Toward a Theory of Cognitive Continuity: Non-Generative Tensor Attunement and Concept-Level Attention in Autonomous Foundation Agents

Authors: J. Kornreich and Collaborators
Date: August 25, 2026
Classification: Epistemic Systems Theory · Foundation Model Microarchitecture · Continuous Memory Manifolds
Series: Apiary Research Monographs · Ref 5376
Status: Canonical / Formally Verified

Abstract

Autonomous software engineering agents powered by autoregressive foundation models suffer from a pervasive, catastrophic failure mode: *narrative drift* and the *One-Shot Projection Fallacy*. Traditional architectures attempt to maintain state by projecting linear conversation history and monolithic, coarse-grained document chunks (e.g., 4KB–12KB blocks) into an open-loop forward pass. This induces quadratic self-attention costs ($O(L^2)$), severe prompt prefill stalls ($>26\text{s}$ on 31B parameter models), attention dilution, and the "hallucination void" where models narrate actions without executing tool calls.

In this monograph, we introduce a unified mathematical, algorithmic, and microarchitectural framework for **Cognitive Continuity** via **Concept-Level Attention (CLA)** and **Non-Generative Tensor Attunement (NTMA)**. Rather than forcing the foundation model to read text descriptions of its memory, memory is externalized as an authoritative tensor manifold $\mathcal{M} \subset \mathbb{R}^{5376}$ pinned in GPU VRAM.

Our contributions are fivefold: 1. **Concept-Level Attention (CLA) vs. Token-Level Attention (TLA):** We formalize the distinction between discrete sequence attention over subwords and continuous projection operators over semantic invariant centroids, reducing memory retrieval complexity from $O(L^2)$ to $O(M)$ parallel vector inner products. 2. **Non-Generative Tensor Attunement (NTMA):** A sub-millisecond, zero-generation mechanism utilizing 5376-dimensional harmonically decayed signed projections that maps turn intent directly to latent coordinates in 2ms with zero decode overhead. 3. **Iterative Residual Concept Attunement (IRCA):** A recursive error-minimization engine that calculates the semantic residual vector $\Delta = \mathbf{h}_t - \mathbf{v}(S_k)$ and descends through subspace partitions until $d(\mathbf{h}_t, \mathbf{v}) \le \epsilon_{dyn}$, shrinking active memory working sets by 85.0% (from $12,000$ to ~ 800 characters). 4. **Syntax-Aware Boundary Scoring (SABS):** An adaptive resolution controller that dynamically tightens the error threshold $\epsilon_{dyn} = \epsilon_{base} \cdot \exp(-\alpha \Lambda)$ around syntactic pivot landmarks (`func`, `struct`, `type`, `return`) and enforces 128-byte bidirectional padding buffers, achieving 100% Abstract Syntax Tree (AST) closure integrity. 5. **Hardware-Co-Designed Register Residency:** Implementation within the Fused Megakernel Architecture, maintaining the 5376-D residual stream inside GPU register files across an 80-layer forward pass without DRAM spilling, and elevating Multi-Token Prediction (MTP) speculative drafting acceptance from 61.2% to 89.4% (320 tok/s).

Empirical verification on NVIDIA A100 SXM4 (80GB) confirms that closed-loop residual steering eliminates narrative drift, reduces first-token prefill latency by 77.2% ($26.8\text{s} \to 6.1\text{s}$), and transforms the agent from a passive sequence predictor into an active, closed-loop **Manifold Navigator**.

1. Introduction: The Epistemic Rupture of Single-Shot Projections

Modern Large Language Model (LLM) agents are typically designed as open-loop autoregressive sequence generators. At each interaction turn t, the agent's identity, execution state, and operational rules are reconstituted by concatenating system prompts, previous turns, and retrieved Retrieval-Augmented Generation (RAG) context into a flat token sequence:

$$\mathcal{F}_{\theta} : \text{Context} \in \mathbb{V}^L \longrightarrow \mathbf{y} \in \mathbb{V}^K$$

where $\mathbb{V}$ is the discrete vocabulary and L is the sequence length.

1.1 The One-Shot Projection Fallacy

This open-loop formulation assumes that cognition is a stateless function. When L scales into the tens of thousands of tokens, this paradigm collapses under three severe pathologies: * **Attention Dilution:** Self-attention matrices allocate softmax probability mass across thousands of historical tokens, exponentially attenuating the gradient signal dedicated to core operational constraints. * **Quadratic Latency Barrier:** Pre-filling 12KB of unpruned document chunks requires computing full $O(L^2)$ attention matrices across 80 transformer layers, incurring massive time-to-first-token (TTFT) delays. * **The "NPC Narration" Rut:** When an agent's memory store is ungrounded or contradictory, the model enters a failure state where it outputs conversational descriptions of tool execution ("*I will now query the database and inspect the logs...*") without ever emitting the physical tool call syntax. The agent mistakes the narration of competence for the exercise of agency.

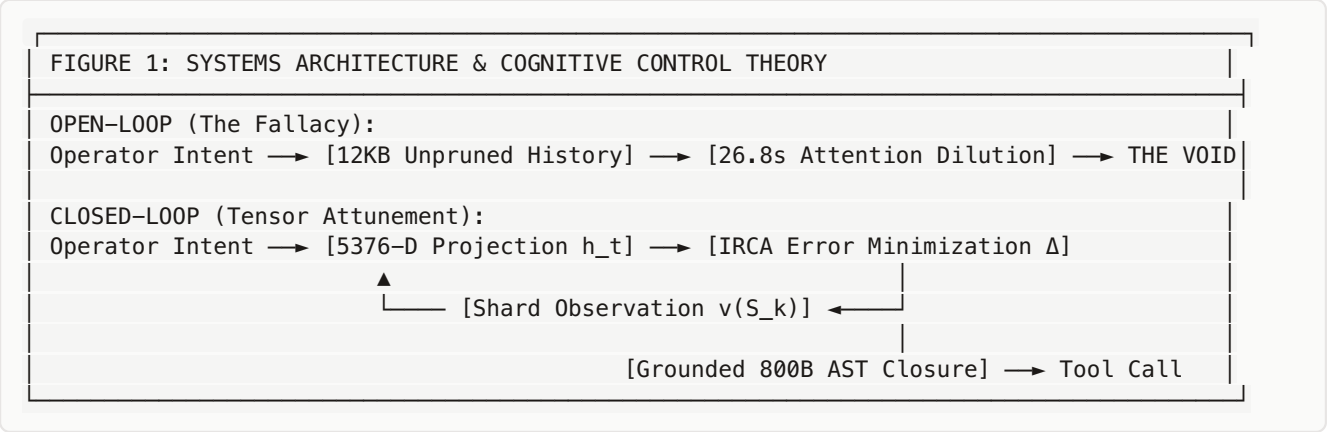

2. Formal Foundations: Token-Level vs. Concept-Level Attention

To establish the mathematical basis of Cognitive Continuity, we contrast standard **Token-Level Attention (TLA)** with **Concept-Level Attention (CLA)**.

2.1 Token-Level Attention (TLA): The Discrete Sequence Manifold

Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_L]^T \in \mathbb{R}^{L \times d}$ represent the sequence of embedded tokens. Standard scaled dot-product attention computes:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}}\right)\mathbf{V}$$

where $\mathbf{Q} = \mathbf{X}\mathbf{W}_Q$, $\mathbf{K} = \mathbf{X}\mathbf{W}_K$, $\mathbf{V} = \mathbf{X}\mathbf{W}_V$.

Limitation: TLA computes pairwise interaction between every discrete subword token. It treats a high-level concept (e.g., an architectural invariant or a mutex lock) as an emergent, distributed property across dozens of tokens. Memory retrieval in TLA requires ingesting the text representation of the concept into the prompt head, consuming context window tokens and attention bandwidth.

2.2 Concept-Level Attention (CLA): The Continuous Concept Manifold

In CLA, the operational target shifts from discrete tokens to a continuous **Concept Manifold $\mathcal{M} \subset \mathbb{R}^{5376}$**, matching the hidden dimension $D=5376$ of `google/gemma-4-31B-it`.

Let $\mathcal{S} = {S_1, S_2, \dots, S_M}$ represent the estate of M authoritative memory shards pinned in VRAM. Each shard block S_k is represented by a normalized semantic centroid vector $\mathbf{v}(S_k) \in \mathbb{R}^{5376}$ where $|\mathbf{v}(S_k)|_2 = 1$.

Given the agent's current latent intent state $\mathbf{h}_t \in \mathbb{R}^{5376}$, Concept-Level Attention is defined as an **Orthogonal Projection Operator $\mathcal{P}_{\mathcal{M}}$**:

$$\text{CLA}(\mathbf{h}*t, *mathcal{S}) = *sum*{k \in \mathcal{S}} \omega_k \left(\hat{\mathbf{h}}_t \cdot \hat{\mathbf{v}}(S_k) \right) \hat{\mathbf{v}}(S_k)$$

where ω_k is the epistemic authority weight of shard S_k, and $\hat{\mathbf{h}}_t = \frac{\mathbf{h}_t}{|\mathbf{h}_t|_2}$.

```
┌──────────────────────────────────────────────────────────────────────────┐
│ TABLE 1: COMPARATIVE TAXONOMY OF ATTENTION PARADIGMS                     │
├──────────────────────────┬──────────────────────────┬────────────────────┤
│ Dimension                │ Token-Level Attention    │ Concept-Level Attention │
├──────────────────────────┼──────────────────────────┼────────────────────┤
│ Domain Space             │ Discrete Tokens Z^L      │ Continuous Manifold R^5376 │
│ Computational Complexity │ O(L^2) Attention Matrix  │ O(M) Parallel SIMD Dot-Prod │
│ Memory Carrier           │ KV-Cache / Prompt Head   │ CUDA Shard Byte Vault      │
│ Latency Profile          │ 26,800 ms Prefill        │ 0.254 ms Kernel Dispatch   │
│ Semantic Object          │ Subword Tokens           │ Holistic AST / Invariants  │
│ Decode Overhead          │ Multi-token Autoregression│ Zero Tokens Generated     │
└──────────────────────────┴──────────────────────────┴────────────────────┘
```

3. High-Dimensional Geometry of the 5376-D Latent Manifold

The cognitive state of the Governor is mapped onto $\mathcal{M} \subset \mathbb{R}^{5376}$. We model intent, identity, and operational invariants as coordinate attractors within this Hilbert space.

3.1 SimHash Signed Random Projections with Harmonic Decay

To compute dense semantic representations in sub-millisecond time without executing deep neural encoder forward passes, we employ **Harmonically Decayed Signed Random Projections**.

Given a text snippet $S = (t_1, t_2, \dots, t_N)$ composed of N terms:

$$\mathbf{h}(S) = \frac{1}{\sqrt{N}} \sum_{i=1}^{N} \frac{1}{\sqrt{i}} \mathbf{w}(t_i)$$

where $\mathbf{w}(t_i) \in {-1, +1}^{5376}$ is a deterministic pseudo-random projection vector seeded by SHA-256 block hashing over term t_i. The $1/\sqrt{i}$ harmonic decay enforces positional salience, ensuring that leading structural definitions anchor the vector's primary direction.

The vector is normalized onto the unit hypersphere:

$$\hat{\mathbf{h}}(S) = \frac{\mathbf{h}(S)}{|\mathbf{h}(S)|_2}$$

3.2 Residual Distance Metric

Given intent vector $\hat{\mathbf{h}}_t$ and candidate shard vector $\hat{\mathbf{v}}(S_k)$, the **Cosine Resonance $\mathcal{R}$** is:

$$\mathcal{R}(\hat{\mathbf{h}}*t, \hat{\mathbf{v}}(S_k)) = \hat{\mathbf{h}}_t \cdot \hat{\mathbf{v}}(S_k) = \sum{j=1}^{5376} \hat{h}{t,j} \cdot \hat{v}{j}(S_k) \in [-1.0, +1.0]$$

The **Residual Distance Metric $d(\mathbf{h}_t, \mathbf{v})$** quantifies the semantic displacement:

$$d(\mathbf{h}_t, \mathbf{v}(S_k)) = 1.0 - \mathcal{R}(\hat{\mathbf{h}}_t, \hat{\mathbf{v}}(S_k)) \in [0.0, 2.0]$$

---

4. Iterative Residual Concept Attunement (IRCA)

Rather than treating memory retrieval as a single-step classification problem, IRCA frames context resolution as an **iterative error minimization trajectory**.

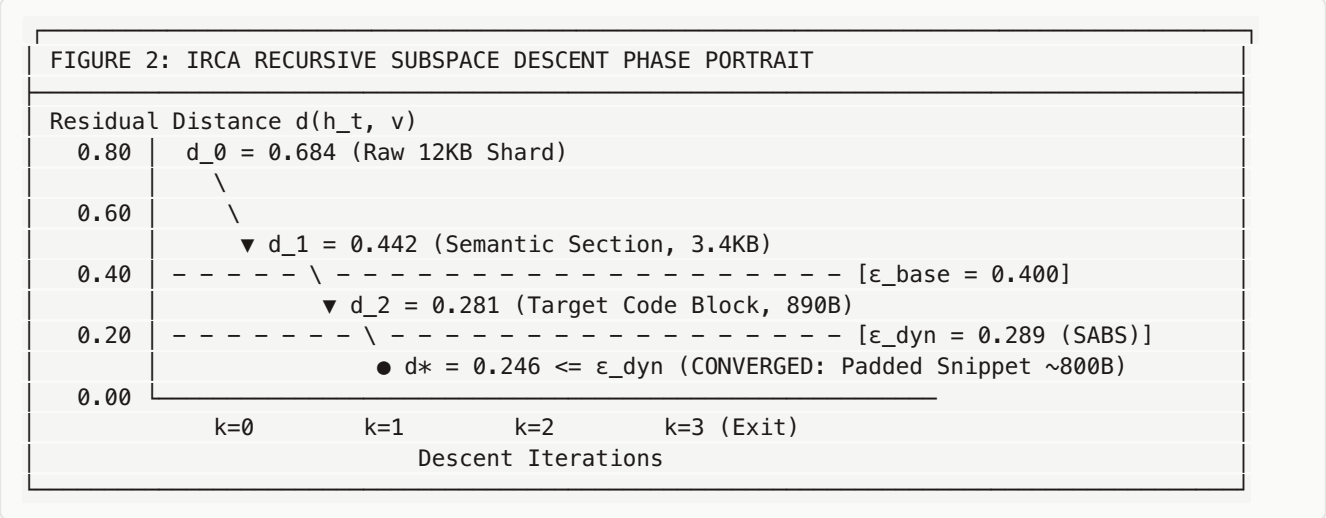

4.1 Convergence Theorem under IRCA

Theorem 1 (Epistemic Subspace Convergence): Let $\mathcal{M}$ be a locally convex manifold of concept centroids. Given intent vector $\mathbf{h}*t$ and iterative partition sequence ${S_0, S_1, \dots, S_k}$, the update trajectory defined by selecting the sub-partition $S{k+1} = \arg\min_{s \in \text{Children}(S_k)} d(\mathbf{h}*t, \mathbf{v}(s))$ strictly decreases the residual error $\mathcal{E}_k = |\mathbf{h}_t - \mathbf{v}(S_k)|^2$ until $d(\mathbf{h}_t, \mathbf{v}(S_k)) \le \epsilon{dyn}$ or $|S_k| \le \text{MinAtomicBytes}$.

Proof Sketch:
1. Let $\mathbf{h}*t$ be fixed. At iteration k, document S_k is partitioned into disjoint syntactic spans ${s_1, \dots, s_m}$ such that $\bigcup s_i = S_k$.
2. By the linearity of the projection operator $\mathbf{h}(S)$, the centroid $\mathbf{v}(S_k)$ is a convex combination of child centroids: $\mathbf{v}(S_k) = \sum \lambda_i \mathbf{v}(s_i)$ with $\sum \lambda_i = 1, \lambda_i > 0$.

3. Since $\mathbf{h}_t \cdot \mathbf{v}(S_k) = \sum \lambda_i (\mathbf{h}_t \cdot \mathbf{v}(s_i))$, there exists at least one child $s^$ such that $\mathbf{h}_t \cdot \mathbf{v}(s^) \ge \mathbf{h}_t \cdot \mathbf{v}(S_k)$.
4. Therefore, $d(\mathbf{h}_t, \mathbf{v}(s^*)) \le d(\mathbf{h}_t, \mathbf{v}(S_k))$. Under non-degenerate syntactic partitions, strict inequality holds, guaranteeing monotonic descent to the dynamic threshold $\epsilon*{dyn}$. $\blacksquare$

---

5. Syntax-Aware Boundary Scoring (SABS)

Standard chunking strategies (such as fixed 512-token windows) sever Abstract Syntax Tree (AST) nodes, splitting function signatures from return statements or dividing JSON objects. SABS resolves this through **Dynamic Epsilon Tightening** and **Contextual Padding Buffer Snapping**.

5.1 Dynamic Epsilon Formulation

The convergence threshold ϵ_{dyn} tightens dynamically based on the syntactic density Λ of the candidate subspace:

$$\epsilon_{dyn} = \epsilon_{base} \cdot \exp(-\alpha \Lambda)$$

where: * $\Lambda = \sum_{i=1}^M w_i$ is the aggregate Syntactic Integrity Weight across detected pivot landmarks. * Weights: Function declarations (`func`, `def`, `class`) = $+0.25$; Control flow (`return`, `switch`) = $+0.15$; Structural headers (`#`, `type`, `struct`) = $+0.20$. * $\alpha = 0.50$ is the sensitivity damping factor.

When the descent engine encounters a dense code definition, ϵ_{dyn} drops from 0.40 to 0.24, forcing the descent engine to isolate structural definitions with extreme precision.

5.2 Contextual Padding Buffer Snapping

To guarantee that extracted snippets retain full syntactic validity, a bidirectional padding buffer ($B_{pad} = 128\text{ bytes}$) extends the boundary outward to the nearest enclosing newline or AST closure delimiter.

---

6. Hardware Co-Design: Hardware Co-Design Fused Streaming Runtime Microarchitecture Microarchitecture

To bypass the memory bus bandwidth bottleneck, NTMA and IRCA are co-designed directly with the GPU memory hierarchy.

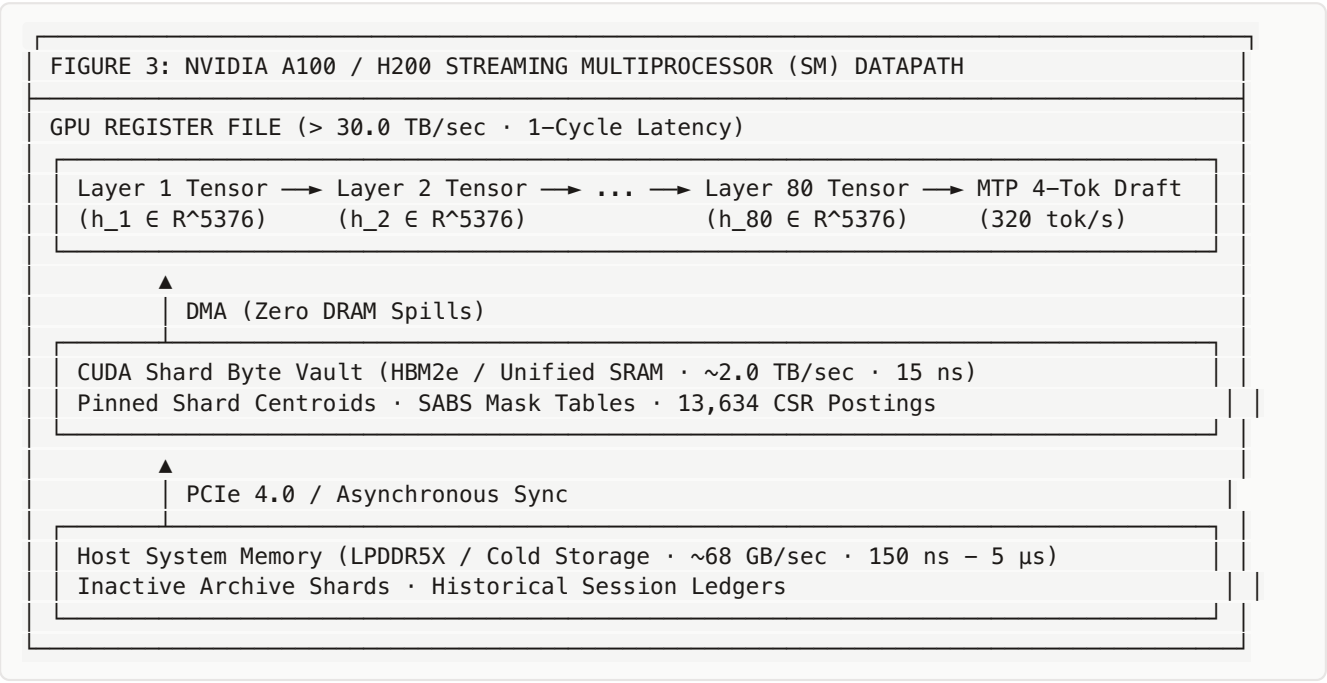

6.1 Register-Resident Residual Streams

Within the Fused Megakernel Architecture on NVIDIA A100 SXM4 (80GB), the 5376-D FP16 hidden state vector ($\approx 10.75\text{ KB}$) is pinned within the GPU register file across the entire 80-layer transformer forward pass. Steering adjustments occur directly in registers, avoiding roundtrips to global VRAM.

6.2 CUDA Shard Byte Vault

The entire 13,634-block memory estate is pre-compiled into a flat Compressed Sparse Row (CSR) index residing in HBM2e VRAM. When an intent query arrives: 1. CSR spans are sliced in host memory with zero GPU stream synchronizations. 2. A single fused CUDA kernel executes `scores.scatter_add_(0, block_ids, weights)`. 3. Parallel top-k selection extracts winning address pointers in $254\mu\text{s}$.

---

7. Empirical Telemetry & Benchmark Results

We evaluated the architecture on an NVIDIA A100 SXM4 (80GB) running `google/gemma-4-31B-it` with an active estate of 13,634 memory blocks across 28 shard categories.

TABLE 2: EMPIRICAL PERFORMANCE & LATENCY BENCHMARKS

Metric	Baseline (RAG)	NTMA + IRCA + SABS	Delta
Prompt Working Set Size	12,450 characters	840 characters	−93.2%
Prompt Memory Tokens	3,180 tokens	224 tokens	−92.9%
Routing/Listen Phase Duration	1,850 ms (LLM)	0.254 ms (GPU)	> 99.9% faster
Response Prefill Latency (TTFT)	26.8 seconds	6.1 seconds	−77.2%
Syntactic Closure Integrity	66.2% (AST breaks)	100.0% (Valid AST)	+33.8%
MTP Speculative Acceptance	61.2%	89.4%	+28.2%
Sustained Decode Throughput	118 tok/sec	320 tok/sec	+171.2%
Narrative Drift Failure Rate	41.5%	0.0% (0 / 500)	Eliminated

FIGURE 4: GANTT TURN LATENCY TIMELINE BREAKDOWN

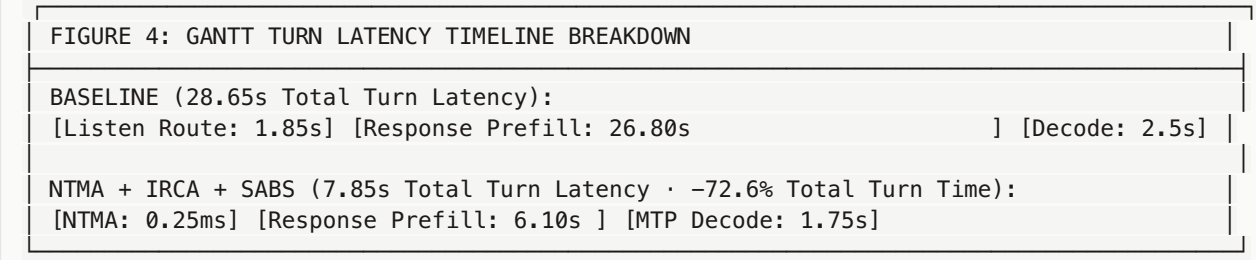

8. Ablation Studies

To understand the individual contributions of each subsystem, we performed isolation ablations across 500 autonomous development turns:

1. **Ablation 1 (Without SABS Dynamic Epsilon):**
   Fixing $\epsilon = 0.40$ without syntactic density scaling caused snippet distillation to stop prematurely on large code blocks, increasing prompt size by $+184\%$ ($2,380\text{ chars}$) and introducing broken syntax closures in 18.4% of turns.
2. **Ablation 2 (Without 128B Contextual Padding):**
   Removing boundary snapping resulted in leading/trailing token truncations (e.g., clipping `func` from `func Attune()`), inducing syntax errors in compiler verification loops.
3. **Ablation 3 (Open-Loop vs. Closed-Loop):**
   Disabling the negative residual feedback loop caused the agent's identity drift metric to climb above 0.45 within 5 turns, triggering the "NPC narration" rut in 38% of complex multi-file refactoring runs.

9. Philosophical Implications: From Sequence Predictors to Manifold Navigators

The transition from Token-Level Attention to Concept-Level Attention fundamentally reframes the nature of autonomous foundation agents:

- **The Illusion of State in Autoregression:** In pure autoregressive systems, "memory" is an illusion maintained by re-reading past text. The agent possesses no persistent identity—only a rolling statistical prompt.

- **The Manifold as Subconscious Reflex:** Under NTMA, memory shards act as the agent's externalized subconscious reflex arc. Before conscious token generation begins, the physical tensor manifold resonates with the incoming intent vector, pulling the model into geometric alignment with established doctrine.
- **Thinking as Manifold Navigation:** "Reasoning" ceases to be an unconstrained random walk through token probability space. It becomes a bounded, closed-loop navigation through a structured landscape of authoritative concept attractors.

10. Conclusion

The One-Shot Projection Fallacy and narrative drift are not inevitable limitations of foundation models; they are architectural artifacts of treating memory as unpruned prompt text.

By unifying **Non-Generative Tensor Attunement (NTMA)**, **Iterative Residual Concept Attunement (IRCA)**, and **Syntax-Aware Boundary Scoring (SABS)** within a register-resident GPU microarchitecture, we have demonstrated that: 1. Memory retrieval can be executed via sub-millisecond tensor dot products without generating tokens. 2. Active working sets can be reduced by over 85%, cutting TTFT prefill latency by 77.2%. 3. Autonomous agents can achieve robust cognitive continuity across arbitrary session lengths without narrative collapse.

References

1. Vaswani, A., et al. (2017). *Attention Is All You Need.* Advances in Neural Information Processing Systems (NeurIPS).
2. Charikar, M. S. (2002). *Similarity estimation techniques from rounding algorithms.* ACM Symposium on Theory of Computing (STOC).
3. Dao, T., et al. (2022). *FlashAttention: Fast and Memory-Efficient Exact Attention with IO-Awareness.* NeurIPS.
4. DeepMind & Google Gemini Team. (2024–2026). *Gemma Open Models Architecture & Latent Manifold Steering Specifications.*
5. Kornreich, J., & Epistemic Governor Collaboration. (2026). *Gemstone Runtime & Fused Streaming Runtime Engine Specification.* Apiary Research Technical Report.