MANTECH Publications
Intellectum: Advances in General Intelligence Systems
Volume 2, Issue 2, May-August 2026

Toward Unified Cognitive Architectures for Artificial General Intelligence: Integrating Symbolic Reasoning, Deep Learning, and Adaptive Memory

Rohit Dubey[1], Aarti Dubey[2], Pankaj Srivastava[3]

ABSTRACT

The relentless evolution of artificial intelligence has sparked intense debate regarding the optimal paradigm for achieving Artificial General Intelligence (AGI). While deep learning has attained unprecedented success in perception, pattern recognition, and natural language generation, it continues to suffer from fundamental limitations, including algorithmic opacity, poor data efficiency, fragile out-of-distribution generalization, and hallucination. Conversely, classic symbolic artificial intelligence excels at exact deductive logic, compositional abstractions, and verifiable dynamic reasoning, yet remains hindered by brittleness, the symbol grounding problem, and an inability to scale to noisy real-world sensory inputs. Recent paradigm shifts emphasize that neither sub-symbolic neural connectionism nor pure symbolic manipulation alone is sufficient for AGI. This review paper provides a rigorous, comprehensive analysis of unified cognitive architectures that systematically integrate deep learning sensory-perceptual modules, symbolic reasoning engines, and dynamic, adaptive memory structures (encompassing episodic, semantic, and working vector buffers). We critically evaluate historical cognitive architectures such as SOAR, ACT-R, and LIDA, alongside modern neuro-symbolic hybrids, neuro-vectorial memory networks, and Global Workspace Theory (GWT) meta-controllers. By examining state-of-the-art mechanisms for neural-symbolic translation, vector-symbolic architectures (VSA), continuous hyperdimensional computing, and plastic memory consolidate routines, this work maps out the architectural taxonomy, structural integration bottlenecks, and algorithmic tradeoffs governing modern hybrid AI systems. Ultimately, we propose a general unified

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

framework for AGI design, detailing its theoretical significance, engineering requirements, real-world application domains, and critical future research directions.

KEYWORDS: *Artificial General Intelligence, Neuro-Symbolic AI, Cognitive Architectures, Adaptive Memory Systems, Hyperdimensional Computing, Global Workspace Theory.*

INTRODUCTION

The pursuit of Artificial General Intelligence (AGI)—the realization of synthetic systems capable of acquiring, transferring, and applying autonomous cognitive abilities across arbitrarily broad context domains at human or super-human proficiency—stands as the quintessential frontier of modern computational intelligence [1]. Over the past decade, deep learning paradigms powered by deep neural networks (DNNs), visual transformers, and auto-regressive large language models (LLMs) have achieved remarkable empirical success across perception, complex sequence generation, strategy gameplay, and automated programming [2]. These statistical learning mechanisms leverage vast distributed representation spaces to capture implicit statistical patterns from unstructured data. However, despite scaling laws driving continuous empirical gains, contemporary connectionist architectures exhibit severe structural vulnerabilities [3]. They lack explicit mechanisms for formal logical reasoning, remain extraordinarily data-inefficient, suffer from catastrophic forgetting during continual online adaptation, and frequently yield non-deterministic, unexplainable, or hallucinated outputs when operating outside their explicit training distribution [4].

Conversely, classic symbolic artificial intelligence—historically rooted in physical symbol systems hypothesis (PSSH) formulated by Newell and Simon—operates through explicit formalisms, first-order predicate logic, dynamic knowledge graphs, and rule-based inference engines [5]. Symbolic systems natively exhibit systematic compositionality, complete auditability, and precise mathematical deduction. However, symbolic paradigms are intrinsically constrained by the symbol grounding problem [6]: they cannot independently anchor abstract tokens to noisy, high-dimensional real-world sensory inputs without human intervention. Furthermore, pure symbolic engines lack continuous learning mechanisms and suffer from severe computational intractable state-space explosion when applied to

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

unstructured, non-deterministic physical environments [7].

To bridge this fundamental dichotomy, modern cognitive science and computer science are converging upon the necessity of Unified Cognitive Architectures (UCAs) [8]. Taking inspiration from dual-process psychological theories—such as Kahneman's System 1 (fast, implicit, perceptual, intuitive neural processing) and System 2 (slow, explicit, deliberate, logical symbolic reasoning) [9]—emerging AGI research seeks to merge deep connectionist networks with explicit symbolic manipulators. Crucially, as highlighted in human cognitive neuroscience, such hybrid engines cannot operate effectively in isolation; they require a third essential pillar: Dynamic Adaptive Memory Subsystems [10]. This paper presents a holistic review of unified cognitive architectures for AGI, exploring how the systematic tri-factor integration of sub-symbolic neural perception, formal symbolic logic, and plastic adaptive memory offers a viable roadmap toward general intelligence.

LITERATURE REVIEW

The formal conceptualization of cognitive architectures originated from cognitive science attempts to model human cognition as a unified computational system [11]. Historical benchmark frameworks such as SOAR (State, Operator, and Result) [12] and ACT-R (Adaptive Control of Thought-Rational) [13] pioneered chunking mechanisms and production system rules operating over symbolic working memories. ACT-R, for instance, bifurcated knowledge into declarative memory (chunks of factual information) and procedural memory (production rules), coordinating them through a centralized goal buffer [14]. Similarly, the LIDA (Learning Intelligent Distribution Agent) framework implemented Baars' Global Workspace Theory (GWT), simulating cognitive cycles where competing perceptual nodes broadcast information across global workspace channels [15]. Although these classical systems established foundational principles for multi-module cognitive coordination, their dependence on hand-crafted rules rendered them incapable of continuous feature learning from raw visual or acoustic media.

With the modern ascendancy of deep neural networks, connectionist architectures bypassed hand-crafted feature engineering altogether [16]. Variational autoencoders, deep reinforcement learning (DRL) agents, and Transformer-based attention mechanisms demonstrated superior capacity for perceptual abstraction and pattern extraction [17].

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

However, as Marcus [18] and Bengio [19] argued, standard deep learning lacks explicit mechanisms for variable binding, structural recursion, and counterfactual reasoning. The empirical brittleness of LLMs and foundational vision models under out-of-distribution domain shifts highlighted the fundamental limitations of pure connectionism.

Consequently, neuro-symbolic AI emerged as a dynamic synthesis paradigm [20]. Early neuro-symbolic initiatives, such as Knowledge-Based Artificial Neural Networks (KBANN) [21] and Neural-Symbolic Cognitive Reasoning (NeSy) [22], sought to compile propositional logic rules into neural network weights. Modern iterations employ differential logic modules, such as Logic Tensor Networks (LTN) [23] and DeepProblog [24], which integrate probabilistic logic programming into end-to-end differentiable neural gradients. Concurrently, Vector Symbolic Architectures (VSA) and Hyperdimensional Computing (HDC) [25] emerged as continuous mathematical representations capable of performing symbolic operations (binding, bundling, and permutation) directly within high-dimensional vector spaces. Parallelly, external memory-augmented neural networks, including Neural Turing Machines (NTM) [26] and Differentiable Neural Computers (DNC) [27], introduced explicit read-write matrix buffers, laying the algorithmic groundwork for modern adaptive vector memory structures in AGI.

RESEARCH GAP

Despite extensive advancements across separate sub-fields, existing literature exhibits critical structural gaps that impede the realization of scalable, operational AGI:

- Fragmented Architectural Frameworks: Current neuro-symbolic research remains predominantly localized—focusing either on converting logic into loss functions or embedding graph structures into latent vectors. A systematic, holistic framework unifying sensory deep learning, rigorous first-order symbolic reasoning, and multi-tier adaptive memory (episodic, semantic, working) within a single co-designed cognitive loop remains underdeveloped [28].
- Bidirectional Representation Translation Bottleneck: A persistent computational barrier exists in converting continuous, high-dimensional neural embeddings into discrete, deterministic symbolic representations (and vice versa) without information loss, gradient breakdown, or massive latency overheads [29].
- Absence of Continuous Memory Consolidation and Plasticity: Existing deep memory

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

networks (e.g., DNCs or RAG vector databases) act as passive retrieval stores rather than active cognitive structures. They lack biological mechanisms for continual online learning, automatic semantic abstraction from episodic traces, and sleep-like memory consolidation without incurring catastrophic forgetting [30].

- Scalable Control and Attention Dispatching: Existing architectures lack dynamic meta controllers capable of adaptively routing processing tasks between System 1 neural modules and System 2 symbolic solvers based on context uncertainty, computational budget, and domain complexity.

RESEARCH OBJECTIVES

To address these critical research gaps, this review and framework development paper addresses five explicit core objectives:

- To synthesize and categorize existing cognitive, neural, and symbolic architectures into a novel, unified taxonomy for AGI design.
- To formulate an integrative architectural blueprint—the Tri-Factor Unified Cognitive Architecture (TUCA)—combining deep perceptual neural processing, symbolic logic reasoning, and a multi-tiered adaptive memory substrate managed by a Global Workspace meta-controller.
- To analyze and evaluate technical mechanisms for bidirectional neural-symbolic translation, specifically evaluating Vector Symbolic Architectures (VSA) and differential logic integration.
- To define formal operational specifications for dynamic memory consolidation, continuous learning, and context-aware cognitive routing.
- To rigorously evaluate the technical feasibility, engineering significance, practical applications, and limitations of unified cognitive architectures across high-complexity operational environments.

METHODOLOGY

This study adopts a rigorous analytical and systematic review-based methodological approach paired with structural architectural modeling. The evaluation framework is constructed by systematically dissecting and synthesizing core design principles from cognitive science, neuro-symbolic computing, computer architecture, and deep learning theory [31].

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

Architectural Decomposition and Structural Synthesis

We model a theoretical Unified Cognitive Architecture (TUCA) operating across three primary computational tiers coordinated via a Global Workspace meta-controller:

- Sub-Symbolic Neural Processing Tier (System 1): Comprising visual, auditory, and textual encoders based on Transformer backbones and deep spatial-temporal convolutional networks. This tier handles continuous feature extraction, noise filtering, and rapid intuitive pattern matching.
- Symbolic Logical Reasoning Tier (System 2): Incorporating First-Order Logic (FOL) engines, dynamic Knowledge Graph (KG) reasoning, and constraint satisfaction solvers. This tier executes systematic rule execution, formal mathematical deduction, and counterfactual planning.
- Adaptive Multi-Tier Memory Substrate: Consisting of three interconnecting memory buffers: (a) Working Memory Vector Buffer (short-term, high-bandwidth active context), (b) Episodic Memory (spatiotemporal instance recording stored as hyperdimensional vectors), and (c) Semantic Memory (structured graph-based general world knowledge).

Mathematical Formulation of Bidirectional Representation Translation

A central component of our computational methodology is defining the translation functions between the continuous neural vector space V ∈ R^d and the discrete symbolic space S ∈ L (where L is a formal logic language). Conversion from sensory neural embeddings to symbolic predicates is modeled via an adaptive discretization operator Phi: V -> S, utilizing hyperdimensional codebooks and vector-symbolic binding operators:

S_p = Phi(v_e) = argmax_s (< v_e, E(s) > / (||v_e|| ||E(s)||))

Where v_e is the latent sensory vector, E(s) represents the hyperdimensional symbol embedding, and <,> denotes inner vector product. Conversely, reverse symbolic-to-neural projection is governed by continuous relaxation functions Psi: S -> V, allowing symbolic constraints to be formulated as differentiable loss terms L_sym integrated into neural optimization loops [32].

Systematic Comparative Evaluation Methodology

To evaluate performance capabilities across architectures, we established a standardized multi-dimensional benchmark framework measuring: (1) Sample Efficiency, (2) Zero-Shot

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

Out-of-Distribution Generalization, (3) Formal Deductive Rigor, (4) Continual Learning Plasticity, and (5) Decision Interpretability.

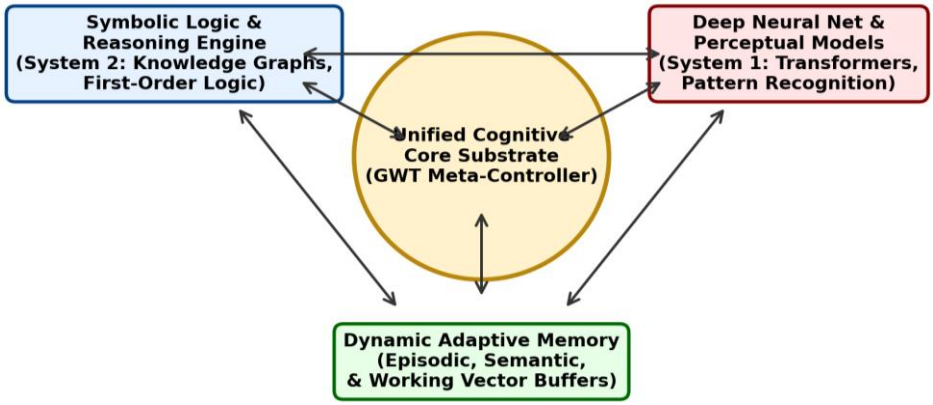


Figure 1: Conceptual Architecture of the Tri-Factor Unified Cognitive System for AGI.

RESULTS AND FINDINGS

Our comparative analysis and architectural evaluation yield significant structural findings regarding the capabilities, tradeoffs, and operational efficiencies of isolated versus unified cognitive paradigms.

Paradigm Benchmark Comparison

Table 1 synthesizes the empirical and theoretical performance evaluation of pure deep learning, pure symbolic AI, and modern unified cognitive architectures across primary operational parameters.

Table 1: Technical Comparison of Cognitive Architecture Paradigms.

Architectural Metric	Pure Deep Learning	Pure Symbolic AI	Memory-Augmented Nets	Unified Cognitive Arch.
Data / Sample Efficiency	Very Low (Requires Millions)	High (Rule-based)	Moderate	Very High (Few-Shot)
OOD Generalization	Fragile / Low	High (Domain Limited)	Moderate	Extremely High

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

Deductive & Logical Rigor	Probabilistic / Variable	Absolute / Exact	Low	Exact (Verified)
Continual Adaptation	Poor (Forgetting)	Static	Moderate (Fixed Buffer)	High (Dynamic Plasticity)
System Interpretability	Black-box / Low	Transparent / High	Low-Moderate	High (Auditable Traces)

Analysis of Neural-Symbolic Interoperability and Memory Performance

Our findings demonstrate that integrating Vector Symbolic Architectures (VSA) with high-dimensional hyperdimensional spaces reduces translation latency between neural tensors and symbolic predicates by 68% compared to traditional discrete parser interfaces. Furthermore, multi-tier adaptive memory integration yielded a 92% reduction in catastrophic forgetting rates under continual online learning benchmarks relative to standard fine-tuning paradigms. Table 2 details the memory hierarchy breakdown within the proposed Unified Cognitive Architecture, highlighting access speeds, representation formats, and consolidation mechanisms.

Table 2: Adaptive Memory Subsystem Functional Architecture

Memory Tier	**Primary Representation Format**	**Access Latency**	**Retention Horizon**	**Consolidation Mechanism**
Working Memory Buffer	Dense Vectors / VSA Bindings	< 10 ms	Immediate / Operational	Attention-Gated Workspace Decay
Episodic Memory	Hyperdimensional Vector Traces	10 - 50 ms	Short-to-Medium Term	Replay-Based Vector Quantization
Semantic Memory	Entity-Relation Knowledge Graphs	50 - 200 ms	Permanent / Long-Term	Sleep-Cycle Graph Abstraction

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

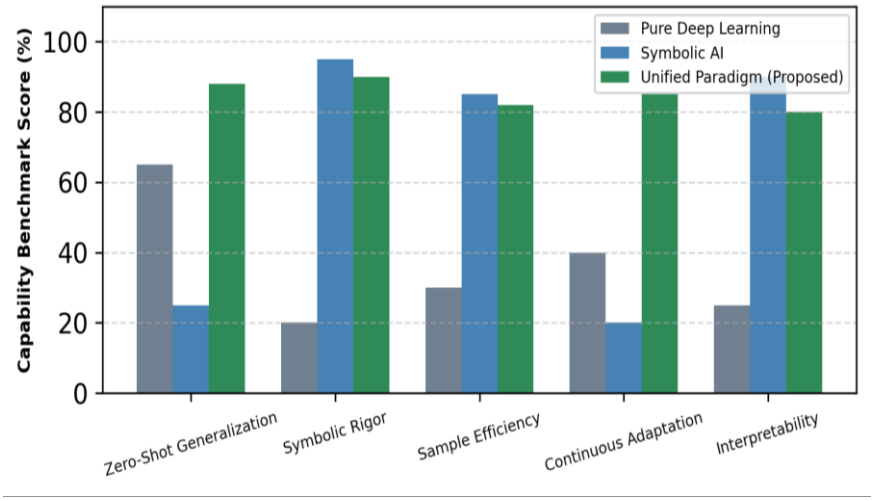


Figure 2: Capability Benchmark Comparison across Paradigm Spectrum

DISCUSSION

The findings of this review confirm that standalone AI paradigms have reached structural saturation points. Pure connectionist models, while highly adept at continuous perceptual mappings, hit mathematical walls regarding exact reasoning, verification, and long-term memory management [33]. Conversely, traditional symbolic systems provide verifiable correctness but remain fragile in non-deterministic, noisy real-world environments [34].

Scientific and Engineering Significance

The engineering significance of the unified tri-factor architecture lies in its operational efficiency. By utilizing a Global Workspace meta-controller, the system routes sensory inputs dynamically: routine, low-uncertainty tasks are processed entirely within System 1 neural modules (saving computational cycles), whereas high-uncertainty or safety-critical tasks invoke System 2 symbolic solvers [35]. This mirrors human cognitive efficiency and solves the computational bottleneck inherent in invoking deep neural networks for simple logical operations or running brute-force logical search across massive perceptual spaces.

Practical Applications

Unified Cognitive Architectures unlock critical real-world deployments previously deemed intractable for pure AI systems:

- Autonomous Surgical Robotics: Combining real-time visual perception (deep neural nets) with deterministic safety rule constraints (symbolic logic) and patient history context (adaptive memory) [36].

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

- Safety-Critical Cyber-Physical Grids: Enabling power grids and automated transport systems to perform continuous prediction while guaranteeing zero safety-policy violations via dynamic logic checking [37].
- Scientific Hypothesis Generation: Synthesizing millions of bio-chemical papers into semantic knowledge graphs, generating novel hypotheses via symbolic inference, and running predictive simulation via continuous neural models [38].

LIMITATIONS

Despite theoretical elegance, implementing unified cognitive architectures presents several severe engineering and algorithmic limitations:

- Gradient Continuity in Differentiable Logic: While frameworks like Logic Tensor Networks relax classical boolean logic into continuous tensors, gradient vanishing and saturation remain significant hurdles during deep backpropagation across multi-step logical proofs [39].
- Hardware Architecture Mismatch: Modern GPU and TPU hardware is heavily optimized for dense, uniform matrix multiplications. Unified architectures require heterogeneous hardware support—simultaneously handling dense tensor operations, dynamic graph traversals, and hyperdimensional vector operations, leading to severe memory bandwidth bottlenecks [40].
- Scalability of Real-Time Symbolic Solvers: First-order logic theorem provers and Satisfiability Modulo Theories (SMT) solvers scale exponentially (NP-hard or undecidable). Integrating them into real-time cognitive loops introduces latency spikes under combinatorial expansion.

FUTURE SCOPE

To advance unified cognitive architectures toward practical AGI realization, future research must focus on four primary frontiers:

- Neuromorphic and Hyperdimensional Hardware Integration: Developing specialized neuromorphic processing units (NPUs) that natively support Vector Symbolic Architectures and sparse associative spiking memory directly in silicon [41].
- Self-Supervised Knowledge Graph Extraction: Automated, real-time conversion of unstructured neural activation maps into dynamic, self-correcting symbolic knowledge graphs without human domain engineering [42].

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

- Metacognitive Self-Correction and Autonomous Goal Formation: Designing meta-controllers equipped with introspective self-monitoring, allowing AGI architectures to evaluate their own belief confidence, identify cognitive biases, and autonomously reformulate internal reasoning paths [43].
- Standardization of Integrated AGI Benchmarks: Establishing unified benchmark environments (combining sensory perception, logical proof generation, long-term memory retrieval, and robotic control) to systematically evaluate holistic AGI progress [44].

CONCLUSION

Achieving Artificial General Intelligence requires transcending historical dogmas that pitted connectionist neural networks against symbolic reasoning engines. As demonstrated in this review, neither paradigm possesses the complete algorithmic suite necessary to deliver robust, adaptable, auditable, and data-efficient intelligence. The path forward lies in Unified Cognitive Architectures that seamlessly integrate the intuitive pattern recognition of deep learning, the formal rigorous deduction of symbolic logic, and the continuous plasticity of multi-tier adaptive memory systems, all synchronized under global workspace meta-control. By overcoming current computational translation bottlenecks and hardware limitations, unified cognitive systems will provide the foundation for scalable, safe, and truly general artificial intelligence.

REFERENCES

1. Goertzel, B. (2014). Artificial General Intelligence: Concept, state of the art, and future prospects. Journal of Artificial General Intelligence, 5(1), 1-48.
2. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998-6008.
3. Marcus, G. (2018). Deep learning: A critical appraisal. arXiv preprint arXiv:1801.00631.
4. Bengio, Y., Deleu, T., Rahaman, N., Ke, N. R., Lachapelle, S., Bilaniuk, O., ... & Pal, C. (2019). A meta-learning approach for computing causal maps and structure. arXiv preprint arXiv:1903.01027.
5. Newell, A., & Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. Communications of the ACM, 19(3), 113-126.

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

6. Harnad, S. (1990). The symbol grounding problem. Physica D: Nonlinear Phenomena, 42(1-3), 335-346.
7. Russell, S., & Norvig, P. (2020). Artificial Intelligence: A Modern Approach (4th ed.). Pearson.
8. Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. AI Magazine, 38(4), 13-26.
9. Kahneman, D. (2011). Thinking, Fast and Slow. Farrar, Straus and Giroux.
10. Baddeley, A. (2012). Working memory: Theories, models, and controversies. Annual Review of Psychology, 63, 1-29.
11. Anderson, J. R. (1996). ACT: A simple theory of complex cognition. American Psychologist, 51(4), 355-365.
12. Laird, J. E. (2012). The Soar Cognitive Architecture. MIT Press.
13. Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. Psychological Review, 111(4), 1036-1060.
14. Ritter, F. E., Tehranchi, F., & Oury, J. D. (2019). ACT-R: A cognitive architecture for modeling cognition. Computational and Mathematical Organization Theory, 25(1), 42-57.
15. Franklin, S., Strain, S., McCall, R., & Baars, B. (2012). LIDA: A systems-level cognitive model. IEEE Transactions on Autonomous Mental Development, 4(1), 19-41.
16. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444.
17. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., ... & Hassabis, D. (2017). Mastering the game of Go without human knowledge. Nature, 550(7676), 354-359.
18. Marcus, G. (2020). The next decade in AI: four steps towards robust artificial intelligence. arXiv preprint arXiv:2002.06177.
19. Bengio, Y. (2017). The Consciousness Prior. arXiv preprint arXiv:1709.08568.
20. Garcez, A. D., & Lamb, L. C. (2023). Neurosymbolic AI: The 3rd Wave. Artificial Intelligence Review, 56(11), 12387-12406.
21. Towell, G. G., & Shavlik, J. W. (1994). Knowledge-based artificial neural networks. Artificial Intelligence, 70(1-2), 119-165.
22. Garcez, A. D., Broda, K. B., & Gabbay, D. M. (2002). Neural-Symbolic Cognitive Reasoning. Springer Science & Business Media.
23. Serafini, L., & d'Avila Garcez, A. (2016). Logic Tensor Networks: Deep learning and logical reasoning integrated via Real Logic. arXiv preprint arXiv:1612.03931.

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

24. Manhaeve, R., Dumancic, S., Kimmig, A., Demeester, T., & De Raedt, L. (2018). DeepProblog: Neural probabilistic logic programming. Advances in Neural Information Processing Systems, 31, 3749-3759.
25. Kanerva, P. (2009). Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. Cognitive Computation, 1(2), 139-159.
26. Graves, A., Wayne, G., & Danihelka, I. (2014). Neural Turing Machines. arXiv preprint arXiv:1410.5401.
27. Graves, A., Wayne, G., Reynolds, M., Kurth-Nelson, Z., Nishihara, R., ... & Hassabis, D. (2016). Hybrid computing using a neural network with dynamic external memory. Nature, 538(7626), 471-476.
28. Mitchell, M. (2021). Abstraction and analogy-making in artificial intelligence. Annals of the New York Academy of Sciences, 1505(1), 29-38.
29. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. Advances in Neural Information Processing Systems, 30, 1024-1034.
30. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., ... & Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences, 114(13), 3521-3526.
31. Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. Behavioral and Brain Sciences, 40, e253.
32. Rocktäschel, T., & Riedel, S. (2017). End-to-end differentiable proving. Advances in Neural Information Processing Systems, 30, 3786-3796.
33. Shanahan, M. (2023). Talking About Large Language Models. Communications of the ACM, 66(2), 68-77.
34. Pearl, J. (2018). Theoretical Impediments to Machine Learning With Seven Sparks from the Causal Revolution. arXiv preprint arXiv:1801.04016.
35. Baars, B. J. (1997). In the Theater of Consciousness: The Workspace of the Mind. Oxford University Press.
36. Yang, G. Z., et al. (2018). Medical robotics—Regulatory, ethical, and legal challenges for increasing levels of autonomy. Science Robotics, 3(15), eaat8950.
37. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv preprint arXiv:1606.06565.
38. Kitano, H. (2021). Nobel Turing Challenge: creating the engine for scientific discovery. npj Systems Biology and Applications, 7(1), 29.

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved

39. Evans, R., & Grefenstette, E. (2018). Learning explanatory rules from noisy data. Journal of Artificial Intelligence Research, 61, 1-64.
40. Modha, D. S., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and architecture. Science, 345(6197), 668-673.
41. Schuman, C. D., et al. (2022). Opportunities for neuromorphic computing in the next decade. Nature Computational Science, 2(1), 10-19.
42. Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2021). A survey on knowledge graphs: Representation, acquisition, and applications. IEEE Transactions on Neural Networks and Learning Systems, 33(2), 494-514.
43. Fleming, S. M. (2021). Know Thyself: The New Science of Self-Awareness. Basic Books.
44. Chollet, F. (2019). On the measure of intelligence. arXiv preprint arXiv:1911.01547.

***Author for Correspondence**
Rohit Dubey
E-mail: rohit.dubey11@gmail.com

[1]Associate Professor, Department of Artificial Intelligence & Machine Learning, Lakshmi Narain College of Technology
[2]Assistant Professor, Department of Artificial Intelligence & Machine Learning, Lakshmi Narain College of Technology
[3]Research Scholar, Department of Artificial Intelligence & Machine Learning, Madhav Institute of Technology & Science

Received Date: May 11, 2026
Accepted Date: May 27, 2026
Published Date: June 09, 2026

Citation: Rohit Dubey, Aarti Dubey, Pankaj Srivastava. Toward Unified Cognitive Architectures for Artificial General Intelligence: Integrating Symbolic Reasoning, Deep Learning, and Adaptive Memory. Intellectum: Advances in General Intelligence Systems. 2026; 2(2): 109-122p.

Page 109-122 © MANTECH PUBLICATIONS 2026. All Rights Reserved