COSINIA

COSINIA RESEARCH SERIES

The Chinese Room Revisited

Intrinsic Memory and the Computational Foundations of Meaning

AUTHOR
Ray Ketabchi
Reza Ketabchi

VERSION
1.0
PUBLISHED
July 2026

Abstract

The Chinese Room Argument has remained one of the most influential discussions in artificial intelligence, cognitive science, and the philosophy of mind. By distinguishing syntactic symbol manipulation from meaning, it challenged the assumption that computation alone is sufficient to produce semantics. More than four decades later, the emergence of Large Language Models has renewed this discussion, yet the underlying computational question remains unresolved.

This paper revisits the Chinese Room from the perspective of the Architecture of Intelligence. Rather than attempting to refute or defend the original argument, it argues that the enduring significance of the Chinese Room lies not in demonstrating the limitations of syntax, but in revealing the absence of the computational conditions under which meaning can emerge.

The central question is therefore not whether syntax can generate semantics, but what computational substrate is required for a continuously evolving Semantic World to exist.

The Architecture of Intelligence proposes that meaning is not an intrinsic property of symbols, statistical associations, or syntactic computation. Rather, meaning is an emergent phenomenon arising from the continuous evolution of a persistent Semantic World whose semantic organization is formed by persistent Identities, immutable Events, and Identity-Centric Persistent Relations.

This paper argues that Intrinsic Memory provides the missing computational foundation by preserving the semantic continuity required for the continuous evolution of the Semantic World. Rather than rejecting the Chinese Room, it extends its central insight by proposing the computational architecture through which meaning can emerge, evolve, and persist.

Scope of this Paper

This paper examines the Chinese Room Argument within the context of the Architecture of Intelligence.

Its purpose is neither to determine whether machines possess consciousness nor to resolve philosophical questions concerning subjective experience, intentionality, or the nature of mind.

Instead, it addresses a different computational question: What architectural conditions are necessary for meaning to emerge from a continuously evolving Semantic World?

Building upon the conceptual foundations established in the preceding papers of this series, it argues that the distinction between syntax and semantics remains incomplete unless it also considers the computational substrate responsible for preserving the semantic continuity upon which a Semantic World continuously evolves.

This paper therefore serves as a conceptual bridge between the Architecture of Intelligence and the broader scientific discussion surrounding meaning, computation, and artificial intelligence.

Introduction

The question of whether computation can produce understanding has occupied artificial intelligence since its earliest days. Among the many arguments that have shaped this discussion, few have had a greater influence than the Chinese Room Argument. By distinguishing the manipulation of symbols from the understanding of their meaning, it challenged one of the central assumptions of symbolic artificial intelligence and continues to influence contemporary discussions surrounding intelligent systems.

Recent advances in Large Language Models have renewed interest in this debate. Modern systems demonstrate remarkable capabilities in language generation, reasoning, translation, and problem solving, leading many to ask whether increasingly sophisticated computation is approaching genuine semantic understanding. Others argue that regardless of their capabilities, such systems remain fundamentally engaged in statistical or syntactic manipulation rather than meaningful comprehension.

The Architecture of Intelligence approaches this discussion from a fundamentally different perspective.

The preceding papers established that Intrinsic Memory preserves semantic continuity and enables the continuous evolution of a persistent Semantic World. They further argued that meaning is not an intrinsic property of language, symbols, or isolated observations, but an emergent phenomenon arising from the accumulated semantic organization of that continuously evolving Semantic World.

From this perspective, the central question is no longer whether syntax alone can produce semantics. Instead, a more fundamental question emerges: Under what computational conditions can meaning emerge at all?

This distinction changes the focus of the discussion. The Chinese Room primarily examines the computational process occurring inside the room—the manipulation of symbols according to formal rules. The Architecture of Intelligence instead examines the computational environment within which such processing occurs. It asks whether semantic understanding requires more than symbol manipulation; specifically, whether it depends upon a continuously evolving Semantic World whose semantic continuity is preserved through time.

This paper argues that the Chinese Room correctly identifies an important limitation of purely syntactic computation, but that it leaves unanswered the architectural conditions required for meaning to emerge. The Architecture of Intelligence accepts the distinction between syntax and semantics while proposing that the missing component is neither a larger rule system nor a more sophisticated statistical model. Rather, it is a computational substrate capable of preserving the Semantic World through which meaning continuously emerges. That substrate is Intrinsic Memory.

Viewed in this way, the Chinese Room is not rejected. It is revisited from the perspective of computational architecture. Its significance lies not in demonstrating the impossibility of artificial intelligence, but in revealing the absence of the computational foundation required for a continuously evolving Semantic World. The remainder of this paper explores this perspective and considers its implications for the future design of intelligent systems.

The Chinese Room

The Chinese Room Argument, first proposed by John Searle in 1980, has become one of the most influential discussions concerning the relationship between computation and meaning. Although originally formulated as a response to symbolic artificial intelligence, its central question remains highly relevant to contemporary intelligent systems.

The thought experiment considers a person who possesses no knowledge of the Chinese language but is enclosed within a room containing an extensive collection of formal rules. Chinese symbols enter the room as input, and by following these rules, the person produces Chinese symbols as output. To an external observer, the responses appear indistinguishable from those of a native speaker, despite the fact that the individual inside the room has no understanding of the language being processed.

The significance of the argument lies in the distinction it draws between syntax and semantics.

> *Syntax concerns the formal manipulation of symbols according to prescribed rules. It specifies how symbols may be arranged, transformed, or combined without reference to what those symbols represent.*

> *Semantics concerns meaning. Within the Architecture of Intelligence, meaning is understood as an emergent phenomenon arising from the continuously evolving organization of a persistent Semantic World.*

Within the Chinese Room, every operation is performed correctly according to syntactic rules. Symbols are received, transformed, and returned according to an established procedure. At no point, however, does the argument attribute meaning to the individual performing those operations.

The Chinese Room therefore distinguishes between the ability to manipulate symbols and the computational conditions required for meaning to emerge.

> *Syntactic symbol manipulation, regardless of its complexity, is not by itself sufficient for meaning to emerge.*

What the Chinese Room Demonstrated

The Chinese Room has generated decades of discussion because it challenges a fundamental assumption underlying many computational theories of intelligence. Regardless of whether one ultimately accepts or rejects every aspect of the argument, it establishes an important conceptual distinction that continues to shape research in artificial intelligence.

The argument demonstrates that successful manipulation of symbols should not automatically be identified with meaning itself. A computational process may operate correctly according to formal rules while remaining entirely independent of the semantic significance those symbols eventually acquire.

This distinction is significant because it separates computational correctness from the computational conditions under which meaning emerges. A system may produce appropriate outputs, satisfy formal rules, and exhibit sophisticated behavior without establishing the persistent semantic organization required for meaning to arise.

The Architecture of Intelligence accepts this conclusion.

> ***Syntax alone does not give rise to meaning.***

This acceptance is not a limitation of the proposed architecture but one of its foundational assumptions. Symbol manipulation is recognized as an important computational capability, but it is not identified as the source of meaning. Meaning emerges only within a continuously evolving Semantic World whose semantic organization has been preserved through time.

From this perspective, the Chinese Room should not be viewed as an argument against artificial intelligence. Rather, it reveals that formal computation alone does not explain the computational conditions required for meaning to emerge.

The question that remains is therefore not whether syntax is sufficient for semantics, but what computational architecture is necessary for a continuously evolving Semantic World to exist.

The Question Left Unanswered

While the Chinese Room provides a compelling argument concerning the limitations of syntactic computation, it does not attempt to identify the computational conditions under which meaning emerges.

The argument explains why syntax alone is insufficient. It does not explain what computational substrate is required for a Semantic World to evolve continuously through time, allowing meaning to emerge from its accumulated semantic organization.

This distinction is fundamental.

The Architecture of Intelligence proposes that meaning is neither an intrinsic property of symbols nor a consequence of increasingly sophisticated rule systems or statistical associations. Rather, meaning is an emergent phenomenon arising from the continuous evolution of a persistent Semantic World.

This shifts the discussion from computational processes to computational architecture.

Rather than asking whether symbol manipulation produces meaning, the Architecture of Intelligence asks a more fundamental question:

> *What must exist before meaning can emerge?*

Syntax

↓

?

↓

Meaning

For more than four decades, the space between syntax and meaning has remained the subject of extensive debate. Numerous theories have proposed richer symbolic systems, statistical learning methods, embodied interaction, or external grounding as possible explanations. While these approaches address important aspects of intelligent behavior, they generally focus on the mechanisms through which symbols acquire or express meaning.

The Architecture of Intelligence approaches the problem from a different direction.

> ***Rather than asking how meaning is generated, it asks what allows a Semantic World to evolve continuously.***

Meaning cannot emerge unless the Semantic World itself evolves continuously through the accumulation of persistent Identities, immutable Events, and Identity-Centric Persistent Relations. Without semantic continuity, the Semantic World fragments into disconnected observations regardless of how sophisticated the computations performed upon those observations may become.

> ***The missing computational question is therefore not one of syntax. It is one of semantic continuity.***

The following sections argue that this missing architectural layer is Intrinsic Memory—a computational substrate responsible for preserving the semantic continuity required for the continuous evolution of the Semantic World.

Meaning Requires Semantic Continuity

The preceding discussion established that syntax alone is insufficient for meaning to emerge. The Architecture of Intelligence accepts this conclusion and proposes that the distinction between syntax and semantics is ultimately a distinction between transient computation and continuously evolving semantic organization.

Meaning is not an intrinsic property of individual symbols or isolated observations. Symbols may represent meaning, communicate meaning, or participate in meaningful computation, yet they derive their semantic significance from the Semantic World into which they are integrated.

> *Meaning emerges from the continuous evolution of a persistent Semantic World.*

Within the Architecture of Intelligence, the Semantic World evolves through the lawful accumulation and organization of persistent Identities, immutable Events, and Identity-Centric Persistent Relations. As observations accumulate, semantic history expands, semantic organization becomes richer, and meaning continuously emerges from that evolving organization.

An Identity preserves semantic reference by allowing observations separated in time to refer to the same semantic entity. Events preserve semantic history without being rewritten. Identity-Centric Persistent Relations preserve the semantic organization connecting those structures throughout the continuous evolution of the Semantic World.

Semantic continuity is therefore indispensable. Without continuity of Identities, Events, and Identity-Centric Persistent Relations, the Semantic World cannot evolve coherently. Individual observations become isolated fragments whose interpretation must continually be reconstructed rather than derived from accumulated semantic organization.

This limitation applies both during observation and during Semantic Retrieval. If semantic continuity is not preserved while observations are admitted, subsequent retrieval cannot reconstruct the intended Semantic World with certainty. It must instead repeatedly infer Identities, semantic organization, and historical continuity from incomplete or ambiguous information, inevitably introducing semantic drift and knowledge loss.

The limitation identified by the Chinese Room therefore extends beyond syntactic computation. Without semantic continuity, no computational process can preserve the continuously evolving Semantic World from which meaning emerges.

Semantic continuity therefore serves two complementary architectural responsibilities.

- During observation, it preserves the continuity of persistent semantic structures, allowing the Semantic World to evolve through accumulated experience.
- During Semantic Retrieval, it enables computation to operate directly upon an already existing Semantic World rather than repeatedly reconstructing semantic context from disconnected historical records.

Semantic continuity is therefore not an implementation strategy or a performance optimization. It is the architectural condition that enables the continuous evolution of a Semantic World. Without that continuity, meaning cannot emerge through time because the semantic organization from which it arises can no longer be preserved.

> ***Syntax governs the manipulation of symbols. Semantic continuity preserves the continuously evolving Semantic World from which meaning emerges.***

Intrinsic Memory as the Missing Computational

Layer

If meaning emerges from the continuous evolution of a Semantic World, then preserving that evolution becomes a computational responsibility.

Conventional computing architectures provide processors to perform computation, storage systems to preserve information, and communication systems to exchange data. Each fulfills a well-defined architectural role. None, however, are defined by the responsibility of preserving semantic continuity.

- Storage preserves records.
- Processors execute instructions.
- Networks transmit information.
- None intrinsically preserve the Semantic World.

The Architecture of Intelligence introduces Intrinsic Memory as a distinct computational substrate whose responsibility is preserving semantic continuity through time. Its purpose is not merely to record information, but to preserve the persistent semantic structures required for the continuous evolution of a Semantic World.

By preserving persistent Identities, immutable Events, and Identity-Centric Persistent Relations, Intrinsic Memory enables semantic history and semantic organization to accumulate continuously rather than being repeatedly reconstructed or discarded.

The architectural dependency established throughout this research series may therefore be summarized as follows.

Intrinsic Memory

↓

Semantic World

↓

Meaning

↓

Intelligent Behaviors

> *Intrinsic Memory preserves the continuously evolving Semantic World from which meaning emerges.*

Viewed from this perspective, Intrinsic Memory is not an enhancement of storage, but a new computational primitive whose responsibility is preserving the computational conditions required for meaning and intelligent behavior to emerge.

Revisiting the Chinese Room

The discussion may now return to the Chinese Room from the perspective established by the Architecture of Intelligence.

The room is capable of manipulating symbols according to formal rules. This is not disputed. The room performs syntactic computation correctly. This is likewise not disputed.

The Architecture of Intelligence accepts these observations while proposing that they do not address the computational conditions required for meaning to emerge.

Within the Chinese Room, symbols enter the room, are transformed according to predefined rules, and leave the room as syntactically appropriate responses. Throughout this process, however, the computational environment contains no continuously evolving Semantic World. There are no persistent Identities accumulating semantic history, no immutable Events preserving that history, and no Identity-Centric Persistent Relations maintaining semantic organization across time.

> ***Without a continuously evolving Semantic World, symbols remain syntactic representations whose apparent meaning exists only from the perspective of an external observer.***

The limitation identified by the Chinese Room therefore extends beyond the distinction between syntax and semantics. It reveals the absence of the computational substrate required for a Semantic World to evolve continuously through time.

Viewed in this way, the room is not merely performing syntactic computation without semantics. It lacks the architectural conditions under which meaning can emerge. The missing component is neither a larger rule system nor a more sophisticated computational procedure, but a persistent Semantic World maintained through Intrinsic Memory.

> *The Chinese Room is therefore revisited not as an argument against artificial intelligence, but as an illustration of the architectural distinction between computation and the continuous evolution of a Semantic World.*

Toward a Computational Architecture of Meaning

The implications of this perspective extend well beyond the Chinese Room itself.

The Chinese Room demonstrates that syntactic computation alone is insufficient for meaning to emerge. The Architecture of Intelligence accepts this conclusion but proposes that the discussion should extend beyond the limitations of syntax toward the computational conditions required for a Semantic World to evolve continuously.

This shift fundamentally changes the architectural focus of artificial intelligence.

Rather than viewing intelligent systems as collections of computational processes that generate meaning through increasingly sophisticated algorithms, the Architecture of Intelligence views intelligent behaviors as computational processes operating upon a continuously evolving Semantic World. The essential question is therefore no longer how symbols are manipulated, but upon what computational foundation those manipulations occur.

Contemporary artificial intelligence systems demonstrate extraordinary capabilities in language processing, statistical prediction, reasoning, planning, and decision making. These achievements represent remarkable advances in computational intelligence. The Architecture of Intelligence does not challenge these capabilities. Instead, it distinguishes computational performance from the computational architecture required for meaning to emerge.

A system may perform increasingly sophisticated computations while relying upon external mechanisms to reconstruct semantic continuity. Such a system may successfully process information, generate language, retrieve knowledge, and perform reasoning without preserving the continuously evolving Semantic World upon which meaning depends.

From this perspective, the defining distinction is architectural rather than algorithmic. Regardless of whether an intelligent system employs symbolic reasoning, statistical learning, neural computation, or future computational paradigms, meaningful intelligent behavior depends upon a continuously evolving Semantic World whose semantic continuity exists independently of the computations performed upon it.

Intrinsic Memory

↓

Semantic World

↓

Meaning

↓

Intelligent Behaviors

Within this framework, meaning is neither an intrinsic property of symbols nor the inevitable consequence of increasingly sophisticated computation. It continuously emerges from the accumulated semantic organization of a persistent Semantic World, while intelligent behaviors operate upon that world rather than creating it independently.

Viewed in this way, the Architecture of Intelligence does not replace existing approaches to artificial intelligence. It introduces a complementary architectural layer beneath them: a computational foundation responsible for preserving the Semantic World from which meaning continuously emerges.

> ***The Chinese Room therefore serves not as the conclusion of this discussion, but as its point of departure. By revealing the distinction between syntax and semantics, it invites a deeper investigation into the computational architecture required for a continuously evolving Semantic World.***

Conclusion

The Chinese Room Argument profoundly influenced artificial intelligence by demonstrating that syntactic symbol manipulation alone is insufficient for meaning to emerge. More than four decades after its introduction, this insight continues to shape discussions concerning the relationship between computation and semantic understanding.

The Architecture of Intelligence accepts this conclusion while extending it in a new direction. Rather than asking only why syntax is insufficient, it asks what computational architecture is required for meaning to emerge in the first place.

This paper has argued that the missing component is neither a richer symbolic system, a larger statistical model, nor a more sophisticated computational procedure. It is a computational substrate capable of preserving the semantic continuity required for the continuous evolution of a Semantic World.

Intrinsic Memory fulfills this responsibility by preserving persistent Identities, immutable Events, and Identity-Centric Persistent Relations. Together, these persistent semantic structures allow semantic history and semantic organization to accumulate continuously through time, enabling meaning to emerge from an evolving Semantic World rather than from isolated symbols or individual computations.

From this perspective, the relationship between syntax and semantics is not bridged by computation alone. It is bridged by the computational architecture that preserves the Semantic World upon which semantic understanding depends.

> ***The central contribution of this work is not a new theory of symbol manipulation, but a computational architecture in which Intrinsic Memory preserves the continuously evolving Semantic World, allowing meaning to emerge and intelligent behaviors to operate upon its accumulated semantic organization.***

References

1. Turing, A. M. (1950). Computing Machinery and Intelligence. Mind, 59(236), 433–460.
2. Searle, J. R. (1980). Minds, Brains, and Programs. Behavioral and Brain Sciences, 3(3), 417–457.
3. Harnad, S. (1990). The Symbol Grounding Problem. Physica D, 42(1–3), 335–346.
4. Newell, A., & Simon, H. A. (1976). Computer Science as Empirical Inquiry: Symbols and Search. Communications of the ACM, 19(3), 113–126.
5. Marr, D. (1982). Vision. W. H. Freeman.
6. Sowa, J. F. (2000). Knowledge Representation: Logical, Philosophical, and Computational Foundations. Brooks/Cole.
7. Floridi, L. (2011). The Philosophy of Information. Oxford University Press.

Citation

Ketabchi, R. The Chinese Room Revisited. Cosinia Research Paper No.4 Version 1.0

← Previous Next →