HEEL: A State-Centric Autobiographical Memory Architecture for Persistent Agents

—A New Paradigm Beyond Session-Centric Ontology

Author Zili Chen
zl.chen.research@outlook.com

https://github.com/L-aaaaaaa/HEEL.git

ArXiv Preprint, April 2026

Abstract

Current agent memory systems are entrenched in a "Session-Centric Ontology"—treating the Session as the sole temporal container and semantic labels as the basis for management. This results in a triple coupling of memory-identity-semantics, leading to systemic pathologies such as cross-session amnesia, personality fragmentation, and cumulative context burden. This paper proposes a "State-Centric Taxonomy," using **Injectability**—the ability to be placed directly into the context window without additional processing—as the sole classification criterion, orthogonal to semantics. From this criterion, a four-part memory structure **HEEL** (History / Experience / External / Logic) is necessarily derived, along with a three-layer identity decoupling model termed **Character-Fenshen-Yinao**. We demonstrate that the "engineering challenges" under the old paradigm are not solved but **naturally dissolved**—cross-session continuity and personality stability become inevitable corollaries under the new axiomatic system. This paper establishes an axiomatic foundation for the agent memory domain for the first time, making the new architecture an endogenous product of theoretical deduction.

This manuscript establishes the theoretical framework and academic priority of the HEEL architecture. The complete metabolic closed-loop engineering instantiation, system implementation, and quantitative

comparative evaluation against LangChain, MemGPT, and other baselines in multi-session scenarios will be reported in a subsequent extended version targeting top-tier conferences. (To comply with double-blind review requirements, an anonymized version with author information removed will be provided alongside this arXiv submission for reviewer use.)

Keywords: Agent Memory; Session-Centric Ontology; Injectability; State-Centric Taxonomy; Autobiographical Memory; Paradigm Shift

1. Introduction: The Pathology of Session-Centric Ontology

Since the rise of conversational agents driven by Large Language Models (LLMs), the **Session** has served as the fundamental design primitive for context management. The Session packages personality settings, dialogue history, and task state into a single container, offering a convenient boundary for resource allocation and isolation. However, this convenience masks a fundamental architectural compromise. As agents evolve from simple chatbots into long-term collaborative partners, three systemic lesions have become unavoidable:

Lesion Type	Clinical Manifestation	User Perception
Cross-Session Amnesia	New sessions cannot access historical memory	"We talked about this yesterday; don't you remember?"
Personality Fragmentation	Personality settings truncated or compressed on context overflow	"Why did you suddenly change personality?"
Context Burden	Large amounts of ineffective redundancy occupy the window	"We've been talking for half an hour, and you've already forgotten the beginning."

These lesions are not bugs in implementation. They are the **structural consequences of a shared ontological commitment**: that Session boundaries should define the scope of memory, and semantic categories should dictate management strategies. All mainstream frameworks today—LangChain, LlamaIndex, MemGPT (Letta), CrewAI, AutoGen—operate under this commitment.

Existing improvement schemes (buffer memory, summary memory, vector retrieval memory, structured memory) all operate within the paradigm of "how to better manage memory **within** a Session." **No one questions whether the Session itself should serve as the fundamental unit of memory management.**

We dissect the lesion of Session-centrism into three interlocked couplings:

- **Semantic-State Coupling**: The "what" of data determines the "how" of management.
- **State-Identity Coupling**: Personality settings and transactional logs coexist in the same Session container; the system cannot differentiate when context overflows.
- **Identity-Semantic Coupling**: The long-term identity marker is tied to a temporary Session ID.

This paper proposes a fundamental paradigm shift: replace the temporal boundary (Session) with the consumable state (Injectability) as the fundamental criterion for memory stratification. We establish a new axiomatic system and deductively derive the corresponding architecture and identity model. We will demonstrate: under this new paradigm, the three major lesions of the old paradigm are no longer "solved," but **cease to be raised**.

✨ This paper focuses on the axiomatic reconstruction of agent memory architecture. In a parallel submission, "Shícè: A Semantic-Driven Time Scheduling Paradigm from Rule Pre-compilation to Runtime Reasoning," we address the agent's **temporal awareness** problem—the Shícè Loop similarly adopts "state-centric" thinking, replacing pre-compiled rules with LLM runtime reasoning. Both theoretical frameworks jointly constitute the cognitive architecture foundation of the JARVIS project.

2. Theoretical Insight: Injectability as an Orthogonal Dimension

2.1 Critique of the Old Axiom

The core assumption of Session-Centric Ontology can be formalized as:

> **Old Axiom (Session Boundary Assumption)**: The agent's memory state is entirely defined by the current Session container. In-Session = online accessible; Out-of-Session = offline inaccessible.

This axiom leads to two irreconcilable internal contradictions:

Contradiction 1: Identity Contradiction. An agent must maintain personality continuity across Sessions, but the old axiom makes persistent identity logically impossible—unless the entire Session is preserved indefinitely (infeasible given finite context windows).

Contradiction 2: Compression Contradiction. When context overflows, memory must be compressed to free space, but the old axiom provides no basis for distinguishing "compressible transactional logs" from "incompressible personality core."

2.2 The New Axiomatic System

We propose three new axioms, forming the cornerstone of State-Centric Ontology.

> **Axiom 0 (Axiom of Existence)**: The Agent possesses a **Lifelong Memory Repository** $\mathcal{M}$, whose existence does not depend on any specific Session. $\mathcal{M}$ is the persistent carrier of the Agent's memory, and its lifecycle is bound to the Agent's identity identifier.

> **Axiom 1 (Axiom of State Orthogonality)**: There exists an injectability criterion $\mathcal{I}(d) \in {0,1}$, indicating whether data item d can, under current technical conditions,

be placed directly into the context window and effectively understood by the LLM without additional processing. $\mathcal{I}(d)$ is **orthogonal** to the semantic content of d and independent of the Session boundary d belongs to.

Operational Definition of Injectability: $\mathcal{I}(d) = 1$ if and only if: (1) the size of d does not exceed a preset token threshold; (2) the format is natively parsable by the LLM; (3) the content is self-contained, requiring no additional context.

Meaning of "Orthogonal": Two data items with identical semantic content (e.g., a raw dialogue transcript vs. a structured fact summary) have fundamentally different injectability. Under semantic classification, both might be grouped together; under state-centric classification, they belong to different layers because their **processing state differs**.

Axiom 2 (Axiom of Metabolic Necessity): The memory repository $\mathcal{M}$ must continuously undergo a metabolic closed loop of **Encoding** → **Consolidation** → **Decay** → **Eviction**; otherwise, infinite growth will lead to dual degradation in retrieval efficiency and context injection quality.

2.3 Formal Definition: Old Paradigm vs. New Paradigm

Under the old paradigm, an Agent is formalized as a function from a Session-Query pair to a response:

$$\text{Agent}{\text{old}} : \text{Session}i \times \text{Query} \rightarrow \text{Response}$$

Under the new paradigm, the Agent is redefined as:

$$\text{Agent}{\text{new}} : \mathcal{M} \times f{\text{concentration}}(\text{Query}, \mathcal{M}) \rightarrow \text{Response}$$

where $\mathcal{M}$ is the lifelong memory repository (Axiom 0), and $f_{\text{concentration}}$ is the **Concentration Function**, dynamically assembling a subset of $\mathcal{M}$ to inject into the current context based on the current Query and the Agent's operating state.

Core of the Paradigm Shift: The Session is **demoted** from a "state container" to a "concentration view"—merely one possible parameterization of $f_{\text{concentration}}$, not the necessary boundary of memory.

3. HEEL Architecture: Deductive Construction from the Axiomatic System

3.1 Theorem 1: The Four-Layer Stratification Theorem

Theorem 1 (Four-Layer Stratification Theorem): Axiom 1 (orthogonality of $\mathcal{I}(d)$) necessarily implies a four-layer logical structure for memory:

- **History** ($\mathcal{I}=0$, source: dialogue stream): Raw transactional logs, append-only, immutable.
- **Experience** ($\mathcal{I}=1$): Ready-to-inject data units, directly consumable.
- **External** ($\mathcal{I}=0$, source: external data): External data sources requiring retrieval/parsing.
- **Logic** (non-data layer): A pure policy layer managing the read/write/evolution/compression of the other three layers.

Proof Sketch (Constructive): Based on the value of $\mathcal{I}(d)$, data naturally falls into two categories. For $\mathcal{I}=0$ data, further distinction is made by source: dialogue stream is History; external data sources are External. Data with $\mathcal{I}=1$ is uniformly categorized as Experience. A policy layer Logic manages the first three layers. This four-part division is exhaustive and mutually exclusive.

Layer	Criterion	Definition	Typical Example
History	$\mathcal{I}=0$, Dialogue Raw Material	Raw transactional logs	User: How's the weather in Beijing? Assistant: Sunny.
Experience	$\mathcal{I}=1$, Ready-to-Use Finished Product	Ready-to-inject data units	Summary: User often works in Beijing; Pref: Prefers concise replies
External	$\mathcal{I}=0$, External Raw Material	External data sources requiring retrieval/parsing	Long PDF documents, SQL databases
Logic	—	Pure policy layer for read/write, eviction, compression	Capacity control, compression triggers

Key Insight: Classification is not determined by what data "means," but by what state it is "in."

3.2 Corollary: The Inevitability of Decoupling

Directly from Axiom 1:

Corollary 1 (Decoupling Corollary): The management strategy for memory should not be determined by semantics, but by $\mathcal{I}(d)$.

Jointly from Axiom 0 and Axiom 1:

Corollary 2 (Session Demotion Corollary): The Session is no longer the necessary boundary for memory management strategy. Memory affiliation is determined by Character ID, and memory processing strategy is determined by $\mathcal{I}(d)$.

3.3 After Decoupling: The Character as the New Anchor

Axiom 0 establishes the persistent existence of $\mathcal{M}$; Axiom 1 liberates data from the dual bondage of semantics and Session. The only stable organizational center of gravity is the **Character**—a persistent personality instance bound to $\mathcal{M}$ that exists across Sessions.

4. Personality Anchoring: The Three-Layer Identity Decoupling Theorem

4.1 Theorem 2: Three-Layer Identity Decoupling Theorem

✨ **Theorem 2 (Three-Layer Identity Decoupling Theorem)**: Jointly derived from Axiom 0 and Axiom 1, the Agent's identity must be decoupled into three orthogonal semantic layers:

- **Character**: The persistent personality identifier bound to $\mathcal{M}$.
- **Fenshen** (分身): A temporary view of the Character under a specific task/focus, i.e., a concrete instantiation of $f_{\text{concentration}}$.
- **Yinao** (义脑): The LLM model driving inference, hot-swappable and independent of memory affiliation.

Proof Sketch: Axiom 0 demands a persistent entity that transcends any individual Session. This is the **Character**. Axiom 1 requires that data in $\mathcal{M}$ be dynamically assembled into context based on $\mathcal{I}(d)$ and current query needs. Each such assembly is a **view** of $\mathcal{M}$—a selection, not the whole. This is the **Fenshen**. The engine executing inference (LLM model) is an independent dimension: switching from GPT-4 to Claude does not change whose memory is being accessed, only the reasoning style. This is the **Yinao**. The three layers are orthogonal because changing one does not necessitate changing the others.

4.2 From Inheritance to Composition: The Paradigm Shift in Memory Injection

Under the old paradigm, a new Session **inherits** the Character's memory—either all or nothing. Under HEEL, a new Fenshen **composes** a subset of the Character's memory—precisely selecting relevant Experience items based on current task, focus, and user intent. Memory injection shifts from a coarse-grained inheritance model to a fine-grained composition model.

5. Paradigm Dissolution

5.1 Dissolution Proposition 1: Cross-Session Continuity

Old Paradigm Problem: Session boundaries cause amnesia. Each new Session starts with blank or manually migrated context.

HEEL Dissolution: By Axiom 0, the lifelong memory repository $\mathcal{M}$ persists across Sessions; by the Concentration Function $f{\text{concentration}}$, each time a new Fenshen (new session view) is created, the system queries all memory items in the Experience layer ($\mathcal{I}=1$) within $\mathcal{M}$ using Character ID as the key, ranks them by composite scoring, and injects them into context. **Thus, Session switching only changes the instantiation parameters of $f{\text{concentration}}$, without affecting the integrity and accessibility of $\mathcal{M}$—cross-session continuity is demoted from a problem requiring manual maintenance to a property automatically guaranteed by the system.**

5.2 Dissolution Proposition 2: Personality Stability

Old Paradigm Problem: Context compression intermixes personality settings and transactional logs; the system cannot differentiate protection priority when forced to prune, causing "personality fragmentation."

HEEL Dissolution: By Theorem 1 (Four-Layer Stratification Theorem), personality settings (such as character descriptions, behavioral constraints) belong to the Experience layer ($\mathcal{I}=1$), while transactional logs belong to the History layer ($\mathcal{I}=0$). The Context Builder (implementation of $f_{\text{concentration}}$) assembles context in a fixed priority order: **System Prompt → Character Experience → Retrieved Relevant Experience → Recent History → Current User Message**. Compression operations begin from the right side (History); Experience is protected by this assembly priority. **Thus, regardless of context window pressure, Experience items defining personality are always the last objects to be pruned—the architectural precondition for personality fragmentation has been removed.**

5.3 Dissolution Proposition 3: Context Management Burden

Old Paradigm Problem: Users must remember "where they said what," manually copying and pasting context across Sessions to maintain continuity.

HEEL Dissolution: By Axiom 0 and $f{\text{concentration}}$, continuity is automatically maintained by the system. When creating a new Fenshen, users need only specify a **focus** (such as project tags, task types); $f{\text{concentration}}$ filters Experience items in $\mathcal{M}$ accordingly. **Thus, the need for manual context migration disappears—users are liberated from the role of "memory administrator" and return to the conversation itself.**

6. Conclusion and Future Work

This paper completes a full paradigm shift from diagnosis to reconstruction: diagnosing Session-Centric Ontology as the root lesion; establishing Axioms 0–2, using Injectability $\mathcal{I}(d)$ as an orthogonal criterion to liberate data from the dual bondage of semantics and Session; deriving the HEEL four-layer architecture and three-layer identity decoupling from the axioms; and demonstrating that the engineering challenges under the old paradigm are naturally dissolved under the new axiomatic system.

The demise of Session-centrism occurs not because a better engineering solution was found, but because its axiomatic premise was proven false. HEEL is not just another memory system, but an axiomatic reconstruction of the agent memory domain.

Future Work: This paper makes public the core theoretical framework of HEEL (Axioms 0–2, Four-Layer Stratification Theorem, Three-Layer Identity Decoupling Theorem) to establish timestamped academic priority. The extended version, planned for submission to top-tier conferences (such as NeurIPS, ICML, or ACL), will build upon this framework with: (1) complete metabolic closed-loop engineering instantiation, including the nine-dimensional MECE data model, incremental compression with two-level capacity control, and concrete implementation of the Summary Pyramid mechanism; (2) quantitative comparative experiments against LangChain (ConversationBufferMemory/SummaryMemory), MemGPT, and other baselines in real multi-session scenarios; (3) user-feedback-driven adaptive weighting strategies for the composite scoring formula.

References

[1] Chase, H. (2022). LangChain: Building applications with LLMs through composability.
[2] Wang, J., et al. MemGPT: Towards LLMs as Operating Systems. NeurIPS 2023 Workshop.
[3] Packer, C., et al. Letta: A Framework for Building Stateful LLM Applications. arXiv:2410.05630, 2024.
[4] CrewAI (2024). CrewAI Documentation: Memory Systems.
[5] Li, G., et al. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. ACM CIKM 2023 Workshop.
[6] Liu, X., et al. Agents: An Open-Source Framework for Autonomous Language Agents. arXiv:2309.07870, 2023.
[7] Tulving, E. Elements of Episodic Memory. Oxford University Press, 1983.
[8] Conway, M. A. Sensory-perceptual episodic memory and its context: Autobiographical memory. Philosophical Transactions of the Royal Society B, 356(1413): 1375–1384, 2001.
[9] Kuhn, T. S. The Structure of Scientific Revolutions (2nd ed.). University of Chicago Press, 1970.

Appendix A: HEEL Core Glossary

Term	Definition
History	Raw transactional logs, $\mathcal{I}=0$.
Experience	Data unit that is ready and can be directly injected into context, $\mathcal{I}=1$.
External	External data source requiring retrieval or parsing, $\mathcal{I}=0$.
Logic	Pure policy layer managing read/write/evolution/compression.
Character	Persistent personality identifier bound to $\mathcal{M}$.
Fenshen	Temporary view of a Character under a specific task.
Yinao	Hot-swappable LLM driving inference.
Injectability	Criterion indicating whether data can be placed directly into the context window without additional processing.

Appendix B: Summary of Formal Symbols

Symbol	Meaning
$\mathcal{M}$	Lifelong memory repository (Axiom 0)
$\mathcal{I}(d) \in {0, 1}$	Injectability criterion for data d (Axiom 1)
$f_{\text{concentration}}$	Concentration function selecting memory subset from $\mathcal{M}$
Session_i	i-th session container
Character_j	j-th Character (memory affiliation subject)

ArXiv Preprint, condensed framework version. Full system implementation and experimental evaluation will be released in a subsequent extended version. Comments and citations welcome.