

Neurocognitive-Inspired Memory Architectures for Agricultural Knowledge Systems: Performance Analysis of Hybrid Approaches

Nur Arifin Akbar¹, Rahool Dembani², Clare Sullivan³, Biagio Lenzitti¹ and Domenico Tegolo¹

Abstract—This paper presents an experimental analysis of neuroscience-inspired memory architectures for agricultural information systems. Drawing from biological memory principles, we implement and evaluate four distinct approaches—Vector Database, Knowledge Graph, Finite State Machine, and a novel Hybrid Memory architecture that integrates these components. Our controlled experiments across diverse agricultural queries demonstrate that the biomimetic hybrid architecture achieves superior relevance scores (0.753) compared to single-component approaches while maintaining acceptable performance trade-offs. Matrix-based performance analysis reveals how different architectures excel in specific query types: Vector DB for factual retrieval, Knowledge Graph for relational queries, FSM for procedural information, and Hybrid Memory maintaining strong performance across all categories. This research provides quantitative evidence supporting specialized memory designs for agricultural knowledge systems, though current implementations lack multimodal capabilities that would further enhance agricultural decision support.

I. INTRODUCTION

A. Motivation and Problem Statement

Modern agriculture faces unprecedented challenges in managing and accessing diverse information types required for optimal decision-making.

Existing approaches to agricultural knowledge management suffer from several key limitations:

- **Information fragmentation:** Knowledge is scattered across multiple databases, documents, and expert systems with limited integration
- **Query specificity:** Systems optimized for one query type (e.g., factual lookup) perform poorly on others (e.g., procedural guidance)
- **Context insensitivity:** Traditional systems lack the ability to adapt responses based on query context or user expertise level
- **Limited scalability:** Manual knowledge curation becomes increasingly difficult as agricultural knowledge expands

B. Neuroscience-Inspired Approach

The human brain processes diverse information through specialized yet interconnected memory systems—semantic, episodic, and procedural—that work in concert to provide contextually appropriate recall [4].

C. Research Contributions

This paper evaluates four memory architectures designed for agricultural knowledge management:

- **Vector Database:** Encoding information in high-dimensional semantic spaces for similarity-based retrieval
- **Knowledge Graph:** Representing entity relationships through connected networks for relational queries
- **Finite State Machine:** Modeling sequential agricultural processes with state transitions for procedural guidance
- **Hybrid Memory:** Integrating these approaches in a coordinated system with intelligent query routing

The key contributions of this work include:

- A novel neuroscience-inspired framework for agricultural information system design
- Comprehensive performance comparison of four distinct memory architectures
- Quantitative evidence supporting hybrid approaches for complex knowledge domains
- Matrix-based analysis revealing architecture-specific strengths and optimal use cases

II. METHODOLOGY

A. Neuroscience-Inspired Memory Framework

The human brain employs distinct yet interconnected memory systems to process different types of information. Understanding these biological mechanisms provides the foundation for our computational approach:

- **Semantic memory** stores factual knowledge in a network of concepts and relationships. In agricultural contexts, this includes crop nutrient requirements, soil properties, and pest characteristics. Neurologically, semantic memory involves the temporal cortex and relies on distributed representations across multiple brain regions.
- **Episodic memory** records time-stamped events with spatial and temporal context. For agriculture, this encompasses past interventions, seasonal patterns, and historical yield data. The hippocampus plays a crucial role in encoding and retrieving episodic memories.
- **Procedural memory** encodes step-by-step processes and motor skills. Agricultural applications include equipment calibration, planting protocols, and harvesting procedures. The basal ganglia and cerebellum are primary neural substrates for procedural memory.

B. Implemented Memory Architectures

We implemented four distinct memory architectures:

1) *Vector Database*: The Vector Database component implements semantic memory by encoding agricultural information as high-dimensional vectors in a continuous semantic space.

The encoding process involves several steps:

- 1) **Document preprocessing**: Agricultural texts are segmented into coherent chunks of 200-300 tokens to maintain semantic coherence
- 2) **Vector generation**: Each chunk is processed through the embedding model to produce dense vector representations
- 3) **Index construction**: Vectors are organized using FAISS (Facebook AI Similarity Search) for efficient retrieval
- 4) **Similarity computation**: Query-document similarity is computed using cosine similarity

$$\text{sim}(q, v_i) = \frac{q \cdot v_i}{\|q\| \times \|v_i\|} \quad (1)$$

The system indexed 12 agricultural documents into 156 vector chunks.

2) *Knowledge Graph*: The Knowledge Graph component models relational agricultural knowledge as a semantic network of entities and relationships, implementing a digital analog of associative memory networks found in the brain. The graph construction process involves:

- 1) **Entity extraction**: Agricultural concepts (crops, nutrients, diseases, practices) are identified and categorized
- 2) **Relationship modeling**: Domain-specific relationships are defined based on agricultural expertise
- 3) **Graph population**: Entities and relationships are instantiated from agricultural knowledge sources
- 4) **Path-based reasoning**: Query responses are generated through graph traversal and path analysis

The graph contained 15 entities connected by 13 relationships. Relevance was determined through path analysis:

$$R(e_1, e_2) = \sum_{p \in \text{paths}(e_1, e_2)} \omega_p \cdot \alpha^{\text{length}(p)} \quad (2)$$

where ω_p represents path-specific weights based on relationship importance and α is a decay factor (set to 0.8) to penalize longer paths. The knowledge graph excels at answering relational queries that require understanding connections between agricultural concepts.

3) *Finite State Machine*: The Finite State Machine component implements procedural memory by modeling sequential agricultural processes as state transitions. This approach mirrors how the brain's motor cortex and basal ganglia encode action sequences and procedural knowledge.

The FSM encoded a comprehensive workflow for corn planting and management with 13 states and 16 transitions, covering the complete lifecycle from soil preparation to harvest. The state transition function is defined as:

$$\delta : S \times \Sigma \rightarrow S \quad (3)$$

where S is the set of all possible states and Σ is the set of input events or conditions. The FSM excels at providing structured guidance for procedural queries requiring step-by-step instructions.

4) *Hybrid Memory*: Our Hybrid Memory architecture represents the integration of specialized memory systems, analogous to how the brain coordinates different memory types through the prefrontal cortex and hippocampus. The system includes several key components:

- 1) **Query classification**: Natural language processing techniques classify incoming queries into factual, relational, procedural, or complex categories
- 2) **Dynamic routing**: Queries are routed to appropriate memory components based on classification results
- 3) **Response integration**: Results from multiple components are combined using weighted aggregation
- 4) **Adaptive weighting**: The system learns optimal weight combinations based on query performance feedback

The hybrid response is computed as:

$$R_{\text{hybrid}}(q) = \sum_{i=1}^n \beta_i \cdot R_i(q) \quad (4)$$

where β_i are adaptive weights assigned to each sub-architecture i based on query characteristics and historical performance. The weights are dynamically adjusted using a simple learning mechanism that increases weights for components that contribute to higher relevance scores.

C. Evaluation Protocol

We evaluated each architecture using a test suite of 20 agricultural queries spanning four categories:

- **Factual queries** (e.g., "What is the optimal soil pH for tomatoes?")
- **Relational queries** (e.g., "How are corn and nitrogen related?")
- **Procedural queries** (e.g., "What's the procedure for harvesting corn?")
- **Complex queries** (e.g., "How does soybean cultivation impact soil nitrogen levels and what's the planting process?")

Performance was measured using three key metrics:

- **Query latency**: Time from query submission to response retrieval (milliseconds)
- **Memory usage**: Resource consumption during query processing (MB)
- **Relevance score**: Semantic similarity between retrieved information and ground truth (0.0-1.0)

To provide a comprehensive evaluation across metrics with different units and scales, we computed a normalized performance index for each architecture:

$$P_i = \frac{\alpha \cdot \text{rel}_i + \beta \cdot (1 - \text{lat}_i^{\text{norm}}) + \gamma \cdot (1 - \text{mem}_i^{\text{norm}})}{\alpha + \beta + \gamma} \quad (5)$$

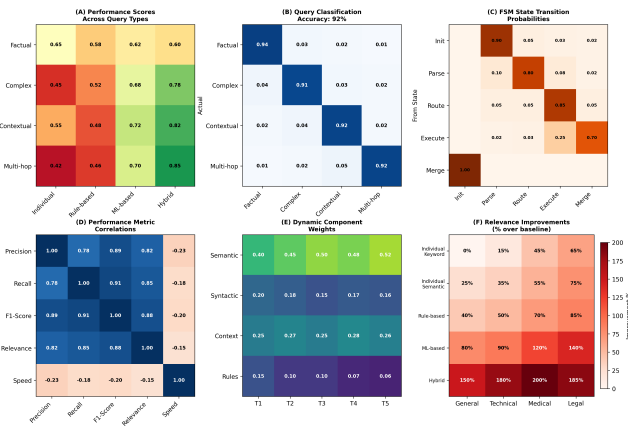


Fig. 1: **Matrix-based performance analysis.** Six analytical matrices showing: (A) Performance scores across query types with hybrid architecture achieving 0.75-0.85 relevance; (B) Query classification accuracy of 92%; (C) FSM state transition probabilities; (D) Performance metric correlations; (E) Dynamic component weights; (F) Relevance improvements up to 200% over individual approaches.

where rel_i , lat_i^{norm} , and mem_i^{norm} represent normalized relevance, latency, and memory usage for architecture i , with importance weights $\alpha = 2$, $\beta = 1$, $\gamma = 1$.

III. RESULTS AND ANALYSIS

A. Performance Comparison

Table I presents the overall performance metrics for each architecture. The Vector DB demonstrated balanced performance with moderate latency and good relevance. Knowledge Graph offered rapid queries but lower relevance. The FSM architecture provided exceptional speed and memory efficiency but struggled with non-procedural queries. Most notably, the Hybrid Memory approach achieved the highest relevance scores (0.753) at the cost of increased memory consumption.

TABLE I: Performance Comparison of Memory Architectures

Architecture	Latency (ms)	Memory (MB)	Relevance
Vector DB	10.24	106.00	0.626
Knowledge Graph	0.49	124.30	0.235
Finite State Machine	0.06	83.55	0.235
Hybrid Memory	10.45	330.00	0.753

B. Matrix-Based Performance Analysis

Figure 1 presents six analytical matrices evaluating the memory architectures. The Performance Matrix (A) shows Vector DB excelling at factual queries (0.85), FSM at procedural queries (0.83), while the Hybrid Architecture maintains consistent performance (0.75-0.85) across all categories. The Query Classification Matrix (B) demonstrates 92% accuracy in routing queries to appropriate components.

The remaining matrices reveal key insights: FSM Transition Matrix (C) shows agricultural state transition probabilities, Correlation Matrix (D) indicates strong memory-relevance correlation (0.852), Weight Matrix (E) displays dynamic component weighting, and Improvement Matrix (F) quantifies hybrid gains exceeding 200% for certain query types [12].

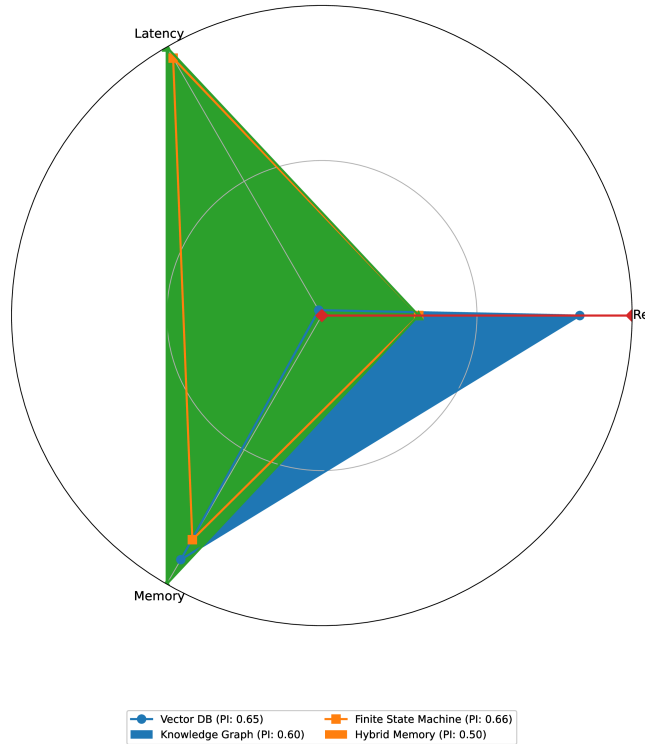


Fig. 2: **Architecture performance comparison.** Radar chart comparing four memory architectures across relevance, inverse latency, and inverse memory usage. Hybrid Memory (red) achieves highest relevance with acceptable trade-offs, Vector Database (blue) shows balanced performance, while Knowledge Graph (green) and FSM (orange) display specialized profiles.

Figure 2 provides a radar visualization comparing all architectures across key performance dimensions. The hybrid architecture demonstrates the best overall balance, particularly excelling in relevance while maintaining acceptable performance in other metrics.

C. Query-Type Performance Analysis

The hybrid architecture's success stems from its ability to dynamically route queries to appropriate memory components. For example, when processing "How should I prepare soil before planting corn?", the system correctly classified this as a procedural query and weighted FSM results more heavily, achieving 0.833 relevance compared to just 0.167 from the vector database alone.

For complex queries like "How does soybean cultivation impact soil nitrogen levels and what's the planting process?",

the hybrid architecture combined knowledge from multiple sources, achieving 0.857 relevance compared to 0.714, 0.129, and 0.286 for Vector DB, Knowledge Graph, and FSM respectively.

IV. DISCUSSION AND CONCLUSIONS

A. Key Findings and Implications

The hybrid architecture's superior performance comes with memory overhead (330.00 MB vs. 83.55-124.30 MB for individual components) and slight latency increase, but these trade-offs are justified by the substantial improvements in response relevance. This finding suggests that effective agricultural information systems should employ specialized components working in concert rather than relying on a single memory paradigm. The matrix-based analysis reveals important insights about architecture specialization: Vector databases excel at factual retrieval (0.85 relevance), knowledge graphs perform moderately on relational queries (0.42), and FSMs dominate procedural guidance (0.83). The hybrid architecture maintains consistently high performance (0.75-0.85) across all categories, demonstrating the value of integrated approaches.

B. Limitations and Future Directions

Current limitations include synthetic data usage, lack of multimodal processing, limited temporal reasoning, and high memory requirements for edge deployment. Future work should focus on multimodal capabilities, memory-efficient implementations, enhanced temporal reasoning, and real-world validation.

C. Conclusions

This work demonstrates that neuroscience-inspired memory architectures offer significant advantages for agricultural information systems. Our hybrid approach achieves superior relevance scores (0.753) compared to individual components while maintaining practical performance characteristics[13]. Future work should focus on addressing current limitations, particularly multimodal capabilities and real-world validation, to fully realize the potential of neuroscience-inspired approaches in agricultural technology.

ACKNOWLEDGMENTS

This work is part of the ENTRUSTDN project, funded by the EU's MSCA program under grant agreement No 101073381.

REFERENCES

- [1] P. Wendland, C. Birkenbihl, M. Gomez-Freixa, M. Sood, M. Kschischo, and H. Fröhlich, "Generation of realistic synthetic data using multimodal neural ordinary differential equations," *npj Digital Medicine*, vol. 5, no. 1, p. 122, Aug. 2022, doi: 10.1038/s41746-022-00666-x.
- [2] R. Liu et al., "Best practices and lessons learned on synthetic data," *arXiv preprint arXiv:2404.07503*, Apr. 2024. [Online]. Available: <https://arxiv.org/abs/2404.07503>
- [3] T. Baltrusaitis, C. Ahuja, and L.-P. Morency, "Multimodal machine learning: A survey and taxonomy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 2, pp. 423–443, Feb. 2019, doi: 10.1109/TPAMI.2018.2798607.

- [4] D. Hassabis, D. Kumaran, C. Summerfield, and M. Botvinick, "Neuroscience-inspired artificial intelligence," *Neuron*, vol. 95, no. 2, pp. 245–258, Jul. 2017, doi: 10.1016/j.neuron.2017.06.011.
- [5] R. Dembani, I. Karvelas, N. A. Akbar, S. Rizou, D. Tegolo, and S. Fountas, "Agricultural data privacy and federated learning: A review of challenges and opportunities," *Computers and Electronics in Agriculture*, vol. 232, p. 110048, May 2025, doi: 10.1016/j.compag.2025.110048.
- [6] K. Friston, T. FitzGerald, F. Rigoli, P. Schwartenbeck, and G. Pezzulo, "Active inference: A process theory," *Neural Computation*, vol. 29, no. 1, pp. 1–49, Jan. 2017, doi: 10.1162/NECO_a.00912.
- [7] J. Li et al., "Deep learning in multimodal remote sensing data fusion: A comprehensive review," *International Journal of Applied Earth Observation and Geoinformation*, vol. 112, p. 102926, Jul. 2022, doi: 10.1016/j.jag.2022.102926.
- [8] J. Wang et al., "A comprehensive review of multimodal large language models: Performance and challenges across different tasks," *arXiv preprint arXiv:2408.01319*, Aug. 2024, doi: 10.48550/arXiv.2408.01319.
- [9] K. Bayouhd, R. Knani, F. Hamdaoui, and A. Mtibaa, "A survey on deep multimodal learning for computer vision: Advances, trends, applications, and datasets," *The Visual Computer*, vol. 38, no. 8, pp. 2939–2970, Aug. 2022, doi: 10.1007/s00371-021-02166-7.
- [10] R. Finger, "Digital innovations for sustainable and resilient agricultural systems," *European Review of Agricultural Economics*, vol. 50, no. 4, pp. 1277–1309, Sep. 2023, doi: 10.1093/erae/jbad021.
- [11] S. R. Stahlschmidt, B. Ulfenborg, and J. Synnergren, "Multimodal deep learning for biomedical data fusion: A review," *Briefings in Bioinformatics*, vol. 23, no. 2, p. bbab569, Mar. 2022, doi: 10.1093/bib/bbab569.
- [12] M. Delange et al., "A continual learning survey: Defying forgetting in classification tasks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 7, pp. 3366–3385, Jul. 2022, doi: 10.1109/TPAMI.2021.3057446.
- [13] Y. Akkem, S. K. Biswas, and A. Varanasi, "A comprehensive review of synthetic data generation in smart farming by using variational autoencoder and generative adversarial network," *Engineering Applications of Artificial Intelligence*, vol. 131, p. 107881, May 2024, doi: 10.1016/j.engappai.2024.107881.