The Imagination Machine 0: Orientation to a Framework for Embedded Epistemic Systems

Mark Tracy

March 2026

Abstract

The Imagination Machine series develops a framework for understanding how knowledge arises in systems embedded within their environment. Because such systems have no access to an external vantage point, knowledge cannot be defined as correspondence with an independently accessible world. Instead, knowledge must be understood operationally as the stabilization of representations through interaction with an environment.

Across the series a common architecture appears repeatedly. Observations generate representations of relational structure; these representations are compressed to retain invariants while discarding detail; extension operations generate predictions of missing structure; and prediction error drives subsequent updates. This recursive cycle of observation, compression, extension, and update forms the core mechanism of the Imagination Machine.

The purpose of the present document is to orient the reader to the series as a whole. The individual papers may be read independently, but they collectively describe a layered architecture for embedded epistemic systems. Early papers establish the epistemic foundations of the framework, later papers explore its structural manifestations across domains such as analogy, institutional learning, and symbolic representation, and subsequent papers develop computational realizations and philosophical implications for scientific knowledge.

1 Introduction

The Imagination Machine series investigates how an epistemic system embedded within the world can construct coherent representations of that world.

Traditional epistemology often treats knowledge as correspondence between a representation and an independently accessible reality. Embedded systems, however, have no such external vantage point. They interact only with their observational surface. Consequently, knowledge must be defined in operational terms: representations are evaluated by their capacity to generate coherent predictions and to remain stable under continued interaction with the environment.

Across the papers of this series a common architectural motif emerges. Learning proceeds through a recursive cycle in which an agent observes its environment, constructs a representation of relational structure, compresses that representation to retain invariant features, extends the compressed structure to predict missing relations, and updates its representation in response to prediction error.

This architecture will be referred to as the *imagination machine.* The individual papers explore how this mechanism manifests across several domains, including cognition, analogy, institutional knowledge production, symbolic representation, computational learning systems, and the methodology of science.

2 The Core Epistemic Loop

The central operation of the framework can be summarized as the following cycle.

1. An agent observes data generated by interaction with an environment.
2. Observations are organized into a representation capturing relational structure.
3. The representation is compressed so that invariant relations are retained while redundant detail is discarded.
4. The compressed representation is extended through prediction of missing relations or future states.
5. Prediction error generated by subsequent observations updates the representation.

Repeated execution of this loop gradually stabilizes representations that capture persistent relational structure in the environment. Such stabilized structures function operationally as knowledge.

This perspective resonates with several research traditions in which learning is understood as a dynamical feedback process. Early cybernetic work emphasized the centrality of feedback loops in adaptive systems [14, 1]. More recent work in neuroscience proposes predictive processing models in which perception and cognition arise through the minimization of prediction error [4, 3]. Reinforcement learning frameworks likewise describe agents that iteratively update internal models based on interaction with their environment [13].

The framework also bears philosophical affinity with Karl Popper's conception of knowledge growth through conjecture and refutation [10, 11, 12]. In Popper's view, scientific theories function as hypotheses that generate testable predictions; empirical feedback then eliminates those that fail to withstand critical testing.

Similarly, evolutionary accounts of knowledge growth have emphasized processes of variation and selective retention. Campbell proposed that scientific and cognitive development proceeds through cycles of hypothesis generation and error elimination [2]. Within the present framework, extension operations generate candidate structural hypotheses, while prediction error functions as a mechanism of selective elimination guiding representational revision.

3 Representation and Closure

A central philosophical challenge for embedded epistemic systems is that representation necessarily involves the imposition of conceptual boundaries upon a world that cannot be accessed independently of those boundaries.

Hilary Lawson has argued that all representation involves acts of closure through which distinctions are drawn and stabilized [8]. In Lawson's account, knowledge arises through the construction of frameworks that impose structure upon experience while remaining open to revision.

The compression operations described in the Imagination Machine framework may be interpreted as formal mechanisms for generating such closures. Observational data are grouped into equivalence classes that preserve selected relational invariants while discarding other distinctions. These closures stabilize representations sufficiently to support prediction and reasoning, while the extension and update stages of the epistemic loop allow those closures to be revised in response to new evidence.

4 A Layered Architecture

Although the papers in the series address diverse domains, they can be viewed as exploring different layers of a single architecture.

- **Epistemic Foundations.** Early papers examine the situation of an embedded observer and introduce the inference–implication loop through which world models stabilize.
- **Dynamical Learning Systems.** Subsequent work develops agent– environment interaction models in which predictive agents recover latent structure from observational data.
- **Structural Reasoning.** Further papers examine mechanisms such as analogy, abstraction, and simplicial completion that enable reasoning systems to generate hypotheses about unseen relations. Classical theories of analogy have long emphasized the role of structure-preserving mappings between domains [6].
- **Institutional Learning.** The framework is extended to communities of interacting agents in which dialogue, compression, and feedback produce evolving institutional knowledge. The dynamics of scientific communities have been studied extensively in the philosophy and sociology of science [7].
- **Symbolic Representation.** Later work shows how conceptual structures may be externalized into symbolic artifacts and interpreted through categorical transformations preserving relational invariants. Category theory provides a natural language for reasoning about structure-preserving mappings between systems [9, 5].
- **Computational Realization.** The architecture is implemented as a learning system whose world model is a dynamically updated knowledge graph interacting with an open textual environment.
- **Philosophy of Science.** Finally, the framework is used to interpret scientific knowledge itself as the stabilization of relational invariants under compression of observational data.

Taken together, these layers suggest that the same epistemic mechanism may operate across multiple levels of organization, from individual cognition to collective scientific practice.

5 Relation to Existing Traditions

The architecture described here bears resemblance to several established research traditions.

Cybernetics emphasized feedback and control as fundamental principles of adaptive systems [14, 1]. Predictive processing models in cognitive science interpret perception and cognition as hierarchical processes in which prediction error drives model revision [4, 3]. Reinforcement learning describes agents that iteratively update policies and value estimates through environmental feedback [13].

Karl Popper's philosophy of science emphasized the iterative interaction between conjecture and refutation as the mechanism by which knowledge grows [10, 11, 12]. The present framework can be viewed as providing a structural and computational interpretation of this dynamic within embedded epistemic systems.

In parallel, category theory has increasingly been used to formalize the structure of learning systems and compositional models of knowledge [5]. These approaches provide mathematical tools for describing how representations preserve relational invariants across transformations.

6 Reading the Series

The papers of the Imagination Machine series may be read independently, but they collectively describe different aspects of the same architecture. Later papers often reinterpret or instantiate principles introduced earlier in the series.

Readers interested primarily in philosophical questions may focus on the early papers concerning epistemic closure and representation. Readers interested in computational architectures may focus on the later papers describing knowledge graph learning systems and experimental environments.

The present document serves simply as an orientation to the series as a whole.

7 Conclusion

The Imagination Machine series explores how embedded systems construct representations capable of supporting prediction, reasoning, and coordinated action. Across the diverse domains examined in the series, a common mechanism appears: knowledge arises through recursive cycles of observation, compression, extension, and update.

The aim of the project is both philosophical and practical. Philosophically, it seeks to clarify how knowledge can arise without appeal to an external vantage point. Practically, it proposes an architectural framework that may guide the construction of artificial systems capable of discovering relational structure within complex environments.

References

[1] W. Ross Ashby. *An Introduction to Cybernetics*. Chapman Hall, 1956.

[2] Donald T. Campbell. Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes. *Psychological Review*, 67(6):380–400, 1960.

[3] Andy Clark. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, 2016.

[4] Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

[5] Brendan Fong and David I. Spivak. *Seven Sketches in Compositionality: An Invitation to Applied Category Theory*. Cambridge University Press, 2019.

[6] Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, 1983.

[7] Thomas S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, 1962.

[8] Hilary Lawson. *Closure: A Story of Everything*. Routledge, 2001.

[9] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 2nd edition, 1998.

[10] Karl R. Popper. *The Logic of Scientific Discovery*. Hutchinson, 1959.

[11] Karl R. Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge*. Routledge, 1963.

[12] Karl R. Popper. *Objective Knowledge: An Evolutionary Approach.* Oxford University Press, 1972.

[13] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction.* MIT Press, 2nd edition, 2018.

[14] Norbert Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine.* MIT Press, 1948.