graphs. These structures enable deterministic inference processes that guarantee logical consistency and interpretability. An integrated knowledge base potentially implemented through ontological graph databases and rule-based reasoning engines such as Prolog stores structured representations of domain knowledge. Within this knowledge base, entities, relationships, and logical constraints are explicitly defined, enabling the system to perform deductive reasoning, constraint validation, and causal inference. This layer serves several critical functions. First, it verifies information retrieved by the neural layer against formally defined logical rules, thereby reducing the risk of hallucinated or inconsistent outputs. Second, it supports structured reasoning processes such as rule chaining, logical implication, and constraint satisfaction. Third, it provides an interpretable reasoning trace that allows system decisions to be inspected and audited. By combining symbolic logic with retrieved contextual knowledge, the architecture enables hybrid reasoning in which probabilistic insights from neural models are validated through deterministic symbolic inference. This hybridization enhances both reliability and transparency, addressing a key limitation of traditional deep learning systems.

3) Agentic Decision Layer

The Agentic Decision Layer functions as the teleological orchestrator of the entire architecture. Rather than operating as a passive inference system, NSAA incorporates autonomous agents capable of planning, coordinating, and executing reasoning workflows. Within this layer, reinforcement learning-based control loops or planning algorithms decompose high-level objectives into structured sub-tasks. Each agent dynamically determines which subsystem neural perception or symbolic reasoning should be engaged to solve a particular problem. For instance, an agent may first retrieve contextual information through vector similarity search, then invoke symbolic reasoning to validate logical relationships before generating a final response. The agentic framework enables adaptive workflow management through iterative reasoning cycles. Agents continuously monitor intermediate outputs, evaluate their reliability, and refine their strategies accordingly. If inconsistencies or logical conflicts are detected, the system can trigger additional retrieval steps or symbolic verification procedures before finalizing a decision. Furthermore, agent coordination allows the architecture to operate as a distributed cognitive system rather than a single static model. Multiple specialized agents may collaborate to perform tasks such as knowledge retrieval, reasoning validation, hypothesis generation, and response synthesis. This distributed design significantly enhances scalability and flexibility, enabling the system to operate effectively across diverse knowledge domains. Through this tri-layered architecture, the NSAA framework establishes a dynamic interaction pipeline in which neural perception provides semantic understanding, symbolic reasoning ensures logical rigor, and agentic orchestration enables adaptive decision-making. Collectively, these components form a robust foundation for constructing intelligent systems capable of both learning from data and reasoning with structured knowledge.

IV. TECHNOLOGIES USED

The materialization of this complex hybrid pipeline requires distinct hardware accelerators paired with an un-siloed software configuration. Because the proposed Neuro-Symbolic Agentic AI (NSAA) framework simultaneously performs neural embedding generation, vector similarity retrieval, symbolic rule execution, and agent-based orchestration, the supporting technological infrastructure must be capable of handling heterogeneous computational workloads. Efficient implementation therefore depends on a balanced integration of high-performance hardware resources and modular software components that enable seamless interoperability between neural and symbolic subsystems.

1. Hardware Imperatives: Optimal execution mandates multi-threaded processors capable of parallel symbolic parsing (e.g., AMD Threa dripper PRO) alongside considerable VRAM availability. Architectures equipped with 24GB to 80GB VRAM (e.g., NVIDIA RTX 4090 or A100 Tensor Cores) are crucial for hosting embedding generators and generation nodes simultaneously. Rapid NVMe SSD configurations are strictly necessary to minimize I/O latency upon querying vast semantic graphs. Beyond raw processing power, the NSAA framework benefits significantly from hardware capable of handling heterogeneous computing tasks. Neural embedding generation and inference rely heavily on GPU acceleration due to the parallelizable nature of tensor operations within Transformer architectures. GPUs equipped with Tensor Cores provide efficient matrix multiplication capabilities required for deep learning workloads such as attention mechanisms and high-dimensional vector encoding. At the same time, symbolic reasoning operations are often CPU-intensive rather than GPU-intensive. Rule evaluation, logical inference, and knowledge graph traversal require strong multi-threaded CPU performance and large system memory to efficiently process symbolic structures. High- core-count processors enable concurrent evaluation of logical predicates and rule chains within the symbolic reasoning engine, ensuring minimal latency when verifying knowledge retrieved by the neural subsystem. Additionally, memory bandwidth and storage performance play a critical role in large-scale knowledge retrieval scenarios. Vector databases and graph-based knowledge bases may contain millions of embeddings and relational nodes. Fast NVMe storage allows rapid indexing, retrieval, and caching of embedding vectors and graph structures. In large deployments, distributed storage architectures or memory-mapped databases may further enhance system scalability by enabling parallel access to semantic data structures. Network infrastructure also becomes relevant when deploying the architecture across distributed systems. High- bandwidth interconnects allow agents, neural modules, and symbolic engines to communicate efficiently across nodes in cluster environments, thereby enabling scalable multi-agent coordination and parallel retrieval operations.

2. Software Stack: At the orchestrator level, Python serves as the unifying medium. Neural layers utilize PyTorch or Tensor Flow environments and fetch localized topological models via Hugging Face. The symbolic reasoning stratum is orchestrated through SWI-Prolog interconnected via the PySwip bridging library. Large-scale structural storage operates utilizing graph database solutions such as Neo4j, while FAISS or ChromaDB facilitates the localized execution of rapid similarity profiling for the Vector RAG segment. The software ecosystem supporting the NSAA frame work is designed around modular interoperability. Python functions as the primary orchestration language due to its extensive ecosystem of machine learning, data processing, and agent- based development libraries. Through Python-based controllers, the system coordinates interactions between neural embedding models, symbolic reasoning engines, and vector retrieval modules. Within the neural computation layer, frameworks such as PyTorch and Tensor Flow provide optimized environments for constructing and deploying deep learning models. These frameworks enable efficient GPU utilization and support a wide range of pre-trained transformer architectures. Model repositories such as Hugging Face further streamline development by providing access to pre-trained language and multimodal models capable of generating high-quality embeddings for semantic retrieval tasks. The symbolic reasoning layer leverages SWI-Prolog as a logical inference engine. Prolog-based systems are well suited for rule-based reasoning due to their ability to represent knowledge through predicates, perform backtracking search, and execute deterministic logical inference. By integrating Prolog with Python through the PySwip interface, the NSAA framework enables neural outputs to be dynamically validated against symbolic rules. This bridge allows the system to translate vector-based semantic results into structured symbolic representations suitable for logical verification. For structured knowledge representation, graph database systems such as Neo4j provide efficient storage and querying of complex relational structures. Graph databases are particularly suitable for representing ontologies, knowledge graphs, and entity relationships, enabling rapid traversal of interconnected data nodes. These structures complement vector-based retrieval by providing explicit semantic relationships that support symbolic reasoning tasks. The Vector RAG component relies on similarity search libraries such as FAISS or ChromaDB to perform high-speed nearest-neighbor retrieval in high-dimensional embedding spaces. These libraries index embedding vectors and allow the system to quickly identify semantically relevant documents, knowledge fragments, or contextual references. By combining vector retrieval with symbolic reasoning verification, the architecture ensures that retrieved information is both semantically relevant and logically consistent. Together, this integrated hardware and software ecosystem forms the operational backbone of the NSAA architecture. The synergy between high-performance computational hardware and modular software frameworks enables the system to efficiently process multimodal inputs, retrieve contextual knowledge, perform logical reasoning, and orchestrate agent-driven decision processes within a unified intelligent pipeline.

V. IMPLEMENTATION AND RESULTS

A. System Modules and Data Flow

The operational workflow of the proposed Neuro-Symbolic Agentic AI (NSAA) frame work is structured as a modular pipeline in which information flows through sequential yet interdependent computational components. Each module performs a specialized function while maintaining bidirectional communication with other modules, enabling the system to dynamically transition between perception, reasoning, integration, and decision-making.

IRJCS: International Research Journal of Computer Science ISSN: 2393-9842
Volume 13, Issue 03, March 2026 https://www.irjcs.com/archives
https://doi.org/10.26562/irjcs.2025.v1303.09
The modular organization not only enhances scalability and interpretability but also ensures that errors in one layer can be detected and corrected through feedback mechanisms present in subsequent layers.

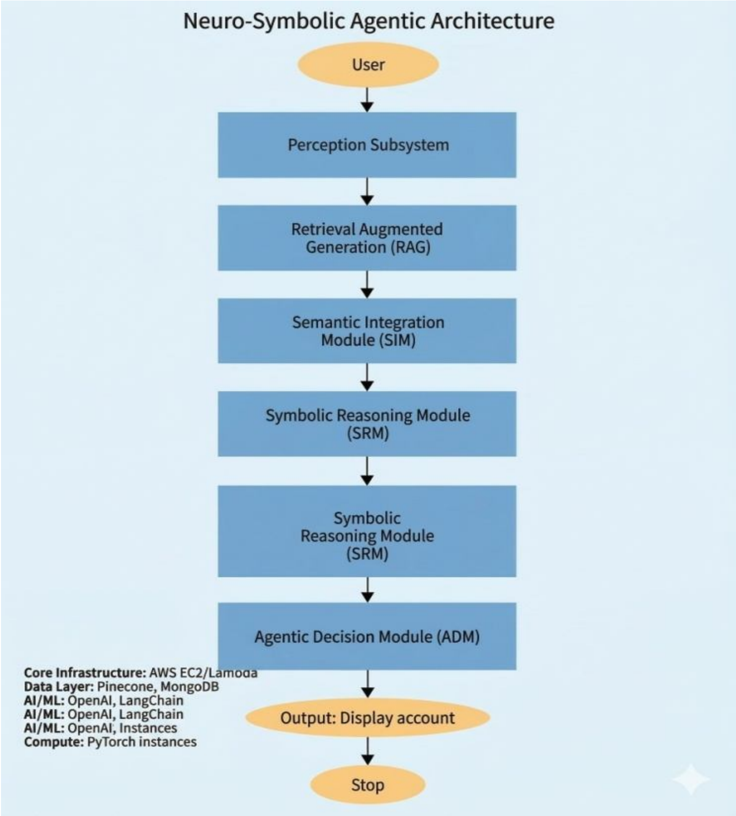

Fig 2: NSAAI of System Implementation

The data flow begins with the ingestion of raw or unstructured input signals, which may originate from textual queries, visual inputs, structured datasets, or multimodal sensory streams. These inputs are progressively transformed into increasingly structured forms as they propagate through the pipeline. Neural representations capture semantic relationships, symbolic reasoning validates logical constraints, and agent-based controllers coordinate the interaction between the two domains. Through this staged transformation process, the system moves from raw perception toward logically validated semantic conclusions.

1. Neural Perception Module (NPM): Dissects unstructured stimulus via distinct modal encoders into normalized latent vectors, computing confidence scores alongside entity bounding predictions. The Neural Perception Module operates as the primary interface between the external data environment and the internal reasoning architecture. Within this module, specialized neural encoders analyze incoming inputs and convert them into dense vector representations that capture semantic and contextual relationships. For textual inputs, transformer-based language encoders extract contextual embeddings that represent syntactic and semantic dependencies across tokens. In multimodal scenarios, vision encoders or graph-based models may be used to extract spatial, relational, or structural features from non-textual inputs. These embeddings are normalized with in a high-dimensional latent space, ensuring that semantically similar entities occupy proximate vector regions. Alongside vector generation, the module computes confidence metrics that estimate the reliability of detected entities, contextual relations, and extracted semantic patterns. Entity recognition models further identify relevant concepts or objects within the input stream, producing structured meta data such as entity boundaries, semantic roles, or relational links. These outputs collectively form a semantic representation layer that is subsequently passed to downstream modules for validation and reasoning.

2. Symbolic Reasoning Module (SRM): Encodes logical invariants. When cross-questioned by the Agent interface, the SRM parses an active memory cluster, firing forward/backward heuristics, yielding definitive logical confirmation or falsification without semantic bleed. The Symbolic Reasoning Module represents the deterministic reasoning core of the architecture. Unlike neural models that rely on probabilistic approximations, this module operates on formally defined logical rules and structured knowledge representations. Logical invariants, ontological hierarchies, and rule-based constraints are encoded within a knowledge base, allowing the system to perform rigorous inference operations. When invoked by the Agentic Decision Module, the SRM analyzes relevant knowledge segments stored within the active reasoning context. Using mechanisms such as forward chaining and backward chaining, the module evaluates logical implications between predicates and determines whether proposed hypotheses are consistent with the encoded rule set. The reasoning process proceeds through explicit rule evaluation steps, ensuring that conclusions derived by the system remain logically coherent and traceable. A key advantage of the SRM is its ability to eliminate semantic ambiguity by enforcing discrete logical boundaries. Because symbolic reasoning operates on structured predicates rather than probabilistic embeddings, the system can definitively confirm or reject hypotheses without introducing semantic drift.

This capability is particularly critical when resolving contradictions between retrieved contextual knowledge and previously established logical constraints.

3. Semantic Integration Module (SIM): This proprietary interface maps continuous continuous valued token embeddings exclusively to discrete symbolic ontology nodes. Inversely, it operationalizes structural nodes back into synthetic multi-dimensional canonical vectors for associative retrieval. The Semantic Integration Module functions as the translation bridge between the neural and symbolic components of the architecture. Neural embeddings and symbolic predicates operate under fundamentally different representational paradigms; embeddings encode meaning in continuous vector spaces, whereas symbolic systems rely on discrete logical structures. The SIM reconciles these differences by establishing a bidirectional mapping mechanism. In the forward direction, the module analyzes neural embeddings generated by the perception layer and aligns them with corresponding ontology nodes within the symbolic knowledge base. This mapping process involves similarity matching, entity linking, and contextual disambiguation procedures that associate vector representations with semantically equivalent symbolic entities. Through this process, continuous semantic signals are transformed into structured symbolic representations suitable for logical reasoning. In the reverse direction, the SIM enables symbolic knowledge structures to be converted back into vector representations that can be utilized by neural models for associative retrieval tasks. By embedding ontology nodes into vector space, the system can integrate symbolic knowledge directly into vector similarity search processes. This bidirectional transformation capability allows neural perception and symbolic reasoning to operate cohesively rather than as isolated computational domains.

4. Agentic Decision Module (ADM): Regulates the execution loop. Leveraging Deep Reinforcement paradigms, the ADM parses terminal objectives, instigates targeted calls to the SIM for embedding alignments, interrogates the SRM when rule conflict manifests, and triggers the generation modules conditionally. The Agentic Decision Module functions as the executive control center of the NSAA frame work. Rather than allowing the system to operate as a static inference pipeline, the ADM introduces autonomous decision-making capabilities through agent-based control mechanisms. Using reinforcement learning strategies and planning algorithms, the ADM continuously evaluates the system's state and determines the optimal sequence of actions required to achieve the desired objective. When at ask is initiated, the ADM decomposes the high-level objective into smaller operational subtasks. These subtasks may involve retrieving contextual knowledge, validating logical relationships, resolving inconsistencies, or synthesizing final outputs. Depending on the nature of the problem, the ADM dynamically determines whether neural perception, symbolic reasoning, or vector retrieval should be invoked. During execution, the ADM monitors intermediate results and detects potential conflicts or uncertainties. If contradictions arise between neural outputs and symbolic rules, the ADM initiates additional verification cycles through the SRM. If contextual knowledge is in sufficient, the ADM triggers retrieval procedures through the vector-based knowledge store. Through this adaptive decision loop, the module ensures that the system remains both context-aware and logically consistent throughout the reasoning process.

5. Feedback and Adaptation Module (FAM): Sustains autonomous alignment. Logical corrections provided by the Symbolic sector against hallucinated Neural derivations automatically spawn verified training tuples, a synchronously fine tuning the NPM's weights to avert compounding errors. The Feedback and Adaptation Module ensures long-term system robustness by enabling continuous learning and error correction. When discrepancies arise between neural predictions and symbolic verification results, the system records these conflicts as structured feedback signals. These signals form validated training tuples that capture both the erroneous sprediction and the corrected logical interpretation. Over time, these tuples are used to refine the neural perception models through incremental fine-tuning procedures. This adaptive learning mechanism gradually reduces the frequency of hallucinated outputs by reinforcing neural representations that align with verified symbolic knowledge. Importantly, this training occurs asynchronously, meaning that the system can continue operating while updates are applied in the background. By integrating symbolic feedback directly into neural learning processes, the NSAA framework establishes a self- correcting intelligence loop. Neural perception improves through experience, symbolic reasoning maintains logical integrity, and agentic orchestration ensures that both components interact effectively within evolving knowledge environments.

b. RAG Implementation

Vector embedding clusters strictly regulate external context anchoring. Upon algorithmic determination that internal neural and symbolic records reflect knowledge deficiency, the ADM triggers the semantic search sub-routine against the Vector store. Utilizing similarity profiling models, relevant chunks are topologically recalled and injected dynamically into a Context Filtering Layer which scrutinizes the payload under the symbolic engine before releasing it into the text synthesis model. The Retrieval-Augmented Generation (RAG) mechanism serves as the external knowledge acquisition component of the NSAA architecture. Rather than relying solely on the knowledge encoded within neural model parameters, the system dynamically retrieves relevant information from external data repositories whenever internal reasoning resources prove insufficient. When the Agentic Decision Module detects a knowledge gap, it initiates a semantic search procedure within the vector database. Embedding vectors representing the current query or reasoning context are compared against stored document embeddings using similarity metrics such as cosine similarity or Euclidean distance. This process identifies semantically related knowledge fragments that may assist in resolving the task at hand. Retrieved content is then routed through a Context Filtering Layer that performs validation and relevance assessment. Within this stage, symbolic reasoning mechanisms examine the retrieved information to ensure that it does not violate established logical constraints or ontological relationships. This verification step prevents the system from incorporating misleading or contradictory external knowledge into the reasoning pipeline. Once validated, the contextual information is integrated into the generation environment where it supports downstream synthesis modules.

ISSN: 2393-9842
https://www.irjcs.com/archives
https://doi.org/10.26562/irjcs.2025.v1303.09

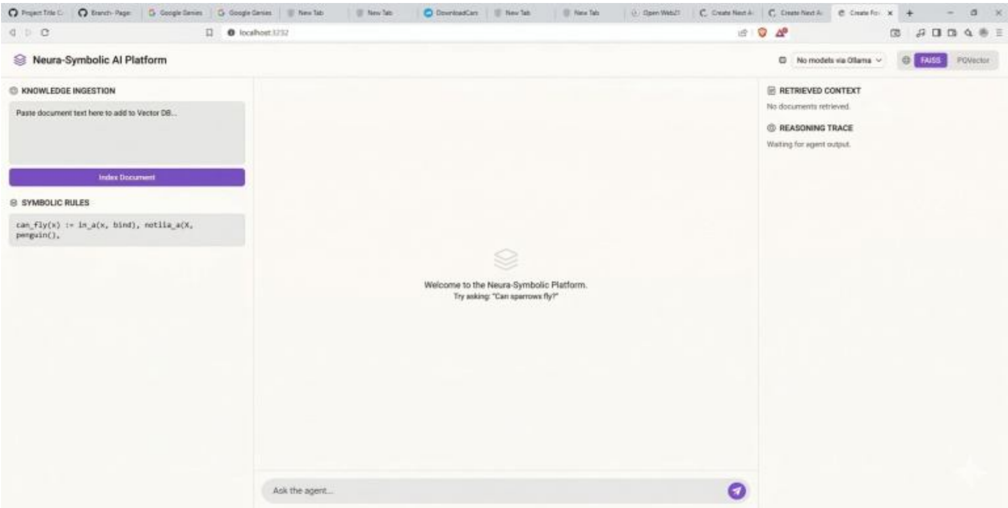

Fig 3: Interface of NSAAI

By grounding the generation process in externally retrieved knowledge, the RAG system significantly enhances factual accuracy and contextual relevance. This architecture ensures that generated outputs are informed by both dynamically retrieved evidence and logically validated reasoning structures.

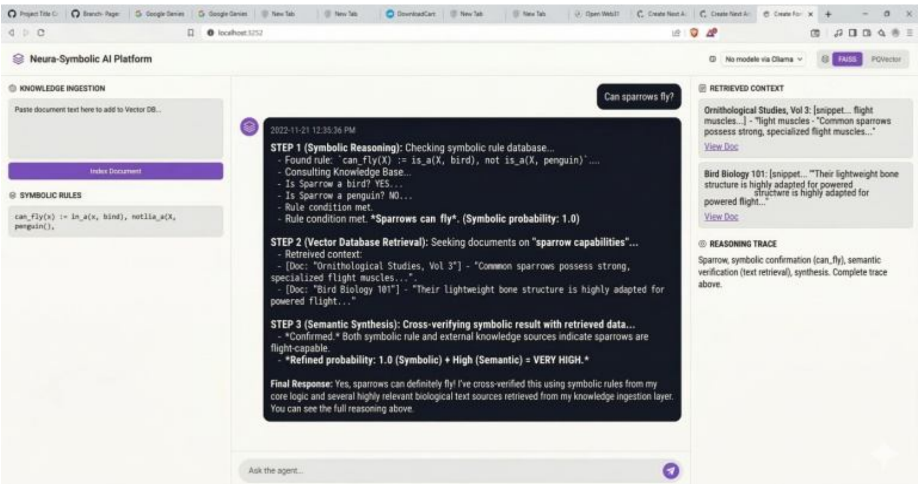

Fig 4: Chat Interface

Through the integration of vector retrieval, symbolic verification, and agent-based orchestration, the RAG subsystem becomes a tightly coupled component of the broader NSAA framework. It enables the system to continuously expand its accessible knowledge while maintaining logical rigor and semantic consistency throughout the reasoning process.

VI. CONCLUSION

This research systematically delineates an architectural resolution bridging the chasm between statistical correlation paradigms and formalized machine reasoning. The inherent algorithmic flaws limiting pure connectionist systems including severe susceptibility to hallucinations, logical ambiguity, and operational capacity are structurally annulled when neural processes are strictly tethered to the deterministic constraints of symbolic computation. The Neura-Symbolic Agentic Artificial Intelligence(NSAA) framework, supplemented by Vector Retrieval-Augmented Generation matrices, succeeds in engineering an ecosystem capable of authentic autonomous deliberation. By distributing cognitive loads across localized modules where neural engines handle abstraction while explicit ontologies manage validation—this framework fosters uninterrupted, continuous learning absent of destructive interference. Conclusively, this methodology constitutes an evolutionary pivot towards highly stable, explicitly explainable, and contextually grounded intellectual synthesis, serving as a comprehensive blueprint for mission- critical and next-generation associative intelligence models.

REFERENCES

1. B.Colelough and W.Regli,"Neuro-SymbolicAIin2024: A Systematic Review,"CEUR Workshop Proceedings,2025.
2. O.Bougzime, S.Jabbar, C.Cruz and F.Demoly, "Unlocking the Potential of Generative AI through Neuro-Symbolic Architectures: Benefits and Limitations," arXiv preprint arXiv:2502.11269, 2025.
3. G.A.Prenosil, M.Huber, M.J.M. fohl et al., "Neuro-Symbolic AI for Auditable Cognitive Information Extraction from Medical Reports," Communications Medicine, Nature Portfolio, 2025.
4. R.J.Tong and X.Hu,"Future of Education with Neuro-Symbolic AI Agents in Self-Improving Adaptive Instructional Systems, "Frontiers of Digital Education,vol.1, no. 1, 2024.