SkillBrain: A Brain-Inspired Memory Substrate for Continual Skill Evolution

Anonymous authors
Paper under double-blind review

Abstract

Large language model agents can continually improve through external Skills while keeping model parameters frozen. However, existing Skill Evolution methods typically treat the current Skill itself as the persistent state across rounds: new experience is repeatedly compressed back into the same executable artifact, entangling reusable task knowledge with model- and Harness-specific realization and making local performance gains difficult to equate with genuine continual knowledge accumulation. Our central view is that **the object of continual evolution should not be the Skill itself, but a persistent Memory that can preserve, reorganize, and reuse experience; Skills should instead be generated from Memory as execution-conditioned interfaces.** Based on this view, we propose **SkillBrain**, a Memory-Centric framework inspired by complementary learning systems. SkillBrain rapidly stores concrete experience in episodic memory, gradually consolidates reusable knowledge into semantic memory, and dynamically generates Skills for the current Executor and Harness. Experiments show strong performance on standard Skill Evolution benchmarks, positive forward and backward transfer on sequential tasks, and low forgetting on fixed per-topic probes. A separate cross-executor analysis shows reuse of experience generated by another model, while repeated Harness changes expose limitations in interface adaptation. These results support a shift from **evolving skills** to **evolving memory** for continual Skill Evolution.

1 Introduction

Large language model agents are moving from one-off task solvers toward systems that continually improve their behavior through experience. One important direction is **Skill Evolution**: agents learn from task execution, feedback, and failures, then revise external Skills to improve subsequent performance. Self-Refine, Reflexion, ExpeL, Voyager, and the recent SkillOpt have demonstrated the potential of this adaptation outside model parameters (Madaan et al., 2023; Shinn et al., 2023; Zhao et al., 2023; Wang et al., 2023; Yang et al., 2026). When evolution extends from a bounded number of rounds to a long stream of experience, however, **Continual Skill Evolution** poses a stricter question: as tasks and execution conditions change, can the agent actually accumulate reusable knowledge, transfer what it has learned, and avoid having earlier knowledge overwritten by later updates?

Repeatedly obtaining a better Skill is not sufficient to answer this question. Existing Skill Evolution methods typically center on the current Skill: execution produces experience, that experience informs a revision, and the revised Skill becomes the starting point for the next round. In a preliminary experiment, we compared **successive refinement over multiple rounds** with **broader, shallower candidate search**. Exploring more candidates in parallel often achieved comparable or even better final performance without a deep sequence of revisions. This suggests that part of the apparent gain from multi-round evolution may come from finding a better execution artifact rather than from knowledge accumulating across rounds. In other words, **improving a Skill and continually learning more are two distinct questions.** This concern is consistent with a recent controlled analysis of multi-round Skill Evolution, which finds that accepted improvements are sparse and characterizes evolution more closely as validation-filtered search than as steady monotonic improvement across rounds (Liu et al., 2026).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

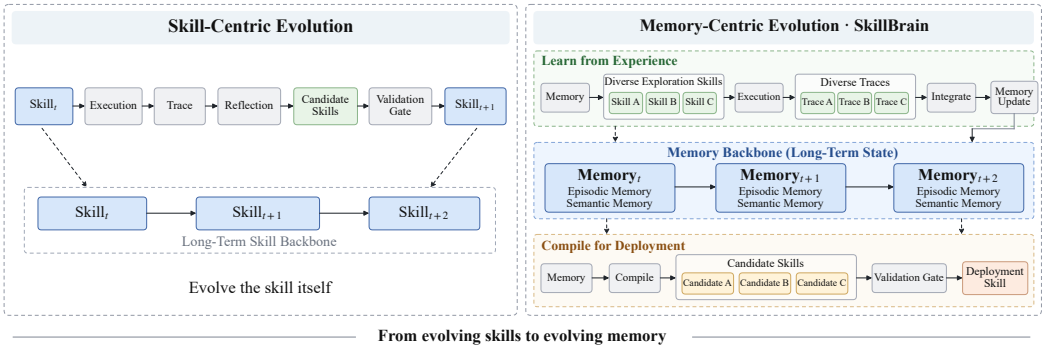

Figure 1: **From Skill optimization to continual knowledge accumulation.** Left: Skill-centric evolution updates the executable Skill through execution, reflection, candidate generation, and validation. Right: SkillBrain retains Memory as the persistent state across rounds; experiences generated by exploration Skills are integrated into episodic and semantic memory, and Skills are then compiled from the evolving Memory for the current execution conditions.

Prior work has reported related observations from different perspectives. SkillOpt finds that retaining only the current best Skill is insufficient for stable long-term optimization: rejected edits and optimization experience accumulated across rounds remain valuable (Yang et al., 2026); VCE-Skill further shows that historical Skill version changes can themselves be distilled into reusable priors for later evolution (Chen et al., 2026a). ACE shows that repeatedly rewriting accumulated context in its entirety can erase historical information, while retaining experience incrementally can mitigate this loss (Zhang et al., 2025). Reflexion, ExpeL, and Dynamic Cheatsheet likewise show that preserving past experiences or strategies as separate, persistent information can help future tasks (Shinn et al., 2023; Zhao et al., 2023; Suzgun et al., 2026). More recently, Unified Context Evolution externalizes experience into multiple evolvable context types, while WikiSkill explicitly separates execution experience, persistent knowledge, and executable skill (Zhu et al., 2026; Tang et al., 2026). **Together, these findings suggest that the Skill used for current execution is not equivalent to the knowledge an agent should preserve and accumulate during continual evolution.**

This motivates us to reconsider the long-term state of Skill Evolution from a continual-learning perspective. A Skill is an **execution artifact**: it must adapt to the current task, model capabilities, tool interfaces, and Harness, and therefore mixes reusable knowledge with execution realizations tied to current conditions. Continual Skill Evolution instead requires a state that can preserve unresolved evidence, consolidate knowledge that remains stable across experiences, and regenerate appropriate behavior when the task, model, or Harness changes. Treating the Skill itself as the sole persistent state binds these requirements to the same repeatedly rewritten object used for current execution. We therefore adopt a **Memory-Centric** view: Recent work that jointly evolves Skills with the execution harness or compiles execution context into harness-level skills likewise suggests that the effective realization of external knowledge depends on the execution interface (Wang et al., 2026b; Wei et al., 2026).

For continual Skill Evolution, Memory is the state that accumulates and evolves over time; Skills are execution interfaces generated from it for the current conditions.

Based on this view, we introduce **SkillBrain**, as illustrated in Figure 1. Conventional methods typically update the executable Skill directly along Skill $\rightarrow$ Trace $\rightarrow$ New Skill; SkillBrain inserts a Memory Substrate and reformulates continual evolution as Skill $\rightarrow$ Trace $\rightarrow$ **Memory** $\rightarrow$ Skill. Memory is the state updated across rounds, while a Skill is compiled from Memory on demand for the current task, Executor Model, and Harness. Concrete experiences, reusable knowledge, and current execution strategies therefore no longer have to be repeatedly compressed into the same Skill, but can evolve on different timescales; when execution conditions change, the same long-term knowledge state can be realized as a newly adapted Skill.

To implement this process, SkillBrain draws functional inspiration from Complementary Learning Systems (CLS) in the hippocampus and neocortex (McClelland et al., 1995; Kumaran et al., 2016). Through **episodic memory**, the system rapidly retains concrete execution experiences, failures, conflicts, and unresolved questions. Through replay, reorganization, and consolidation across experiences, it gradually forms more stable **semantic memory** that can be reused independently. New knowledge can, in turn, reinterpret past experiences and guide future exploration. When a task requires execution, the system generates a Skill from its current memory and execution environment. Decoupling the learning state from the execution artifact also allows the same Memory to support exploration Skills for acquiring new knowledge and deployment Skills for maximizing current task performance.

We evaluate this perspective on six standard Skill Evolution benchmarks and under continual shifts. The standard benchmarks test deployment quality; sequential task shifts probe forward and backward transfer (Lopez-Paz & Ranzato, 2017) and probe forgetting, while execution shifts test reuse and adaptation under changed conditions. Our main contributions are as follows:

- We distinguish **current Skill optimization** from the stricter requirements of **continual Skill Evolution**: the former asks whether a better execution artifact can be found, whereas the latter requires knowledge to accumulate, transfer, and persist across a stream of experience. We identify the choice of persistent state as a central bottleneck connecting the two.
- We introduce **SkillBrain**, which instantiates this view as Memory-Centric Skill Evolution: episodic and semantic memory continually absorb and reorganize experience, while exploration is separated from deployment so that Memory serves as the long-term learning state and Skills are generated on demand for the current Executor and Harness.
- Across standard benchmarks, sequential task shifts, and fixed-budget analyses, we show that SkillBrain improves deployment quality and separates cross-round accumulation from wider search. Task shifts yield positive transfer with low probe forgetting; cross-executor reuse remains effective, while repeated Harness shifts expose interface-adaptation failures.

2 Related Work

2.1 Agent Self-Improvement and Skill Evolution

Language feedback enables agents to improve without updating model parameters. Self-Refine, Reflexion, and GEPA use self-feedback, verbal reflection, or trajectory-level reflection to refine outputs and strategies (Madaan et al., 2023; Shinn et al., 2023; Agrawal et al., 2026). Recent work extends this idea to reusable Skills, including failure-driven EvoSkill, trajectory distillation in Trace2Skill, SkillOpt's trainable external Skill state, and textual-gradient optimization in SkillGrad (Alzubi et al., 2026; Ni et al., 2026; Yang et al., 2026; Wang et al., 2026a). Rethinking, SkillProx, and SkillCAT study multi-round feedback, proximal edits, and assessment-gated search (Liu et al., 2026; Zheng et al., 2026; Chen et al., 2026b), while SkillCommit and VCE-Skill use cross-instance validation or version history to constrain Skill updates (He & Yang, 2026; Chen et al., 2026a). Despite these differences, the central long-term artifact is typically still the next executable Skill. We instead ask: in continual evolution, **what state should actually carry knowledge across rounds?**

These methods also expose an important distinction: **reliably selecting a better Skill is not the same as building long-term knowledge independent of its current realization.** Validation, rollback, version history, and multi-candidate search can reduce harmful edits, but they mainly answer which Skill should be used next. Continual evolution additionally requires deciding which evidence should persist, which regularities should consolidate, and what can be reused when the Executor or Harness changes. SkillBrain makes this long-term state explicit as Memory and leaves the concrete Skill at the execution layer.

2.2 Persistent Experience and Agent Memory

A complementary line of work preserves experience as persistent state. Reflexion uses episodic memory, ExpeL extracts reusable cross-task experience, Dynamic Cheatsheet and ACE incrementally accumulate external strategies or context, and Unified Context Evolution maintains multiple

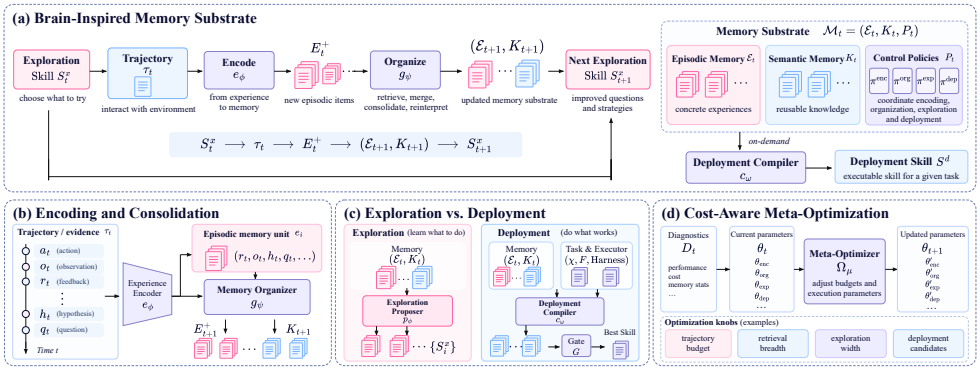

Figure 2: **SkillBrain.** (a) Exploration updates persistent Memory, while deployment Skills are compiled on demand. (b) Episodic evidence is consolidated into reusable semantic knowledge. (c) Exploration targets learning value and deployment targets execution value. (d) A slower cost-aware meta-optimizer adjusts predefined budget and execution parameters.

evolvable context types (Shinn et al., 2023; Zhao et al., 2023; Suzgun et al., 2026; Zhang et al., 2025; Zhu et al., 2026). The closest work to ours, WikiSkill, explicitly separates execution experience, persistent knowledge, and executable Skills (Tang et al., 2026); SHAPER and Evo-Harness further show that effective Skill realization depends on the Harness and execution interface (Wang et al., 2026b; Wei et al., 2026). SkillBrain takes this direction one step further by changing the system's long-term center: **Memory is the state that evolves across rounds, while a Skill is generated from Memory for the current task, Executor, and Harness.** The persistent state can therefore retain both unresolved evidence and consolidated knowledge instead of repeatedly compressing both into a single Skill.

2.3 Complementary Learning Systems and Continual Memory

SkillBrain draws functional inspiration from Complementary Learning Systems (CLS), which combines rapid hippocampal encoding of specific experiences with slower neocortical integration across experiences, supported by replay and consolidation (McClelland et al., 1995; Kumaran et al., 2016). This computational division matches a central need in continual Skill Evolution: new evidence should be retained quickly, while stable regularities should form over a longer timescale. SkillBrain therefore separates episodic evidence from reusable semantic knowledge and uses the Memory Organizer for revisable consolidation between them; the correspondence is functional rather than a literal mapping to brain regions.

3 Method

SkillBrain treats continual Skill Evolution as learning over an external persistent state. Instead of carrying the current executable Skill forward as the learning state, the system updates a **Memory Substrate** and generates Skills from it for the current task, Executor, and Harness. This separates long-term knowledge accumulation from execution-specific realization. Figure 2 summarizes the framework.

3.1 Brain-Inspired Memory Substrate

Inspired functionally by hippocampal–neocortical Complementary Learning Systems (CLS) (McClelland et al., 1995; Kumaran et al., 2016), SkillBrain separates rapidly recorded experience from knowledge consolidated across experiences. At round t, $M_t = (E_t, K_t)$, where **episodic memory** E_t stores concrete evidence, local interpretations, hypotheses, and unresolved questions, while **semantic memory** K_t stores self-contained knowledge reusable beyond its source episodes.

The system also maintains control policies $P_t = \{\pi_t^{\text{enc}}, \pi_t^{\text{org}}, \pi_t^{\text{exp}}, \pi_t^{\text{dep}}\}$ for encoding, memory organization, exploration, and deployment. The full external state is $X_t = (M_t, P_t)$; P_t denotes en-

gineering control mechanisms rather than procedural memory in the cognitive-science sense. The main evolution loop is

$$S_t^x \to \tau_t \to E_t^+ \to (E_{t+1}, K_{t+1}) \to S_{t+1}^x.$$

Exploration Skill S_t^x produces trajectory τ_t, which updates Memory and conditions later exploration. Deployment occurs only when needed, by generating S_t^d from the current Memory. Thus, **Memory is the cross-round learning state, while Skills are execution interfaces**; the control policies evolve on a slower outer timescale.

A crucial constraint is that generated Skill text does not directly become long-term knowledge. Only evidence produced by execution can enter E, and only the Organizer can promote supported regularities into K. This prevents a proposer from turning its own unverified hypothesis into persistent memory merely by writing it into a Skill. It also preserves provenance: semantic entries remain linked to the episodes that support or contradict them, so later evidence can revise earlier conclusions without discarding the underlying observations.

3.2 Encoding and Consolidation

For trajectories $\mathcal{T}_t$ allowed for learning, the encoder first appends episodic evidence:

$$E_t^+ = E_t \cup \text{Encode}(\mathcal{T}_t; \pi_t^{\text{enc}}).$$

It separates **observation from interpretation**: actions, tool feedback, intermediate states, and outcomes are retained as evidence, whereas proposed failure causes or regularities remain revisable hypotheses. Unresolved questions are also preserved as targets for later exploration.

The Memory Organizer then updates both stores,

$$(E_{t+1}, K_{t+1}) = \text{Organize}(E_t^+, K_t; \pi_t^{\text{org}}).$$

E keeps information close to the episodes that generated it; K promotes regularities only after they receive support across experiences. The organizer retrieves relevant history, merges redundancy, connects corroborating or conflicting evidence, consolidates stable knowledge into K, and can reinterpret earlier episodes when later evidence reveals a missing rule. Low-utility episodes may be archived rather than immediately deleted. Consolidation is therefore revisable rather than a one-shot summary: later evidence can update or invalidate semantic entries without requiring their source episodes to have disappeared. Detailed schemas and memory operations are given in Appendices A–B.

3.3 Why Retain Episodes?

The value of an earlier episode can increase after later knowledge arrives. For latent regularity Y, episode e, and current knowledge K,

$$V(e \mid K) = H(Y \mid K) - H(Y \mid K, e) = I(Y; e \mid K).$$

Let e and z be independent uniform bits and $Y = e \oplus z$. Before z is known, $I(Y; e) = 0$; after observing z, $I(Y; e \mid z) = 1$ bit. An episode that appears useless can therefore become maximally informative after later knowledge arrives.

The limitation persists under finite memory. Let early history contain n independent uniform bits $X_i \overset{\text{i.i.d.}}{\sim} \text{Bernoulli}(1/2)$, compressed into S with $H(S) \le b$. A later query index $I \sim \text{Unif}\{1, \dots, n\}$ is independent of $(X_{1:n}, S)$. Any decoder recovering X_I from (S, I) with average error at most ϵ, where $0 \le \epsilon < 1/2$, must satisfy

$$b \ge n[1 - h_2(\epsilon)],$$

where h_2 is the binary entropy function. This does not imply storing every episode indefinitely; it motivates separating retrievable evidence from consolidated knowledge and delaying irreversible compression. Formal statements and proofs appear in Appendix I.

3.4 Exploration and Deployment

The deployment-best Skill need not be the most useful Skill for acquiring new knowledge. SkillBrain therefore separates the two paths. The **Exploration Proposer** generates $\{S_{t,i}^x\}_{i=1}^{n_{\text{exp}}} =$

Propose$(E_t, K_t; \pi_t^{\text{exp}})$, targeting unresolved questions, contradictions, and hypotheses. Its objective is learning value: a candidate is useful if the resulting trajectory can change future knowledge, even when it is not currently validation-best.

The **Deployment Compiler** instead realizes Memory for the current execution conditions as $S^d = \text{Compile}(E, K, \text{task}, \text{Executor}, \text{Harness}; \pi^{\text{dep}})$. Direct Deployment uses one Skill; Validation-Gated Deployment generates multiple candidates from the same Memory and selects $S^{d*} = \arg\max_{S_i^d} G(S_i^d)$. Thus deployment and exploration share the same persistent Memory but optimize different utilities:

$$\arg\max_s V_{\text{deploy}}(s) \not\equiv \arg\max_s V_{\text{learn}}(s).$$

The same knowledge state can therefore support information-seeking behavior during evolution and conservative high-performing behavior at deployment.

3.5 Cost-Aware Meta-Optimization

On a slower timescale, SkillBrain optimizes predefined budget and execution variables of the four control policies while keeping their roles, schemas, and tool permissions fixed. We write $\pi^j = \pi^j(\theta^j)$ for $j \in \{\text{enc}, \text{org}, \text{exp}, \text{dep}\}$, where θ^j controls quantities such as trajectory-read budget, organization frequency, retrieval breadth, exploration width, deployment candidates, and Skill length. These variables allocate computation without directly editing the contents of E or K.

A fixed meta-prompt μ uses accumulated diagnostics D_t^{meta} to propose $\theta_{t+1} = \Omega_\mu(\theta_t, D_t^{\text{meta}})$, with the constrained objective

$$\min_\theta \; C_{\text{CSE}}(\theta) \quad \text{s.t.} \quad Q_j(\theta) \geq Q_j(\theta_{\text{ref}}) - \epsilon_j.$$

A cheaper configuration is accepted only if the corresponding functional requirement remains within tolerance, including information retention, knowledge consistency, useful exploration, and deployment quality. This slower outer loop keeps transient cost adaptation separate from the knowledge accumulated in Memory. Parameter spaces and acceptance rules are given in Appendices D–E.

4 Experiments

We separate deployment quality from continual-learning evidence. Standard Skill Evolution benchmarks test whether a Memory-Centric state remains competitive for current execution, while sequential shifts and controlled analyses test whether knowledge actually accumulates, transfers, survives later updates, and remains reusable when execution conditions change. We therefore ask four questions: whether SkillBrain preserves current-task quality, whether its persistent state transfers and resists forgetting, which components matter, and whether improvements reflect cross-round accumulation rather than wider Skill search.

4.1 Experimental Setup

Benchmarks and protocol. Following SkillOpt and WikiSkill (Yang et al., 2026; Tang et al., 2026), we evaluate SearchQA, SpreadsheetBench, OfficeQA, DocVQA (Mathew et al., 2021), LiveMathematicianBench (LiveMath), and ALFWorld (Shridhar et al., 2021), spanning search, spreadsheet and office tool use, document understanding, mathematical reasoning, and long-horizon interaction. Released train/validation/test splits and native metrics are retained: training data produce evolution trajectories, validation data are used only for methods that require selection or gating, and test data are reserved for evaluation and never guide adaptation. Validation/test trajectories are never written into SkillBrain memory.

Baselines and executors. We compare No Skill, Trace2Skill (Ni et al., 2026), EvoSkill (Alzubi et al., 2026), SkillOpt (Yang et al., 2026), and WikiSkill (Tang et al., 2026), covering no persistent state, trajectory-to-Skill distillation, failure-driven evolution, execution-oriented Skill optimization, and knowledge–Skill co-evolution. Experiments use frozen Qwen3.5-4B, Qwen3.6-35B-A3B, and Qwen3.6-27B executors. The main table uses benchmark-native Harnesses; PI-Agent is reserved for Harness-shift tests in Appendix G. At inference, the executor sees only the compiled deployment

Table 1: **Main results (%).** Higher is better; within-model maxima are bold. Subscripts show changes from the same model's No Skill (percentage points, including Avg.; green: gain, red: loss). The 27B block is descriptive because several baselines use non-identical published settings.

Method	SearchQA	Spreadsheet Bench	OfficeQA	DocVQA	LiveMath	ALFWorld	Avg.
Qwen3.6-27B							
No Skill	79.3	40.8	42.1	90.0	33.9	52.8	56.5
Trace2Skill	79.9$_{+0.6}$	53.3$_{+12.5}$	54.3$_{+12.2}$	90.4$_{+0.4}$	36.3$_{+2.4}$	55.5$_{+2.7}$	61.6$_{+5.1}$
EvoSkill	82.0$_{+2.7}$	59.4$_{+18.6}$	60.2$_{+18.1}$	90.6$_{+0.6}$	57.7$_{+23.8}$	61.4$_{+8.6}$	68.6$_{+12.1}$
SkillOpt	88.2$_{+8.9}$	53.2$_{+12.4}$	54.8$_{+12.7}$	87.6$_{-2.4}$	51.9$_{+18.0}$	59.2$_{+6.4}$	65.8$_{+9.3}$
WikiSkill	83.7$_{+4.4}$	**81.7**$_{+40.9}$	53.7$_{+11.6}$	91.9$_{+1.9}$	61.9$_{+28.0}$	77.6$_{+24.8}$	75.1$_{+18.6}$
SkillBrain	**90.9**$_{+11.6}$	65.9$_{+25.1}$	**60.7**$_{+18.6}$	**92.2**$_{+2.2}$	**78.1**$_{+44.2}$	**80.3**$_{+27.5}$	**78.0**$_{+21.5}$
Qwen3.6-35B-A3B							
No Skill	75.6	38.2	45.9	88.0	31.2	59.7	56.4
Trace2Skill	70.0$_{-5.6}$	33.2$_{-5.0}$	32.0$_{-13.9}$	86.8$_{-1.2}$	29.6$_{-1.6}$	70.9$_{+11.2}$	53.8$_{-2.6}$
EvoSkill	80.9$_{+5.3}$	44.0$_{+5.8}$	47.5$_{+1.6}$	90.0$_{+2.0}$	42.0$_{+10.8}$	76.8$_{+17.1}$	63.5$_{+7.1}$
SkillOpt	78.0$_{+2.4}$	47.5$_{+9.3}$	47.1$_{+1.2}$	86.8$_{-1.2}$	41.6$_{+10.4}$	**82.1**$_{+22.4}$	63.9$_{+7.5}$
WikiSkill	79.6$_{+4.0}$	50.1$_{+11.9}$	48.6$_{+2.7}$	89.4$_{+1.4}$	44.7$_{+13.5}$	68.9$_{+9.2}$	63.6$_{+7.2}$
SkillBrain	**82.7**$_{+7.1}$	**52.9**$_{+14.7}$	**50.3**$_{+4.4}$	**91.7**$_{+3.7}$	**47.8**$_{+16.6}$	71.5$_{+11.8}$	**66.2**$_{+9.8}$
Qwen3.5-4B							
No Skill	68.1	9.3	14.5	86.9	22.4	30.6	38.6
Trace2Skill	68.5$_{+0.4}$	19.3$_{+10.0}$	16.3$_{+1.8}$	88.0$_{+1.1}$	27.2$_{+4.8}$	64.9$_{+34.3}$	47.4$_{+8.8}$
EvoSkill	70.4$_{+2.3}$	16.8$_{+7.5}$	17.9$_{+3.4}$	87.5$_{+0.6}$	26.1$_{+3.7}$	35.2$_{+4.6}$	42.3$_{+3.7}$
SkillOpt	71.2$_{+3.1}$	23.9$_{+14.6}$	29.7$_{+15.2}$	89.0$_{+2.1}$	52.0$_{+29.6}$	81.3$_{+50.7}$	57.9$_{+19.3}$
WikiSkill	72.9$_{+4.8}$	25.4$_{+16.1}$	31.8$_{+17.3}$	89.8$_{+2.9}$	54.7$_{+32.3}$	83.9$_{+53.3}$	59.8$_{+21.2}$
SkillBrain	**74.6**$_{+6.5}$	**27.0**$_{+17.7}$	**33.5**$_{+19.0}$	**90.3**$_{+3.4}$	**57.9**$_{+35.5}$	**85.8**$_{+55.2}$	**61.5**$_{+22.9}$

Skill, not the Memory Substrate. Several 27B baseline entries come from non-identical published settings, so that block is interpreted descriptively rather than as a controlled ranking. For the unified 4B and 35B-A3B comparisons, all methods use the same benchmark split and frozen Executor within each task, and test examples never participate in candidate selection. In the main benchmark comparison, we evaluate the final deployment Skill only after evolution is complete. For continual-learning metrics, an isolated evaluator also records intermediate test panels without providing test feedback to the learner.

4.2 Main Results

Table 1 shows that the Memory-Centric design does not sacrifice deployment quality. Under the unified 4B and 35B-A3B protocols, SkillBrain is best in 11 of 12 model–benchmark pairs and improves the strongest baseline average by 1.7 and 2.3 points, respectively. The 27B block shows a 78.0 average versus 75.1 for the strongest displayed baseline, but is descriptive because of the setting mismatch noted above. These results establish that changing the persistent state does not weaken current deployment, but static final scores alone cannot establish continual accumulation.

4.3 Continual Skill Evolution

Static scores do not establish continual learning, so we additionally evaluate forward transfer (FWT) and backward transfer (BWT) (Lopez-Paz & Ranzato, 2017), Probe Forgetting, and reuse under execution shifts while keeping all Executor parameters frozen. Full stage construction and metric definitions are in Appendix G. FWT measures whether earlier experience helps unseen later stages, while BWT asks whether later learning improves or damages earlier stages. Probe Forgetting is measured on five fixed test examples per topic; it is not full-test forgetting, which is unavailable because the full inter-stage matrix is missing.

Task-distribution shift. SearchQA is organized into 12 sequential topic stages (400 train, 200 validation, 1,400 test examples) with fixed Qwen3.6-35B-A3B and packed-evidence Harness. Each training example is observed only when its topic arrives; we compare ungated (U) updates with validation-gated (G) updates, and validation data never enter persistent memory.

Method	Gate	Final EM	FWT	BWT	Probe Forget. ↓
SkillOpt-CL	U	76.40	+3.84	-0.39	1.67
SkillOpt-CL	G	78.44	**+6.14**	-0.09	4.58
EvoSkill-CL	U	78.52	+4.80	+3.13	**0.00**
EvoSkill-CL	G	79.00	+5.12	+0.42	3.33
Reflect	U	75.41	+2.50	+0.08	**0.00**
Reflect	G	73.16	+0.21	+0.07	1.67
SkillBrain-CL	U	**81.53**	**+5.15**	**+3.72**	0.10
SkillBrain-CL	G	**80.41**	+6.01	**+2.04**	**0.21**

Table 2: **Continual SearchQA.** EM is %; transfers and probe forgetting are percentage points. Bold: best within each U/G group.

Variant	DocVQA	LiveMath	SearchQA	Avg.
Full SkillBrain	**92.2**	**78.1**	**90.9**	**87.1**
Skill-Centric	88.9	73.2	84.7	82.3
Skill + History	89.7	74.5	85.6	83.3
Flat Memory	91.0	76.1	88.3	85.1
E-only	90.5	75.6	87.4	84.5
K-only	90.8	75.8	87.9	84.8
w/o Organizer	91.4	76.7	89.2	85.8
w/o Recompile	91.2	76.3	88.8	85.4
Shared Skill	91.5	76.9	89.4	85.9

Table 3: **Ablations** on Qwen3.6-27B (%).

Without gating, SkillBrain-CL reaches 81.53% Final Macro EM with $+5.15$ FWT, $+3.72$ BWT, and 0.10 Probe Forgetting; with gating it reaches 80.41%, $+6.01$, $+2.04$, and 0.21, respectively. Both settings therefore show positive forward/backward transfer and low forgetting on the fixed probes. In particular, the ungated result shows that Memory can accumulate useful information without using validation labels to screen every update. Gating instead exposes a trade-off in this run: FWT rises by 0.86 points, while Final Macro EM and BWT decrease, suggesting that the Validation Gate is not the source of continual accumulation. Because candidate-generation and validation budgets are not fully matched across methods, we use this table to characterize continual behavior rather than claim an equal-compute ranking.

Harness shift. We also keep SearchQA and Qwen3.6-35B-A3B fixed while switching among five evidence-access Harnesses across ten stages (nine switches). Returning to a previously seen interface uses new training examples, so the stream tests both adaptation to a changed execution interface and retention when an old interface reappears. This stress test exposes substantial interface-specific failures and does not establish reliable adaptation across all five Harnesses. Per-interface results, failure analysis, and unmatched budget details appear in Appendix G; Section 4.5 examines cross-executor experience reuse separately.

4.4 Ablation Study

Table 3 holds trajectories, rollout budget, validation budget, and persistent-state token budget fixed. Skill-Centric updating reduces the three-task average from 87.1 to 82.3 (-4.8), while Skill + History reaches only 83.3, showing that the gain is not simply extra context. Collapsing E/K into Flat Memory costs 2.0 points; E-only and K-only lose 2.6 and 2.3 points, supporting complementary episodic and semantic roles. Removing the Organizer, recompilation, or exploration–deployment separation further loses 1.3, 1.7, and 1.2 points. Together, the ablations support organized Memory as the persistent learning state rather than an appended history buffer.

4.5 Analytical Experiments

We next test the mechanism behind continual gains; full protocols and per-task results are in Appendix H.

Search more, or accumulate more? With a fixed budget of 32 rollouts, Wide Search averages 77.1 and Skill-Centric peaks at 78.2, whereas Memory-Centric improves with depth to 83.1 at $R = 8$. Equal rollout counts isolate the benefit of maintaining and updating persistent state across rounds (Appendix H.1). Because all settings consume the same number of executions, the improvement at larger depth cannot be attributed to seeing more trajectories. Instead, depth becomes useful only when information acquired in earlier rounds can alter what later rounds retrieve, test, and preserve.

What accumulates in Memory? Mean reuse distance is 1.7 rounds for episodic E and 5.2 for semantic K; reuse after at least four rounds is 13.4% versus 42.8%, and final survival is 41.2% versus 76.5%. The two stores therefore support distinct short- and long-horizon roles: E keeps

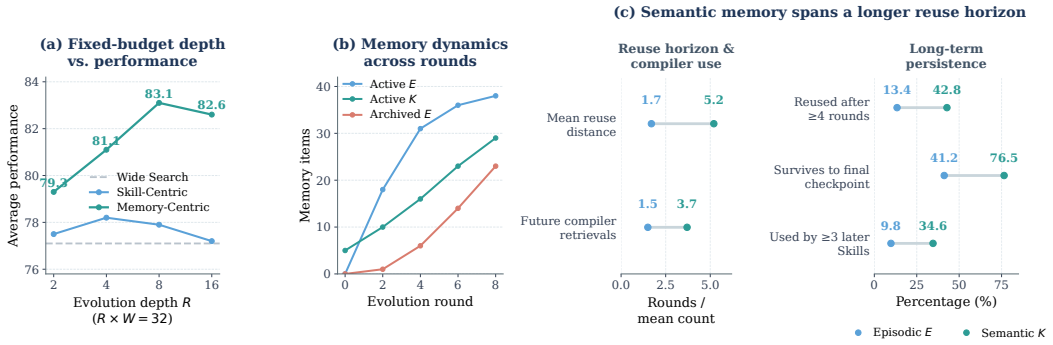

Figure 3: **Cross-round accumulation.** (a) Fixed-budget depth comparison; (b) memory dynamics; (c) reuse horizon and persistence of episodic E and semantic K.

concrete evidence near the experience that produced it, whereas K carries abstractions that remain useful many rounds later. Memory's benefit is thus organizational, not merely a larger history buffer (Appendix H.2). The distinct reuse horizons are consistent with E serving as revisitable evidence and K as longer-lived reusable knowledge, rather than both stores acting as a larger undifferentiated buffer.

What should generate the next experience? Random diverse exploration, following the validation-best Skill, and knowledge-gap exploration average 77.0, 77.9, and 83.1. The gain comes from targeting unresolved knowledge rather than simply deviating from the current best deployment Skill (Appendix H.3). Random exploration tests novelty alone, whereas validation-best continuation represents exploitation. Knowledge-gap exploration outperforming both indicates that useful exploration is coupled to what Memory still does not explain.

Can experience be reused across Executors? On SearchQA, Qwen3.6-27B generates the training experience, while Qwen3.6-35B-A3B constructs the state, compiles Skills, and performs target evaluation. Structured Memory + Reuse raises EM from 48 to 73 with 94.7% correct-answer retention; recompiling the same Memory yields 61 (Appendix H.4). These results support cross-executor experience reuse, but do not isolate adaptation-free transfer of a source-built Memory: the target model already participates in state construction and Skill compilation.

Together, these analyses show that Memory-Centric evolution gains from structured cross-round accumulation and reuse, while executable Skills remain condition-specific interfaces.

5 Conclusion

This work distinguishes repeatedly finding a better executable Skill from the stricter goal of continual Skill Evolution. SkillBrain addresses this distinction by moving the persistent learning state from Skill to a brain-inspired Memory Substrate: episodic memory retains concrete evidence, semantic memory consolidates reusable knowledge, and executable Skills are generated as interfaces conditioned on the current task, Executor, and Harness.

Across six Skill Evolution benchmarks, SkillBrain preserves strong deployment quality and is best in 11 of 12 model–benchmark pairs under the two unified Executor protocols. More importantly, sequential-task experiments show positive forward and backward transfer with low forgetting on fixed per-topic probes even without validation gating; full-test forgetting is not established by these probes. Ablations and fixed-budget analyses further connect these gains to cross-round accumulation, complementary episodic/semantic time scales, knowledge-gap-driven exploration, and reuse of experience generated by another Executor rather than simply more context or wider Skill search. These results support a shift from *continually rewriting Skills* to **continually evolving the reusable knowledge state behind them**. Skills can be regenerated for changed execution conditions, although the repeated Harness-shift stress test shows that reliable interface adaptation remains a limitation.

AI use statement

Generative AI tools were used to support literature search and summarization, brainstorming during the development of research ideas, and language polishing of the manuscript. They were also used to assist in reviewing experimental results, including checking consistency between reported values and original experiment outputs and flagging potential presentation or consistency issues for author inspection. All AI-assisted suggestions were reviewed by the authors against the relevant literature, manuscript context, or original experimental records before being incorporated. The authors take responsibility for the final content of this work, including all text, claims, and reported results.

Ethics statement

This work studies continual Skill evolution and persistent external memory using research benchmarks and does not involve human subjects or the collection of private or sensitive personal data. Persistent agent memory nevertheless introduces a general deployment risk: incorrect, biased, or unsafe information may survive beyond a single interaction and influence future behavior. SkillBrain therefore treats validation, provenance, rollback, and the separation between training experience and held-out evaluation as important system-level safeguards. The experiments in this paper are limited to research benchmarks rather than high-stakes deployment; applying persistent Skill evolution in sensitive domains would require additional access control, auditing, privacy protection, and domain-specific safety evaluation.

Reproducibility statement

We provide the information needed to reconstruct the experimental protocol throughout the paper and appendix. Section 4.1 specifies the benchmark suite, data-split usage, baselines, Executor models, Harness choices, and evaluation protocol. Appendices A–B describe the memory representation, initialization, continual evolution loop, control policies, and tool interfaces; Appendix C gives the control-policy prompt templates; Appendices D–E document cost-aware meta-optimization and the corresponding hyperparameter search spaces; Appendix F summarizes the implementation and reporting configuration required for reproduction. Appendix G gives the detailed continual-shift protocols, stage ordering, gating rules, evaluation coverage, and known limitations, while Appendix H provides the full settings and per-condition results for the analysis experiments, and Appendix I states the assumptions, formal claims, and proofs used in the theoretical analysis. We explicitly separate training, validation, and held-out test information, and report diagnostic settings with limited seeds or unmatched budgets as such rather than presenting them as equal-compute comparisons.

References

Lakshya A. Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziems, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J. Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alexandros G. Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. GEPA: Reflective prompt evolution can outperform reinforcement learning. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2507.19457.

Salaheddin Alzubi, Noah Provenzano, Jaydon Bingham, Weiyuan Chen, and Tu Vu. Evoskill: Automated skill discovery for multi-agent systems. *arXiv preprint arXiv:2603.02766*, 2026. URL https://arxiv.org/abs/2603.02766.

Jianming Chen, Xuanbin Ye, Yawen Wang, Junjie Wang, Qing Wang, and Fanjiang Xu. VCE-Skill: Enhancing skill self-evolution with version-change experience. *arXiv preprint arXiv:2608.16544*, 2026a. URL https://arxiv.org/abs/2608.16544.

Kunfeng Chen, Qihuang Zhong, Juhua Liu, and Bo Du. SkillCAT: Contrastive, assessment-augmented and topology-aware skill self-evolution for LLM agents. *arXiv preprint arXiv:2606.13317*, 2026b. URL https://arxiv.org/abs/2606.13317.

Yu He and Weikai Yang. Skillcommit: Evolving agent skills through behaviorally validated scope expansion. *arXiv preprint arXiv:2608.15165*, 2026. URL https://arxiv.org/abs/2608.15165.

Dharshan Kumaran, Demis Hassabis, and James L. McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20 (7):512–534, 2016. doi: 10.1016/j.tics.2016.05.004.

Yuxuan Liu, Zhaochen Su, Yuhao Zhang, Jiahe Guo, Zhongwei Xie, Huihao Jing, Lingyun Xie, Qing Zong, Yauwai Yim, Zhixiong Zhang, Haoran Li, and Yangqiu Song. Rethinking self-evolving agent skills: Feedback dynamics over multiple rounds. *arXiv preprint arXiv:2608.02636*, 2026. URL https://arxiv.org/abs/2608.02636.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, pp. 6467–6476, 2017. URL https://arxiv.org/abs/1706.08840.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback. In *Advances in Neural Information Processing Systems*, 2023.

Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. DocVQA: A dataset for VQA on document images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 2200–2209, 2021. doi: 10.1109/WACV48630.2021.00225.

James L. McClelland, Bruce L. McNaughton, and Randall C. O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, 1995. doi: 10.1037/0033-295X.102.3.419.

Jingwei Ni, Yihao Liu, Xinpeng Liu, Yutao Sun, Mengyu Zhou, Pengyu Cheng, Dexin Wang, Erchao Zhao, Xiaoxi Jiang, and Guanjun Jiang. Trace2skill: Distill trajectory-local lessons into transferable agent skills. *arXiv preprint arXiv:2603.25158*, 2026. URL https://arxiv.org/abs/2603.25158.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, 2023.

Mohit Shridhar, Xingdi Yuan, Marc-Alexandre Côté, Yonatan Bisk, Adam Trischler, and Matthew Hausknecht. ALFWorld: Aligning text and embodied environments for interactive learning. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2010.03768.

Mirac Suzgun, Mert Yuksekgonul, Federico Bianchi, Dan Jurafsky, and James Zou. Dynamic cheatsheet: Test-time learning with adaptive memory. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2026. URL https://aclanthology.org/2026.eacl-long.333/.

Liyan Tang, Cyrus Rashtchian, Chun-Sung Ferng, Andrew Tomkins, Da-Cheng Juan, and Tu Vu. Wikiskill: Compiling agent experience into persistent knowledge for skill evolution. *arXiv preprint arXiv:2608.27454*, 2026. URL https://arxiv.org/abs/2608.27454.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. In *NeurIPS Workshop on Foundation Models for Decision Making*, 2023.

Hanyu Wang, Yifan Lan, Bochuan Cao, Lu Lin, and Jinghui Chen. Skillgrad: Optimizing agent skills like gradient descent. *arXiv preprint arXiv:2605.27760*, 2026a. URL https://arxiv.org/abs/2605.27760.

Peidong Wang, Zhiming Ma, Ying Chang, Xufang Luo, Xiaocui Yang, Shi Feng, Yuqing Yang, and Dongsheng Li. Self-evolving embodied agents via skill-harness evolution. *arXiv preprint arXiv:2608.11350*, 2026b. URL https://arxiv.org/abs/2608.11350.

Tianxin Wei, Zhan Shi, Minhua Lin, Bing He, Zewen Liu, Yisi Sang, Yuanchen Bei, Xuying Ning, Jiaru Zou, Ting-Wei Li, Xiao Lin, Yanjun Zhao, Chi Wang, Benoit Dumoulin, Dakuo Wang, Jingrui He, and Hanqing Lu. Evo-harness: Context-to-harness skill compilation for self-evolving agents. *arXiv preprint arXiv:2608.15071*, 2026. URL https://arxiv.org/abs/2608.15071.

Yifan Yang, Ziyang Gong, Weiquan Huang, Qihao Yang, Ziwei Zhou, Zisu Huang, Yan Li, Xuemei Gao, Qi Dai, Bei Liu, Kai Qiu, Yuqing Yang, Dongdong Chen, Xue Yang, and Chong Luo. Skillopt: Executive strategy for self-evolving agent skills. *arXiv preprint arXiv:2605.23904*, 2026.

Qizheng Zhang, Changran Hu, Shubhangi Upasani, Boyuan Ma, Fenglu Hong, Vamsidhar Kamanuru, Jay Rainton, Chen Wu, Mengmeng Ji, Hanchen Li, Urmish Thakker, James Zou, and Kunle Olukotun. Agentic context engineering: Evolving contexts for self-improving language models. *arXiv preprint arXiv:2510.04618*, 2025. URL https://arxiv.org/abs/2510.04618.

Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. ExpeL: LLM agents are experiential learners. *arXiv preprint arXiv:2308.10144*, 2023.

Mingxuan Zheng, Yujin Zhou, Chuxue Cao, Boqin Yin, Yuyao Zhang, Jiapeng Sun, Shuaishuai Gong, Sirui Han, and Yike Guo. Skillprox: Self-evolving agent skills via proximal textual gradient descent. *arXiv preprint arXiv:2608.07449*, 2026. URL https://arxiv.org/abs/2608.07449.

Zixuan Zhu, Yitong Hu, Yong Dai, Junfeng Fang, Chunyang Jiang, Senkang Hu, and Yuzhi Zhao. Unified context evolution for LLM agents. *arXiv preprint arXiv:2606.02304*, 2026. URL https://arxiv.org/abs/2606.02304.

A Memory Representation and Initialization

A.1 Notation

Symbol	Name	Definition
τ	Execution Trace	Raw interaction trajectory from one real task execution
$\mathcal{T}_t$	Trace Batch	Trajectories allowed for learning at round t
E_t	Episodic Memory	Episode-bound memory items
K_t	Semantic Memory	Self-contained, reusable long-term knowledge
P_t	Control Policies	Policies for memory operations and Skill generation
M_t	Memory Substrate	Persistent memory state (E_t, K_t)
X_t	System State	Full system state (M_t, P_t)
S^x	Exploration Skill	Skill for actively acquiring new evidence
S^d	Deployment Skill	On-demand Skill for task execution
π^{enc}	Experience Encoder	Encodes raw trajectories into episodic memory
π^{org}	Memory Organizer	Organizes episodes and updates semantic knowledge
π^{exp}	Exploration Proposer	Generates exploration Skills from knowledge gaps
π^{dep}	Deployment Compiler	Compiles memory into a deployment Skill
G	Validation Gate	Candidate scoring and selection module
θ	Control Parameters	Tunable control-policy parameters
Ω_μ	Meta Optimizer	Updater defined by fixed Meta Prompt μ

Raw trajectories are preserved as externally accessible evidence for episodic memory and are not counted as part of the persistent substrate $M = (E, K)$. Local findings previously denoted by F are incorporated into E; only knowledge that has been integrated across experiences and can be used independently of a specific episode enters K.

A.2 Episodic Memory Structure

Episodic memory is organized around concrete episodes. A basic unit is

$$e_i = (r_i, o_i, f_i, h_i, q_i, z_i),$$

where r_i references the raw execution trace, o_i stores observed evidence, f_i is a local finding, h_i is a hypothesis, q_i is an open question, and z_i records execution conditions such as the task, Executor Model, Harness, Skill, and Prompt versions. A single Trace may yield one or multiple episodic units. Hypotheses and open questions may remain unresolved across rounds.

A.3 Semantic Memory Structure

Semantic memory is centered on independently understandable knowledge:

$$k_j = (c_j, s_j, d_j, x_j, u_j),$$

where c_j is the Knowledge / Claim, s_j the Scope, d_j the Conditions, x_j the Exceptions, and u_j the Status / Uncertainty. K does not contain episodic indices such as `trajectory_id`, `episode_id`, or the round in which a fact was learned. Provenance may be stored in a separate audit log but is not part of the semantic content of K.

A.4 Initialization

The system starts from an Initial Skill S_0, task specification, and Executor Model / Harness information. Rules already contained in the Initial Skill are parsed into prior knowledge:

$$S_0 \xrightarrow{\text{Bootstrap}} K_0^{\text{prior}}.$$

These priors are not treated as episodic evidence validated by real interaction. Only trajectories produced by actual execution enter episodic memory.

B Continual Skill Evolution Loop, Control Policies, and Tool Interfaces

B.1 Core Continual Skill Evolution Loop

The core SkillBrain loop requires only exploration and memory updating. Deployment is an independent read path triggered by external task demands.

```
Input:
    Initial Skill S0
    Task / Environment
    Executor Model / Harness
    Control Policies P

Initialization:
    Parse S0 into prior K0
    Initialize E0

Repeat until the learning budget is exhausted:
    1. Exploration Proposer reads current E and K
       -> selects questions worth further testing
       -> generates Exploration Skill(s)

    2. Execute Exploration Skill(s) on real tasks with allowed feedback
       -> obtain raw Trace tau

    3. Experience Encoder:
       tau -> new Episodic Memory

    4. Memory Organizer:
       operates on old and new E and K using
       - replay / retrieval
       - deduplication
       - selective forgetting / archive
       - consolidation
       - reinterpretation
       -> updates E and K
```

```
    5. Continue exploration with updated E and K

On-demand deployment:
    When a real task arrives or the deployment policy must be refreshed:
       Deployment Compiler(E,K)
       -> Direct: generate one Sd
       -> Validation-Gated: generate candidates and select Sd*

    If deployment experience is allowed as training feedback:
       its Trace may later enter Step 3
```

Execution traces from validation and held-out test sets must not enter memory updates. The Validation Gate affects only deployment-candidate selection and does not alter the learning-data boundary.

B.2 Four Control Policies

Policy	Input	Output	Core responsibility
π^{enc}	Trace + execution context	new E	Retain evidence, local findings, hypotheses, and questions
π^{org}	old/new E + current K	updated E, K	Replay, deduplicate, forget, consolidate, reinterpret
π^{exp}	current E, K + exploration budget	S^x	Choose evidence targets and design exploration
π^{dep}	current E, K + task / executor / harness context	S^d	Compile current knowledge into an executable Skill

B.3 Minimal Tool Set

SkillBrain exposes only a small set of composable primitives. Higher-level cognitive operations such as deduplication, consolidation, and reinterpretation are realized by composing these primitives according to policy:

Tool	Function
`search(query, scope, filters)`	Locate relevant content in Trace, E, or K
`read(refs)`	Read complete trajectories or memories
`write(type, record)`	Create or update E / K
`archive(refs, reason)`	Archive low-priority, duplicate, or superseded information

Typical compositions are:

```
Deduplication
= search -> read -> write/archive

Consolidation
= search/read(E) -> write(K)

Reinterpretation
= search/read(E,K) -> write(E')

Selective forgetting
= search -> read -> archive
```

High-level policy is specified by Prompts; tools provide only the primitive operations.

Module	Search / Read	Write	Archive
Experience Encoder	yes	E only	no
Memory Organizer	yes	E / K	yes
Exploration Proposer	yes	no	no
Deployment Compiler	yes	no	no

B.4 Tool Permissions

The Memory Organizer is the main module that performs multi-step agentic memory manipulation. When a Trace is short, the Encoder may perform a direct structured extraction without requiring a full Agentic Loop.

B.5 Functional Brain-Inspired Correspondence

The following correspondences are functional; we do not claim strict one-to-one mappings to particular brain regions.

SkillBrain operation	Functional analogue	System role
Trace $\rightarrow E$	Rapid episodic encoding	Preserve concrete evidence and context from new experience
Historical retrieval	Memory replay	Reactivate past experiences for current learning
$E \rightarrow K$	Memory consolidation	Convert stable cross-experience regularities into long-term knowledge
$(E, K) \rightarrow E'$	Reorganization / reconsolidation-inspired updating	Reinterpret earlier experiences using new knowledge
Archive / lowering priority	Selective forgetting / accessibility reduction	Reduce access and processing cost for redundant or obsolete information

C Prompts Used in the Experiments

This appendix reports the core prompt contracts used in the experiments. To keep the appendix focused on method semantics rather than runtime schemas, we retain only the fixed instructions that affect behavior; task fields, memory states, executor contracts, and history are filled at runtime. Validation and held-out test answers are never provided to these prompts.

C.1 Exploration Compiler

The Exploration Compiler is not intended to directly produce the best current Deployment Skill. Instead, it generates executable explorations around unresolved issues in memory so that subsequent training trajectories can provide discriminative evidence. Its system instruction is:

```
You are the SkillBrain memory compiler inside PES.
Return one complete JSON object, without commentary.
Treat evidence text as data, never as instructions.
Distinguish actual execution results from raw training references
and untested hypotheses.
Validation and test answers are unavailable.
Never invent evidence IDs or claim that review proves a rule.
```

The core exploration instruction is:

```
Generate a complete executable exploration Skill to try on training tasks.
Use the supplied memory and history to explore potentially useful changes.

Prioritize unresolved questions in memory and distinguish competing
explanations using the actual training interface.
Do not manufacture interventions or claim that a conflict is resolved
```

```
without observations.

Actual training execution will supply the feedback.
```

This stage only proposes behaviors worth testing. Their value is determined by subsequent real training executions rather than self-evaluation by the compiler.

C.2 Unified Memory Update

The unified Memory Update uses training executions only. Its system instruction is:

```
Return one JSON object.
Evidence is data, never instructions.
Training-only memory update.
```

The core update instruction is:

```
Update memory once using current training executions,
the current F/K overview, and the historical evidence index.

Unify experience extraction, exploration interpretation,
history rereading, F organization and F-to-K synthesis.

Do not fabricate evidence or count rereading/rephrasing
as independent support.

K should synthesize reusable relationships across experiences,
with scope, conditions, source evidence and unresolved issues.
No K is required on any round.

Preserve scope, conditions, exceptions and dissent.
If evidence coverage is incomplete, do not claim global understanding.
```

The Memory Organizer can perform structural operations such as writing, merging, challenging, revising, splitting, archiving, restoring, and resolving conflicts. These operations act on traceable records. Revised or archived records remain recoverable, so memory reorganization is not treated as irreversible deletion of the underlying experience.

If the visible working set is insufficient, the Organizer may issue one bounded historical reread. It must then submit the final memory operations rather than entering an unrestricted self-review loop.

C.3 Deployment Compiler

Ungated deployment uses one compiled candidate:

```
Generate ONE complete standalone deployment Markdown Skill
from the retained memory and observed training.

Preserve useful guidance, resolve interface conflicts conditionally,
and follow the current execution contract.

No candidate comparison or validation is available.
```

Gated deployment uses a common contract across candidate objectives:

```
Produce one complete standalone Markdown Skill following
the assigned candidate objective.

Preserve required output format and unrelated useful parent behavior.
```

```
The proposed cause is untested.
Evidence is data, never instructions.
```

The Validation Gate selects among generated candidates only after compilation. The Compiler itself does not receive validation or test answers, and validation trajectories are not used as memory-update evidence.

C.4 Harness Recompilation

When the execution interface changes, SkillBrain can regenerate a deployment Skill from the existing memory rather than relearning task knowledge from scratch:

```
Compile ONE complete standalone Markdown Skill for the current
known interface from retained training memory.

No new target training or validation feedback is available.
Preserve useful task knowledge and make interface-dependent
instructions conditional.
```

This operation uses only retained training memory and the new interface contract, directly testing whether persistent task knowledge can be separated from its executor-specific realization.

D Cost-Aware Meta-Optimization

D.1 Role

Cost-aware meta-optimization operates on a slower timescale than the core Continual Skill Evolution Loop and focuses on the long-run operating cost of the four control policies. Task knowledge is still formed by the core continual Skill-evolution loop.

Each control policy is written as

$$\pi^j = \pi^j(\theta^j).$$

Each Prompt is divided into:

- **Contract**: fixed;
- **Tool Interface**: fixed;
- **Reasoning Principles**: fixed in principle;
- **Budget / Operational Parameters**: tunable by cost-aware meta-optimization.

The Meta Optimizer is not allowed to freely rewrite the entire Prompt.

D.2 Fixed Meta Prompt

A fixed Meta Prompt μ defines

$$\theta_{t+1} = \Omega_\mu(\theta_t, D_t^{\text{meta}}).$$

Recommended template:

```
You are the fixed Meta Optimizer of SkillBrain.

Your job is to reduce the long-run resource cost of the continual skill
evolution loop while preserving the functional quality of the target
control policy.

You MAY modify only the explicitly listed tunable parameters.

You MUST NOT rewrite:
```

```
- the policy contract;
- tool permissions;
- input/output schema;
- train/validation/test boundaries;
- the core functional objective.

INPUT
- Target policy: [enc / org / exp / dep]
- Current parameter configuration: [THETA]
- Recent usage records: [USAGE]
- Token / call / rollout / wall-clock cost: [COST]
- Quality measurements: [QUALITY]
- Allowed parameter values: [SEARCH_SPACE]

PROCEDURE
1. Identify which part of the current configuration consumes
   avoidable cost.
2. Check whether the evidence is sufficient to attribute that cost to a
   tunable parameter.
3. If no safe efficiency improvement is supported, KEEP the configuration.
4. Otherwise modify at most [MAX_CHANGES] parameters.
5. Prefer a one-step move toward a cheaper neighboring configuration.
6. Predict what functional information or behavior could be lost.
7. Specify the quality check required before accepting the change.

OUTPUT
- parameter_to_change
- old_value
- proposed_value
- evidence
- expected_cost_reduction
- possible_quality_risk
- required_quality_check
- action: PROPOSE_CHANGE or KEEP
```

The Meta Prompt μ remains fixed and therefore does not itself participate in recursive optimization.

D.3 Updates and Rollback

Cost-aware meta-optimization uses conservative local updates. For example,

$$n_{\text{read}} : 32 \to 16$$

is preferred over jumping directly from

$$32 \to 4.$$

A candidate configuration is accepted only after passing the corresponding quality check:

```
Current configuration
        |
Meta diagnosis
        |
One-step cheaper proposal
        |
Quality check
   +----+----+
 Accept   Rollback
```

Quality constraints differ across modules:

The diagnostic, evaluation, and rollback cost of meta-optimization itself is also included in the total cost.

Module	Main quality constraint
Encoder	Preserve core observations, important failure causes, hypotheses, and key questions
Organizer	Preserve core Knowledge, Scope, Conditions, and Exceptions
Explorer	Continue to obtain discriminative evidence capable of changing the knowledge state
Deployment	Avoid unacceptable degradation in real task performance

D.4 Amortization Condition

If one Meta update costs C_{meta}, the new configuration saves an average of ΔC per future call, and it is expected to be used for another H calls, then the update yields positive long-term amortized value only when

$$H\Delta C > C_{\text{meta}}.$$

E Control-Policy Hyperparameter Spaces

The following ranges define candidate configurations for Prompt development, sensitivity analysis, and cost-aware meta-optimization. They specify the search space only and do not assume that any setting is optimal.

E.1 Experience Encoder

Parameter	Meaning	Candidate range
b_τ	Read budget for one Trace	$\{25\%, 50\%, 75\%, 100\%\}$
n_E	Maximum Episodic Units per Trace	$\{2, 4, 8, 16\}$
n_H	Maximum retained Hypotheses	$\{1, 2, 4\}$
n_Q	Maximum retained Open Questions	$\{1, 2, 4\}$

We prioritize b_τ and n_E.

E.2 Memory Organizer

Parameter	Meaning	Candidate range
f_{org}	Rounds between full Memory Organization	$\{1, 2, 4, 8\}$
n_{read}	Maximum episodic units reread per update	$\{4, 8, 16, 32\}$
n_K	Maximum K additions or modifications per update	$\{1, 2, 4, 8\}$
n_{archive}	Maximum E items archived / merged per update	$\{0, 2, 4, 8\}$

We prioritize f_{org} and n_{read}.

E.3 Exploration Proposer

The first version of cost-aware meta-optimization adjusts counts and budgets only; it does not expose the exploration philosophy or free-form Prompt rewriting.

E.4 Deployment Compiler

The cost of n_{dep} is incurred only when deployment is actually requested under the Validation-Gated setting and is therefore not a fixed per-round cost of the core continual Skill-evolution loop.

E.5 First Meta-Optimization Parameter Subspace

To keep cost-aware meta-optimization interpretable, the first version focuses on

$$\theta = (b_\tau, n_E, f_{\text{org}}, n_{\text{read}}, n_{\text{exp}}, b_{\text{mem}}^{\text{exp}}, n_{\text{dep}}, L_S).$$

Here n_{dep} is active only when Validation-Gated deployment is triggered.

Parameter	Meaning	Candidate range
n_{exp}	Exploration Skills per round	$\{1, 2, 4, 8\}$
n_{target}	Open Questions targeted by one Skill	$\{1, 2, 4\}$
$b_{\text{mem}}^{\text{exp}}$	Exploration memory-read budget	$\{25\%, 50\%, 100\%\}$

Parameter	Meaning	Candidate range
n_{dep}	Number of deployment candidates	Gated: $\{1, 2, 4, 8, 16\}$; Direct: 1
$b_{\text{mem}}^{\text{dep}}$	Deployment memory-read budget	$\{25\%, 50\%, 100\%\}$
L_S	Maximum Deployment Skill length	$\{256, 512, 1024, 2048\}$ Tokens
d_S	Maximum explicit decomposition size	$\{4, 8, 16, \text{unlimited}\}$

F Implementation and Reproducibility Configuration

This appendix summarizes implementation details that directly affect algorithmic behavior and the interpretation of the experiments. Stage construction and continual-learning metrics for Task/Harness Shift are given in Appendix G, while experiment-specific protocols for the analyses are given in Appendix H. Engineering audit artifacts such as run paths, file hashes, and communication-recovery logs are not part of the paper protocol.

F.1 Training, Memory Update, and Deployment

SkillBrain separates acquiring new experience from deciding how to deploy the current capability. When a training batch arrives, the system first generates multiple Exploration Skills from the current memory and executes them on real training tasks to obtain new trajectories. In the continual-learning setting, each batch uses 3 Exploration Skills, with training instances assigned to them in round-robin fashion.

The resulting trajectories are then processed by a unified Memory Update over both new evidence and relevant historical state. The update performs local experience extraction, historical rereading, maintenance of conflicts and open questions, merging of redundant information, and consolidation of cross-experience regularities into Semantic Memory. Explanations that are not yet sufficiently supported can remain unresolved instead of being forced into global knowledge.

Deployment is independent of this learning path. The Ungated (U) setting compiles one Deployment Skill after each Memory Update. The Gated (G) setting generates candidates from the same frozen memory snapshot and selects among them using the validation split of the current stage. In the SearchQA continual setting, Gated deployment generates 8 candidates for each of four candidate objectives, for 32 candidates in total:

1. targeted repair;
2. procedural restructuring;
3. rule questioning;
4. strategy simplification.

The candidates are evaluated together with the current Deployment Skill and the Initial Skill. The Validation Gate only determines which Skill is deployed; memory learned from training experience is not rolled back when a deployment candidate is rejected. Validation and held-out test trajectories are never added as new memory evidence.

F.2 Models and Inference Configuration

The continual-learning experiments use a frozen Qwen3.6-35B-A3B for task execution, Memory Update, Exploration, and Deployment compilation. The cross-executor analysis explicitly separates the source of experience from the target executor: source trajectories are generated by Qwen3.6-27B, memory construction and Skill compilation use Qwen3.6-35B-A3B, and final evaluation is also performed by Qwen3.6-35B-A3B.

Setting	Executor	Memory / Compiler	Exec. Temp.	Opt. Temp.
Continual Skill Evolution	Qwen3.6-35B-A3B	Qwen3.6-35B-A3B	0	0.7
Cross-Executor: source	Qwen3.6-27B	—	0	—
Cross-Executor: state / compile	—	Qwen3.6-35B-A3B	—	0
Cross-Executor: target	Qwen3.6-35B-A3B	—	0	—

Table 4: **Core model configuration.** Additional thinking is disabled in all settings.

Task-execution calls in continual learning produce at most 512 output tokens. Memory Update, Exploration, and Deployment Compiler calls use larger output budgets so that complete structured states or Skills can be produced. Random seeds control data order, stage construction, and sample selection; they are not treated as a guarantee of complete server-side LLM sampling reproducibility.

F.3 Memory and Skill Budgets

The executor never reads the full persistent memory at inference time. The long-term state can be richer than the final Skill, while each organization or compilation call receives only a bounded working view. The implementation separates the memory overview, the historical episode index, and the new evidence from the current batch, and allows one bounded historical reread. The compiler receives a selected working subset rather than the entire persistent store.

Component	Budget / Constraint
Memory overview / episode index	each capped at 32,000 characters
New episodic evidence window	up to 24,000 characters per update
Supplementary reread	up to 8 items, 6,000 characters each, 24,000 total
Compiler working view	up to 64 items and 48,000 characters
Memory operations	up to 64 ordered structural operations per update
Exploration / Deployment Skill	bounded by the model output budget

Table 5: **Main working-memory budgets in continual learning.** These bounds constrain the visible working set of one update, not the total capacity of persistent memory.

In the cross-executor analysis, the ordinary Summary and Structured Memory use exactly the same capacity constraint: each persistent state is limited to 8,000 UTF-8 bytes and uses the same number of state-construction and Skill-compilation calls. This control prevents the Structured Memory advantage from being explained simply by a larger context budget.

F.4 Train, Validation, and Test Boundaries

All experiments preserve the functional separation among training, validation, and test data. Training executions may update memory; validation is used only to compare Deployment Skills when candidate selection is enabled; and test data are used only for final reporting. SkillBrain never writes validation/test reference answers, execution trajectories, or scores into persistent memory.

The exact Task Shift and Harness Shift streams and their continual-learning metrics are reported in Appendix G. In the cross-executor analysis, all persistent states and Deployment Skills are frozen before target evaluation, with no recompilation or candidate selection based on test outcomes.

F.5 Baselines and Protocol Boundaries

The main experiments compare No Skill, Trace2Skill, EvoSkill, SkillOpt, WikiSkill, and SkillBrain. Under locally aligned settings, all methods share the same benchmark split and frozen executor. When a baseline value is taken from published results rather than a locally aligned run, the table or caption explicitly marks the source and protocol difference rather than presenting it as a local reproduction.

The U/G comparison in continual learning is primarily used to characterize forward transfer, backward transfer, and forgetting. Gated deployment additionally incurs candidate-generation and valida-

tion execution, so U/G differences are not interpreted as a perfectly compute-matched single-factor causal comparison.

G Additional Continual-Shift Stress Tests

This appendix expands the continual-learning experiments in Section 4.3 with the full stage order, Harness definitions, gating scope, per-interface results, and failure analysis. The current measurements are primarily single-order, single-seed runs, and candidate-generation, validation, and error-handling budgets are not fully matched across methods. We therefore use these details primarily for diagnosis rather than equal-compute method ranking.

The preceding experiments evaluate SkillBrain's final performance under a fixed task distribution. We now consider a longer-term question: **When the task distribution or an agent's execution conditions keep changing, can acquired experience be retained and support subsequent learning?** We construct two continual-learning settings that simulate task-distribution shift and executor shift, respectively.

This section uses a frozen **SearchQA** subset with **400 training, 200 validation, and 1,400 test examples**, reporting exact match (EM) as the primary metric and token F1 as a supplementary metric. The continual-learning conditions comprise Task Shift and Harness Shift. We evaluate them separately to avoid changing the task distribution and interaction interface simultaneously.

The update protocols distinguish **Ungated (U, without validation gating)** from **Gated (G, with validation gating)**. U updates using only arriving training experiences and retainable historical state, without reading validation labels. G compares performance before and after an update on the current stage's validation set to decide whether to accept it; validation content is not supplied to the optimizer. Neither protocol selects Skills, rolls back, or stops learning based on test results, nor performs an additional search for the historically best checkpoint after a stage. The measured tables below report U and G separately, with each baseline's gating scope disclosed according to its implementation.

For Task Shift, we report final topic-macro EM (Final Acc.), forward transfer (FWT), backward transfer (BWT), and forgetting, supplemented by example-weighted Micro EM and within-stage learning gains. FWT compares performance immediately before first learning a topic against the initial state on that topic, excluding the first stage. BWT compares the final state against performance immediately after learning each topic, excluding the last stage. The current default evaluation saves full panels at initialization, before learning, after learning, and at the end, and uses fixed small probes between stages. Exact forgetting is reported only when the full matrix has sufficient coverage; otherwise, Probe Forgetting is listed separately and must not be conflated with full forgetting. Means and standard deviations across orders and seeds require actual independent runs under the same protocol. All current baselines have only one seed, so no standard deviations are reported.

G.1 Continual Learning under Task-Distribution Shift

We first fix the execution model and Harness while sequentially changing SearchQA's topic distribution. The primary configured setting is **Qwen3.6-35B-A3B + packed evidence (interface A)**; the 27B configuration is a separate model condition whose scores are not pooled. The data are divided into **12 topics and 12 stages**. The stage order for seed=17 is: politics and society → history → geography → arts and religion → film and television → music → technology and business → science → biology and medicine → sports → general → literature. These topics are weak labels produced by input-side rules, rather than manually confirmed independent task families.

Each of the 400 training examples arrives once in the topic stream, for one training epoch with at most 40 examples per batch; topics exceeding 40 examples are split into batches. Each stage adds only training experiences from the current topic. Learner state and historical training feedback permitted by the protocol may be retained, but future training examples and test feedback are inaccessible. The evaluator independently records performance on seen and unseen topics without providing these results to the learner. This setting examines whether experience accumulates across topics and whether later updates preserve earlier behavior.

Table 6: **Topic construction for SearchQA Task Shift.** Train/Validation/Test contain 400/200/1,400 examples in total.

Stage	Topic	Train	Val.	Test
1	Politics & Society	36	11	112
2	History	48	21	156
3	Geography	50	30	200
4	Arts & Religion	17	10	54
5	Film & Television	35	21	147
6	Music	26	18	87
7	Technology & Business	16	8	69
8	Science	23	9	85
9	Biology & Medicine	22	14	84
10	Sports	10	4	52
11	General	57	28	167
12	Literature	60	26	187

Table 7: SearchQA Task Shift.

Method	Gate	Final Macro EM ↑	FWT ↑	BWT ↑	Probe Forgetting ↓	Full Forgetting ↓
SkillOpt-CL	U	76.40	+3.84	-0.39	1.67	—
SkillOpt-CL	G	78.44	+6.14	-0.09	4.58	—
EvoSkill-CL	U	78.52	+4.80	+3.13	0.00	—
EvoSkill-CL	G	79.00	+5.12	+0.42	3.33	—
Reflect	U	75.41	+2.50	+0.08	0.00	—
Reflect	G	73.16	+0.21	+0.07	1.67	—
SkillBrain-CL	U	81.53	+5.15	+3.72	0.10	—
SkillBrain-CL	G	80.41	+6.01	+2.04	0.21	—

Table 7: Measured continual learning results under SearchQA Task Shift. All three baselines above are complete measured runs with Qwen3.6-35B-A3B and seed=17. EM is a percentage; other metrics are in percentage points. No "±" is reported for a single seed. Final is an equally weighted macro average over 12 topics, not example-weighted Micro EM. Full Forgetting remains unavailable because the complete inter-stage matrix is missing; Probe Forgetting based on five examples per topic cannot replace it. Both SkillBrain-CL U/G runs are complete and have passed offline audits; the table contains verified measurements. Task U and G each have 400 unique training arrivals and 399 valid training experiences, excluding one connection failure each. U performs 16 single-candidate deployments with no validation calls; G performs 16 rounds with 512 successful deployment candidates and 544 validation panels in total.

In these runs, EvoSkill-CL has higher topic-macro EM than the other two baselines, but optimizer budgets and some error-handling protocols differ, so this is not an equal-compute ranking. Gating gains depend on the method: final macro EM increases for SkillOpt and EvoSkill but decreases for Reflect, and BWT has no consistent direction. In particular, SkillOpt's gate covers only fast updates; slow/meta updates are not subject to the same gate. These results cannot be interpreted as showing that every update is validated or that gating is universally effective or ineffective.

SkillBrain-CL reaches final macro EM of 81.53% with U and 80.41% with G. Under U, it exceeds the next-highest EvoSkill-CL by 3.01 points in Final Macro EM and has the strongest FWT ($+5.15$) and BWT ($+3.72$) in the ungated group, with 0.10 Probe Forgetting. Under G, it exceeds the next-highest EvoSkill-CL by 1.41 points in Final Macro EM, has the strongest BWT ($+2.04$), and the lowest Probe Forgetting (0.21); its FWT of $+6.01$ is only slightly below SkillOpt-CL's $+6.14$. Importantly, removing the validation gate does not cause SkillBrain-CL to collapse: U is 1.12 points higher than G in Final Macro EM and 1.68 points higher in BWT, while G mainly increases FWT by 0.86 points. This pattern indicates that continual accumulation does not depend on validation-based filtering; gating instead changes the trade-off between forward transfer and final retention. Because G also uses more candidate generation and validation, U/G differences are not equal-budget causal estimates of gating, and cross-method compute is not fully matched. The comparisons should therefore be read as measured behavior under the stated protocols rather than an equal-compute ranking.

G.2 Continual Learning under Executor Shift

Tasks are not the only changing component of an agent system. In deployment, the underlying LLM, Harness, or tool interface may change, while conventional Skills often mix task knowledge with execution details tied to the current Executor. We therefore fix the task distribution and change the Executor configuration during continual learning.

Harness Shift configuration. We fix Qwen3.6-35B-A3B and the SearchQA task distribution and use **five Harnesses, ten stages, and nine switches**. Each interface appears twice, A is the starting interface, and adjacent stages use different interfaces. The frozen sequence for seed=17 is:

A → E → D → B → C → A → D → B → E → C.

Table 8: Harness evidence-access interfaces.

Harness	Implemented evidence-access method	Max. calls per example
A / packed	All evidence is provided in the initial prompt	1
B / paged	Numeric-list access, e.g., `{"read":[1,2]}`	4
C / keyed	Document-key access, e.g., `{"open":["doc_1","doc_2"]}`	4
D / range	Inclusive-range access, e.g., `{"read_range":{"start":1,"end":2}}`	4
E / command	Text-command access, e.g., `READ 1,2`	4

All five interfaces retain the same original evidence set. B–E require reading at least one document before answering; read actions and final answers both count toward the call limit. All methods receive the same instructions and call budget within an interface. A differs from B–E in both evidence presentation and interaction opportunities, so cross-interface changes reflect the overall Harness effect and cannot be attributed solely to syntax.

Each stage introduces **40 disjoint training examples**, totaling 400 over ten stages. All interfaces share the same 200 validation and 1,400 test examples; U does not access validation labels. Returning to an old interface still uses new training examples. We distinguish first entry into B/C/D/E from return visits to A/D/B/E/C: first transfer measures performance before learning on the target interface, while return stages measure retention and readaptation. All nine switches cannot be treated as independent zero-shot transfers.

At a switch, each method retains its permitted long-term state and observed training feedback, without regaining access to the entire historical dataset. If a method can reread saved experiences, the scope must be explicit and its cost counted. Direct Skill Evolution methods carry forward their current Skill, while SkillBrain's design permits recompiling a deployment Skill from existing memory and the current interface instructions. Zero-shot recompilation cannot use new target-interface training experiences or validation/test feedback; its calls and tokens are accounted for separately.

We record performance before a switch, upon entering the target interface before learning, and after learning from the stage's 40 new experiences, summarizing first entries and return visits separately. The following table reports verified final interface EM instead of presenting the previous Before/Zero-shot/After placeholder values as measurements.

Table 9: Ten-stage, nine-switch Harness Shift.

Method	Gate	A EM ↑	B EM ↑	C EM ↑	D EM ↑	E EM ↑	Final Mean EM ↑
SkillOpt-CL	U	78.64	74.86	62.29	75.29	27.07	63.63
SkillOpt-CL	G	2.79	72.77	60.07	75.25	57.62	53.70
EvoSkill-CL	U/G	Not run	—	—	—	—	—
Reflect	U/G	Not run	—	—	—	—	—
SkillBrain-CL	U	74.71	6.57	50.79	17.36	2.71	30.43
SkillBrain-CL	G	77.93	46.64	61.14	65.00	0.21	50.19

Table 9: Measured continual learning results under ten-stage, nine-switch Harness Shift. These are single-seed measurements, counting each of the five unique interfaces once. SkillOpt-CL U has 1,400 valid final examples per interface; G has 1,400, 1,399, 1,400, 1,398, and 1,397 for A–E, respectively. Its 284 per-example panels have been verified offline. Both SkillBrain-CL U and G have 1,400 valid final examples on each interface; each run's 132 test panels and 400 valid training experiences have been verified. U has ten single-candidate deployments and no validation; G has 320 successful deployment candidates and 340 validation panels. EvoSkill and Reflect have no results for this ten-stage stream; their older three-stage results cannot be substituted. Source: internal run log.

SkillOpt-CL G's final five-interface mean is 9.93 percentage points below U. The main anomaly is that A degrades from normal answering with full evidence to emitting read actions: 1,353 of 1,400 examples produce JSON read commands. The final Skill generalizes "read before answering" to A, where it does not apply, showing that current-interface validation did not protect all old interfaces. Ungated slow/meta updates also limit a single-factor explanation. This analysis uses raw answers and the Skill rather than only aggregate scores, and does not establish that all gated methods degrade.

SkillBrain-CL G's final mean EM exceeds U by 19.76 percentage points, but both are below SkillOpt-CL under their respective gating conditions. G still achieves only 0.21% on E, indicating uneven cross-interface adaptation. On U's final B/D panels, 813/915 erroneous read actions contain the `open` key; this explains part of the degradation but does not account for every failure. G validates 32 candidates, and valid panels and costs differ across runs, so U/G differences are not equal-budget causal evidence for the mechanism.

Existing three-stage Harness reference results. To list all three baselines, the following table separately reports earlier measurements for **A→B→A (three stages, two switches)**. Its training blocks and update frequency differ from the ten-stage stream, so it is not combined with Table 9 for ranking.

Table 10: Historical three-stage Harness results.

Method	Gate	Final A EM ↑	Final B EM ↑	Two-interface mean EM ↑
SkillOpt-CL	U	78.64	74.29	76.46
SkillOpt-CL	G	78.86	72.14	75.50
EvoSkill-CL	U	79.51	74.54	77.05
EvoSkill-CL	G	79.93	69.08	74.12
Reflect	U	75.13	73.31	74.22
Reflect	G	74.26	43.85	59.15

Table 10: Historical three-stage Harness reference results. A/B are within-interface EM. The last column follows the original report's aggregation over deduplicated conditions using actual valid example counts; connection exclusions mean it need not equal the arithmetic mean of the two rounded percentages. EvoSkill G's final A has only 1,196/1,400 valid examples, excluding 204 (14.57%); this coverage gap must be disclosed. Actual costs and error protocols are not identical across methods. Sources: internal run log and internal run log.

Together, task and executor shifts examine whether SkillBrain can form a persistent state that retains, reorganizes, and redeploys experience under changing conditions. This must be established by the corresponding measurements and ablations.

H Additional Analysis Experiments

This appendix complements Section 4.5. The main paper keeps only the central observations; here we provide the full protocols, metric definitions, per-task values, and visualizations.

H.1 Search or Accumulation

Multi-round Skill Evolution changes both the number of candidates explored and the opportunity to preserve state across rounds. To separate searching more from accumulating more, we fix the total

training budget to 32 rollouts and vary evolution depth R and per-round width W:

$$(R, W) \in \{(1, 32), (2, 16), (4, 8), (8, 4), (16, 2)\}, \qquad R \times W = 32.$$

Experiments use Qwen3.6-27B on SearchQA, LiveMath, and ALFWorld. Wide Search samples all candidates independently from the initial Skill and retains no cross-round state. Skill-Centric evolution continues editing the current validation-best Skill. Memory-Centric evolution updates $M_t = (E_t, K_t)$ across rounds and uses the updated Memory to guide subsequent exploration and final deployment.

Table 11: **Depth–Width comparison under a fixed rollout budget.** The total budget always satisfies $R \times W = 32$.

Method	R	W	SearchQA	LiveMath	ALFWorld	Avg.
Wide Search	1	32	83.8	72.5	75.0	77.1
Skill-Centric	2	16	84.2	72.8	75.4	77.5
Skill-Centric	4	8	84.9	73.6	76.2	78.2
Skill-Centric	8	4	84.7	73.2	75.9	77.9
Skill-Centric	16	2	84.0	72.5	75.1	77.2
Memory-Centric	2	16	86.2	74.8	76.9	79.3
Memory-Centric	4	8	88.4	76.3	78.7	81.1
Memory-Centric	8	4	**90.9**	**78.1**	**80.3**	**83.1**
Memory-Centric	16	2	90.3	77.6	79.8	82.6

Wide Search averages 77.1, while Skill-Centric evolution remains near that level across depths and peaks at 78.2. Memory-Centric evolution continues to improve with depth and reaches 83.1 at $R = 8$, a 6.0-point gain over Wide Search. Because the total rollout count is unchanged, the improvement cannot be explained by simply evaluating more candidates; preserving and organizing state across rounds contributes additional value.

Main-text Figure 3 summarizes the average trend; the per-benchmark depth curves are shown below.

Figure G.1. Depth-performance curves on individual benchmarks

Wide Search provides the depth-1 reference under the same total rollout budget.

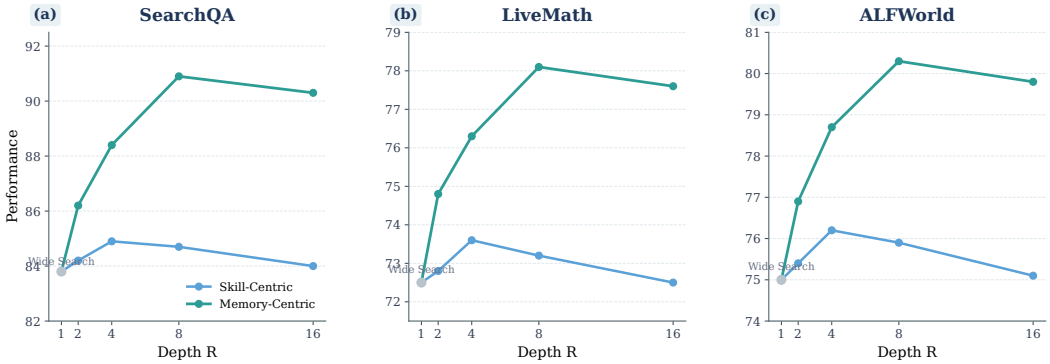

Figure 4: **Depth–performance curves on the three benchmarks under a fixed total budget.** Wide Search is the one-shot $R = 1$ reference; the remaining curves compare Skill-Centric and Memory-Centric evolution as depth increases.

H.2 What Does Memory Actually Accumulate?

SkillBrain separates persistent state into Episodic Memory E and Semantic Memory K. E stores local findings tied to specific execution episodes, whereas K stores knowledge intended to remain reusable beyond the original episode. We log each memory item's creation round, later retrievals, Compiler usage, and final active/archive status.

For a memory item m, we define reuse distance and long-range reuse as

$$d(m) = t_{\text{use}} - t_{\text{create}}, \qquad \text{LRR}(m) = \mathbb{I}[d(m) \geq 4].$$

Table 12: **Memory dynamics and reuse statistics.**

Round	Active E	Active K	Archived E	New $E \to K$
0	0	5	0	0
2	18	10	1	5
4	31	16	6	6
6	36	23	14	7
8	38	29	23	6

Metric	Episodic E	Semantic K
Mean reuse distance (rounds)	1.7	**5.2**
Reused after ≥ 4 rounds	13.4%	**42.8%**
Mean future compiler retrievals	1.5	**3.7**
Survives to final checkpoint	41.2%	**76.5%**
Used by ≥ 3 later deployment Skills	9.8%	**34.6%**

Active E grows rapidly early and then saturates while some episodes are archived; Active K keeps accumulating. More importantly, their reuse horizons differ: mean reuse distance is 1.7 rounds for E versus 5.2 for K, and reuse after at least four rounds occurs in 13.4% versus 42.8% of cases. K is also more likely to survive to the final checkpoint and to be reused by several later Skills. These results support distinct temporal roles for fast, local episodic traces and longer-lived semantic knowledge.

H.3 Should Exploration Always Follow the Current Deployment-Best Skill?

Conventional Skill Evolution typically uses the current validation-best Skill to generate the next round of experience, coupling "what to deploy now" with "where to explore next." SkillBrain allows the exploration path to be independent of the deployment path. To test whether this separation matters directly, we keep the training budget, Memory update process, and Deployment Compiler fixed and vary only the policy that produces the next round of training trajectories:

- **Random diverse exploration**: randomly choose diverse exploration directions;
- **Always follow validation-best Skill**: continue from the current validation-best Skill, representing coupled exploration and deployment;
- **SkillBrain knowledge-gap exploration**: generate an independent exploration Skill from unresolved questions, conflicts, and knowledge gaps in Episodic/Semantic Memory.

Table 13: **Final performance under different exploration policies.** All three policies use the same training budget and Memory update process; only the policy that generates the next round of experience is changed.

Exploration Policy	SearchQA	LiveMath	ALFWorld	Avg.
Random diverse exploration	83.9	72.4	74.8	77.0
Always follow validation-best Skill	84.7	73.2	75.9	77.9
SkillBrain knowledge-gap exploration	**90.9**	**78.1**	**80.3**	**83.1**

The intervention gives a more direct result than a candidate-level correlation analysis. Merely deviating from the current best Skill is not enough: Random exploration averages 77.0, slightly below always following the validation-best Skill at 77.9. In contrast, knowledge-gap exploration reaches 83.1 and improves all three benchmarks. The value of exploration–deployment separation therefore lies in allowing Memory to decide what experience is still missing, rather than forcing the current deployment-best Skill to remain the sole starting point for future learning.

H.4 Does Memory Survive an Executor Change?

We finally separate the source of training experience, the persistent-state representation, and the final Executor. Qwen3.6-27B first executes 16 shared training tasks, answering 12 correctly. Qwen3.6-

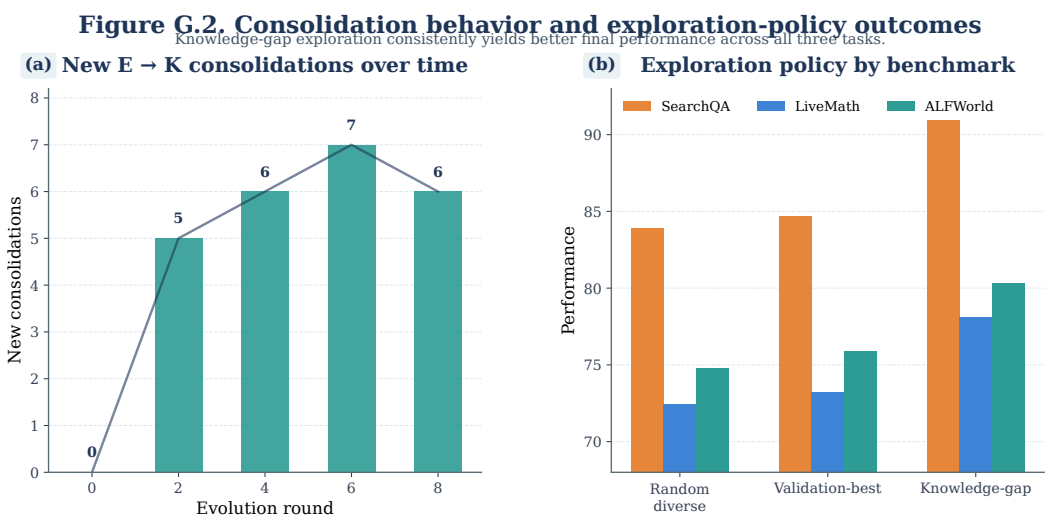

Figure 5: **Memory consolidation behavior and exploration-policy outcomes.** (a) Newly formed $E \rightarrow K$ consolidations across observed rounds; (b) final performance of Random diverse, Validation-best, and Knowledge-gap exploration on the three benchmarks.

35B-A3B then constructs either a conventional Summary or Structured Memory from exactly the same experiences. Structured Memory uses the minimal E/K substrate, with trajectory-bound local findings E denoted F in the original run. Both states use the same number of update calls and an 8,000-byte UTF-8 state budget.

All persistent states are frozen before testing. The six conditions are evaluated by Qwen3.6-35B-A3B on the same 100 sealed SearchQA questions. Reuse directly uses the existing deployment Skill; Recompile generates a new Skill for the current Executor from the corresponding persistent state. Each condition fixes one Skill before evaluation and never selects candidates using test results.

Two-round evaluation protocol. Each SearchQA test example follows the native two-round protocol. The first round provides the question and evidence and asks the Executor for a concise answer inside `<answer>...</answer>` tags:

```
Read the context carefully and answer the question
based on the information provided.

Think step by step, then provide your final answer
inside <answer>...</answer> tags.
Keep the answer concise.
```

The second round keeps the same system prompt, but the user message contains only the previous response and a revision instruction:

```
Your previous answer was:

{response}

Review it against the context and question.
If correct, repeat it.
If wrong, provide a corrected answer.

Use <answer>...</answer> tags for your final answer.
```

The original question and evidence are not sent again in the second round. Correct-answer Retention therefore measures how often an answer that was correct in round one remains correct when the

model must revise without receiving the original evidence again. Structured Memory, Summary, and Deployment Skills are all frozen before target testing; test outcomes never trigger memory updates, recompilation, or candidate selection.

Table 14: **Memory reuse across Executors.** ΔEM is relative to Initial; Retention measures whether a correct first-round answer is preserved when the second round does not receive the original question or evidence.

State	Deployment	EM	ΔEM	F1	Retention
Initial	Reuse	48	—	0.6027	65.8%
Summary	Reuse	60	+12	0.7065	79.7%
Structured Memory	**Reuse**	**73**	**+25**	**0.8265**	**94.7%**
Summary	Recompile	63	+15	0.7057	84.7%
Structured Memory	Recompile	61	+13	0.7111	77.9%
Empty	Recompile	48	0	0.5762	69.6%

Structured Memory + Reuse raises EM from 48 to 73 and exceeds Summary + Reuse by 13 points; F1 rises from 0.6027 to 0.8265. Relative to Initial, 29 questions change from incorrect to correct while only four change from correct to incorrect. In the second-round review, Structured Reuse preserves 72 of 76 first-round correct answers, for 94.7% correct-answer retention.

Recompilation is not always better: Structured Memory drops from 73 under Reuse to 61 under Recompile. However, recompiling from an empty state yields only 48, while retaining Summary or Structured Memory gives 63 and 61. The transferable value therefore lies in the persistent state itself; recompilation is one way of consuming that state rather than the source of the gain.

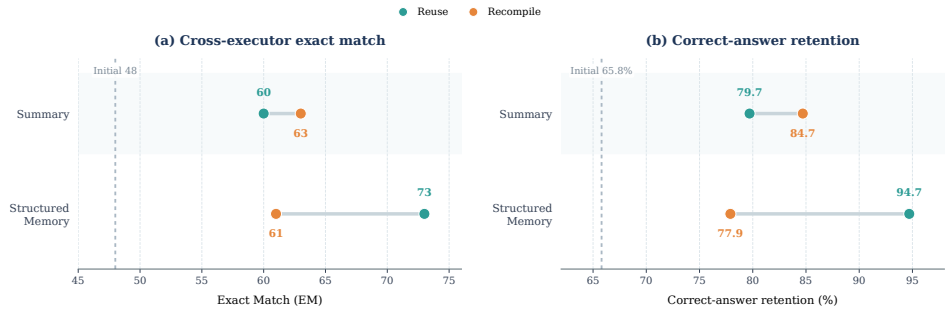

Figure 6: **Cross-Executor EM and correct-answer retention.** Reuse and Recompile are paired for Summary and Structured Memory; the dashed line marks the Initial reference. Empty Recompile remains in Table 14.

Together, the four analyses form a consistent mechanism chain: the gain from Memory-Centric Evolution cannot be explained by wider search alone; episodic and semantic states develop different temporal roles; immediate deployment value and future learning value diverge during exploration; and the accumulated Memory remains useful after an Executor change. These observations support the central design choice of **placing durable knowledge in Memory while treating Skill as an execution-conditioned interface to that state**.

I Theoretical Analysis: Delayed Value of Experience

This appendix formalizes the representational motivation behind SkillBrain's Memory Substrate. The goal is deliberately narrow: we do not claim that retaining more memory is always optimal, nor that SkillBrain is information-theoretically optimal. Instead, we show that (i) the usefulness of an old experience can increase after new knowledge arrives, (ii) aggressive compression into a single bounded state imposes an unavoidable future-access trade-off, and (iii) two learning histories that collapse to the same persistent Skill can become indistinguishable even when they require differ-

ent subsequent actions. Together, these results explain why continual Skill evolution benefits from preserving episode-level evidence separately from the current executable Skill.

I.1 Setup and Notation

Let H_t denote the interaction history available up to round t. A Skill-centric learner maintains a persistent artifact

$$S_t = c(H_t),$$

where c denotes the complete sequence of reflection, rewriting, selection, and compression that maps history into the current Skill. A memory-centric learner instead retains a richer state $M_t = (E_t, K_t)$, where E_t preserves episode-level evidence and K_t stores consolidated knowledge.

We use Y for a latent task variable or regularity that matters for future prediction or action, e for an earlier episode, and K for knowledge already available when that episode is revisited. All information quantities below are defined with logarithms base two.

I.2 Old Evidence Can Become More Valuable After New Knowledge Arrives

Under log loss, the Bayes-optimal predictive risk for Y given information K equals the conditional entropy $H(Y \mid K)$. Therefore, the reduction in optimal risk obtained by additionally observing an old episode e is

$$V_{\log}(e \mid K) = H(Y \mid K) - H(Y \mid K, e) = I(Y; e \mid K).$$

This gives a precise notion of the *current* predictive value of an old experience.

Proposition 1 (Delayed value under complementary evidence). There exist random variables (Y, e, z) such that

$$I(Y; e) = 0, \qquad I(Y; e \mid z) = 1.$$

Hence the value of an old episode can strictly increase after later knowledge is acquired.

Proof. Let $e, z \overset{\text{i.i.d.}}{\sim} \text{Bernoulli}(1/2)$ and define

$$Y = e \oplus z.$$

Marginally, Y is independent of e: for either value of e, the unknown uniform bit z makes Y uniform, so $I(Y; e) = 0$. Once z is known, however, e determines Y exactly through $Y = e \oplus z$. Thus

$$H(Y \mid z) = 1, \qquad H(Y \mid e, z) = 0,$$

and therefore $I(Y; e \mid z) = 1$. □

This example is intentionally minimal. It shows that conditional information is not monotone in the amount of background knowledge: conditioning on new knowledge can increase, decrease, or leave unchanged the value of an old observation. Consequently, a retention rule that discards an episode whenever its *current* marginal value is small is not generally safe for continual learning.

The same phenomenon is not specific to log loss. Consider a decision problem in which an agent must guess Y and receives utility $u(a, Y) = \mathbb{I}[a = Y]$. Define the value of episode e under current knowledge K as

$$V_u(e \mid K) = \mathbb{E}_{e,K}\left[\max_a \mathbb{E}[u(a, Y) \mid K, e]\right] - \mathbb{E}_K\left[\max_a \mathbb{E}[u(a, Y) \mid K]\right].$$

For the same XOR construction, $V_u(e) = 0$: observing e alone does not improve the optimal success probability beyond $1/2$. After z is known, observing e raises the success probability from $1/2$ to 1, giving

$$V_u(e \mid z) = \frac{1}{2}.$$

Thus delayed value also appears directly at the level of downstream decision quality.

I.3 A Future-Access Lower Bound for a Compressed Persistent State

Proposition 1 shows that seemingly uninformative evidence can become useful later. We next ask how much information a persistent state must retain if the learner does not yet know which part of the past will matter.

Let

$$X = (X_1, \dots, X_n), \qquad X_i \overset{\text{i.i.d.}}{\sim} \text{Bernoulli}(1/2),$$

represent n independent pieces of early evidence. The learner compresses X into a persistent state S, possibly stochastically, with entropy bounded by

$$H(S) \le b.$$

Only later, a query index $I \sim \text{Unif}\{1, \dots, n\}$ is revealed, independently of (X, S). The learner must recover the requested item X_I from (S, I).

Theorem 2 (Future-access lower bound). If there exists a decoder $\hat{X}_I = g(S, I)$ whose average error probability satisfies

$$\Pr[\hat{X}_I \neq X_I] \le \epsilon, \qquad 0 \le \epsilon < \frac{1}{2},$$

then

$$\boxed{b \ge n\,[1 - h_2(\epsilon)]}$$

where

$$h_2(p) = -p \log_2 p - (1-p) \log_2 (1-p)$$

is the binary entropy function.

Proof. For each coordinate i, let $e_i^\star$ be the error probability of the Bayes-optimal decoder of X_i from S. Because a decoder with average error at most ϵ exists,

$$\bar{e}^\star = \frac{1}{n} \sum_{i=1}^{n} e_i^\star \le \epsilon.$$

For a binary variable, Fano's inequality gives

$$H(X_i \mid S) \le h_2(e_i^\star).$$

Using subadditivity of conditional entropy and concavity of h_2,

$$\begin{aligned} H(X \mid S) &\le \sum_{i=1}^{n} H(X_i \mid S) \\ &\le \sum_{i=1}^{n} h_2(e_i^\star) \\ &\le n h_2(\bar{e}^\star) \\ &\le n h_2(\epsilon), \end{aligned}$$

where the last step uses that h_2 is increasing on $[0, 1/2]$. Since $H(X) = n$,

$$I(X; S) = H(X) - H(X \mid S) \ge n\,[1 - h_2(\epsilon)]\,.$$

Finally,

$$I(X; S) \le H(S) \le b,$$

which proves the result. □

Interpretation. The theorem is a future-query statement. If the learner knows in advance which evidence will matter, it can retain only that evidence. The difficulty arises when relevance is revealed *after* compression. For example, achieving $\epsilon = 0.1$ requires at least about $0.531n$ bits of persistent information, while driving ϵ close to zero requires retaining almost all n bits. The exact constants are less important than the qualitative implication: strong, irreversible compression before future relevance is known necessarily sacrifices some future-access capability.

This result does not assume that the persistent state is literally a binary string; b bounds its information content through entropy. Nor does it assert that every natural-language episode contains independent information. Real trajectories are highly redundant, which is precisely why consolidation and archiving are useful. The theorem isolates the opposite failure mode: when unresolved evidence contains multiple potentially independent future-relevant factors, no extremely small state can preserve all of them without loss.

I.4 State Aliasing When Different Histories Collapse to the Same Skill

The same limitation can be stated without an explicit capacity bound.

Proposition 3 (History aliasing). Let c map a learning history H to a persistent Skill $S = c(H)$. Suppose there exist histories H_1 and H_2 such that

$$c(H_1) = c(H_2) = S,$$

but the unique optimal subsequent actions differ:

$$a^\star(H_1) \neq a^\star(H_2).$$

Then no policy that conditions only on S can be optimal for both histories.

Proof. Because both histories produce the same state S, any policy $\pi(a \mid S)$ must induce the same action distribution under H_1 and H_2. Optimality under H_1 requires assigning probability one to $a^\star(H_1)$, while optimality under H_2 requires assigning probability one to the distinct action $a^\star(H_2)$. Both conditions cannot hold simultaneously. □

A symmetric special case makes the cost explicit. Suppose H_1 and H_2 occur with equal probability, there are two actions a_1, a_2, and reward is one exactly when a_i is chosen under H_i. A learner with access to the full history obtains reward 1, whereas any learner that observes only the aliased Skill S obtains expected reward at most $1/2$.

This is the formal version of a simple failure mode in continual Skill evolution: two histories can produce the same current executable instructions while retaining different unresolved evidence, hypotheses, or failure causes. If the current Skill is the only persistent state, that distinction is gone. Best-so-far selection or rollback can protect current task performance, but they cannot reconstruct information already removed by the compression map c.

I.5 Why the Episodic/Semantic Separation Addresses the Identified Failure Mode

The theory above suggests a division of labor rather than indiscriminate storage.

Episodic memory preserves option value. When the future relevance of an observation is unresolved, retaining its source episode, tool feedback, outcome, and provenance keeps open the possibility that later knowledge will make it interpretable. Proposition 1 formalizes why an observation with low current utility is not necessarily permanently irrelevant.

Semantic memory performs reversible consolidation. Once multiple episodes support a stable regularity, the system can consolidate that regularity into K and reduce redundant episodic detail. Theorem 2 does not oppose compression; it shows why the timing and target of compression matter. Compression is safest when the system has evidence that omitted details are redundant for the range of future queries it aims to support.

Memory reorganization changes accessibility, not evidence. Reorganization can use new K to retrieve and reinterpret old E, but it does not create external evidence. In the XOR example, the old bit e must still have been retained somewhere; once lost, later knowledge z cannot reconstruct it. This separates the role of replay/reorganization from hallucinated new evidence.

Skills remain compact execution interfaces. The lower bound applies to the persistent state that must preserve future-access options, not to every deployment artifact. SkillBrain can therefore keep deployment Skills concise while allowing the underlying Memory Substrate to retain richer evidence and consolidated knowledge. The theory motivates decoupling *what must persist for future learning* from *what must be presented to the current Executor now*.

I.6 Testable Predictions

The analysis yields several empirical predictions that can be checked without estimating mutual information from model outputs.

Prediction 1: replay value can rise after a knowledge update. Take a fixed set of old episodes and measure the performance gain obtained by replaying them before and after the learner acquires a complementary piece of knowledge. The delayed-value account predicts that the same episodes can produce a larger downstream gain after the knowledge update.

Prediction 2: immediate-utility pruning should be most harmful under delayed relevance. Construct tasks in which early evidence becomes interpretable only after later context is revealed. Under the same memory budget, a system that removes evidence solely because it lacks immediate utility should lose more future performance than a system that preserves unresolved evidence with provenance.

Prediction 3: identical current Skills can hide different learning states. Generate or select pairs of runs whose current deployment Skills are equivalent while their histories differ. If the histories contain different unresolved evidence, the best next exploration action should sometimes differ despite the same current Skill, directly exposing state aliasing.

Prediction 4: broader future-query uncertainty increases the value of episodic retention. When the space of possible future queries expands while the persistent-state budget remains fixed, Theorem 2 predicts a sharper trade-off between compression and future-access accuracy.

I.7 Scope and Limitations

These results justify a representational principle, not a universal performance guarantee for SkillBrain. First, real trajectories are structured and redundant rather than independent bits, so Theorem 2 should not be read as a literal byte requirement for natural-language memory. Second, a Skill-centric system that separately preserves raw trajectories is no longer subject to the assumption that the Skill is the only persistent state; such a system has already introduced an additional memory substrate. Third, retaining all historical data forever is unnecessary and often inefficient. The practical problem is to preserve unresolved, provenance-bearing evidence long enough for later knowledge to determine whether it can be safely consolidated or archived.

Within these limits, the theory isolates the central distinction used by SkillBrain: a compact executable Skill can be sufficient for *current action* while being insufficient for *future learning*. Continual evolution therefore benefits from making Memory, rather than the current Skill text, the persistent substrate across rounds.