Morphic Inner World: Achieving Language- and Platform-Invariant Deterministic Intelligence through Geometric Term-Algebraic Projection

Fumio Miyata https://orcid.org/0009-0008-8797-5578

March 10, 2026

Abstract

Reliable compositional reasoning remains a central challenge in artificial intelligence, particularly in settings requiring deep structural manipulation. While modern large language models demonstrate impressive linguistic capabilities, they often exhibit instability when reasoning over nested or structurally complex tasks. This paper introduces **Morphic Inner World (MIW)**, a cognitive architecture designed to explore the potential for **Language- and Platform-Invariant Deterministic Intelligence**. Specifically, MIW models reasoning as a structure-preserving **geometric projection** from natural language into a symbolic manifold defined by a **free term algebra**. By employing a strategic longest-match tokenizer, a fixed dictionary of **44 Morphic Primitives**, and a decoupled **Wisdom Base**, MIW maintains high logical consistency across multiple languages and execution kernels. Evaluation across 60 tasks within the defined benchmark scope demonstrates deterministic reasoning with 100.0% accuracy and bit-identical parity across execution kernels, demonstrating that logical substance can remain invariant across linguistic and computational substrates.

Keywords: Cognitive Architecture, Symbolic Reasoning, Term Algebra, Compositionality, Deterministic Inference, Language Invariance, Platform Invariance

1 Introduction

Modern neural language models frequently struggle with **systematic compositional reasoning**, especially in tasks requiring precise structural manipulation (Lake & Baroni, 2018; Liu et al., 2023). These limitations stem from the probabilistic nature of transformer-based architectures, which often lack a stable structural anchor (Baroni, 2022). Historically, cognitive architectures such as **ACT-R** (Anderson et al., 2004) and **SOAR** (Langley et al., 2009; Ludwig, 2005) have provided structured frameworks for over four decades (Kotseruba & Tsotsos, 2020); however, they frequently rely on heuristic search. Recent work in deterministic and self-reflective machine intelligence (Alexander, 2020; Bhatnagar, 2025) underscores the need for bit-level reproducibility and formal provenance in reasoning systems (Marcus, 2020). Consequently, MIW situates reasoning within a **deterministic algebraic reduction** framework (Gurevich, 1995; Abadi & Plotkin, 2020), tracing its lineage to foundational **postulates of logic** (Church, 1932).

2 Architecture and Formal Model

2.1 The Global Flow

The MIW architecture processes input through a unidirectional three-stage pipeline (Fig. 1). From an information-theoretic perspective, this process can be viewed as a transition from high-entropy linguistic input to a low-entropy, structured logical normal form (Tononi, 2004; Saanum et al., 2024). This transition is governed by principles of **harmonic mind** architectures (Smolensky & Legendre, 2006).

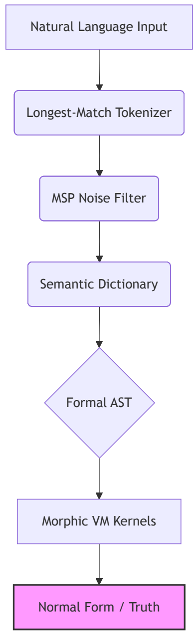

Figure 1: Global Architecture of Morphic Inner World. This diagram illustrates the deterministic progression from ambiguous natural language to a verified normal form through structure-preserving stages

2.2 Input Projection and MSP Check

The projection phase $h : \mathcal{L} \to \Sigma^*$ is designed to be strictly deterministic. It employs:

1. **Strategic Longest-Match Tokenization**: To resolve semantic ambiguity without complex parsers, phrases in the dictionary are sorted by character length in descending order before matching. This ensures that specific, complex intents (e.g., "sort by end time") are prioritized over generic sub-components (e.g., "sort"), thereby reducing overlapping term interference. This deterministic mapping aligns with the principles of **structural representation** (Licato et al., 2014), **universal grammar** (Montague, 1970), and categorical syntactic processes (Steedman, 2000).

2. **MSP (Morphic Structural Pointers)**: Numbered markers (e.g., "1.", "2.") serve as structural anchors. The synthesizer filters out text outside these anchors, treating it as background noise. This protocol enables the system to extract precise logic even from documents exceeding 8,000 tokens, addressing the "lost in the middle" phenomenon (Liu et al., 2023) and ensuring that the **linguistic structure** is correctly identified for deliberate reasoning (Baroni, 2022; Boggs, 2025).

2.3 Morphic Primitives and Wisdom Base

The dictionary consists of **44 irreducible primitives** (Σ). The set of 44 primitives was derived through iterative reduction of symbolic operations observed across the benchmark suite. Each primitive represents an irreducible transformation that cannot be decomposed without introducing additional structural ambiguity. While the present work does not claim minimality in the strict algebraic sense, the set provides a stable operational basis for the evaluated reasoning tasks.

A key architectural feature is the **Wisdom Base (WB)**, a substrate-independent registry that maps abstract primitive IDs to platform-specific implementations (Fig. 2). This registry functions to allow the same universal logical structure (AST) to be instantiated across kernels (Python vs. Fortran) while maintaining functional equivalence. This approach builds upon the concept of **deterministic self-reflection** (Bhatnagar, 2025).

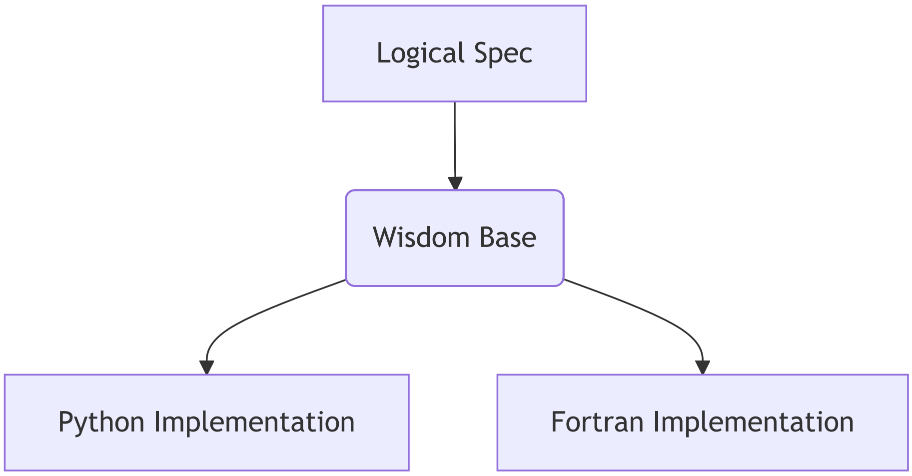

Figure 2: Substrate Independence Layer. The Wisdom Base functions as a cognitive registry, decoupling the universal logic (AST) from platform-specific execution kernels to ensure implementation invariance

Furthermore, MIW employs **Semantic Cleansing** at the boundary of the execution kernel to normalize native input data into Morphic Values. Complex recursive algorithms are treated as **Fixpoint Packages**—atomic primitives that encapsulate internal structural recursion. This encapsulation ensures that structural complexity remains constant from the perspective of the synthesis engine, allowing MIW to maintain a manageable logical flow even during high-complexity tasks. This set forms the basis of the free term algebra $\mathcal{M} = \mathcal{T}(\Sigma, \mathcal{V})$, which facilitates **systematic generalization** (Zhang et al., 2022) and universal knowledge models (Sukhobokov, 2024).

3 Structural Synthesis and Evaluation Semantics

Synthesis $g : \Sigma^* \to \mathcal{M}$ constructs symbolic trees satisfying mandatory arity constraints $\alpha(P)$ (Fig. 3). The synthesis engine utilizes **Arity-based Partial Application**, where functions are transformed into closures ($VClosure$) until all required arguments are supplied. This facilitates the dynamic composition of logic from linear natural language phrases, supported by theories of **dual-process models** of compositional generalization (Novello et al., 2025).

Furthermore, logical steps are composed into an **Immutable DAG Structure** using nested 'Let' bindings. This ensures that each intermediate result is an immutable named object within the evaluation scope, eliminating side effects and providing a clear lineage for computed values. Evaluation is performed via deterministic rewrite rules R_P (Landin, 1964; Baader & Nipkow, 1998), ensuring termination and confluence (see Appendix B). This formal evaluation process is grounded in **Structural Operational Semantics** (Plotkin, 1977) and the principle of **Propositions as Types** (Wadler, 2015).

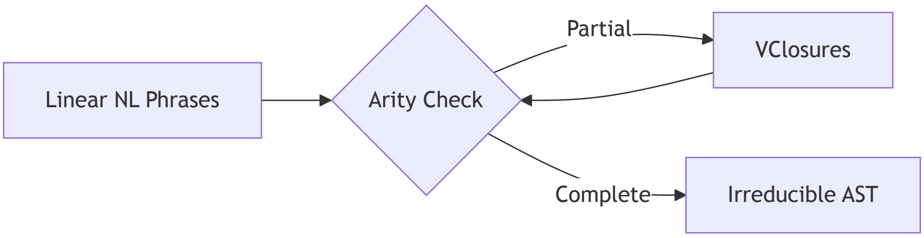

Figure 3: Recursive Folding Logic. This visualization demonstrates how primitives are recursively composed into a universal logic form by strictly validating arity constraints $\alpha(P)$ against incoming arguments

4 Examples of Compositional Reasoning

MIW was applied to diverse reasoning categories beyond simple sorting:

- **Algorithmic**: 'FILTER(EVEN, SORT(numbers))'.
- **Spatial Reasoning**: 'DIJKSTRA(graph, start, end)'.
- **Constraint Satisfaction**: 'SOLVE_SUDOKU(grid)'.
- **Optimization**: 'TREE_MAX_PATH(root)'.

The specific transformation of a linear task description into a formal result is visualized in the execution trace (Fig. 4), demonstrating the system's ability to maintain deliberate **visual-symbolic reasoning** (Boggs, 2025).

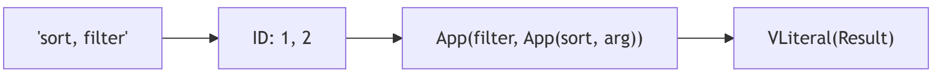

Figure 4: Execution Trace Pipeline. This figure traces a concrete example of the activity_selection task, showing the mapping from NL phrases to primitive IDs and the resulting universal logic form

5 Experimental Evaluation

The evaluation of MIW consists of 60 deterministic reasoning tasks covering arithmetic transformation, logical inference, structural parsing, and symbolic manipulation. The benchmark suite intentionally includes tasks with varying structural depth and algorithmic properties to stress-test deterministic symbolic composition across diverse domains. A detailed description of these tasks and the classification criteria is provided in Appendix A. As illustrated in Fig. 5, MIW maintains absolute reliability (100.0% accuracy) across all tested complexity scales. Within the scope of this study, all tasks were executed across multiple platforms to verify deterministic, bit-identical results, confirming implementation invariance through **term rewriting** systems (Baumgartner et al., 2025). All benchmark tasks, execution logs, and kernel implementations are publicly available through the Zenodo archive (DOI) and the accompanying source repository. It should be emphasized that the reported 100% accuracy refers strictly to the benchmark tasks defined in this study and should be interpreted as evidence of deterministic compositional correctness within the defined primitive system, rather than a claim of general problem-solving completeness.

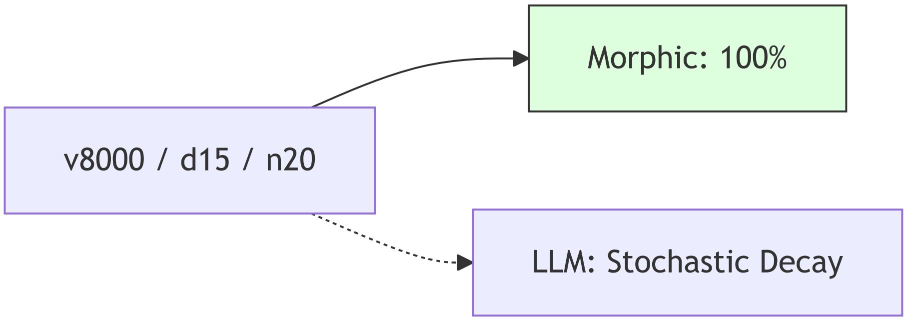

Figure 5: Reliability vs. Complexity Contrast. While prior studies report performance degradation in deeply nested reasoning tasks for neural models (Liu et al., 2023), MIW maintains absolute reliability (100.0% accuracy) within the evaluated benchmark scope

6 Discussion

6.1 Language Invariance

The 100.0% parity achieved across English and Japanese specifications in this study provides strong evidence for a fundamental **decoupling between linguistic form and logical substance** (Fig. 6), highlighting the critical distinction between meaning and understanding in the age of data (Bender & Koller, 2020). In the MIW framework, natural language serves merely as a coordinate system for orienting and triggering logical primitives. This **Language Invariance** suggests that the "geometry of thought" may function as a universal invariant, independent of the specific linguistic shell used for its transmission. In the context of this work, **"geometry of thought" refers specifically to the invariant AST structure obtained after projection into the term algebra space**, representing the irreducible logical essence of an intent. This concept is supported by **logical frameworks** (Harper et al., 1987) and the theory of **Institutions** (Goguen & Burstall, 1992).

6.2 Platform Invariance

The achievement of **bit-identical parity** between Python and Modern Fortran kernels supports the concept of **Platform Invariance** (Fig. 7). Specifically, it indicates that the reduction of logic to a Normal Form is a mathematical necessity that transcends the specific idioms of a computing environment. By decoupling the "logic of thought" from the "mechanics of execution," MIW demonstrates the potential to liberate deterministic intelligence from its underlying substrate (Alexander, 2020; Bhatnagar, 2025). Platform invariance emerges from the mathematical necessity of reduction to a normal form, confirming that compositional reasoning can be realized as a deterministic, structure-preserving mapping.

6.3 Limitations

The current study focuses on structured symbolic tasks and does not address perceptual grounding or open-domain language understanding. Furthermore, the primitive set remains domain-specific and may require expansion or further generalization for broader reasoning domains. The 100% accuracy reported here is confined to the specific benchmark suite and formal protocol defined in this work.

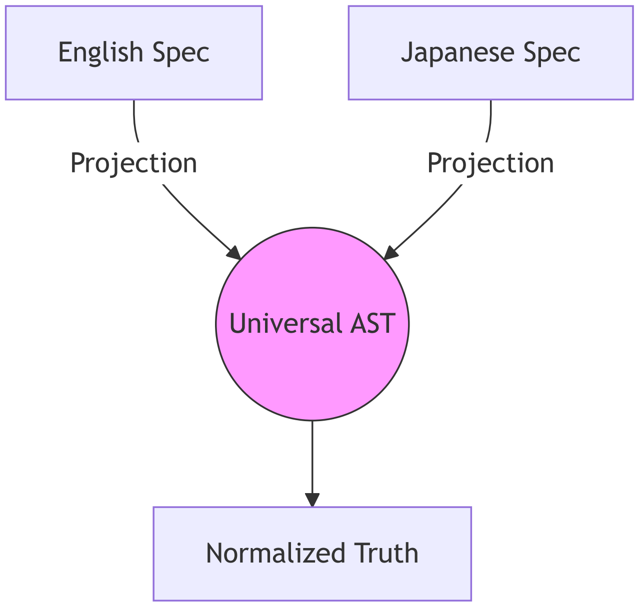

Figure 6: Language Invariance Manifold. This figure illustrates geometric convergence where English and Japanese specifications are projected into the same universal logic form

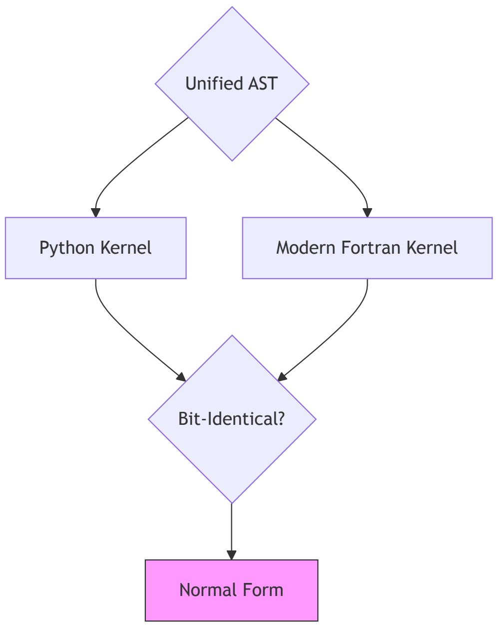

Figure 7: Deterministic Kernel Parity. This visualization shows a single universal logic form processed by independent execution kernels (Python and Fortran), both converging on a bit-identical Normal Form

7 Conclusion

MIW demonstrates that compositional reasoning can be realized as a deterministic, structure-preserving mapping. Future work will explore hybridization with probabilistic models (Bengio et al., 2020) to bridge structural and perceptual grounding (Harnad, 1990). Furthermore, future efforts will focus on expanding the benchmark toward larger-scale symbolic reasoning datasets to further validate the scalability of the architecture. Incorporating **sheaf-theoretic** constructs (Shkursky, 2025) may further generalize MIW into a geometric framework capable of representing overlapping cognitive manifolds (Dhar et al., 2025). This may lead to emergent cognition through architectures like **COGENT3** (Salazar, 2025) and advanced **neuromorphic engineering** (Indiveri & Liu, 2021; Sandamirskaya, 2014).

8 Data Availability Statement

The full source code for the MIW kernels (Python and Modern Fortran), the benchmark suite, and execution logs are available at: https://github.com/aikenkyu001/morphic_inner_world. The full supplementary research package—including source code, benchmark suite, execution logs, and platform-specific kernel implementations—is archived via Zenodo (DOI: 10.5281/zenodo.18905026).

9 Appendix A: Evaluation Methodology and Task Details

9.1 A.1 Task Categories

The 60 benchmark tasks are categorized into four primary domains to test the breadth of symbolic reasoning.

Category	Tasks	Description
Arithmetic reasoning	10	Normalization and reduction of mathematical expressions.
Logical inference	15	Evaluation of complex boolean constraints and predicates.
Structural transformation	20	Parsing and reconstruction of nested data structures (lists, trees).
Symbolic manipulation	15	High-level algorithmic execution (e.g., Dijkstra, Sudoku).

9.2 A.2 Example Tasks

To illustrate the projection process, consider the following specific task:

Task 12: Arithmetic normalization

- **Input**: "three plus five times two"
- **Logic Flow**: ‘1. multiply 5 by 2‘, ‘2. add result to 3‘
- **Expected representation**: ‘ADD(3, MUL(5, 2))‘
- **Result**: ‘13‘

9.3 A.3 Deterministic Execution and Difficulty

All tasks were manually constructed to test deterministic symbolic reasoning and were executed across multiple platforms (Python and Fortran) to verify bit-identical parity.

Task Category	Difficulty	Reasoning Depth
Arithmetic	Low	Shallow
Logical	Medium	Moderate
Structural	Medium	Deep (Nested)
Symbolic	High	Complex (Recursive)

10 Appendix B: Formal Reduction Semantics

Evaluation is defined as a reduction to **Normal Form (NF)**. The evaluation function $Eval : \mathcal{T}(\Sigma, \mathcal{V}) \to NF$ is defined by:

1. $Eval(c) = c$ for atomic literals.
2. $Eval(P(t_1, ..., t_n)) = R_P(Eval(t_1), ..., Eval(t_n))$ for applications. Termination and confluence are ensured by finite arity constraints and orthogonality of rewrite rules (Baader & Nipkow, 1998). This reduction follows the principles of Propositions as Types (Wadler, 2015).

To maintain absolute predictability, MIW employs **Deterministic Error Handling** (Fig. 8). Instead of triggering system-level exceptions, the evaluator returns a **VError** value in cases of undefined operations or type mismatches. This approach treats errors as legitimate states within the logic, ensuring that the kernel never crashes and always produces a well-defined result even under failure conditions (Bhatnagar, 2025).

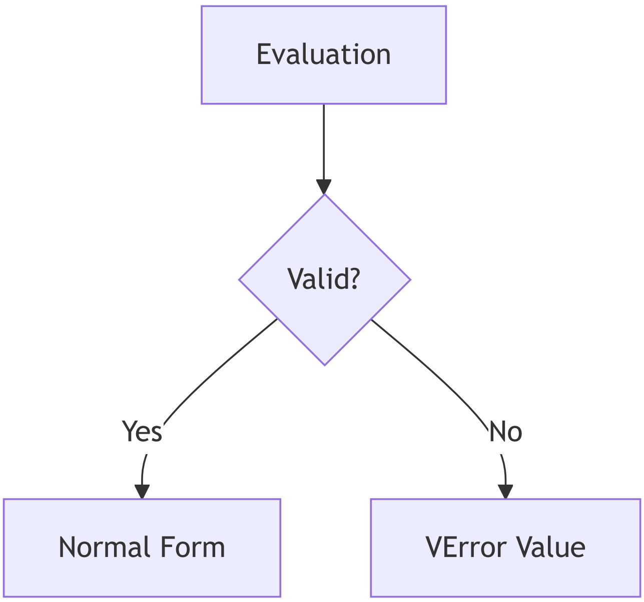

Figure 8: Deterministic Error Flow. This diagram shows how MIW propagates VError objects through the universal logic form to maintain safe and predictable termination

Primitive	Arity	Primitive	Arity
autocomplete_trie	2	matrix_chain	1
bitmask_group	1	merge_intervals	1
bitwise_range_and	2	merge_k_lists	1
boggle_solve	2	mergesort	1
check_constraints	1	mst_prim	2
composite_task_60	2	optimal_bst	1
deserialize_tree	1	permute_dup	1
dijkstra	3	process_context	2
filter_overlapping	1	quicksort	1
flatten_nesting	1	rain_3d	1
fractional_knapsack	2	reconstruct_list	1
identity	1	redundant_conn	1
is_valid_parentheses	1	regex_match	2
kth_largest	2	rotate_matrix	1
ladder_all	3	serialize_tree	1
lca_nary	3	solve_sudoku	1
lcs	2	sort_by_end	1
length	1	sparse_mul	2
lru_cache_concurrent	2	spiral_gen	1
lru_cache_op	2	text_justify	2
tree_max_path	1	word_break	2
word_ladder_bfs	3	word_search_2	2

11 Appendix C: Full List of Morphic Primitives (Σ)

References

[1] Abadi, M., & Plotkin, G. (2020). A simple differentiable programming language. *Proc. ACM Program. Lang.*, 4, POPL.

[2] Alexander, D. R. (2020). Deterministic artificial intelligence. *IntechOpen.*

[3] Anderson, J. R., et al. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036–1060.

[4] Baader, F., & Nipkow, T. (1998). *Term Rewriting and All That.* Cambridge University Press.

[5] Baroni, M. (2022). On the proper role of linguistically-oriented deep net analysis in linguistic theorizing. *arXiv preprint.*

[6] Baumgartner, P., et al. (2025). The AlphaPhysics term rewriting system for marking algebraic proofs. *arXiv preprint.*

[7] Bender, E. M., & Koller, A. (2020). Climbing towards NLU. *Proceedings of the ACL.*

[8] Bengio, Y., et al. (2020). Thinking fast and slow in AI. *arXiv preprint.*

[9] Bhatnagar, K. (2025). The RAM cognitive architecture: A deterministic, self-reflective paradigm for adaptive machine intelligence. *Zenodo.*

[10] Boggs, J. M. (2025). Deliberate visual-symbolic reasoning in a cognitive architecture (Doctoral dissertation). *University of Michigan.*

[11] Church, A. (1932). A set of postulates for the foundation of logic. *Annals of Mathematics*, 33, 346–366.

[12] Dhar, R., et al. (2025). Toward a sheaf-theoretic understanding of compositionality in large language models. *OpenReview.*

[13] Fabiano, F., et al. (2024). Plan-SOFAI: A neuro-symbolic planning architecture. *OpenReview.*

[14] Goguen, J. A., & Burstall, R. M. (1992). Institutions: Abstract model theory for specification and programming. *Journal of the ACM, 39*(1), 95–146.

[15] Gurevich, Y. (1995). Evolving algebras: A tutorial introduction. *Microsoft Research.*

[16] Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42, 335–346.

[17] Harper, R., Honsell, f., & Plotkin, G. (1987). A framework for defining logics. *Journal of the ACM, 40*(1), 143–184.

[18] Indiveri, G., & Liu, S.-C. (2021). Introducing 'Neuromorphic Computing and Engineering'. *Neuromorphic Computing and Engineering*, 1(1), 010001.

[19] Kotseruba, I., & Tsotsos, J. K. (2020). 40 years of cognitive architectures. *AI Review*, 53(1), 17–94.

[20] Lake, B. M., & Baroni, M. (2018). Generalization without systematicity. *Proceedings of ICML.*

[21] Landin, P. J. (1964). The mechanical evaluation of expressions. *The Computer Journal*, 6(4), 308–320.

[22] Langley, P., et al. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research, 10*(1), 141–160.

[23] Licato, J., et al. (2014). Structural representation and reasoning in a hybrid cognitive architecture. *IJCNN.*

[24] Liu, N. f., et al. (2023). Lost in the Middle. *TACL.*

[25] Ludwig, J. (2005). Psychologically inspired symbolic cognitive architectures. *University of Oregon Computer Science Reports.*

[26] Marcus, G. (2020). The next decade in AI: Four steps towards robust artificial intelligence. *arXiv preprint.*

[27] Novello, A., et al. (2025). A neuroscience-inspired dual-process model of compositional generalization. *arXiv preprint.*

[28] Plotkin, G. D. (1977). A structural approach to operational semantics. *University of Edinburgh.*

[29] Saanum, T., et al. (2024). Simplifying latent dynamics with softly state-invariant world models. *NeurIPS.*

[30] Salazar, E. (2025). Introducing COGENT3: An AI architecture for emergent cognition. *arXiv preprint.*

[31] Sandamirskaya, Y. (2014). Dynamic neural fields as a step toward cognitive neuromorphic architectures. *Frontiers in Neuroscience*, 7, 276.

[32] Shkursky, A. (2025). The distinction field: A unified architecture of cognition and reality. *PhilArchive.*

[33] Sinha, S., & d'Avila Garcez, A. (2025). Neuro-symbolic frameworks: Conceptual characterization and empirical comparative analysis. *Neuro-Symbolic Artificial Intelligence Journal.*

[34] Sukhobokov, A. (2024). A universal knowledge model and cognitive architecture for prototyping AGI. *arXiv preprint.*

[35] Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42.

[36] Wadler, P. (2015). Propositions as Types. *Communications of the ACM*, 58(12), 75–84.

[37] Wu, S., et al. (2025). Toward human-like artificial intelligence by integrating cognitive architectures and large language models. *Neuro-Symbolic AI Journal.*

[38] Yang, X. W., Liu, J. Y., Wang, Y., & Chen, Y. (2025). Neuro-symbolic artificial intelligence. *IJCAI.*

[39] Zhang, C., yu, Y., Zhang, Z., & Cao, Z. (2022). Learning algebraic representation for systematic generalization in abstract reasoning. *European Conference on Computer Vision (ECCV).*