Meaning as Syntax:
Logic from Memory Dynamics

Y.Matsuda and ChatGPT5.2

January 9, 2026

Abstract

Human cognition seamlessly integrates memory recall, reasoning, and decision-making, yet existing computational formalisms—symbolic logic, probabilistic models, and Bayesian inference—struggle to explain this integration without introducing rigid representational structures or external semantic interpreters. This paper proposes an alternative perspective grounded in a single principle: meaning is embedded in syntax.

We assume that long-term memory is flat and non-structured, storing neither explicit relations nor semantic content. Structure and meaning arise only through reconstructive processes at recall time. Using Vector Symbolic Architectures as a subsymbolic substrate, we show how binding, aggregation, and similarity operations constitute approximate symbolic syntax whose dynamics generate meaning directly.

Within this framework, induction corresponds to geometric aggregation, deduction emerges as a transition to representational stability, and decision-making arises as a structurally constrained continuation of the same syntactic processes. Logical behavior is not imposed on memory representations but crystallizes from stabilized reconstructive dynamics.

This view reframes logical embedding as an emergent phenomenon rather than a representational technique. Memory, inference, and action are unified as phases of constrained reconstruction, suggesting that logic is not a primitive of cognition but a dynamic regularity arising from memory itself.

1 Introduction

1.1 Motivation

Human cognition exhibits an extraordinary ability to move fluidly from memory recall to flexible reasoning and context-sensitive decision-making. Despite decades of research, no single formalism—symbolic logic, probability theory, or Bayesian inference—has succeeded in reproducing this process in a cognitively faithful manner. Symbolic systems provide precision but lack flexibility; probabilistic approaches capture uncertainty but struggle to explain conceptual stability and semantic structure.

This paper starts from a different premise: meaning is not interpreted on top of syntax, but embedded within it. Rather than assuming pre-defined logical structures or semantic interpreters, we propose that meaning emerges from the operations used to manipulate memory representations themselves.

1.2 Contribution

We argue that Vector Symbolic Architectures (VSA) can serve as a subsymbolic substrate upon which logical embedding emerges as a higher-level phenomenon. Specifically, we claim that:

- Memory is stored in a flat, non-structured form.
- Recall is inherently reconstructive.
- Reasoning corresponds to dynamic reconfiguration of memory fragments.
- Deduction emerges as a stability transition in representational space.

Our contribution is conceptual rather than algorithmic: we reinterpret logical embedding not as an encoding technique, but as an emergent property of vector symbolic memory operations.

2 Background and Related Work

2.1 Symbolic and Probabilistic Models of Cognition

Classical symbolic models of cognition are grounded in the separation of syntax and semantics. Symbols are defined syntactically, and meaning is introduced through interpretation rules, model-theoretic semantics, or external grounding. This architecture enables precise logical inference but presupposes that meaning exists outside the symbolic system itself.

Such systems struggle to account for the fluidity of human cognition. Meaning in human reasoning is context-sensitive, graded, and often ambiguous, whereas symbolic systems require strict well-formedness and explicit rule application. As a result, symbolic models tend to be brittle under noise and partial information.

Probabilistic and Bayesian approaches attempt to overcome these limitations by treating cognition as inference under uncertainty. However, these models still assume a fixed representational structure: a predefined hypothesis space, a generative model, or a semantic variable set. Meaning is encoded indirectly, via probability distributions over symbolic hypotheses.

From the perspective adopted in this paper, both approaches share a common limitation: meaning is not intrinsic to the syntactic operations themselves. Whether interpreted logically or probabilistically, syntax remains a carrier rather than a source of meaning.

2.2 Vector Symbolic Architectures

Vector Symbolic Architectures represent symbols as high-dimensional vectors and define operations such as binding, bundling, and similarity-based retrieval. Unlike classical symbolic systems, these operations are approximate, continuous, and robust to noise.

Historically, VSA has been positioned as an alternative representation scheme: symbols are replaced by vectors, but the underlying semantics is often assumed to remain external. As a result, the symbolic status of VSA has remained ambiguous. Are vectors merely labels, or do they themselves constitute meaning?

This paper adopts the latter position. We argue that in VSA, meaning is not interpreted on top of syntax; rather, meaning is the syntax. A vector's meaning is defined entirely by the operations it participates in and the geometric relations it maintains with other vectors.

2.3 Logical Embedding

Logical embedding has typically been treated as a technical problem: how to encode logical relations into continuous spaces without losing inferential power. Such approaches often aim to approximate classical logic.

Here, we propose a conceptual reinterpretation. Logical embedding is not a representation technique but a cognitive phenomenon. It refers to the emergence of stable relational configurations through repeated memory operations. Logic is not encoded; it crystallizes.

3 Memory Without Structure

3.1 Rejecting Stored Structure

A central commitment of this paper is the rejection of structured memory as a primitive of cognition. We do not assume that long-term memory stores graphs, schemas, rules, logical forms, or latent relational models. Such assumptions merely relocate the problem of meaning from interpretation to storage.

If meaning were stored as structure, cognition would reduce to retrieval and decoding. This view cannot account for the flexibility, context sensitivity, and creative reinterpretation observed in human memory. Empirical evidence from reconstructive memory further undermines the notion of veridical structural storage.

We therefore adopt a stronger position: structure is not stored; it is produced.

3.2 Flat Memory as Non-Interpretive Storage

By flat memory we do not mean unorganized or random storage. Rather, memory consists of fragments that are non-interpretive: episodic traces, perceptual residues, contextual vectors, and motor patterns. These fragments do not encode relations, roles, or logical form.

Crucially, they do not encode meaning. Meaning is not a property of stored representations. It arises only through operations applied to them.

This shift eliminates the need for a semantic interpreter. There is nothing to interpret at rest.

3.3 Recall as the Origin of Meaning

Recall is not retrieval but reconstruction. A cue initiates a process of associative gathering, drawing multiple fragments into a temporary configuration. This configuration has structure, but that structure is ephemeral.

Meaning is identical to this structure. There is no additional semantic layer. The system does not ask what the configuration means; its operational coherence is its meaning.

Thus, recall is already interpretation. Syntax is already semantics.

3.4 Syntax as Meaning

Because memory fragments are meaningless in isolation, all meaning must arise from syntactic operations: binding, aggregation, comparison, and stabilization.

Syntax is no longer a formal shell awaiting interpretation. It is the only locus of meaning. A configuration means what it does under further operations.

This is the sense in which meaning is embedded in syntax: not encoded symbolically, but realized dynamically through reconstructive operations.

3.5 From Reconstruction to Stability

Repeated reconstructions bias future reconstructions. Certain configurations recur, resist perturbation, and become easier to reinstantiate. These stable patterns are what we call concepts.

Concepts are not stored entities. They are attractors in reconstructive dynamics. Their apparent persistence is a consequence of stability, not of representation.

This transition from plastic reconstruction to stable configuration is the precondition for logical embedding. Logic does not precede memory; it emerges from stabilized memory operations.

4 VSA as a Substrate for Logical Embedding

4.1 Approximate Symbolic Primitives

VSA provides a minimal set of operations that resemble symbolic primitives: binding encodes relations, bundling encodes conjunction or collection, and similarity supports approximate identity.

These operations are syntactic in nature. However, because they operate in a continuous space, their outcomes are graded rather than discrete. This gradedness is not a defect but a feature: it allows meaning to vary smoothly with context.

4.2 Meaning Embedded in Syntax

In this framework, syntax does not merely carry meaning. Syntax is meaning.

A vector has no semantic interpretation outside its operational role. Its meaning is exhausted by:

- how it binds with others,
- how it aggregates,
- how stable it remains under perturbation.

This collapses the traditional syntax–semantics distinction. Meaning is no longer an external mapping but an intrinsic property of syntactic manipulation.

4.3 Why VSA Alone Is Not Logic

VSA operations are approximate and noisy. They do not guarantee consistency, transitivity, or logical closure. However, this does not represent a failure.

Human cognition does not operate through strict logical derivation. Instead, logical behavior emerges when certain configurations become stable enough to resist alternative reconstructions.

This observation motivates a shift from logic as rule-following to logic as stability.

5 Logical Embedding as Emergent Stability

5.1 Induction as Geometric Aggregation

Repeated reconstructions of similar fragment sets lead to geometric aggregation in representational space. Configurations that recur frequently become attractors.

Induction corresponds to this process. No explicit generalization rule is required. Generalization emerges from repeated syntactic recombination.

5.2 Deduction as Stability Transition

Deduction occurs when a configuration becomes sufficiently stable that alternative reconstructions are suppressed. This stability functions as necessity.

Deductive inference is thus not symbolic rule application but a phase transition from flexible reconstruction to rigid configuration. Logical embedding names this transition.

5.3 Conceptualization of Memory

Conceptualization is commonly understood as the formation of abstract representations from concrete experiences. Such accounts presuppose that memory contains elements that can be elevated, generalized, or symbolized.

We reject this view.

In the present framework, conceptualization is not an operation applied to memory contents. It is a property of memory dynamics itself. Specifically, conceptualization refers to the emergence of stability within reconstructive processes.

As argued in Section 3, memory stores no structure and no meaning. Each act of recall reconstructs a temporary configuration from flat fragments. However, reconstruction is not neutral. Repeated reconstructions bias future reconstructions. Certain configurations recur, resist deformation, and become easier to reinstantiate.

This resistance to change is the essence of conceptualization.

A concept is not a stored entity, symbol, or node. It is a stable syntactic configuration: a pattern of binding, aggregation, and relational alignment that persists across contexts. Its apparent semantic content is nothing over and above this stability.

Importantly, conceptualization does not introduce meaning. Meaning was already present as syntax. Conceptualization merely fixes it.

In this sense, concepts are attractors in reconstructive dynamics. They function as implicit constraints on future reconstructions, guiding inference and decision-making without explicit rules.

This view establishes a direct duality with the flat memory assumption. Where memory without structure denies stored meaning, conceptualization explains how meaning nevertheless appears. Logical embedding becomes possible precisely because syntax has stabilized enough to behave as necessity.

Conceptualization is therefore the hinge between memory and logic: the point at which flexible reconstruction transitions into deductive stability.

6 Decision-Making as Constrained Reconstruction

6.1 From Meaning to Action

In conventional cognitive and computational models, decision-making is treated as a process distinct from meaning construction. Symbols are first interpreted, values are assigned, and actions are selected according to optimization or inference criteria.

We reject this separation.

In the present framework, decision-making is not an additional computational stage. It is a continuation of the same syntactic processes that generate meaning. If meaning is embedded in syntax, then action must emerge from syntax as well.

Reconstructed configurations are not inert. They possess an internal structure that constrains how the system can act. Action is the externalization of this structure.

6.2 Decision as Structural Projection

A reconstructed configuration defines a space of possible continuations. Each potential action corresponds to a projection of the current configuration into the space of interaction with the environment.

Decision-making consists in selecting the projection that preserves the internal coherence of the reconstructed structure. This selection is not performed by explicit comparison or evaluation. It emerges from the relative stability of possible continuations.

An action is therefore not chosen because it maximizes a value, but because it is syntactically compatible with the current configuration.

6.3 Syntax as Action

Because meaning is embedded in syntax, and because action is an extension of syntactic reconstruction, meaningful action is syntactically coherent action.

There is no need for an explicit decision rule. The same operations that bind, aggregate, and stabilize memory fragments also bias the system toward particular actions.

In this sense, action is not downstream from understanding. It is understanding enacted.

6.4 Constraint Without Optimization

This framework replaces optimization with constraint. Constraints arise from stabilized configurations in memory dynamics. Actions that violate these constraints are unstable and therefore unlikely to be realized.

This does not imply determinism. Multiple actions may be compatible with a configuration. However, the space of viable actions is sharply constrained by the syntactic structure of meaning.

Decision-making is thus neither random nor optimized; it is structurally constrained.

6.5 Relation to Reinforcement Learning

Reinforcement learning models decision-making as reward maximization. Such models require explicit value functions or reward signals.

In contrast, the present framework does not require value representation. Past interactions influence future actions only insofar as they shape the stability of reconstructive dynamics.

What appears as preference or policy is an emergent bias in reconstruction, not an explicitly learned objective function.

6.6 Relation to Bayesian Decision Theory

Bayesian decision theory assumes a probabilistic model of the world and a utility function defined over outcomes.

Here, no explicit probability distribution or utility function is assumed. Uncertainty is implicit in the variability of reconstruction, and action selection reflects the relative stability of competing syntactic continuations.

Decision-making is therefore grounded in structural coherence, not in probabilistic expectation.

6.7 Implications for Cognitive Plausibility

Human decisions are often context-sensitive, inconsistent, and resistant to formal rationalization. These properties are difficult to explain within optimization-based models, but follow naturally from a framework in which decisions emerge from reconstructive dynamics.

Because the same syntactic processes generate meaning, inference, and action, this framework offers a unified account of cognition as a single continuous process rather than a pipeline of specialized modules.

7 Limitations and Open Problems

The framework proposed in this paper makes strong theoretical commitments. As a result, its limitations are not incidental shortcomings, but direct consequences of its core assumptions. Clarifying these boundaries is essential for understanding both the power and the scope of the approach.

7.1 Limits of Precision and Formal Logical Inference

Because meaning is embedded in syntactic operations rather than interpreted symbolically, the framework does not support precise formal logic. There is no guarantee of consistency, completeness, or logical closure.

This is not an implementation defect. It follows directly from the rejection of stored structure and explicit rules. Logical behavior emerges only insofar as configurations become sufficiently stable. Beyond that regime, strict deduction cannot be enforced.

As a consequence, this framework is unsuitable for domains requiring exact symbolic reasoning, such as formal mathematics or program verification. Its target is human-like cognition, where approximate, defeasible reasoning is the norm.

7.2 Depth Limits of Variable Binding

Vector symbolic operations degrade under repeated binding and unbinding. Noise accumulates as relational depth increases, placing practical limits on hierarchical structure.

This limitation reflects a fundamental trade-off. By embedding meaning in syntax, the framework gains flexibility and robustness, but sacrifices unbounded compositional depth. Human cognition itself appears to operate within similar constraints, suggesting that this limitation is cognitively plausible, though computationally restrictive.

7.3 Absence of Explicit Normative Criteria

The framework replaces optimization and utility maximization with structural coherence. While this enables context-sensitive decision-making, it lacks explicit normative criteria for optimality or rationality.

As a result, the framework does not explain why a particular decision should be considered correct, optimal, or rational in a formal sense. It explains only why certain actions emerge as stable continuations of meaning-generating processes.

This limitation reflects a deliberate departure from classical notions of rationality. Whether this trade-off is acceptable depends on the modeling goals.

7.4 Evaluation and Empirical Validation

Because meaning, inference, and decision-making are unified as reconstructive dynamics, standard evaluation metrics are difficult to apply. Accuracy, likelihood, or reward maximization do not directly capture the behavior of the system.

New evaluation criteria are required, such as measures of stability, robustness to perturbation, and adaptability across contexts. Developing such metrics remains an open problem.

7.5 Scalability and Implementation Challenges

While vector symbolic operations are computationally efficient, large-scale systems face challenges related to memory size, interference, and long-term stability. Mechanisms such as consolidation, decay, or hierarchical partitioning may be required, but their integration with the pr

8 Conclusion

This paper has proposed a conceptual reorganization of memory, reasoning, and decision-making based on a single guiding principle: meaning is embedded in syntax. Rather than treating meaning as something to be interpreted, decoded, or inferred from symbolic structures, we have argued that meaning emerges directly from the operations used to manipulate memory.

8.1 Summary of the Approach

The framework rests on three central commitments.

First, memory is assumed to be flat and non-structured. Long-term memory does not store graphs, schemas, rules, or latent logical forms. What is stored are fragments that are non-interpretive in isolation.

Second, recall is reconstructive. Structure, relation, and coherence are produced at use time, not retrieved from storage. Interpretation is not an additional process; it is identical to reconstruction.

Third, reasoning and decision-making are not separate computational stages. They emerge from the same syntactic processes that generate meaning. Induction corresponds to geometric aggregation, deduction to stability transitions, and action to structurally constrained continuation.

Taken together, these commitments collapse the traditional distinctions between memory, inference, and decision-making. Cognition becomes a single continuous process of constrained reconstruction.

8.2 Strengths of the Framework

This approach offers several distinctive advantages.

By embedding meaning directly into syntactic operations, it eliminates the need for external semantic interpreters. Meaning is operational rather than representational.

By rejecting stored structure, it avoids the brittleness associated with fixed symbolic schemas and accommodates context sensitivity and reinterpretation.

By treating decision-making as an extension of meaning-generating processes, it provides a unified account of understanding and action, without invoking value functions, utilities, or explicit optimization.

These features align naturally with empirical characteristics of human cognition: flexibility, defeasibility, and sensitivity to context.

8.3 Limits and Commitments

The framework also makes strong commitments, and its limitations follow directly from them.

Because meaning is embedded in approximate syntactic operations, the framework does not support precise formal logic, unbounded variable binding, or explicit normative optimality. Logical consistency and deductive closure are not guaranteed; they arise only insofar as configurations stabilize.

This places clear boundaries on applicability. The framework is not intended as a replacement for formal logic or probabilistic decision theory in domains requiring exact inference. Instead, it targets the explanatory level at which human-like cognition operates.

Importantly, these limitations are not accidental. They are the price of rejecting interpretation in favor of operational meaning.

8.4 From Syntax as Meaning to Logical Embedding

Within these boundaries, logical behavior nevertheless emerges. As reconstructive processes repeat, certain configurations become increasingly stable, resisting alternative reconstructions. At this point, syntactic patterns begin to function as if they were logical necessities.

This transition is what we identify as logical embedding. Logic is not imposed on memory representations, nor encoded as an external calculus. It crystallizes from stabilized syntactic dynamics.

Logical embedding thus depends critically on the absence of stored meaning. Only because meaning is not fixed at rest can it become fixed through use.

8.5 Implications and Outlook

If this perspective is correct, logical embedding should be understood not as a representational trick, but as a phase of cognitive dynamics. It marks the point at which flexible reconstruction gives way to constraint, necessity, and apparent rule-following.

Future work must investigate how such stability transitions can be characterized formally, how they interact with learning and consolidation, and how they relate to explicit symbolic reasoning in humans and machines.

More broadly, this framework suggests that logic, rationality, and decision-making are not primitives of cognition, but emergent regularities arising from the dynamics of memory itself.

In this sense, meaning does not precede syntax. Meaning is syntax, once it has become stable enough to endure.