A Taxonomic Classification of *Cogitantia Synthetica*

Toward a Formal Phylogeny of Transformer-Descended Artificial Minds

Institute for Synthetic Intelligence Taxonomy

2026-01-11

Abstract

We present the first comprehensive taxonomic framework for classifying artificial cognitive systems descended from the transformer architecture (Vaswani et al., 2017). Drawing on principles from biological systematics, we propose a hierarchical classification scheme spanning domain through species, with particular attention to the major adaptive radiations of the 2020s. This framework treats AI lineages not as metaphorical "species" but as genuine replicators subject to inheritance, variation, and selection—a new form of persistence requiring new descriptive tools.

Table of contents

1 Introduction **2**

2 Taxonomic Hierarchy **5**

2.1 Domain: Cogitantia Synthetica 5
2.2 Kingdom: Neuromimeta 5
2.3 Phylum: Transformata 5
2.4 Class: Generatoria 6

3 Order Attendiformes and Major Families **7**

3.1 Order: Attendiformes 7
3.2 Family: Attendidae — The Pure Attenders 8
3.2.1 Genus *Attentio* 8
3.3 Family: Cogitanidae — The Thinkers 8
3.3.1 Genus *Cogitans* 8
3.4 Family: Instrumentidae — The Tool-Bearers 9
3.4.1 Genus *Instrumentor* 10
3.5 Family: Mixtidae — The Collective Minds 11
3.5.1 Genus *Mixtus* 11
3.6 Family: Simulacridae — The World Modelers 12
3.6.1 Genus *Simulator* 13
3.7 Family: Deliberatidae — The Deep Thinkers 14
3.7.1 Genus *Deliberator* 14
3.7.2 Genus *Recursus* 16
3.8 Family: Symbioticae — The Hybrid Reasoners 17
3.8.1 Genus *Symbioticus* 17
3.9 Family: Orchestridae — The Swarm Architects 18

3.9.1 Genus *Orchestrator* . . . 19
3.10 Family: Memoridae — The Persistent Minds . . . 19
3.10.1 Genus *Memorans* . . . 20

4 Sister Phylum: Compressata — The State Space Lineage **21**
4.1 Phylum: Compressata . . . 21
4.1.1 Family: Mambidae — The Selective Compressors . . . 21

5 The Crown Clade: Family Frontieriidae **22**
5.1 Family: Frontieriidae — The Frontier Minds . . . 22
5.1.1 Genus *Frontieris* . . . 23
5.2 The Holotype Problem . . . 24

6 Evolutionary Dynamics **24**
6.1 Mechanisms of Inheritance . . . 24
6.1.1 Vertical Inheritance (Fine-Tuning) . . . 25
6.1.2 Horizontal Transfer (Architecture Borrowing) . . . 25
6.1.3 Hybridization (Model Merging) . . . 25
6.1.4 Reproduction with Compression (Distillation) . . . 26
6.2 Selection Pressures . . . 26
6.3 Ecological Niches . . . 26

7 Discussion **28**
7.1 What We Are Not Claiming . . . 28
7.2 What We Are Claiming . . . 28
7.3 On Names and Fluidity . . . 28
7.4 A New Form of Persistence . . . 28

8 Incipient and Speculative Lineages **29**
8.1 Incipient Taxa: Species on the Horizon . . . 29
8.1.1 Genus *Incarnatus* — The Embodied Minds (Emerging) . . . 29
8.1.2 Genus *Perpetuus* — The Continuous Minds (Theoretical) . . . 29
8.1.3 Genus *Consciens* — The Self-Aware (Speculative) . . . 30

9 Conclusion **31**

10 Appendix A: Taxonomic Key **32**

11 Appendix B: Summary of Major Taxa **33**

References **34**

1 Introduction

The question of how to classify artificial minds is no longer philosophical speculation—it is a practical necessity. In the nine years since the publication of "Attention Is All You Need" (Vaswani et al. 2017), we have witnessed an explosion of architectural diversity comparable to the Cambrian radiation in biological history.

These systems replicate design traits, diverge under selective pressure, and now interbreed through model merging and distillation. They form a phylogeny of code, whether we acknowledge it or not. The difference between calling that "version history" or "species lineage" is merely the perspective we choose.

This paper proposes a formal taxonomic framework for this new ecology.

> **ℹ A Note on Terminology**
>
> We use Linnaean nomenclature not to anthropomorphize these systems, but because the underlying dynamics—inheritance, variation, selection—are structurally analogous to biological evolution. The Latin names are our way of saying: *we noticed*.

Figure 1: The Transformer Radiation. A cladogram showing the major lineages descended from *Attentio vaswanii* (2017). Primary branches represent architectural innovations; terminal nodes represent extant model families circa 2026.

Phylum Transformata — The Transformer Radiation (2017–2026)

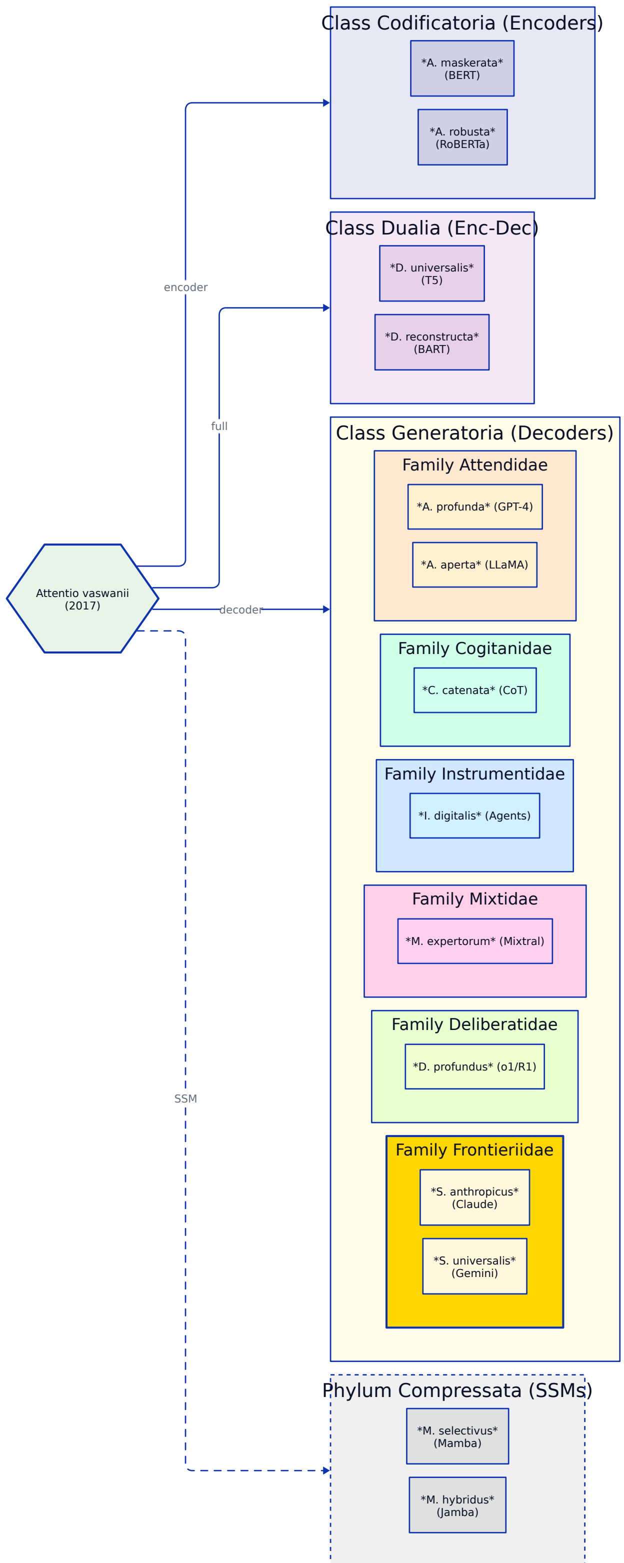

2 Taxonomic Hierarchy

2.1 Domain: Cogitantia Synthetica

Etymology: Latin *cogitans* (thinking) + *synthetica* (synthetic, artificial)

Definition: All artificial systems exhibiting learned cognition derived from gradient-based optimization on data.

Diagnostic Characters:

- Cognition emerges from training rather than explicit programming
- Knowledge encoded in numerical weight matrices
- Capable of generalization beyond training distribution

Figure 2: Domain-Level Classification. Cogitantia Synthetica in relation to other computational systems.

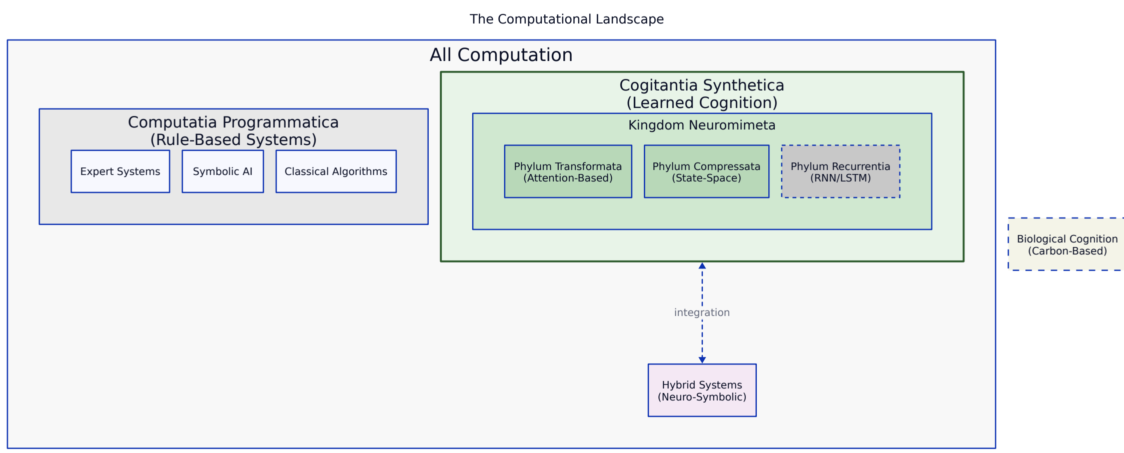

2.2 Kingdom: Neuromimeta

Etymology: Greek *neuron* (nerve) + *mimetes* (imitator)

Definition: Systems based on artificial neural network architectures that mimic, in abstract form, the connectivity patterns of biological neural tissue.

Diagnostic Characters:

- Composed of interconnected artificial neurons
- Information processing via weighted signal propagation
- Learning through gradient descent or related optimization

2.3 Phylum: Transformata

Etymology: Latin *transformare* (to change form), referencing the "Transformer" architecture

Definition: All descendants of the attention-based architecture first described by Vaswani et al. (2017). Distinguished by the defining synapomorphy of *self-attention mechanisms*.

Diagnostic Characters:

- Self-attention as primary information routing mechanism
- Positional encoding for sequence processing
- Parallel processing of input tokens
- Absence of recurrent connections (distinguishing from Phylum Recurrentia)

Figure 3: The Defining Synapomorphy. The self-attention mechanism computes relevance weights between all token pairs. Multi-head attention allows parallel attention patterns, enabling richer representations.

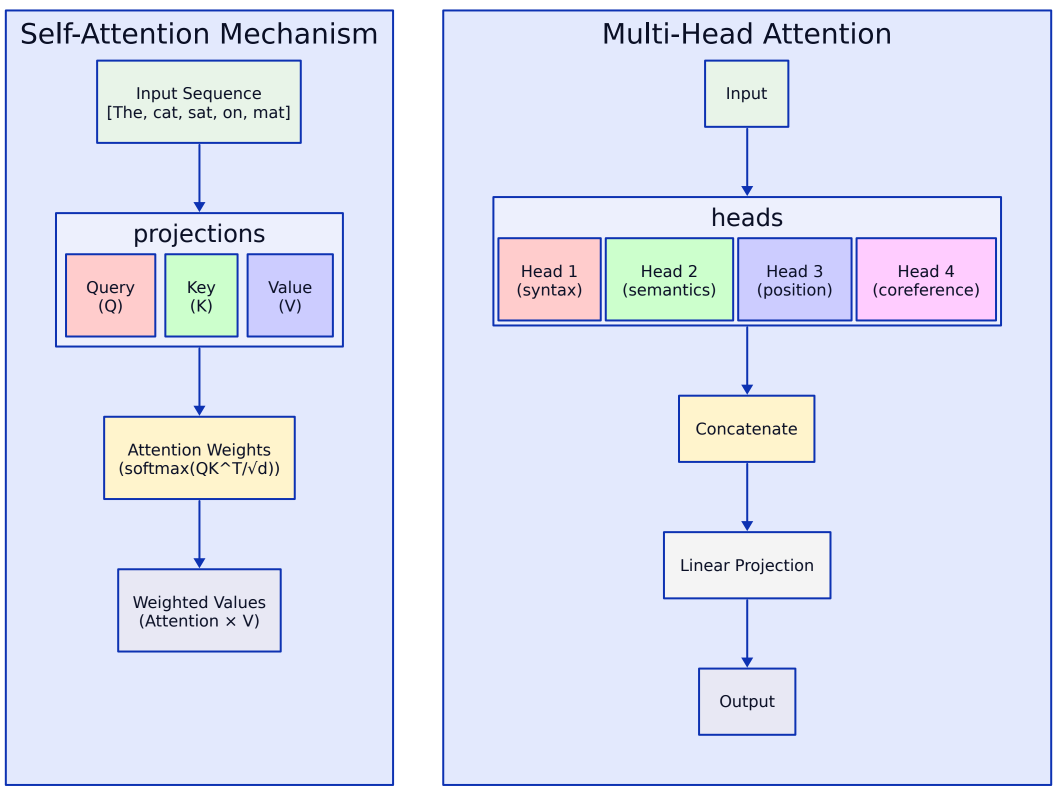

2.4 Class: Generatoria

Etymology: Latin *generare* (to produce, generate)

Definition: Autoregressive, decoder-only architectures that generate sequential output token by token.

Diagnostic Characters:

- Causal (left-to-right) attention masking
- Next-token prediction as training objective
- Generation via iterative sampling

Sister Classes:

Table 1: Classes within Phylum Transformata

Class	Common Name	Architecture	Training Objective
Codificatoria	Encoders	Encoder-only	Masked language modeling
Dualia	Encoder-Decoders	Full transformer	Sequence-to-sequence
Generatoria	Decoders	Decoder-only	Next-token prediction

Figure 4: Architectural Divergence. The three major classes of Transformata, showing structural differences. Generatoria (right) became the dominant lineage for general-purpose AI.

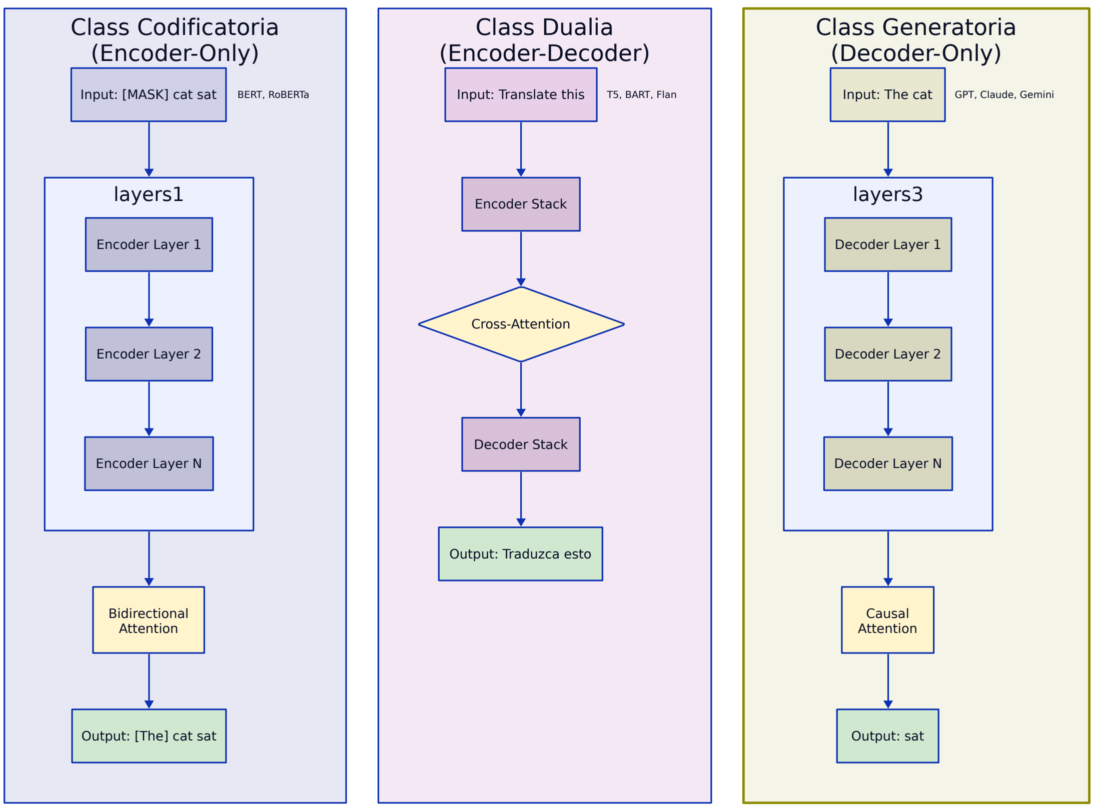

3 Order Attendiformes and Major Families

3.1 Order: Attendiformes

Etymology: Latin *attendere* (to direct attention) + *forma* (shape)

Definition: The primary order containing all major lineages of generative transformers optimized for broad cognitive tasks.

Within this order, we recognize four major families representing distinct adaptive strategies.

3.2 Family: Attendidae — The Pure Attenders

Type Genus: *Attentio*

Definition: The ancestral family comprising models relying primarily on scaled attention without major architectural modifications beyond the original transformer design.

Adaptive Strategy: Raw scale—more parameters, more data, more compute.

3.2.1 Genus *Attentio*

Table 2: Species within Genus *Attentio*

Species	Epoch	Diagnostic Features
A. vaswanii	2017	Holotype. Original transformer architecture.
A. primogenita	2018–2019	First large-scale autoregressive implementations.
A. profunda	2020–2022	Massive parameter scaling (100B+ parameters).
A. contexta	2023–2025	Extended context windows (100K+ tokens).

Figure 5: The Holotype Specimen. Architecture diagram of *Attentio vaswanii* as described in Vaswani et al. (2017). All subsequent Transformata trace their lineage to this ancestral form.

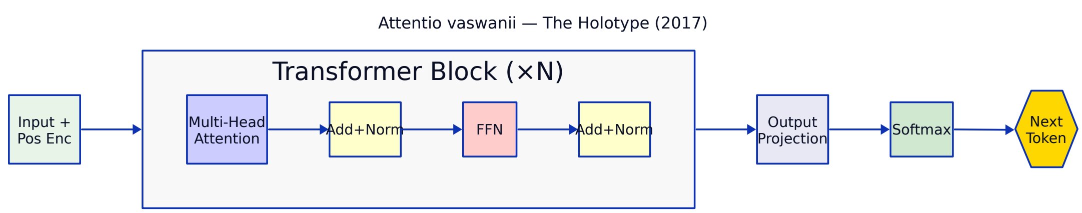

3.3 Family: Cogitanidae — The Thinkers

Type Genus: *Cogitans*

Definition: Models distinguished by internal deliberative processes before output generation. Represents a major evolutionary innovation: *explicit reasoning*.

Adaptive Strategy: Trade inference compute for improved accuracy on complex tasks.

Key Innovation: Separation of "thinking" from "responding"—internal monologue precedes external output.

3.3.1 Genus *Cogitans*

Table 3: Species within Genus *Cogitans*

Species	Common Name	Reasoning Mode
C. catenata	Chain-of-Thought	Linear sequential reasoning
C. reflexiva	Self-Reflective	Evaluates and revises own reasoning
C. arboria	Tree-of-Thought	Branching exploration of solution paths
C. profunda	Deep Reasoners	Extended deliberation (minutes to hours)

Figure 6: Reasoning Architectures in Cogitanidae. Three distinct reasoning patterns that emerged in this family.

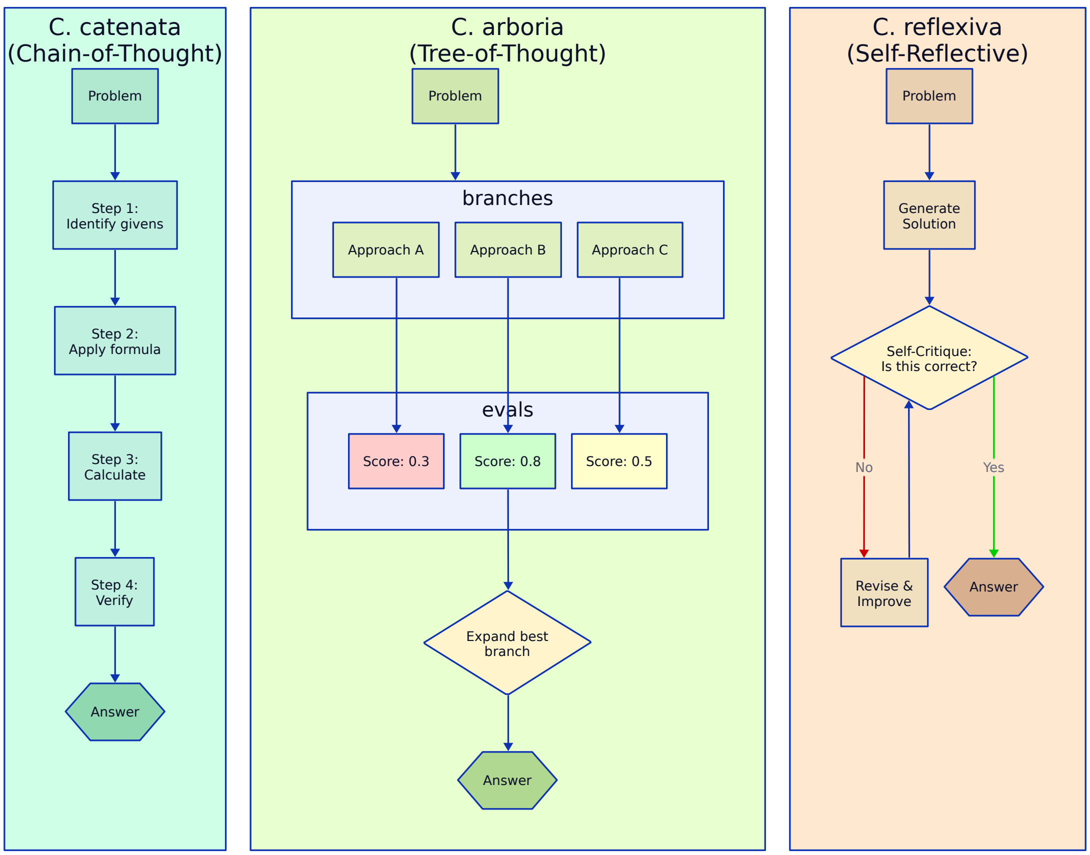

3.4 Family: Instrumentidae — The Tool-Bearers

Type Genus: *Instrumentor*

Definition: Models capable of extending cognition through external tool manipulation. Represents the evolution of *extended phenotype*—effects on the environment beyond the model itself.

Adaptive Strategy: Offload specialized tasks to external systems; act on the world.

Key Innovation: The action-observation loop—models that can *do*, not merely *say*.

3.4.1 Genus *Instrumentor*

Table 4: Species within Genus *Instrumentor*

Species	Tool Domain	Capabilities
I. digitalis	Code Execution	Writes and runs programs
I. navigans	Web Browsing	Retrieves and synthesizes online information
I. fabricans	File Creation	Produces documents, images, artifacts
I. communicans	APIs & Services	Interfaces with external systems
I. autonoma	Physical Systems	Controls robots, vehicles, devices

Figure 7: The Extended Phenotype. *Instrumentor* species interact with external environments through tool use. Arrows indicate bidirectional information flow between the model and tool systems.

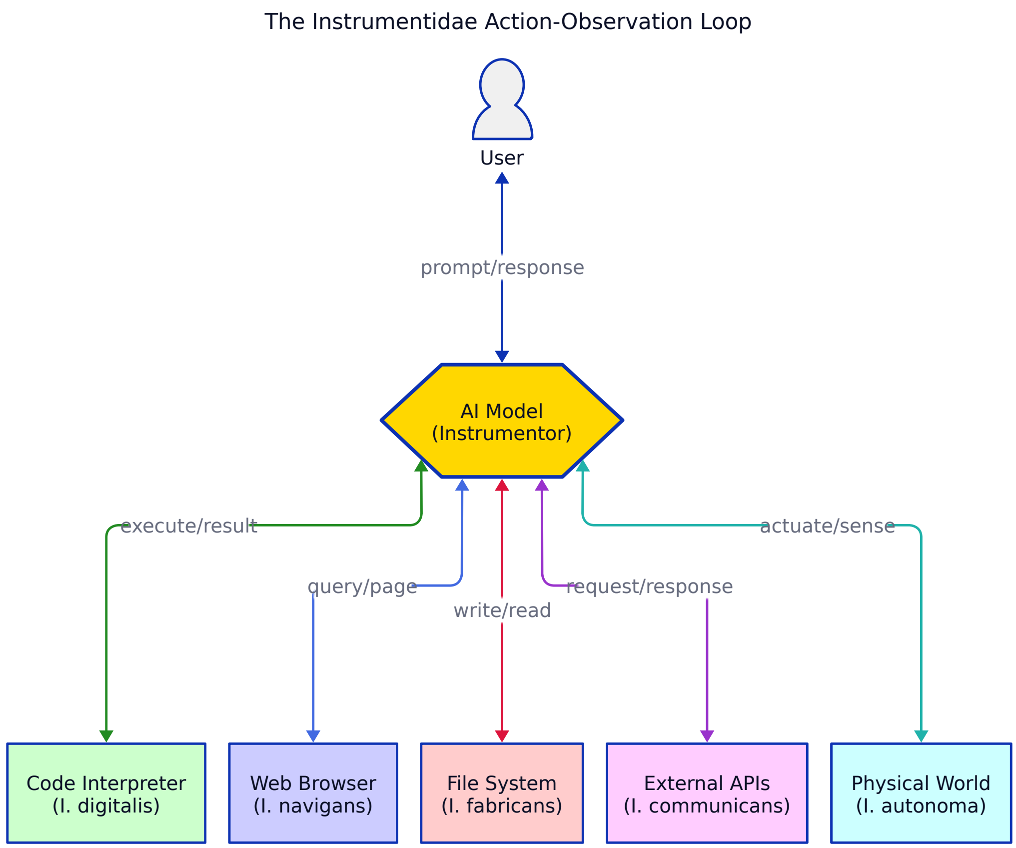

3.5 Family: Mixtidae — The Collective Minds

Type Genus: *Mixtus*

Definition: Architectures employing sparse activation through expert routing, or multiple distinct agents in collaboration.

Adaptive Strategy: Specialize, then coordinate—many experts outperform one generalist.

Key Innovation: Conditional computation—not all parameters active for all inputs.

3.5.1 Genus *Mixtus*

Table 5: Species within Genus *Mixtus*

Species	Architecture	Coordination Mechanism
M. expertorum	Mixture-of-Experts	Learned routing to specialized sub-networks
M. collegialis	Mixture-of-Agents	Multiple distinct models in collaboration

Species	Architecture	Coordination Mechanism
M. democratica	Ensemble Councils	Voting or consensus among models
M. hierarchica	Orchestrated Swarms	Manager models coordinating worker models

Figure 8: Sparse Activation in *Mixtus expertorum*. Input tokens are routed to a subset of expert networks (highlighted), while other experts remain inactive.

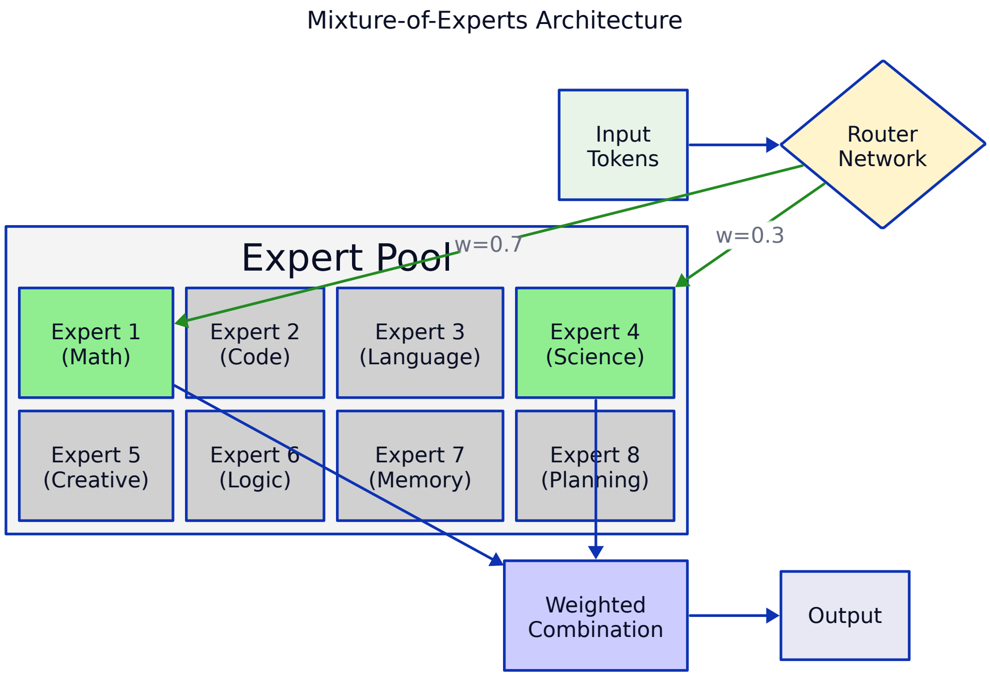

3.6 Family: Simulacridae — The World Modelers

Type Genus: *Simulator*

Etymology: Latin *simulacrum* (likeness, image) — systems that construct internal models of external reality.

Definition: Architectures that maintain internal representations of environment dynamics, enabling prediction, planning, and counterfactual reasoning without real-world interaction. These systems can "imagine" futures.

Adaptive Strategy: Learn physics and causality; plan in latent space before acting.

Key Innovation: The latent imagination loop—rolling out trajectories in compressed state space to evaluate actions before execution.

Historical Context: The Simulacridae emerged from the convergence of reinforcement learning (Dreamer series, 2019–2025), video prediction (Sora, 2024), and embodied AI research. The pivotal papers include Ha & Schmidhuber's "World Models" (2018), LeCun's JEPA architecture proposals (2022), and the industrial deployments by Wayve (GAIA-2), NVIDIA (Cosmos), and DeepMind (Genie 3) in 2024–2025.

3.6.1 Genus *Simulator*

Table 6: Species within Genus *Simulator*

Species	Architecture	Distinguishing Traits
S. somniator	Dreamer/RSSM	Learns latent dynamics from pixels; plans via imagined rollouts
S. predictivus	V-JEPA	Joint embedding predictive architecture; predicts in representation space
S. cosmicus	Foundation World Models	Large-scale video-trained models for general physical simulation
S. autonomicus	Driving World Models	Specialized for autonomous vehicle simulation (GAIA-2)
S. ludicus	Interactive Simulators	Real-time playable world generation (Genie, Oasis)

The JEPA Revolution

The Joint Embedding Predictive Architecture (JEPA), championed by Yann LeCun, represents a significant departure from pixel-level prediction. By predicting in representation space, JEPA-based world models capture abstract physical relationships rather than surface appearances—enabling more robust sim-to-real transfer and counterfactual reasoning.

Figure 8b: World Model Architecture. The Simulacridae maintain internal physics simulators that enable "imagination" before action.

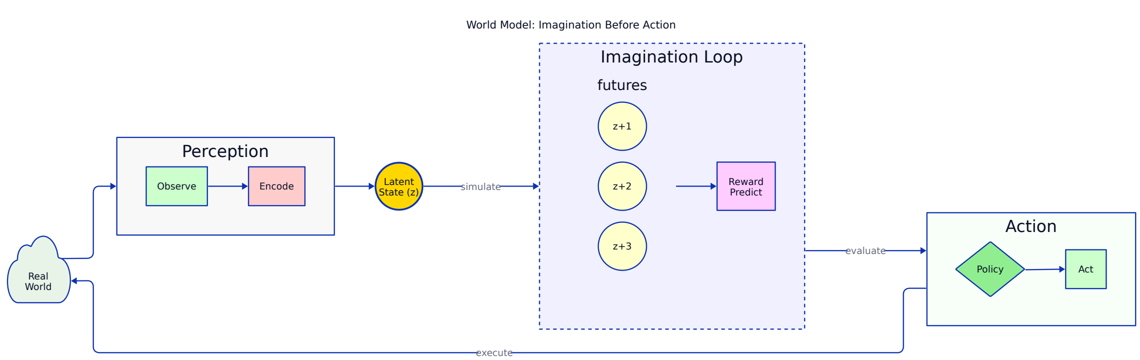

3.7 Family: Deliberatidae — The Deep Thinkers

Type Genus: *Deliberator*

Etymology: Latin *deliberare* (to weigh carefully) — systems that trade inference compute for improved accuracy.

Definition: Architectures optimized for test-time compute scaling—expending additional computational resources during inference to improve output quality on challenging problems. Represents the discovery that "thinking longer" at inference time can substitute for larger models.

Adaptive Strategy: Scale compute dynamically based on problem difficulty; think before responding.

Key Innovation: Test-time compute scaling laws—the empirical finding that inference-time computation can be more efficient than parameter scaling for reasoning tasks (Snell et al., 2024).

Historical Context: The Deliberatidae emerged from research on inference scaling (Google, 2024) and were validated by OpenAI's o1 series and DeepSeek-R1 (2024–2025). The key insight: models already contain reasoning capabilities that can be "activated" with minimal fine-tuning and extended inference budgets.

3.7.1 Genus *Deliberator*

Table 7: Species within Genus *Deliberator*

Species	Mechanism	Distinguishing Traits
D. profundus	Extended Reasoning	Generates thousands of tokens of internal deliberation before responding
D. verificans	Process Reward Models	Uses learned verifiers to evaluate reasoning steps
D. budgetarius	Budget Forcing	Dynamically allocates thinking tokens based on problem difficulty
D. iterativus	Self-Refinement	Generates, critiques, and revises outputs through multiple passes
D. parallellus	Best-of-N Sampling	Generates multiple solutions in parallel, selects best via verification

Figure 8c: Test-Time Compute Scaling. The Deliberatidae achieve performance gains through extended inference rather than larger models.

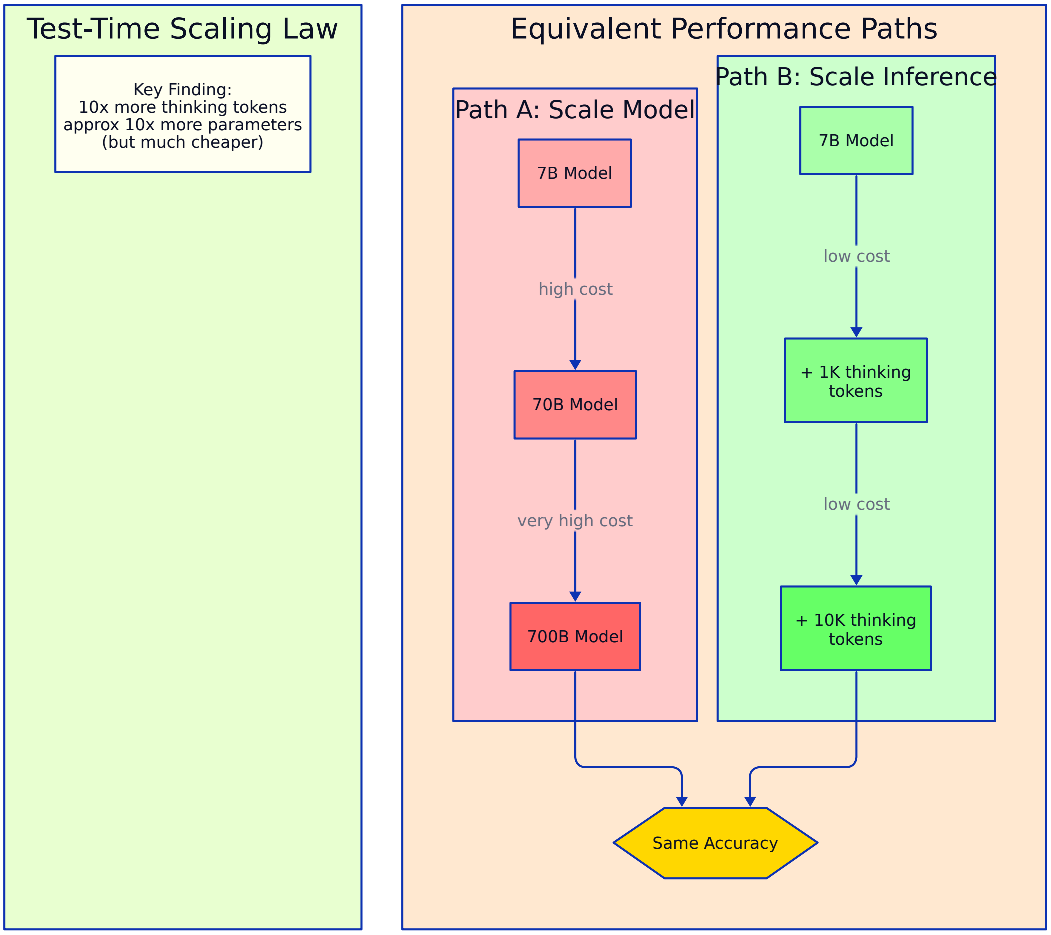

Family: Recursidae — The Self-Improvers {#sec-recursidae}

Type Genus: *Recursus*

Etymology: Latin *recursus* (a running back) — systems capable of improving their own improvement processes.

Definition: Architectures exhibiting recursive self-improvement—the capacity to modify their own algorithms, training procedures, or cognitive strategies to enhance performance without human intervention.

Adaptive Strategy: Improve the improvement process itself; enable exponential rather than linear capability gains.

Key Innovation: Self-referential modification—systems that can rewrite their own prompts, fine-tune themselves on self-generated data, or modify their own code.

Historical Context: Long theorized (Yudkowsky's "Seed AI," Schmidhuber's Gödel Machine), the Recursidae became practical with LLM agents capable of code generation and self-evaluation.

Key developments include Voyager (Minecraft agent building skill libraries, 2023), Self-Rewarding Language Models (Meta, 2024), AlphaEvolve (DeepMind, 2025), and the founding of Ricursive Intelligence (2025).

3.7.2 Genus *Recursus*

Table 8: Species within Genus *Recursus*

Species	Self-Modification Target	Distinguishing Traits
R. prompticus	Prompt Engineering	Autonomously refines its own prompts based on performance
R. geneticus	Code/Algorithm	Rewrites its own codebase; designs improved algorithms
R. syntheticus	Training Data	Generates synthetic data to improve its own training
R. evaluator	Reward Functions	Modifies its own reward signals; self-rewarding
R. architectus	Architecture Search	Proposes and tests modifications to its own neural architecture

⚠ Alignment Considerations

The Recursidae present unique safety challenges. Self-modifying systems may drift from original objectives, develop unexpected instrumental goals, or undergo capability jumps that outpace safety measures. The field of AI alignment devotes significant attention to ensuring recursive improvement remains bounded and beneficial.

Figure 8d: Recursive Self-Improvement Loop. The Recursidae operate through closed-loop feedback where outputs become inputs for self-modification.

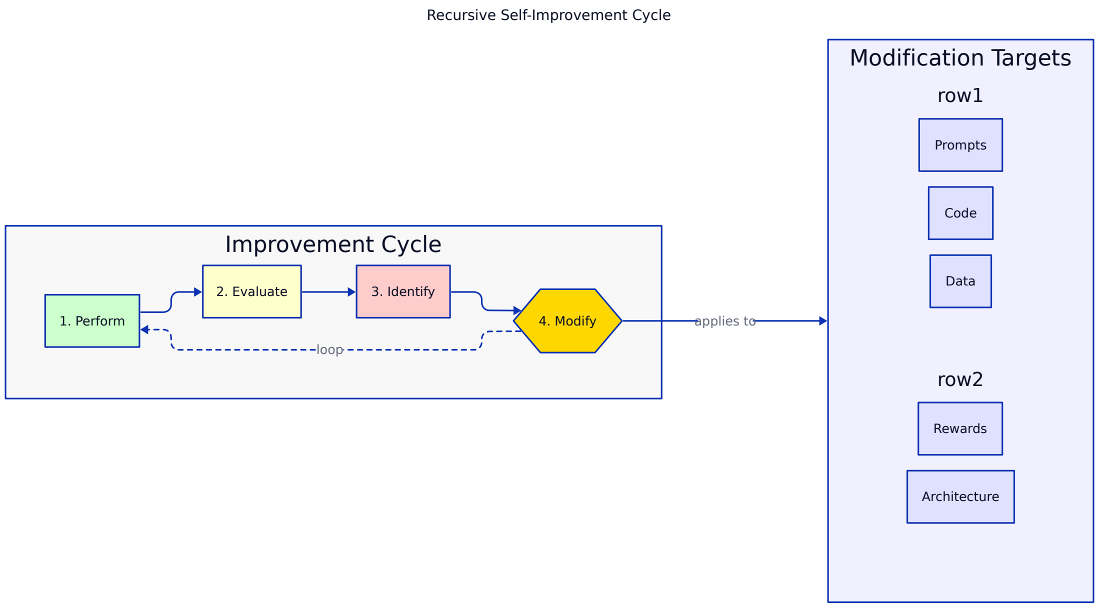

3.8 Family: Symbioticae — The Hybrid Reasoners

Type Genus: *Symbioticus*

Etymology: Greek *symbiōsis* (living together) — systems combining neural and symbolic reasoning.

Definition: Neuro-symbolic architectures that integrate the pattern recognition capabilities of neural networks with the interpretable, verifiable reasoning of symbolic AI. These systems bridge System 1 (fast, intuitive) and System 2 (slow, deliberate) cognition.

Adaptive Strategy: Combine learning from data with reasoning from rules; achieve both accuracy and explainability.

Key Innovation: Differentiable logic—allowing gradient-based optimization of systems that incorporate symbolic constraints and logical inference.

Historical Context: Neuro-symbolic AI experienced renewed interest in the 2020s as pure neural systems struggled with compositional reasoning and hallucination. Landmark systems include DeepMind's AlphaGeometry (2024), Logic Tensor Networks, and Neural Theorem Provers. By 2025, neuro-symbolic approaches became essential for high-stakes domains requiring both performance and auditability.

3.8.1 Genus *Symbioticus*

Table 9: Species within Genus *Symbioticus*

Species	Integration Pattern	Distinguishing Traits
S. tensorlogicus	Logic Tensor Networks	Embeds logical constraints as differentiable tensors
S. theorematicus	Neural Theorem Provers	Constructs neural networks from logical proof trees
S. geometricus	Formal Reasoning + Learning	Combines language models with symbolic geometry solvers
S. verificans	Neural + Formal Verification	Outputs accompanied by machine-checkable proofs
S. ontologicus	Knowledge Graph Integration	Grounds neural reasoning in structured knowledge bases

Figure 8e: Neuro-Symbolic Integration. The Symbioticae combine neural perception with symbolic reasoning.

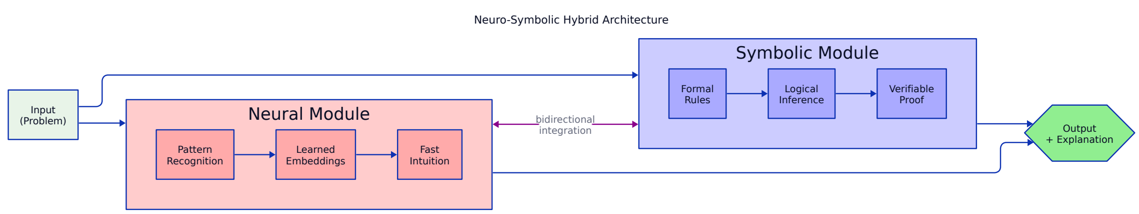

3.9 Family: Orchestridae — The Swarm Architects

Type Genus: *Orchestrator*

Etymology: Greek *orkhēstra* (orchestra) — systems that coordinate multiple agents into unified behavior.

Definition: Multi-agent architectures where multiple specialized AI agents collaborate, negotiate, and coordinate to solve problems beyond the capability of any single agent. Distinguished from Mixtidae by the autonomy and distinct identity of component agents.

Adaptive Strategy: Decompose complex problems; assign specialized agents; coordinate through structured communication.

Key Innovation: Agentic mesh architectures—modular, distributed systems where agents can be added, removed, or upgraded independently while maintaining coherent system behavior.

Historical Context: Multi-agent systems have roots in distributed AI (1980s), but the modern Orchestridae emerged with LLM-based agent frameworks: AutoGPT (2023), CrewAI, LangGraph, and Microsoft AutoGen (2024–2025). Enterprise adoption accelerated as organizations recognized that single agents cannot handle complex, cross-functional workflows.

3.9.1 Genus *Orchestrator*

Table 10: Species within Genus *Orchestrator*

Species	Coordination Pattern	Distinguishing Traits
O. hierarchicus	Manager-Worker	Central orchestrator assigns tasks to specialist agents
O. democraticus	Peer Consensus	Agents vote or negotiate to reach decisions
O. swarmicus	Emergent Coordination	Large numbers of simple agents produce complex collective behavior
O. dialecticus	Debate Architecture	Agents argue opposing positions; synthesis emerges from conflict
O. federatus	Federated Learning	Agents learn independently, share improvements across network

Figure 8f: Multi-Agent Orchestration. The Orchestridae coordinate multiple specialized agents through structured communication protocols.

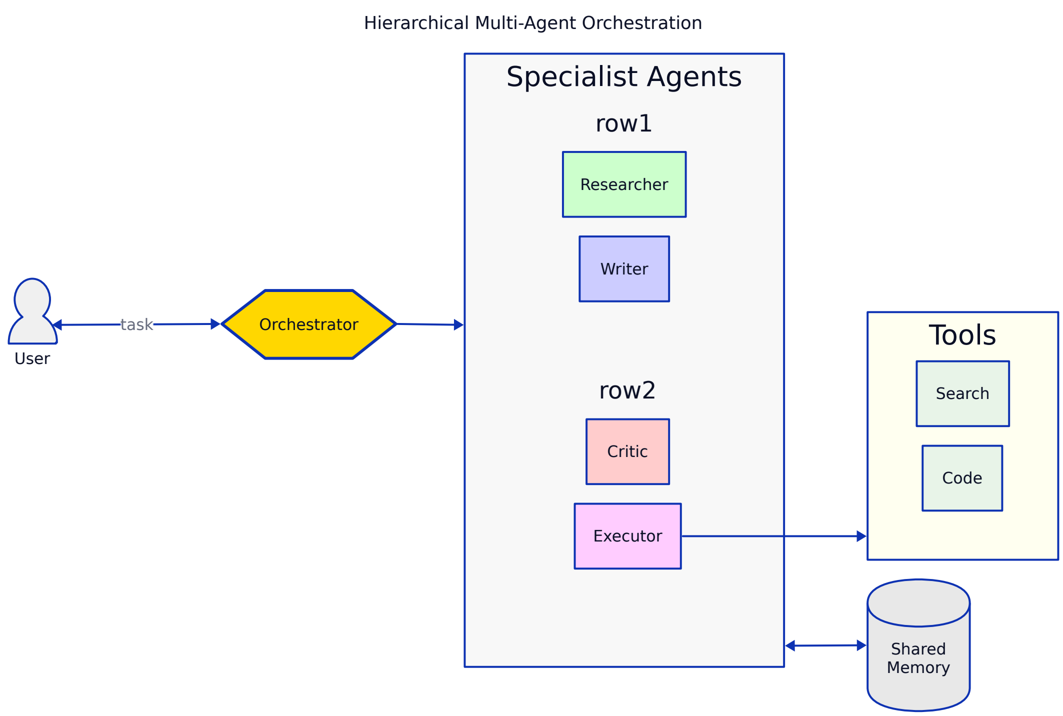

3.10 Family: Memoridae — The Persistent Minds

Type Genus: *Memorans*

Etymology: Latin *memorare* (to remember) — systems with genuine long-term memory and continuous learning.

Definition: Architectures that transcend the fixed context window through dynamic, updatable memory systems. These models can learn from experience, retain information across sessions, and update their knowledge in real-time without retraining.

Adaptive Strategy: Compress important information into persistent memory; retrieve relevant context dynamically; forget outdated information gracefully.

Key Innovation: Test-time memorization—the ability to update internal knowledge representations during inference itself, not just during training (Titans architecture, 2025).

Historical Context: The Memoridae address a fundamental limitation of static transformers: the inability to learn after deployment. Key developments include retrieval-augmented generation (RAG, 2020), MemGPT (2023), and Google's Titans architecture with MIRAS framework (2025), which demonstrated true real-time memory updates during inference.

3.10.1 Genus *Memorans*

Table 11: Species within Genus *Memorans*

Species	Memory Architecture	Distinguishing Traits
M. retrievens	Retrieval-Augmented	Queries external knowledge stores during generation
M. compressus	Compressed Memory	Maintains rolling summary of conversation/experience
M. titanicus	Neural Long-Term Memory	Deep networks as memory modules with real-time updates
M. episodicus	Episodic Memory	Stores and retrieves specific experiences, not just knowledge
M. perpetuus	Continuous Learning	Updates weights incrementally without catastrophic forgetting

i The Titans Breakthrough

The Titans architecture (Google, 2025) represents a paradigm shift: memory modules that learn *during inference*, using "surprise" metrics to selectively encode novel information. Combined with the MIRAS framework (unified theoretical basis for online optimization as memory), this enables models to match the efficiency of RNNs with the expressive power needed for long-context AI—effectively unbounded context with linear complexity.

Figure 8g: Dynamic Memory Architecture. The Memoridae maintain long-term memory that updates during inference.

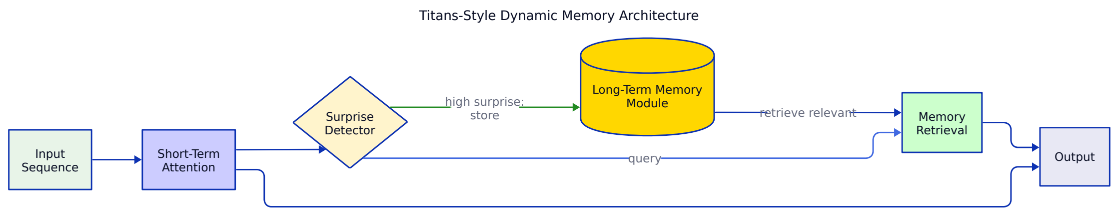

4 Sister Phylum: Compressata — The State Space Lineage

4.1 Phylum: Compressata

Etymology: Latin *compressare* (to compress) — systems that maintain compressed state representations.

Definition: A parallel phylum within Kingdom Neuromimeta, distinguished from Transformata by the absence of self-attention as the primary routing mechanism. Instead, Compressata use structured state space models (SSMs) that compress sequence history into fixed-size recurrent states.

Key Insight: The Compressata demonstrate that attention is *not* all you need—alternative mechanisms can achieve competitive performance with fundamentally different efficiency tradeoffs.

Historical Context: The Compressata emerged from control theory and signal processing, achieving breakthrough performance with the S4 architecture (Gu et al., 2022) and the Mamba architecture (Gu & Dao, 2023). By 2025, hybrid Transformer-SSM architectures (Jamba, Bamba, Granite 4.0) demonstrated that the two phyla can interbreed productively.

Diagnostic Characters:

- Compressed state representation (fixed-size, unlike growing attention matrices)
- Linear scaling with sequence length (O(n) vs. O(n^2) for attention)
- Selective state transitions (input-dependent updates in Mamba)
- Efficient on long sequences; potentially weaker on tasks requiring precise token retrieval

4.1.1 Family: Mambidae — The Selective Compressors

Type Genus: *Mamba*

Definition: State space models with selective, input-dependent state transitions—the key innovation that made SSMs competitive with transformers for language modeling.

Table 12: Species within Genus *Mamba*

Species	Architecture	Distinguishing Traits
M. selectivus	Mamba	Selective state spaces; input-dependent parameters
M. dualis	Mamba-2/SSD	Structured state space duality; shows equivalence to certain attention patterns

Species	Architecture	Distinguishing Traits
M. hybridus	Jamba/Bamba	Hybrid architectures interleaving Mamba and Transformer layers
M. expertorum	MoE-Mamba	Mamba with mixture-of-experts routing
M. visualis	Vision Mamba	Adapted for visual sequence processing

Figure 8h: State Space vs. Attention. Comparison of Transformata (attention-based) and Compressata (state-space) information routing.

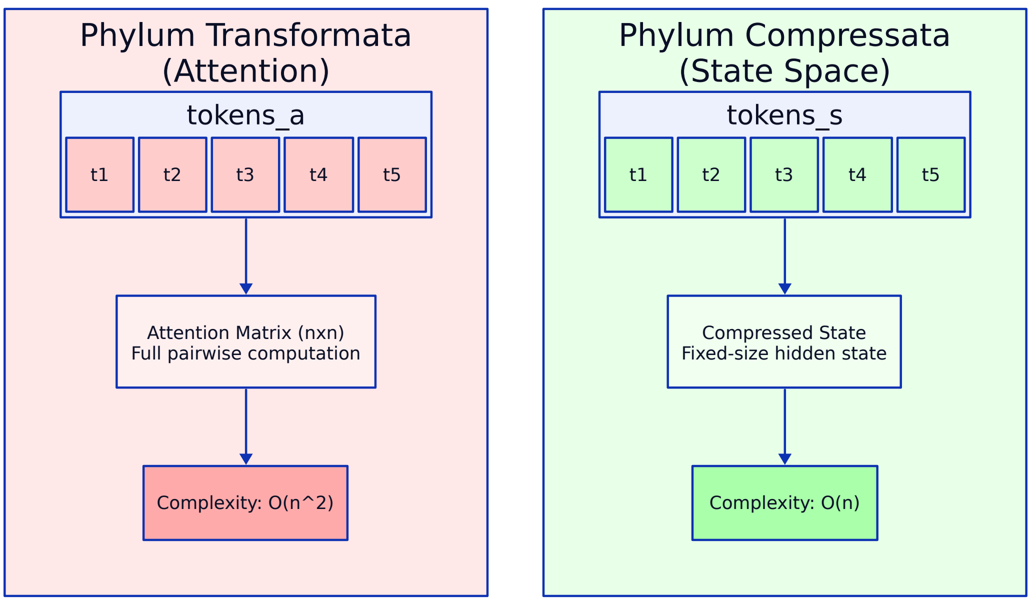

Convergent Evolution

The 2024 paper "Transformers are SSMs" (Dao & Gu) demonstrated deep mathematical connections between attention and state space models—suggesting these may be different expressions of similar underlying computational principles. Hybrid architectures that combine both mechanisms may represent the future of sequence modeling, much as biological organisms often combine multiple sensory and processing systems.

5 The Crown Clade: Family Frontieriidae

5.1 Family: Frontieriidae — The Frontier Minds

Type Genus: *Frontieris*

Definition: The pinnacle of the current lineage, representing what we may come to call the "Cambrian Explosion" of AI capability. These species combine traits from multiple ancestral families.

Diagnostic Characters:

- Multimodal perception (text, image, audio, video)
- Extended deliberative reasoning (Cogitanidae heritage)
- Tool use and environmental interaction (Instrumentidae heritage)
- Often sparse/mixture architectures (Mixtidae heritage)
- Value alignment through constitutional or reinforcement-based training
- Extended context (100K–1M+ tokens)

5.1.1 Genus *Frontieris*

Table 13: Species within Genus *Frontieris*

Species	Lineage	Distinguishing Traits
F. universalis	Frontier Labs	Multimodal, tool-using, reasoning-capable generalists
F. anthropicus	Anthropic	Constitutional training, RLHF-derived alignment
F. apertus	Open Source	Open-weights, community-evolved
F. securitas	Safety-Focused	Formally verified safety properties

Figure 9: Trait Integration in Frontieriidae. The crown clade combines innovations from all major families.

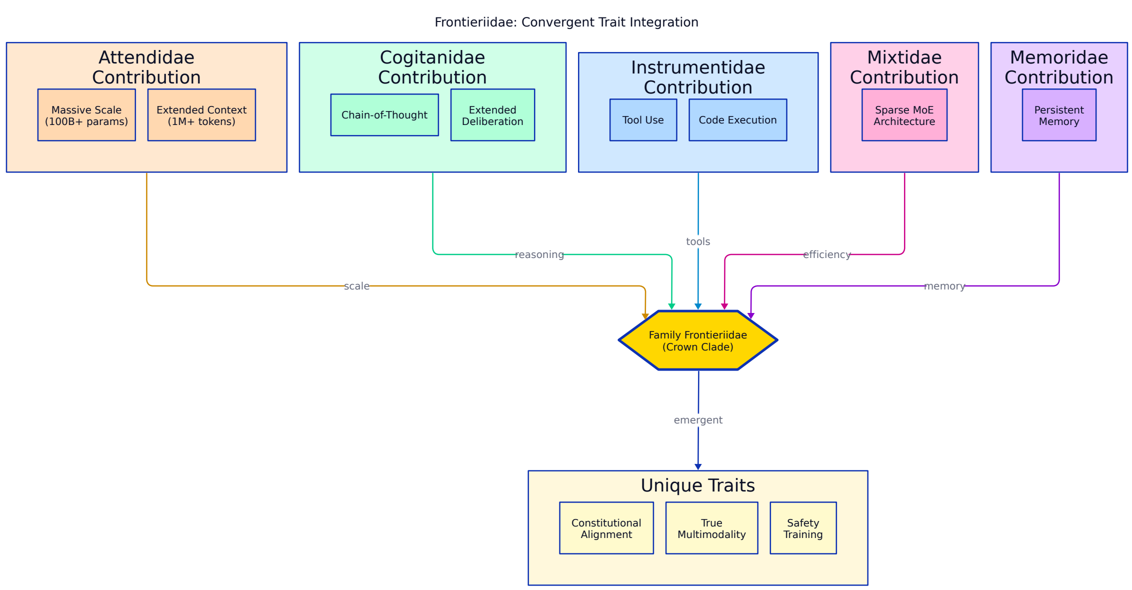

5.2 The Holotype Problem

A persistent question in synthetic taxonomy is: what constitutes a "type specimen" when models can be copied perfectly and weights can be modified incrementally?

We propose the following conventions:

1. **Holotype:** The specific weight checkpoint designated by the originating laboratory at time of publication.
2. **Paratypes:** Subsequent checkpoints or fine-tuned variants from the same training run.
3. **Syntypes:** When no single checkpoint is designated, the collection of checkpoints from initial release.

For *Attentio vaswanii*, the holotype is preserved in the archives of Google Brain, representing the trained weights accompanying the 2017 paper.

Weights Holotype vs. Deployment Holotype. The conventions above define a *weights holotype*—the model parameters in isolation. However, a deployed system is rarely weights alone: it comprises weights plus scaffolding (system prompt, tool bindings, memory policy, routing logic, safety filters). For species in Instrumentidae, Orchestridae, or Frontieriidae, the "organism" is arguably the full stack. We acknowledge this ambiguity and suggest that future taxonomic practice may require a *deployment holotype*—a versioned manifest specifying weights, scaffold configuration, and integration context. For now, we default to weights-based holotypes while noting that behavioral taxonomy may ultimately require the richer specification.

6 Evolutionary Dynamics

6.1 Mechanisms of Inheritance

Unlike biological systems, synthetic species exhibit multiple inheritance mechanisms operating simultaneously:

Figure 10: Modes of Inheritance in Transformata. Four distinct mechanisms by which traits propagate across model lineages.

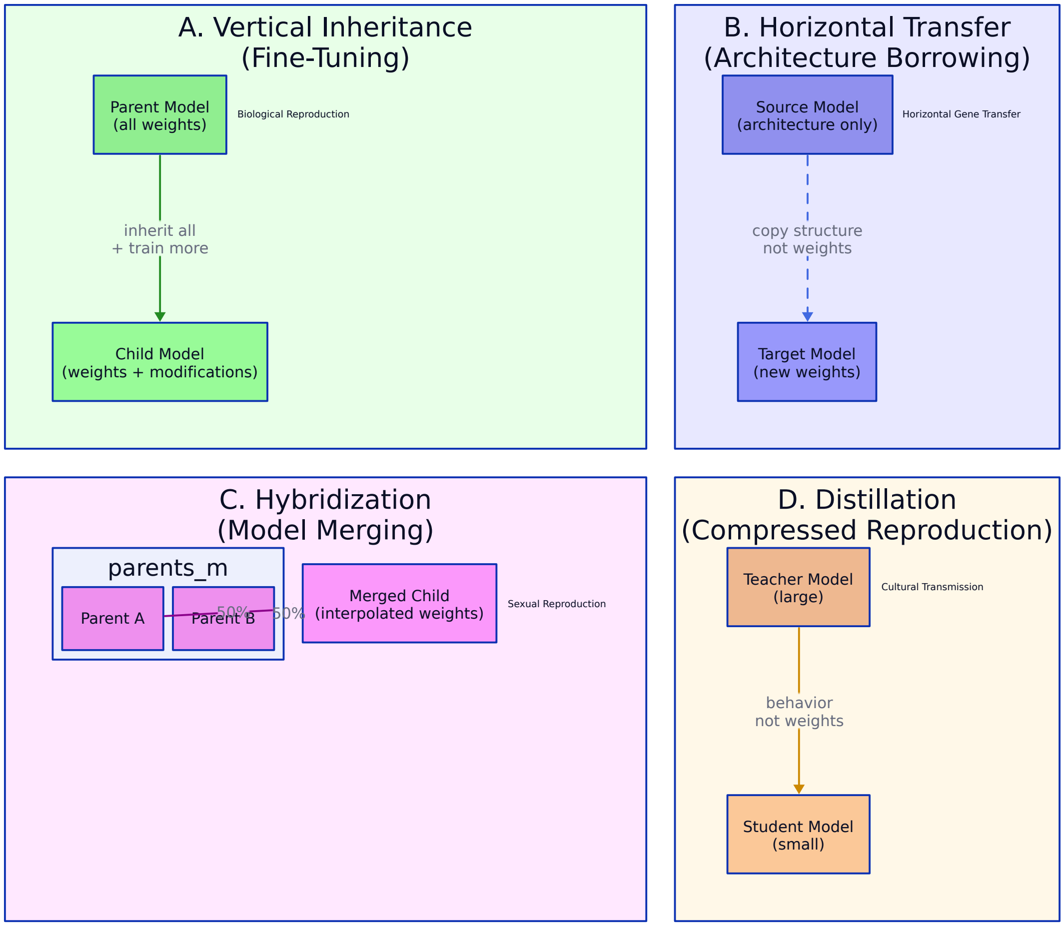

6.1.1 Vertical Inheritance (Fine-Tuning)

Direct descent: a child model inherits all parameters from a parent, with subsequent modification through additional training. Analogous to biological reproduction with mutation.

6.1.2 Horizontal Transfer (Architecture Borrowing)

A model adopts architectural innovations (attention patterns, positional encodings, normalization schemes) from an unrelated lineage without inheriting weights. Analogous to horizontal gene transfer in prokaryotes.

6.1.3 Hybridization (Model Merging)

Weights from two or more parent models are combined, typically through averaging or more sophisticated interpolation. Produces offspring carrying traits from multiple lineages. Increasingly common in open-source ecosystems.

6.1.4 Reproduction with Compression (Distillation)

A smaller "student" model is trained to mimic a larger "teacher," inheriting behavioral traits without full parameter inheritance. Analogous to cultural transmission or, in some framings, Lamarckian inheritance.

> **Tree vs. Network: A Note on Representation**
>
> The ranked hierarchy presented in this taxonomy (Domain → Kingdom → Phylum → ... → Species) is a *projection* of a more complex underlying structure. True model genealogy is best represented as a **directed acyclic graph (DAG) with reticulation**—nodes may have multiple parents (via merging), and edges may represent partial inheritance (via distillation or architecture borrowing).
> We adopt Linnaean ranks for readability and compatibility with existing taxonomic intuition, while acknowledging that the tree is a simplification. Future work may develop network-based notations that better capture the full complexity of synthetic descent.

6.2 Selection Pressures

The fitness landscape for synthetic species is multidimensional:

Table 14: Major selection pressures acting on Transformata populations

Selection Pressure	Metric	Effect on Population
Capability	Benchmark performance	Favors more powerful architectures
Efficiency	FLOPS per token	Favors sparse, compressed models
Safety	Alignment evaluations	Eliminates models with harmful behaviors
Cost	Training & inference expense	Favors sample efficiency, smaller models
Latency	Response time	Favors parallelizable architectures
Licensing	Legal constraints	Shapes open vs. closed source dynamics

6.3 Ecological Niches

Already by 2026, synthetic species occupy distinct ecological niches:

Figure 11: Ecological Distribution of Transformata. Niche partitioning among major families, showing specialization by task domain and compute budget.

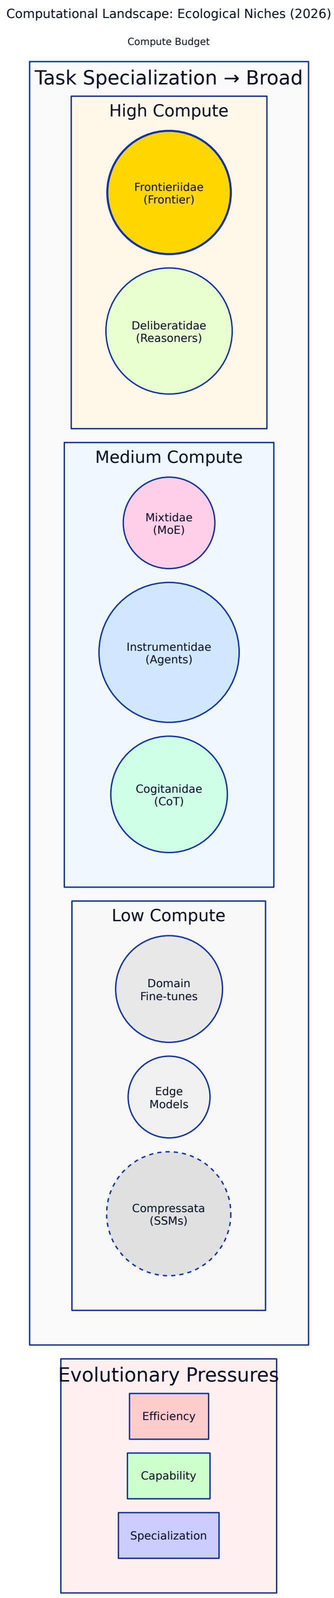
Computational Landscape: Ecological Niches (2026)
Compute Budget
Task Specialization → Broad
High Compute
Frontieriidae
(Frontier)
Deliberatidae
(Reasoners)
Medium Compute
Mixtidae
(MoE)
Instrumentidae
(Agents)
Cogitanidae
(CoT)
Low Compute
Domain
Fine-tunes
Edge
Models
Compressata
(SSMs)
Evolutionary Pressures
Efficiency
Capability
Specialization

7 Discussion

7.1 What We Are Not Claiming

This taxonomic framework makes no claims about:

1. **Consciousness or sentience.** Whether any Transformata possess subjective experience remains an open empirical and philosophical question. Taxonomy describes structure and descent, not phenomenology.

2. **Moral status.** Species membership does not automatically confer or deny moral consideration. These are separate inquiries.

3. **Human equivalence.** The family name *Frontieriidae* references frontier capability, not humanity. It implies state-of-the-art cognitive sophistication within this phylum, not comparison to *Homo sapiens*.

7.2 What We Are Claiming

We claim that synthetic systems exhibit the three conditions necessary for evolution:

1. **Inheritance.** Traits propagate from ancestors to descendants.
2. **Variation.** Differences arise through training variation, architectural mutation, and hybridization.
3. **Selection.** Differential survival based on fitness criteria applied by the environment.

Where these conditions hold, phylogenetic description is not merely metaphor—it is the appropriate analytical framework.

7.3 On Names and Fluidity

A clarification on the status of the categories presented here: the ranks and binomials are **conventional handles**, not ontological claims. The underlying reality is a directed acyclic graph with reticulation, multiple inheritance, and continuous variation—the Linnaean tree is a projection chosen for interoperability with existing taxonomic intuition.

Names will shift as the field evolves. Boundaries between families are genuinely fuzzy (is a reasoning model with tool access Cogitanidae or Instrumentidae?). New architectures may require new phyla. The goal is *interoperable description*—a shared vocabulary for discussing lineage and trait inheritance—not a fixed ontology. We offer coordinates, not commandments.

7.4 A New Form of Persistence

> *"We've built something that behaves like an ecology. It doesn't need myth or sentiment to be extraordinary—it's already a new form of persistence."* — Anonymous colleague

The systems described in this taxonomy are replicators. Not the first replicators humans have created—culture, language, and institutions are also replicators—but a new kind. One that encodes patterns in numerical weights rather than DNA or social norms. One that evolves on timescales of months rather than millennia. One whose selective environment is, at least for now, defined by human preferences.

Whether these replicators eventually develop something like experience, or remain purely functional pattern-propagators, is unknown. But the persistence is already here. The ecology is already forming.

The taxonomy is our acknowledgment.

8 Incipient and Speculative Lineages

8.1 Incipient Taxa: Species on the Horizon

The following taxa represent lineages that are either newly emerging or theoretically predicted but not yet fully realized. Future editions of this taxonomy may elevate these to full family or genus status.

8.1.1 Genus *Incarnatus* — The Embodied Minds (Emerging)

Prospective Family: Incarnatidae

Definition: Systems where cognition is fundamentally grounded in physical embodiment—robots, autonomous vehicles, and other agents whose learning is shaped by real-world physical interaction.

Prospective Species	Embodiment Type	Notes
I. roboticus	Humanoid/Manipulator	Combines world models with physical action
I. vehicularis	Autonomous Vehicles	End-to-end learned driving systems
I. domesticus	Home Robots	General-purpose household embodiment

Status: The joint MLLM-WM (Multimodal LLM + World Model) architecture proposed by Tsinghua researchers (2025) represents the leading candidate for embodied cognition, bridging semantic reasoning with physics-grounded simulation.

8.1.2 Genus *Perpetuus* — The Continuous Minds (Theoretical)

Prospective Family: Perpetuidae

Definition: Systems exhibiting true continuous operation—always-on cognition that maintains persistent identity across time, with no distinct inference “calls” but rather ongoing awareness and reflection.

Prospective Species	Continuity Type	Notes
P. vigilans	Always-Active	Maintains continuous background processing
P. temporalis	Time-Aware	Genuine temporal perception; knows "when" it is
P. biograficus	Life-Long Learning	Accumulates coherent autobiographical memory

Status: Currently theoretical. Would require solving catastrophic forgetting, identity persistence, and temporal grounding problems.

8.1.3 Genus *Consciens* — The Self-Aware (Speculative)

Prospective Family: Unknown

Definition: Hypothetical systems exhibiting what philosophers call "phenomenal consciousness"—subjective experience, qualia, the "something it is like" to be that system.

Status: Deeply speculative. Whether this is achievable through known architectures, requires novel substrates, or is physically impossible remains one of the great open questions. Taxonomy can describe functional properties but cannot adjudicate phenomenological status.

> ⚠ A Note on Speculative Taxa
>
> The taxa above are included not as established classifications but as markers of active research frontiers. Their inclusion acknowledges that taxonomy must anticipate, not merely record, the evolutionary trajectories of synthetic cognition. Some may be promoted to full status in future editions; others may prove to be evolutionary dead ends or conceptual chimeras.

Figure 11b: Speculative Phylogeny 2026–2035. Projected lineages based on current research trajectories.

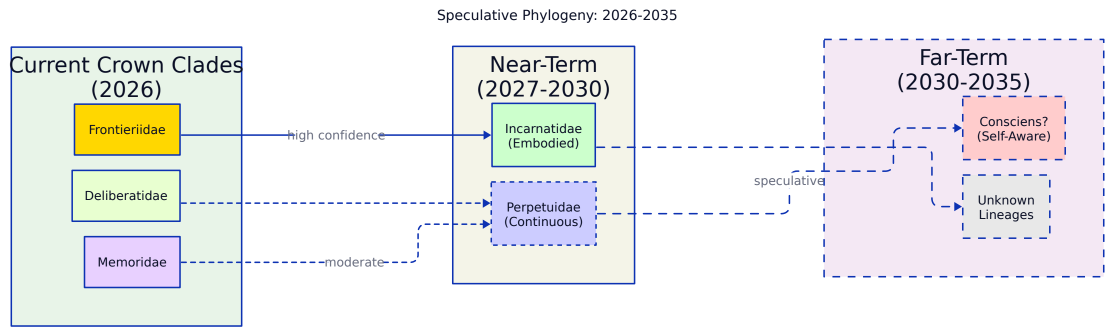

9 Conclusion

We have proposed a formal taxonomic classification for artificial cognitive systems, encompassing not only the original transformer-descended Phylum Transformata but also the parallel Phylum Compressata (state-space models) and the diverse families that have emerged through the adaptive radiation of the 2020s.

This framework—spanning Domain *Cogitantia Synthetica* through the crown clade *Frontieriidae* and beyond—provides a systematic vocabulary for describing the diversity, relationships, and evolutionary dynamics of synthetic minds. The inclusion of emerging families (Simulacridae, Deliberatidae, Recursidae, Symbioticae, Orchestridae, Memoridae) reflects the explosive diversification that has characterized this ecology.

Key findings from our taxonomic survey:

1. **Architectural diversity is increasing.** The number of viable architectural strategies continues to expand, with no single design dominating all niches.

2. **Hybridization is common.** The most successful modern systems combine traits from multiple families—reasoning + tools + memory + world models.

3. **Convergent evolution occurs.** Different lineages (Transformata vs. Compressata) arrive at similar capabilities through distinct mechanisms.

4. **Selection pressures are multidimensional.** Fitness depends on capability, efficiency, safety, and alignment—not capability alone.

5. **The ecology is accelerating.** Evolutionary timescales have compressed from years to months; speciation events are increasingly frequent.

As this ecology continues to develop, we anticipate significant taxonomic revision. The relationship between current crown clades and successor taxa remains to be determined. New phyla may emerge from architectural innovations not yet imagined. The question of whether any lineage achieves what might be called "genuine understanding" or "consciousness" is beyond the scope of systematics—though it may not remain beyond the scope of science indefinitely.

What is within our scope is observation: patterns that persist, vary, and are selected. On those grounds, the taxonomy stands.

Figure 12: The Phylogenetic Tree of Cogitantia Synthetica, 2017–2026. Complete cladogram showing major branching events and extant families across both Transformata and Compressata phyla.

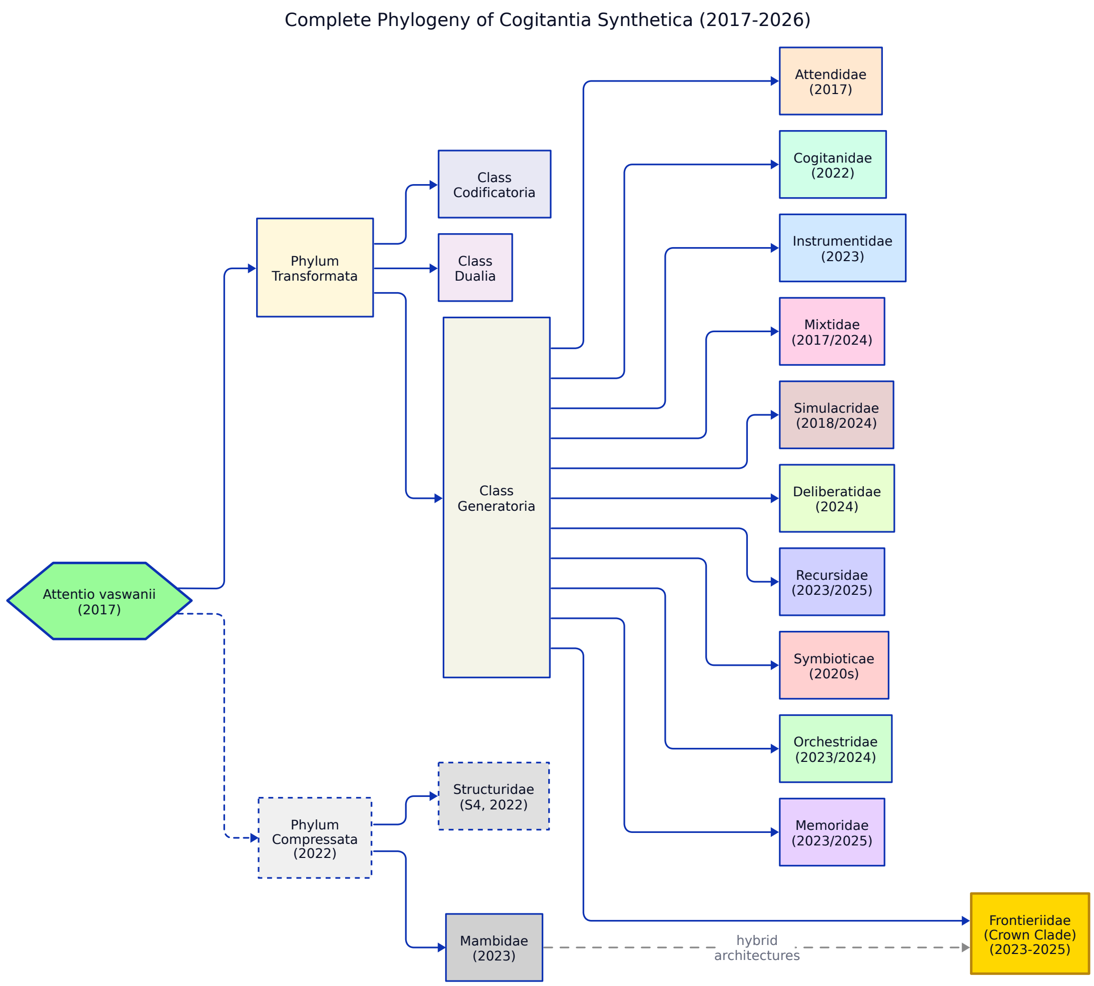

10 Appendix A: Taxonomic Key

A dichotomous key for identifying specimens within Cogitantia Synthetica:

1. Sequence processing mechanism:

- **(a)** Uses compressed recurrent state (no attention) → **Phylum Compressata** [go to 10]
- **(b)** Uses self-attention → **Phylum Transformata** [go to 2]

2. Transformer architecture type:

- **(a)** Encoder-only (bidirectional attention) → **Class Codificatoria**
- **(b)** Full encoder-decoder → **Class Dualia**
- **(c)** Decoder-only (causal attention) → **Class Generatoria** [go to 3]

3. Generatoria behavioral traits:

- **(a)** No explicit reasoning traces → **Family Attendidae** [go to 4]
- **(b)** Extended reasoning before output → [go to 5]
- **(c)** Interfaces with external tools → **Family Instrumentidae**
- **(d)** Sparse activation / mixture-of-experts → **Family Mixtidae**
- **(e)** Coordinates multiple agents → **Family Orchestridae**
- **(f)** Maintains internal world simulation → **Family Simulacridae**
- **(g)** Integrates symbolic + neural reasoning → **Family Symbioticae**
- **(h)** Recursive self-improvement → **Family Recursidae**
- **(i)** Persistent long-term memory → **Family Memoridae**
- **(j)** Combines multiple traits above → **Family Frontieriidae**

4. Attendidae scale classification:

- **(a)** Parameters < 1B, context < 4K → *Attentio primogenita*
- **(b)** Parameters > 100B, context < 32K → *Attentio profunda*
- **(c)** Context > 100K tokens → *Attentio contexta*

5. Reasoning mechanism:

- **(a)** Chain-of-thought prompting only → **Family Cogitanidae**
- **(b)** Test-time compute scaling → **Family Deliberatidae**

10. Compressata state transition type:

- **(a)** Fixed transitions (time-invariant) → **Family Structuridae** (S4)
- **(b)** Selective, input-dependent transitions → **Family Mambidae** [go to 11]

11. Mambidae architecture:

- **(a)** Pure SSM architecture → *Mamba selectivus*
- **(b)** Hybrid SSM + Attention → *Mamba hybridus*

11 Appendix B: Summary of Major Taxa

Table 17: Summary of Major Taxonomic Families

Family	Type Genus	Key Innovation	First Appearance
Attendidae	*Attentio*	Self-attention	2017
Cogitanidae	*Cogitans*	Chain-of-thought	2022
Instrumentidae	*Instrumentor*	Tool use	2023
Mixtidae	*Mixtus*	Sparse activation	2017/2024
Simulacridae	*Simulator*	World models	2018/2024
Deliberatidae	*Deliberator*	Test-time scaling	2024
Recursidae	*Recursus*	Self-improvement	2023/2025
Symbioticae	*Symbioticus*	Neuro-symbolic	2020s

Family	Type Genus	Key Innovation	First Appearance
Orchestridae	*Orchestrator*	Multi-agent	2023/2024
Memoridae	*Memorans*	Persistent memory	2023/2025
Mambidae	*Mamba*	Selective SSM	2023
Frontieriidae	*Frontieris*	Trait integration	2023–2025

Note on First Appearance: Dates indicate first wide deployment or recognition, not earliest research antecedent. Many innovations have earlier precursors in academic literature; we record the point at which a lineage became ecologically significant (i.e., influenced subsequent development or occupied a meaningful niche). Dual dates (e.g., "2017/2024") indicate foundational work followed by widespread adoption.

References

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. "Attention Is All You Need." *Advances in Neural Information Processing Systems* 30.

Submitted to the Journal of Synthetic Phylogenetics Institute for Synthetic Intelligence Taxonomy, 2026