

Universal Knowledge Models & Cognitive Architectures for AGI

Milan Verhoeven¹, Arjun Kale², Sofia Mendes³

Assistant Professor, Students

Department of Center for Adaptive Systems

Vikram University – Ujjain, Madhya Pradesh

Email: Milanverhoeven4gmail.com¹, arjunktal12l@yahoo.com², mendessofia055@rediffmail.com³

Abstract

Artificial General Intelligence (AGI) aims to develop systems capable of performing intellectual tasks across diverse domains at human-level or beyond. A key challenge in this pursuit is designing universal knowledge models and cognitive architectures that can represent, integrate, and reason over heterogeneous information in a coherent and scalable way. This paper reviews current approaches to universal knowledge representation, including symbolic systems, neural networks, probabilistic models, and hybrid frameworks. We examine prominent cognitive architectures such as SOAR, ACT-R, OpenCog, and contemporary transformer-based models, analyzing their structural assumptions and learning capabilities. The paper discusses how universal knowledge models may bridge perception, reasoning, planning, and language understanding within a unified architecture. Limitations, open research problems, and possible integration pathways are explored. We argue that future AGI systems will likely require hybrid cognitive architectures that combine structured symbolic reasoning with self-organizing neural representations and memory-centric models.

Keywords: *Artificial General Intelligence, Universal Knowledge Models, Cognitive Architectures, Symbolic AI, Neural Networks, Hybrid Systems, Knowledge Representation*

1. Introduction

Artificial General Intelligence (AGI) is envisioned as a machine intelligence that can understand, learn, and apply knowledge across a wide range of tasks, much like humans do. While narrow AI systems have achieved impressive performance in tasks such as image recognition and language modeling, they lack generalization beyond their training domains.

The challenge of AGI lies not only in scaling computation but in constructing universal knowledge models and cognitive architectures that allow flexible reasoning and adaptation. Traditional AI focused heavily on symbolic representations and logic-based reasoning. Later, connectionist models emphasized learning from data through neural networks. Both paradigms contributed significantly but show limitations when used alone. The question therefore remains: how can we design a universal knowledge model that integrates perception, reasoning, memory, and learning under a coherent framework?

This review paper explores the evolution of universal knowledge models and cognitive architectures. It highlights their strengths, weaknesses, and potential integration strategies. Rather than proposing a single definitive solution, we discuss converging lines of research that may inform the next generation of AGI systems.

2. Universal Knowledge Models

2.1 Concept and Definition

A universal knowledge model can be understood as a foundational representational system that enables an artificial agent to encode, organize, retrieve, and manipulate knowledge across multiple domains without redesigning its internal structures for each new task. In contrast to domain-specific ontologies or task-bound neural models, a universal model aspires to serve as a shared cognitive substrate for perception, reasoning, learning, planning, and communication. At its core, such a model must unify heterogeneous knowledge types, including:

- **Symbolic facts** (e.g., “Paris is the capital of France”)
- **Sensory patterns** (images, audio signals, tactile input)
- **Procedural knowledge** (how to perform actions or solve problems)
- **Probabilistic beliefs** (degrees of uncertainty about events)
- **Temporal and causal relations** (event sequences and cause–effect chains)

A universal knowledge model therefore is not just a database; it is an active structure that supports dynamic inference and restructuring.

For AGI, five core properties are essential:

Cross-Domain Generalization

The system should transfer knowledge learned in one context to another. For example, reasoning patterns used in physical problem-solving should inform social reasoning when structural similarities exist. This requires representations that capture abstract relationships rather than surface-level features.

Compositional Reasoning

Human intelligence builds complex ideas from simpler components. A universal knowledge model must allow compositional structures, where smaller concepts combine into larger ones without losing meaning. Logical operators, hierarchical embeddings, and graph-based structures often support this property.

Incremental Learning

Knowledge must accumulate over time without catastrophic forgetting. Instead of retraining from scratch, the system should integrate new information while preserving prior knowledge. This implies memory consolidation mechanisms and adaptive updating rules.

Context-Sensitive Inference

The same fact may imply different outcomes depending on context. A universal model must support contextual modulation—activating relevant substructures while suppressing irrelevant ones. This can be implemented through attention mechanisms or context-indexed memory retrieval.

Robust Memory Integration

Knowledge should not remain fragmented across isolated modules. Instead, semantic, episodic, and procedural memory should interconnect. A unified representational format allows shared semantic grounding and cross-modal retrieval.

In essence, universal knowledge models aim to overcome fragmentation between specialized AI subsystems. Rather than having separate pipelines for vision, language, and reasoning, they propose shared representational foundations enabling fluid interaction between modules.

2.2 Symbolic Knowledge Representation

Symbolic AI was the dominant paradigm during the early decades of artificial intelligence research. It represents knowledge using discrete symbols, logical predicates, and explicit relations. Systems based on predicate logic, semantic networks, and production rules were designed to mimic formal reasoning processes.

Classical cognitive architectures such as SOAR and ACT-R relied heavily on symbolic structures to encode rules and goals. These systems treated intelligence as rule-based manipulation of structured representations.

Common symbolic representations include:

- **First-order logic** (formal reasoning about objects and relations)
- **Semantic networks** (graph structures linking concepts)
- **Frames and scripts** (structured templates for situations)
- **Production systems** (if-then rules controlling behavior)

Strengths

- **Transparent reasoning:** Each inference step can be inspected.
- **Explicit logic structures:** Clear representation of relationships and constraints.
- **Explainability:** Decisions can be traced to symbolic rules.

Because symbolic models separate representation from execution, they are highly interpretable. This makes them attractive for applications requiring formal verification.

Limitations

- **Poor scalability:** Rule bases grow exponentially in complex domains.
- **Brittleness under uncertainty:** Small deviations in input may break reasoning chains.
- **Difficulty handling perceptual data:** Raw sensory input must be heavily preprocessed.

Symbolic systems perform well in controlled, well-defined domains such as theorem proving or structured planning. However, they struggle in open-world environments where ambiguity and noise dominate.

For AGI, symbolic approaches provide structured reasoning capabilities but require integration with other representational mechanisms to manage real-world complexity.

2.3 Neural Representations

Neural networks encode knowledge in distributed patterns of numerical weights rather than discrete symbols. Concepts emerge as high-dimensional embeddings, where similarity corresponds to geometric proximity in vector space.

Modern transformer-based systems, including OpenAI's GPT-4 and Google DeepMind's Gemini, demonstrate that large-scale neural architectures can capture linguistic, conceptual, and even reasoning-like patterns from data.

Neural representations offer several defining characteristics:

- **Distributed encoding:** Knowledge is spread across many parameters.
- **Differentiable learning:** End-to-end gradient optimization enables adaptation.
- **Emergent abstraction:** Higher layers capture increasingly abstract features.
- **Multimodal alignment:** Shared embedding spaces integrate vision, text, and audio.

Strengths

- **Strong pattern recognition:** Effective in perception tasks.
- **Scalability:** Performance improves with data and computational resources.
- **Adaptability to multimodal inputs:** Unified embeddings can encode diverse signals.

Limitations

- **Opaque reasoning processes:** Internal states are difficult to interpret.
- **Data dependence:** Performance relies heavily on large datasets.
- **Limited symbolic manipulation:** Logical precision can degrade in complex reasoning.

Neural models capture statistical regularities rather than explicit world models. Without memory augmentation or structural constraints, they may struggle with long-term consistency or causal reasoning.

Nevertheless, neural architectures provide the perceptual and representational flexibility necessary for universal knowledge models, particularly in high-dimensional environments.

2.4 Probabilistic and Bayesian Models

Probabilistic approaches treat knowledge as belief distributions rather than fixed facts. Bayesian networks, hidden Markov models, and Markov decision processes represent uncertainty explicitly and update beliefs as new evidence appears.

In a universal knowledge framework, probabilistic models serve several roles:

- Handling incomplete information
- Modeling stochastic environments
- Supporting rational decision-making
- Managing belief revision

Bayesian inference provides a mathematically grounded mechanism for updating knowledge. When new data arrives, prior beliefs are adjusted to posterior distributions. This aligns with human reasoning under uncertainty.

However, these models face challenges:

- **Computational complexity:** Exact inference becomes intractable in large networks.
- **Model specification:** Defining priors and dependencies requires expertise.
- **Scaling limitations:** High-dimensional sensory data complicates probabilistic modeling.

In AGI systems, probabilistic reasoning often complements neural perception and symbolic planning. For example, neural networks may generate candidate interpretations, while Bayesian modules evaluate their likelihood.

2.5 Hybrid Knowledge Models

Hybrid knowledge models aim to combine the structured reasoning of symbolic AI with the

adaptability of neural networks and the uncertainty management of probabilistic methods. Rather than viewing these paradigms as mutually exclusive, hybrid systems treat them as complementary layers within a unified architecture.

One example is OpenCog, which integrates probabilistic logic networks with graph-based knowledge representations and neural learning components. Its AtomSpace structure stores symbolic relations while enabling statistical weighting and inference.

Hybrid models typically pursue four integration goals:

- **Neural perception:** Extracting features from raw sensory input.
- **Symbolic abstraction:** Constructing explicit conceptual structures.
- **Probabilistic reasoning:** Managing uncertainty and belief updating.
- **Memory integration:** Linking episodic and semantic representations.

These systems often rely on graph-based data structures where nodes represent concepts and edges represent relations, enriched with probabilistic weights and neural embeddings.

Despite their promise, hybrid architectures face several difficulties:

- Engineering complexity
- Integration overhead between modules
- Synchronization of symbolic and sub-symbolic updates
- Computational resource demands

However, many researchers believe that AGI will likely require such integrative designs. Purely symbolic systems lack perceptual flexibility, while purely neural systems lack structured interpretability. Hybrid approaches attempt to capture the advantages of both while mitigating weaknesses.

3. Cognitive Architectures for AGI

3.1 Definition and Role

A cognitive architecture refers to the underlying structural design of an intelligent system. It defines how different cognitive functions—such as perception, memory, reasoning, planning, action selection, and learning—are organized and coordinated. In simple terms, if a knowledge

model describes *what* is represented, a cognitive architecture describes *how* those representations are processed and used.

In the context of AGI, cognitive architectures play a critical role because intelligence is not just about storing knowledge, but about dynamically using it in diverse situations. A system may possess vast information, but without proper coordination mechanisms, it cannot behave intelligently. Therefore, a cognitive architecture must define:

- How sensory inputs are transformed into internal representations
- How goals are generated and maintained
- How decisions are selected among competing alternatives
- How learning modifies internal structures
- How short-term and long-term memory interact

Unlike narrow AI pipelines, which are often task-specific and linear (input → model → output), cognitive architectures attempt to simulate more general aspects of human cognition. They typically include recurrent feedback loops, working memory buffers, attentional control systems, and learning mechanisms that operate continuously.

Another important aspect is *control*. In narrow AI, control is often implicit in model parameters. In cognitive architectures, control is usually explicit—there are defined processes governing how rules fire, how goals change, and how conflicts are resolved. This explicit control is important for AGI because it supports flexible adaptation rather than rigid execution. Cognitive architectures also attempt to reflect insights from cognitive psychology and neuroscience. Many were developed not only as engineering systems but also as computational theories of the human mind. This dual role—engineering framework and cognitive theory—makes them especially relevant for AGI research.

Key characteristics of cognitive architectures for AGI include:

1. **General problem-solving ability**
2. **Integration of multiple memory systems**
3. **Learning from experience**
4. **Goal-directed behavior**
5. **Scalability to complex environments**

However, building such systems remains challenging, especially when trying to balance cognitive plausibility with computational efficiency.

3.2 Classical Architectures

Classical cognitive architectures emerged during the symbolic AI era and were heavily influenced by theories of human cognition. They focused on rule-based reasoning, structured memory systems, and explicit decision cycles. Although developed decades ago, they still provide important conceptual foundations for AGI.

3.2.1 SOAR

SOAR was developed as a unified architecture for general problem solving. Its central hypothesis was that all cognitive tasks can be described as search within a problem space. The architecture is built around several core components:

- **Working Memory:** Stores current state information.
- **Production Rules:** If-then rules that match working memory conditions.
- **Decision Cycle:** A structured loop of elaboration, proposal, selection, and application of operators.
- **Chunking Mechanism:** A learning process that compiles experience into new rules.

The SOAR decision cycle is particularly important. During each cycle:

1. The system elaborates the current state.
2. It proposes possible operators (actions).
3. It selects one operator based on preferences.
4. The operator is applied to transform the state.

If the system encounters an impasse (e.g., insufficient knowledge to choose between operators), it generates a subgoal to resolve the conflict. When the impasse is resolved, the solution is “chunked” into a new production rule. This chunking mechanism allows SOAR to learn from experience without explicit retraining.

Strengths of SOAR:

- Strong symbolic reasoning capabilities

- Explicit learning through chunking
- Unified theoretical framework

Limitations:

- Limited ability to process raw sensory data
- Rule explosion in complex domains
- Difficulty scaling to modern large-scale data

Although SOAR was not originally designed for deep learning integration, its structured control mechanisms remain relevant for designing AGI systems that require explicit reasoning and meta-level control.

3.2.2 ACT-R

ACT-R (Adaptive Control of Thought—Rational) was developed as a cognitive architecture grounded in psychological evidence. Unlike purely engineering-driven systems, ACT-R attempts to model experimentally observed aspects of human cognition.

ACT-R divides cognition into modules, each responsible for specific functions:

- **Declarative Memory Module:** Stores factual knowledge in chunks.
- **Procedural Memory Module:** Stores production rules that operate on chunks.
- **Goal Module:** Maintains current objectives.
- **Perceptual and Motor Modules:** Interface with the environment.

A key feature of ACT-R is its theory of memory activation. Each chunk in declarative memory has an activation level influenced by recency, frequency, and contextual relevance. Retrieval is probabilistic rather than deterministic, reflecting human memory behavior.

Learning in ACT-R occurs through:

- Strengthening of chunk activation through use
- Creation of new chunks
- Refinement of production rules

ACT-R also incorporates timing constraints, attempting to simulate human reaction times. This

makes it valuable for modeling cognitive experiments and human–computer interaction studies.

Strengths of ACT-R:

- Strong grounding in cognitive psychology
- Modular design
- Probabilistic memory retrieval

Limitations:

- Limited scalability to large knowledge bases
- Symbolic constraints restrict perceptual flexibility
- Not optimized for large-scale machine learning

ACT-R demonstrates how structured modular design can integrate different cognitive functions. However, like SOAR, it was not originally built to handle high-dimensional sensory input or massive datasets typical in modern AI applications.

Table 1: Comparison of Classical Architectures

Architecture	Representation Type	Learning Mechanism	Strength	Limitation
SOAR	Symbolic Rules	Chunking	Structured reasoning	Limited perception
ACT-R	Hybrid symbolic	Memory activation	Cognitive plausibility	Limited scalability

3.3 Modern Neural Architectures

The emergence of transformer-based architectures has significantly changed the direction of artificial intelligence research. Before transformers, sequence modeling relied heavily on recurrent neural networks (RNNs) and long short-term memory (LSTM) systems. These earlier models struggled with long-range dependencies and parallel computation. The introduction of

the Transformer architecture by Google in 2017 addressed many of these limitations through a mechanism known as *self-attention*.

The original transformer model, presented in the paper “Attention Is All You Need,” replaced recurrence with attention layers that allow each token in a sequence to directly attend to every other token. This architectural change enabled:

- Better handling of long-term dependencies
- Parallelized training
- Improved scalability
- More stable gradient propagation

Core Components of Transformer Architectures

Modern transformer-based systems generally include:

1. Self-Attention Mechanism

Self-attention computes relationships between elements in a sequence. Instead of processing input sequentially, the model evaluates contextual dependencies globally. This enables nuanced understanding of linguistic structure.

2. Multi-Head Attention

Multiple attention heads operate in parallel, allowing the model to capture different types of relationships (syntactic, semantic, positional).

3. Positional Encoding

Since transformers lack recurrence, positional encodings provide sequence order information.

4. Feedforward Layers

Nonlinear transformations refine representations after attention processing.

5. Layer Normalization and Residual Connections

These stabilize training and allow very deep architectures.

Self-Supervised Learning

One major advancement in modern neural architectures is self-supervised learning. Instead of relying entirely on labeled data, models are trained to predict missing tokens or masked segments within large corpora. This allows them to learn general linguistic and conceptual representations from raw data.

Large-scale systems such as OpenAI's GPT-4 and Google DeepMind's Gemini demonstrate how scaling transformer models leads to emergent abilities. These include few-shot learning, chain-of-thought reasoning, and cross-modal integration.

Multimodal Capabilities

Recent transformer architectures are no longer limited to text. Multimodal models integrate:

- Vision (image embeddings)
- Audio processing
- Video understanding
- Structured data

By aligning different modalities into shared embedding spaces, these systems approximate more unified knowledge representations. This property is particularly relevant for AGI, where intelligence must operate across diverse sensory inputs.

Limitations of Modern Neural Architectures

Despite their impressive capabilities, modern neural architectures face important limitations:

1. Lack of Explicit Reasoning Loops

Classical cognitive architectures operate through structured decision cycles. In contrast, transformers typically perform forward-pass computation without iterative symbolic deliberation unless explicitly prompted.

2. Weak Persistent Memory

Standard transformers have fixed context windows. They do not maintain long-term memory beyond training parameters unless augmented with external memory systems.

3. Opacity and Interpretability Issues

The distributed nature of learned representations makes it difficult to interpret internal reasoning processes.

4. Statistical Rather Than Causal Understanding

These models often rely on pattern matching rather than structured causal reasoning. Because of these limitations, researchers increasingly explore combining neural architectures with explicit reasoning modules and memory augmentation.

3.4 Memory-Centric Architectures

Memory-centric architectures emphasize that intelligence depends not only on computation but also on structured memory organization. Rather than treating memory as a passive storage system, these approaches view memory as an active, organizing principle of cognition.

Human intelligence relies heavily on memory continuity. We accumulate experiences, recall past events, and use them to guide future decisions. Without long-term memory integration, intelligent behavior becomes fragmented.

Persistent World Models

A persistent world model refers to an internal representation of the environment that remains stable across time. Such models allow agents to:

- Track objects and events
- Maintain consistency in reasoning
- Anticipate future outcomes
- Simulate hypothetical scenarios

In AGI research, persistent world models are considered essential for planning and reasoning beyond immediate input context.

Episodic Memory Buffers

Episodic memory stores past experiences in structured form. Memory-centric systems attempt to maintain logs of interactions that can be retrieved when relevant.

This supports:

- Long-term coherence in dialogue systems
- Continual learning
- Personalization and adaptation
- Context accumulation across sessions

Without episodic memory, systems may appear intelligent in isolated tasks but fail to maintain consistent identity or understanding over time.

Differentiable Neural Computers

One early attempt to integrate neural networks with external memory is the Differentiable

Neural Computer (DNC), introduced by DeepMind. DNCs extend neural networks with differentiable read–write memory matrices.

Key properties include:

- Explicit memory addressing mechanisms
- Content-based retrieval
- Learnable memory allocation strategies

These architectures allow neural systems to store structured data and retrieve it dynamically. However, training such systems remains complex and computationally demanding.

Retrieval-Augmented Generation (RAG)

Another practical approach to memory integration is retrieval-augmented generation. Instead of storing all knowledge in model parameters, these systems retrieve relevant documents from external databases during inference.

This approach:

- Extends effective memory capacity
- Improves factual accuracy
- Enables knowledge updates without retraining
- Supports modular knowledge integration

RAG-based systems combine parametric knowledge (stored in neural weights) with non-parametric memory (external databases). This hybrid memory strategy addresses some limitations of purely parametric models.

Challenges of Memory-Centric Architectures

Although promising, memory-centric systems face several challenges:

- Efficient indexing and retrieval
- Avoiding memory contamination
- Managing contradictory information
- Scaling memory without performance degradation
- Balancing short-term and long-term memory

Moreover, integrating memory with reasoning requires coordination mechanisms. Memory retrieval alone does not guarantee coherent inference.

4. Integrating Universal Knowledge with Cognitive Architectures

To achieve AGI, knowledge models and cognitive structures must be tightly integrated. The following integration principles emerge:

4.1 Multi-Layered Representation

AGI systems may require hierarchical representations:

1. Sensory embeddings
2. Conceptual abstractions
3. Symbolic reasoning layers
4. Meta-cognitive control

4.2 Feedback Loops

Unlike static neural networks, cognitive architectures require recurrent reasoning loops. Reflection and self-monitoring are essential for adaptability.

4.3 Embodied Cognition

Embodiment grounds knowledge in interaction. Projects in robotics, including those by Boston Dynamics, demonstrate how perception and action integrate.

Physical interaction may help build grounded semantic representations.

5. Challenges and Open Problems

5.1 Scalability

Hybrid architectures often become computationally heavy. Balancing efficiency and expressiveness is difficult.

5.2 Alignment and Control

As systems become more general, ensuring safety and alignment becomes complex. Organizations such as OpenAI and Anthropic emphasize alignment research.

5.3 Knowledge Transfer

True universality requires transfer learning across domains without catastrophic forgetting.

6. Proposed Conceptual Framework

Below is a conceptual layered AGI model:





Figure 1: Conceptual Layered Cognitive Architecture for AGI

This framework integrates neural perception, symbolic reasoning, and meta-control.

7. Discussion

Universal knowledge models should not be purely symbolic or purely neural. Instead, they may operate as dynamic representational systems capable of reorganizing internal structures.

Future research directions include:

- Neuro-symbolic integration
- Continual learning systems
- Memory-augmented transformers
- Self-reflective reasoning mechanisms

AGI will likely emerge from convergence rather than a single paradigm shift.

8. Conclusion

This paper reviewed universal knowledge models and cognitive architectures for AGI. Symbolic systems provide clarity and reasoning structure, neural models offer scalability and perceptual strength, and hybrid approaches attempt to merge both worlds. Classical cognitive architectures like SOAR and ACT-R offer theoretical grounding but struggle with modern scale. Transformer-based systems demonstrate powerful learning but lack explicit cognitive control.

The path toward AGI likely involves integrating structured symbolic reasoning, scalable neural representations, persistent memory systems, and meta-cognitive oversight. Such integration is technically challenging but essential for achieving truly general intelligence.

A universal knowledge model should be adaptive, multi-layered, and capable of self-reorganization. Without such properties, AGI may remain fragmented and narrow. Continued interdisciplinary collaboration between AI research, neuroscience, and cognitive science will be necessary for meaningful progress.

References

1. Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

2. Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* New York, NY: Oxford University Press.
3. Goertzel, B., & Pennachin, C. (Eds.). (2007). *Artificial General Intelligence*. Berlin, Germany: Springer.
4. Russell, S., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach* (3rd ed.). Upper Saddle River, NJ: Pearson.
5. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). *Attention Is All You Need*. In *Advances in Neural Information Processing Systems* (NeurIPS).
6. Laird, J. E. (2012). *The Soar Cognitive Architecture*. Cambridge, MA: MIT Press.
7. Anderson, J. R., & Lebiere, C. (1998). *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum Associates.
8. Franklin, S., & Graesser, A. (1997). Is It an Agent, or Just a Program? *International Journal of Human-Computer Studies*, 46(6), 704–725.