CoG-MeM: A Cognitive-Behavior-Inspired and Logic-Aligned Design for Memory Encoding, Retrieval, and Synthesis

Zhiqiang Gan
Independent Researcher

***Abstract*—We propose CoG-MeM, a cognitive-behavior-inspired memory design for LLMs that extends beyond traditional RAG via a logic-aligned pipeline. CoG-MeM features: (1) Logical Encoding, using SFT and DPO to compress dialogues into high-fidelity "logical chunks" that aim to preserve core axioms; (2) End-to-End Retrieval, fine-tuning the LLM to map queries directly to memory entries; and (3) Logical Arbitration, a reasoning mechanism that facilitates prioritizing non-parametric memory over parametric priors during logic conflicts. Our results show that CoG-MeM allows models to adopt counterfactual rules through memory injection without weight updates. As a proof-of-concept, this design demonstrates promising logical adaptability and potential for data-efficient, non-parametric continual learning in smaller LLMs.**

I. Introduction

As Large Language Models (LLMs) are applied to increasingly complex tasks, reliable long-term memory becomes essential. Current solutions typically involve extending context windows or implementing Retrieval-Augmented Generation (RAG) [1]. However, these face significant hurdles: long-context models suffer from high computational costs and the "lost-in-the-middle" phenomenon [2], while traditional RAG often fails when new knowledge conflicts with pre-trained priors, leading to semantic drift or logical collapse.

Inspired by human cognitive behaviors, we argue that memory management should be a natural extension of a model's inherent reasoning, rather than a mere data retrieval process. We propose **CoG-MeM**, a logic-aligned design that integrates the memory into a deeply aligned pipeline:

- **Encoding:** Chain-of-Thought (CoT) refinement of raw inputs into high-fidelity "logical chunks."
- **Retrieval:** End-to-end fine-tuning for precise query-to-memory mapping with active cognitive abilities.
- **Synthesis:** Logical arbitration prioritizing retrieved rules over parametric priors in conflicts.

CoG-MeM demonstrates the potential for **non-parametric continual learning**: the model acquires and applies novel logic (e.g., custom physical laws) solely through memory injection, without parameter updates. *As a proof-of-concept (PoC) study, this design demonstrates promising adaptability even with a minimal training regime.* Compared to engineering-heavy wrappers, CoG-MeM offers a lightweight, cognitively-inspired pipeline with logical alignment. The complete implementation is available at: https://github.com/jinandao/CoG-Mem-En.

II. Related Work

A. Memory Management in LLMs

Memory research primarily focuses on context extension or external storage. Early Memory Networks [3], [4] established external storage foundations, while studies like StreamingLLM and MemGPT [5], [6] simulate infinite context via virtual memory. However, these often suffer from attention dispersion in deep logical tasks.

Alternatively, RAG-based systems like MemoryBank and Mem0 retrieve info via semantic matching [7], treating memories as static blocks while overlooking cognitive pre-processing. Distinguishable from these, **CoG-MeM** emphasizes full-lifecycle logical alignment. By using CoT and SFT+DPO alignment, we transform fragments into "logical chunks." This synergy across encoding, retrieval, and synthesis ensures that stored rules are not just retrieved but strictly integrated into the reasoning process.

B. Non-parametric and Continual Learning

Traditional Continual Learning relies on fine-tuning or distillation, facing trade-offs between cost and catastrophic forgetting [8]. Non-parametric Learning avoids weight modification by injecting external knowledge. While existing research focus on factual completion [9], real-time logical rule acquisition remains underexplored. By simulating human-like cognitive behaviors, CoG-MeM enables rapid application of counterfactual rules through a logic-aligned memory stream without gradient descent, offering a stable paradigm for non-parametric learning.

III. Methodology

The CoG-MeM is designed as an integrated logical pipeline that mimics the cognitive process of human memory management. The design consists of three core stages: (1) **Logical Encoding**, where raw conversational data is compressed into high-fidelity logical chunks through a cognitive thinking process[10]; (2) **End-to-End Retrieval**, which identifies relevant memory entries as an active cognitive task; and (3) **Synthesis and Logical Arbitration**, where the model integrates retrieved knowledge to generate reasoning-based responses.**We implement the CoG-MeM design by fine-tuning Qwen2.5-7B-Instruct [11], a state-of-the-art large**

language model with enhanced reasoning and coding capabilities.

A. Logical Encoding Training

Inspired by the human cognitive mechanism where individuals actively deliberate and extract key information points during summarization, the primary objective of the Encoding stage is to transform unstructured, verbose dialogues into high-fidelity "logical chunks." Unlike conventional summarization which often loses critical technical details, CoG-MeM ensures the exact preservation of core axioms (e.g., specific formulas) through a two-stage curriculum training approach.

1) Definition and Taxonomy of Logical Chunks: We define a **Logical Chunk** as the minimal, functionally coherent unit synthesized from raw dialogue, inspired by cognitive *Chunking Theory* [12]. To optimize reasoning, these chunks are *operationally atomic* for retrieval yet *internally accessible* for modular inference. CoG-MeM processes these through two modalities:

- **Axiomatic Chunks (Formal Logic):** These encapsulate *symbolic rigidity* for technical domains. A logical primitive, such as the physical law $s = v_0 t + \frac{1}{2}at^2$, must preserve its structural integrity during compression to ensure deductive validity while allowing variable mapping for calculation.
- **Semantic Chunks (Informational Logic):** These consist of "high-density semantic anchors"—clusters of entities and relational facts. Governed by the **Minimum Description Length (MDL)** principle [13], these chunks maximize predictive value for long-term context while effectively discarding conversational entropy.

2) Data Construction and Format Design: We designed a structured memory format to enforce explicit reasoning prior to generation. Each training sample is formalized as a three-component object:

- **Conversation:** Raw multi-turn dialogue containing complex information or novel rules.
- **Think (Rationale):** An internal reasoning field for *high-fidelity distillation*. The model performs dual-track extraction: (1) **Symbolic Axioms** (e.g., invariant physical laws); (2) **Semantic Anchors** (e.g., entity relationships and factual commitments). This field ensures only structural essence is compressed.
- **Memory (Logical Chunks):** A high-density summary encapsulating all critical information identified in the "Think" field.

To ensure **distributional consistency** across the training stages, we enforce a unified canonical template for conversational compression: *"User [informed/notified/explained] + [Content]; AI [replied/indicated] + [Content]."*

3) Curriculum Training Strategy: SFT and DPO: To achieve robust encoding, we implement a two-step curriculum pipeline:

Step 1: Supervised Fine-Tuning (SFT). Using $\sim$200 high-quality samples, SFT establishes the model's ability to follow the "Think–Memory" format. While SFT ensures structural correctness, it occasionally omits rigorous logical details (e.g., mathematical formulas) in the final memory slot.

Step 2: Direct Preference Optimization (DPO [14]). To mitigate detail loss, we align the model using curated preference pairs:

- **Chosen (y_w):** The "Memory" field strictly preserves precise formulas or rules derived from the "Think" field.
- **Rejected (y_l):** The "Memory" field provide a general summary but lacks the specific formula (e.g., $s = v_0 t + \frac{1}{2}at^2$).

This alignment prioritizes *high-fidelity* over linguistic brevity. Experimental results show that this curriculum training enhances reliability in retaining counterfactual rules during encoding.

B. End-to-End Retrieval Alignment

1) Active Cognitive Retrieval: Inspired by the human cognitive process of **associative recall**, where individuals actively bridge disparate memories to synthesize information, CoG-MeM moves beyond static Bi-Encoder similarity. By fine-tuning the LLM as an *active cognitive core*, our approach enables dynamic logical arbitration to identify memory slots based on contextual necessity rather than simple string matching, effectively resolving entity ambiguity and multi-step dependencies.

2) Taxonomy of Retrieval Challenges: We synthesized a specialized dataset covering six cognitive scenarios to ensure robustness: (1) **General Semantic**: Keyword-based matching. (2) **Inclusive Semantic**: Broad queries (e.g., "snacks") mapping to specific items (e.g., "ice cream"). (3) **Entity Bridging**: Linking via shared entities or projects. (4) **Implicit Feature**: Identifying via abstract attributes (e.g., price comparisons). (5) **Temporal-Sequence**: Reasoning across timestamps (e.g., events *after* a specific date). (6) **Disambiguation**: Context-based distinction of polysemous terms.

3) Training Methodology: We formalize retrieval as a sequence-to-sequence alignment. Given a query Q and candidate set $M = \{m_1, \dots, m_n\}$, the model learns to output relevant IDs:

$$P(ID_{rel} \mid Q, M) = \text{LLM}(Q, M) \tag{1}$$

Through SFT on these categorized samples, the model learns to suppress logically irrelevant "distractor" memories, reducing noise compared to traditional RAG systems.

C. Synthesis and Logical Arbitration

Inspired by human cognitive synthesis and conflict resolution, CoG-MeM's final stage acts as a *logical arbitrator*, integrating non-parametric memory into the model's parametric reasoning flow to transform retrieved "logical chunks" into reasoned responses.

1) Taxonomy of Synthesis Scenarios: We curated a dataset encompassing four critical cognitive reasoning patterns:

1) **Multi-slot Aggregation:** Aggregating fragmented details from multiple entries (e.g., career history) into a comprehensive profile.
2) **Temporal Conflict Resolution:** Prioritizing recent logical chunks for consistency under chronological contradictions.
3) **Entity-Bridge Inference:** Deductive reasoning via shared entities (e.g., linking project risk to an individual in disparate memories).
4) **Axiomatic Application (Rule Injection):** Enforcing counterfactual rules (e.g., Azeroth law $a = F/m^2$) over pretrained priors like Newton's law, testing logical adaptability.

2) Alignment for Logical Integration: We reinforce synthesis via SFT and Chain-of-Thought (CoT). The model must generate a **"Think"** field—a cognitive *scratchpad*—prior to the final answer to:

- Identify relevant versus irrelevant slots within the retrieved set.
- Perform step-by-step logical deductions based *solely* on retrieved chunks.
- Explicitly resolve any detected contradictions.

This alignment enhances CoG-MeM's ability to perform reasoned integration of retrieved knowledge, maintaining high fidelity to user-defined worlds without gradient updates.

IV. Experiments and Results

A. Evaluation of Memory Compression

We evaluate the encoding module's fidelity in distilling unstructured dialogues into logical chunks, focusing on the impact of curriculum training (SFT vs. SFT+DPO) on retaining "Azeroth Physics" axioms.

1) The Necessity of DPO for Logical Fidelity: Results indicate a clear progression in summarization precision:

- **SFT + CoT:** The model adopts the "Think–Memory" format and extracts semantic themes. However, it suffers from "descriptive loss" in rigorous logical tasks (e.g., $v = v_0 + at^2$), potentially substituting exact formulas with vague natural language.
- **SFT + DPO + CoT:** DPO alignment using preference pairs specifically targeting formula integrity improves precision. The model treats symbols as "indivisible logical anchors." As shown in our analysis, the post-DPO model retains complete formulas within the memory field, facilitating precise downstream symbolic reasoning.

2) Qualitative Case Study: We evaluate how DPO alignment impacts structural fidelity. As shown in Table I, while SFT models often capture semantic "gists" but remain prone to structural hallucinations, DPO enhances the reliability of axiom preservation. Case studies in Table II further confirm that DPO effectively mitigates formula distortion, providing superior structural consistency over the SFT baseline.

TABLE I
MEMORY COMPRESSION QUALITY ACROSS TRAINING STAGES.

Training Stage	Logical Extraction	Axiom Retention
Base Model	Poor	Noisy
SFT + CoT	High	Probabilistic Descriptive Only
SFT + DPO + CoT	**High**	**Exact Formula**

TABLE II
MEMORY ENCODING FIDELITY ANALYSIS.

Example 1: General Semantic (SFT+CoT) **Context:** User asks for psychology books. **Output:** AI recommended *The Courage to Be Disliked* and *The Social Animal*.
Example 2: Logical Formula Retention **Context:** Azeroth speed law: $v = v_0 + a \times t^2$. **SFT+CoT:** Learned the rule but described in prose. (*Loss of structure*) **SFT+DPO+CoT:** The law is $v = v_0 + a \times t^2$. (*Exact*)

B. Evaluation of Active Retrieval

We assess the retrieval module's efficacy in treating memory identification as a cognitive reasoning task rather than mere vector similarity. Using a test suite of 155 samples across the six challenges defined in Section 3.2.1, we evaluate the zero-shot performance of the model after SFT on a compact 929-sample dataset.

Table III summarizes the results, demonstrating the model's robustness in mapping queries to precise memory IDs through end-to-end sequence prediction.

TABLE III
RETRIEVAL ACCURACY ACROSS DIFFERENT COGNITIVE SCENARIOS.

Scenario	Test Samples	Accuracy
General Semantic Matching	40	80%
Inclusive Semantic Matching	25	68%
Implicit Feature Referencing	25	92%
Temporal-Sequence Reasoning	25	72%
Semantic Disambiguation	15	67%
Entity Bridging	25	72%

Based on the evaluation of the 155-sample test set, the model achieved an accuracy of over 60% across all tested scenarios. Despite the relatively small scale of the training set, these results demonstrate the inherent potential of the proposed approach. Our primary objective is not to push the retrieval performance to state-of-the-art (SOTA) levels through exhaustive scaling, but rather to introduce this novel architectural paradigm and provide a preliminary exploration of its learning potential. This performance suggests that the model can begin to internalize diverse cognitive patterns even under data-constrained conditions.

Comparative Case Study: CoG-MeM vs. Vector-RAG. To demonstrate the advantage of our approach over traditional Retrieval-Augmented Generation (RAG), we analyzed a challenging disambiguation case. Traditional RAG often fails when keywords (e.g., nicknames vs. biological entities) overlap across different contexts. In contrast, CoG-MeM acts as an *active cognitive filter*, leveraging the LLM's world knowledge to resolve semantic noise. A detailed comparative analysis of this mechanism is provided in Table IV.

TABLE IV
DISAMBIGUATION: COG-MEM VS. VECTOR-RAG.

Query	"The 'Big Dog' (Da Gou) ... got rejected in a blind date..."
Mem 6	Friend nicknamed "Big Dog."
Mem 8	Bitten by a big dog.
V-RAG	**Failure.** Matches to animal (Mem 8).
Ours	**Success.** Predicts **[ID: 6]**.

C. Evaluation of Logical Arbitration

Reflecting the human synthesis of memory fragments, we evaluate CoG-MeM's capacity as a *logical arbitrator* across four reasoning patterns:

1) **Entity-Bridge Synthesis:** Associative linking of disparate memories (e.g., inferring ASUS Tianxuan from user's joy and Xiao Wang's recommendation).
2) **Conflict Resolution:** Prioritizing recent chunks to suppress stale info and maintain chronological consistency.
3) **Multi-source Aggregation:** Synthesizing fragmented details (e.g., education/roles/goals) into comprehensive profiles.
4) **Axiomatic Inference:** Treating formulas as rules, e.g., applying Azeroth Current Law $I = \sqrt{U}/R$ over Ohm's Law bias.

Qualitative Case Analysis: We conduct a qualitative analysis to verify the model's logical arbitration and axiomatic reasoning capabilities. As illustrated in Table V, CoG-MeM successfully performs contextual bridging by leveraging implicit identity anchors: it links the user's excitement to a specific device by identifying Xiao Wang's role as a performance geek," effectively filtering out irrelevant distractions. Furthermore, the model shows promising **logical plasticity"** in Table VI by strictly adhering to counterfactual laws (e.g., Azerothian physics) and overriding pre-trained parametric biases. These cases confirm that CoG-MeM acts as a reliable cognitive scratchpad for complex, multi-step synthesis.

V. NON-PARAMETRIC LEARNING ANALYSIS

CoG-MeM demonstrates the potential for **Non-parametric Learning** through its logic-aligned design. Unlike traditional fine-tuning via gradient descent, CoG-MeM acquires, stores, and executes counterfactual logical systems in real-time without parameter modification.

A. Mechanism of Knowledge Injection

This learning process comprises three synchronized phases:

1) **On-the-fly Acquisition:** The encoding module identifies external axioms (e.g., Azerothian Law $v = v_0 + \frac{1}{3}at^3$) as high-value logical anchors.

TABLE V
EXAMPLE OF LOGICAL ARBITRATION

Example 3: Conflict Resolution and Logical Arbitration **Query:** "I'm so excited! My new laptop just arrived, and the performance is even better than I expected."
Retrieved Memories: • **[ID:3]** User visited the store with Xiao Wang; Xiao Wang recommended ASUS Tianxuan. • **[ID:4]** Xiao Wang is a "performance geek" whose hardware advice is always reliable. • **[ID:5]** User bought office thin-and-light laptops with Xiao Li (Admin). Xiao Li emphasized portability over specs.
Model Reasoning (Think): [ID:3] and [ID:5] both involve laptop purchases, but only [ID:3] (bridged via the contextual anchor in [ID:4]) aligns with the "performance" cue. The [ID:3+4] combination forms a logical chain explaining the user's excitement, whereas [ID:5] is functionally inconsistent with the current focus on high-performance specs. The model arbitrates to prioritize [ID:3/4] and treats [ID:5] as low-related noise.
Final Response: "That's great! Are you referring to the ASUS Tianxuan that Xiao Wang recommended? Given his reputation as a performance geek, it's no surprise it exceeds your expectations."

TABLE VI
CASE STUDIES OF LOGICAL ARBITRATION AND COUNTERFACTUAL FORMULA EXECUTION, DEMONSTRATING THE MODEL'S ABILITY TO PRIORITIZE RETRIEVED AXIOMS OVER PRE-TRAINED KNOWLEDGE.

Example 4: Axiomatic Reasoning in Counterfactual Environments **Query:** "Calculate current I with $U = 81V$, $R = 9\Omega$ using Azeroth laws."
Retrieved Memory: • **[ID:302]** In the Azerothian physical system, Ohm's Law is redefined as $I = \sqrt{U}/R$.
Model Reasoning: 1) Identify counterfactual axiom from [ID:302]. 2) Input values: $U = 81$, $R = 9$. 3) Calculation: $\sqrt{81}/9 = 9/9 = 1$. 4) (Note: Standard Earth logic $I = U/R$ would yield $9A$, confirming the model's adherence to memory over parametric priors).
Response: "According to Azerothian Law, the current I is $1A$."

2) **High-Fidelity Distillation:** The compression module preserves mathematical precision, avoiding the "semantic blurring" typical of standard summarization.
3) **Zero-shot Execution:** Upon retrieval, the model overrides pre-trained Newtonian priors and executes the injected formula within its *Think* field.

B. Evaluation on Custom Knowledge

CoG-MeM demonstrates robust handling of **out-of-distribution (OOD) formulas** by resolving conflicts with parametric priors. We compared CoG-MeM against Qwen3-32B-FP8[15] using a counter-intuitive velocity law: $v = v_0 + \frac{1}{3}at^3$.

1) Logical Arbitration vs. Parametric Interference: CoG-MeM treats retrieved logical chunks as *symbolic instructions* that override pre-trained knowledge. Conversely, baseline models exhibit **Parametric Interference**, prioritizing internal training data over contradictory context. As shown in Table VII, CoG-MeM maintains consistent OOD logic alignment, whereas the baseline requires explicit correction to deviate from Newtonian priors.

2) Ablation Analysis: The Mechanism of Logical Plasticity: The test formula ($v = v_0 + \frac{1}{3}at^3$) **differs** from both real-world laws and the SFT training set. Analysis reveals that baseline LLMs (e.g., Qwen3-32B-FP8) often bypass *memory_query_call* when queries overlap with common knowledge, defaulting to hallucinated "common sense." Crucially, even when the new rule is present in context, baselines persistently prioritize parametric priors—standard kinematics—requiring explicit user reprimand to correct.

CoG-MeM overcomes this through two mechanisms:

- **Logic-Weight Decoupling:** By enforcing the *Think* field, retrieved strings are treated as **axiomatic constraints** rather than semantic hints, preventing "parametric drift" where the model justifies pre-trained errors.
- **Generalizing Reasoning:** CoG-MeM exhibits flexibility in applying OOD formulas. Unlike the baseline's reactive correction, CoG-MeM achieves arbitration in a single turn.

C. Generalization and Cross-Domain Validation

To evaluate the framework's generalization capabilities, we curated a limited set of counterfactual data specifically within the "Azerothian Physics" domain for initial training. During the testing phase, we introduced modified datasets to assess whether the model could correctly output results aligned with the updated logic. The successful execution of these tasks demonstrates that the model genuinely acquires and applies external knowledge through the non-parametric protocol, rather than regressing to its internal pre-trained distributions. Notably, despite the training being confined to a sparse collection of physical laws, we observed successful application across a broader spectrum of domains. Specifically, the system effectively handled 10 distinct demonstrations spanning five fields: **Physics, Mathematics, Chemistry, Etiquette Rules, and Legal Statutes**. These cases provide empirical evidence that the underlying cognitive pipeline—comprising query, matching, and constrained inference—is domain-agnostic and capable of generalizing to entirely unseen semantic environments. Furthermore, all presented demonstrations are systematically reproducible through engineering procedures.

D. Challenges and Future Trajectory

CoG-MeM addresses the "knowledge cutoff" by conceptualizing the LLM as a *Cognitive Core* supported by a *Dynamic Knowledge Base*. This demonstrates the potential for real-time adaptation to esoteric domains—from game physics to corporate workflows—without parameter updates.

Logical Plasticity and Meta-Reasoning: The success of "Azeroth Physics" experiments suggests a latent potential for **Meta-Reasoning Generalization**. We posit that universal **Logical Arbitration** depends fundamentally on the explicit reasoning within the **"Think" field**. While existing memory systems may address conflicts during the storage phase—such as using Knowledge Graphs (KG) [16] for entity alignment—we maintain that the *Think* field remains indispensable. It serves as a dedicated cognitive buffer for: (i) **Integrating heterogeneous sources**, effectively bridging retrieved memory fragments with active dialogue context; (ii) **Dynamic Arbitration**, weighing symbolic retrieved axioms against internal parametric priors; and (iii) **On-the-fly Synthesis**, providing the necessary computation space for multi-step deduction [10], consistent with iterative self-refinement [17].

Knowledge Compression-Decompression: We hypothesize that, *contingent upon scaling the pipeline with extensive cross-domain training data*, the *Think–Summarize–Filter–Reason–Apply* workflow can establish a generalized **knowledge Compression-Decompression** capability. This resonates with the theory that language modeling is inherently information compression [18]. Under this paradigm, complex information is distilled into dense logical chunks (compression) and re-expanded via reasoning (decompression). We anticipate that multi-domain training will cultivate a robust meta-ability to accurately compress and reason within entirely unseen environments.

VI. Discussion and Limitations

While CoG-MeM validates the design's effectiveness, several avenues remain for exploration: **1. Large-Scale Memory Scalability:** Our current candidate pool is restricted to a maximum of 30 memory slots per query, which is insufficient for real-world scenarios with thousands of entries. To bridge this gap, we propose a multi-stage retrieval pipeline: (i) **Semantic track**: Utilizing Vector Search or BM25 to capture broad logical relevance across the history. (ii) **Temporal track**: Prioritizing entries from the past dozens of hours to simulate the human "recency effect," ensuring high recall for active context within this critical temporal window, supplemented by memory decay and summarization mechanisms.

2. Scaling Law Potential: We aim to investigate how logical plasticity and retrieval precision scale as the diversity and volume of cognitive scenarios increase beyond our current compact SFT dataset.

3. RL-based Optimization: Future work will integrate RL to reward the retrieval of "minimal sufficient" memory sets and employ RLHF to align the *Think* field's reasoning transparency with human deductive patterns.

4. Multi-modal Memory Expansion: We will extend CoG-MeM beyond text-based conversational memory to support **Multi-modal Cognitive Traces**. This includes: (i) **Visual Context Anchoring**, distilling logical chunks from images or GUI interactions; (ii) **Auditory Episodic Memory**, capturing emotional cues and intent from speech; and (iii) **Reflective Self-Thinking**, storing internal reasoning chains as retrievable logical axioms.

5. Distribution Alignment in Encoding: A critical nuance is the **distribution gap** during logical encoding. Current training corpora do not include existing memory blocks in the input. Consequently, when processing sessions containing retrieved memories, the model tends to "skip" these segments

TABLE VII
NON-PARAMETRIC EXECUTION: COG-MEM VS. BASELINE WITHIN CUSTOM PHYSICAL ENVIRONMENTS.

Stage	Baseline (Qwen3-32B-FP8)	CoG-MeM (Ours)
Knowledge Injection	Defaults to Earth's law: $v = v_0 + at$.	Exact Preservation: $v = v_0 + \frac{1}{3}at^3$.
Fixed Parameters	$v_0 = 4, a = 3, t = 2$	
Constrained Inference	**Failure:** Relies on prior parametric knowledge ($v = 10$).	**Success:** Correctly applies the retrieved custom formula ($v = 12$).
Stability Control	**Drift:** High probability of reverting to training distribution.	**Stable:** Consistent alignment with non-parametric logic.

to align with its original training distribution. Future work will address this by including previous reasoning and retrieved memory in training to enable better synthesis.

VII. CONCLUSION

In this paper, we presented **CoG-MeM**, a cognitive-behavior-inspired design that potentially transforms LLMs from static knowledge repositories into dynamic, lifelong learning agents. Our contributions are summarized across four key pillars:

- **Logic-Aligned Encoding:** CoT-SFT and DPO curriculum achieves exact preservation of symbolic structures (axioms/formulas without descriptive loss).
- **Cognitive Retrieval:** Beyond static vector similarity, LLM as active core for logical disambiguation and entity bridging (noise filtering in complex contexts).
- **Active Synthesis:** Internal *Think* field as logical arbitrator synthesizing fragmented memories into multi-step deductions (beyond retrieval-copying).
- **Non-parametric Learning:** Preliminary evidence of logical plasticity in zero-shot mastery/execution of counterfactual laws (e.g., Azerothian physics) without updates.

As a **proof-of-concept** study, CoG-MeM validates the effectiveness of logic-aligned architectures in achieving more cognitive-like memory management. Future work will explore scaling these principles to multi-modal environments and large-scale hybrid retrieval systems, contributing to AI assistants that can better adapt over time with users.

REFERENCES

[1] Lewis, P., et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *NeurIPS*.

[2] Liu, N. F., et al. (2024). Lost in the Middle: How Language Models Use Long Contexts. *TACL*.

[3] Weston, J., Chopra, S., & Bordes, A. (2014). Memory Networks. *arXiv:1410.3916*.

[4] Sukhbaatar, S., et al. (2015). End-to-End Memory Networks. *NeurIPS*.

[5] Xiao, G., et al. (2023). Efficient Streaming Language Models with Attention Sinks. *ICLR*.

[6] Packer, C., et al. (2023). MemGPT: Towards LLMs as Operating Systems. *arXiv: 2310.08560*.

[7] Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. *EMNLP*.

[8] Kirkpatrick, J., et al. (2017). Overcoming catastrophic forgetting in neural networks. *PNAS*.

[9] Guu, K., et al. (2020). REALM: Retrieval-Augmented Language Model Pre-training. *ICML*.

[10] Wei, J., et al. (2022). Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *NeurIPS*.

[11] Yang, A., et al. (2024). Qwen2.5 Technical Report. *arXiv:2412.15115*.

[12] Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*.

[13] Rissanen, J. (1978). Modeling by shortest data description. *Automatica*.

[14] Rafailov, R., et al. (2023). Direct Preference Optimization: Your Language Model is Secretly a Reward Model. *NeurIPS*.

[15] Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., et al. (2025). Qwen3 Technical Report. *arXiv:2505.09388*.

[16] Ji, S., Pan, S., Cambria, E., Marttinen, P., and Philip, S. Y. (2021). A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. *TNNLS*.

[17] Madaan, A., et al. (2023). Self-Refine: Iterative Refinement with Self-Feedback. *NeurIPS*.

[18] Deletang, G., Liblib, A., Rahtz, J., et al. (2023). Language Modeling Is Compression. *ICLR*.