

EvoAgentBench: Benchmarking Agent Self-Evolution via Ability Transfer

Xingze Gao^{1,2} Chuanrui Hu^{2†} Hongda Chen^{2,3} Pengfei Yao² Zhao Wang² Yi Bai²
Zhengwei Wu² Yunyun Han² Xiaofeng Cong³ Jie Gui³ Yafeng Deng^{2*} Teng Li^{1*}

¹Anhui University ²EverMind, Shanda Group ³Southeast University
{xingze.gao, chuanrui.hu, hongda.chen, yaopengfei}@shanda.com
{zhao.wang, baiyi, hanyunyun, dengyafeng}@shanda.com
zhengwei.wu@evermind.ai tenglwy@gmail.com
cxf_svip@163.com guijie@ustc.edu

Abstract

Agent self-evolution in long-horizon LLM systems is largely procedural: useful experience is not merely stored information, but reusable procedures for searching, debugging, and verification. Yet current evaluations do not isolate this form of transfer. Agent benchmarks test single-episode task solving; memory benchmarks target information retention rather than procedural reuse. We introduce **EvoAgentBench**, a benchmark for agent self-evolution with Ability-guided transfer evaluated separately in four agentic domains: web research, algorithmic reasoning, software engineering, and knowledge work. EvoAgentBench extracts trace-grounded Abilities from agent executions, canonicalizes them into operational units, and builds domain-specific Ability Graphs linking tasks that share procedural overlap. By design, every test task is backed by verified training-side Ability support. Across a 528/267 train/test split, two scaffolds, and three backbones, curated Ability content yields a positive aggregate effect and positive gains in every evaluated configuration, while gains from current automatic methods remain inconsistent across configurations. EvoAgentBench shifts self-evolution evaluation from aggregate accuracy comparison to fine-grained diagnosis of experience encoding, routing, exposure, and uptake. The benchmark is publicly available at <https://huggingface.co/datasets/EverMind-AI/EvoAgentBench>.

1 Introduction

A useful agent should not solve every task from scratch. After prior executions, it should retain reusable procedures for searching, debugging, and verification. As LLM agents increasingly master isolated long-horizon tasks (Patwardhan et al., 2026; Chen et al., 2026; Jain et al., 2025; Jimenez et al., 2024), the evaluation frontier is shifting to-

ward exactly this capability: improvement from experience.

We study skill-based agent self-evolution, where past trajectories are distilled into external artifacts such as skills, cases, workflows, or evolved prompts, and made available when solving future tasks. This paradigm offers a lightweight, model-agnostic path to improving agent reliability (Shinn et al., 2023; Zhao et al., 2024; Wang et al., 2025, 2024; Fernando et al., 2024).

Current evaluations fail to isolate trajectory-to-procedure transfer. Task-centric benchmarks measure whether an agent can solve a held-out task, but they never test whether prior experience converts into reusable procedural knowledge (Patwardhan et al., 2026; Chen et al., 2026; Jain et al., 2025; Jimenez et al., 2024; Mialon et al., 2024). Long-context and memory benchmarks evaluate recall, retrieval, and personalization over stored information, but stop at information retention, never reaching procedural reuse (Bai et al., 2024; Maharana et al., 2024; Wu et al., 2025).

Benchmarks for self-evolving agents come closer but still fall short. These evaluate how skill libraries grow over sequential tasks, yet they entangle artifact quality with task order, retrieval policy, and scaffold dynamics (Cai et al., 2025; Wu et al., 2024; Jiang et al., 2026). Skill-oriented benchmarks probe procedure quality, but evaluation operates within individual tasks rather than measuring cross-task transfer (Li et al., 2026; Zhong et al., 2026). The result is a shared blind spot: **no current benchmark explicitly controls for ability-level support while measuring whether reusable procedures transfer to unseen tasks with related but non-identical demands.**

Evaluating trajectory-to-procedure transfer introduces a tension absent from ordinary held-out evaluation. Evaluation tasks must be related enough to prior experience that transferable procedures exist, yet distinct enough that improvement cannot

*Co-corresponding authors.

†Project leader.

stem from memorized answers or near-duplicate exposure. *How can a benchmark guarantee both?* We resolve this with a single structural principle: a benchmark for agent self-evolution must be ability-supported yet instance-disjoint.

We introduce **EvoAgentBench**, a multi-domain benchmark built around this principle. It provides a training split on which self-evolution methods collect experience and construct reusable artifacts, together with held-out tasks requiring related capabilities without duplicating instances. The benchmark spans four long-horizon agentic domains: web research, algorithmic reasoning, software engineering, and knowledge work. Rather than using benchmark or domain labels as transfer units, EvoAgentBench extracts trace-grounded Abilities from agent executions to build domain-specific Ability Graphs whose splits measure procedure-level transfer separately within each domain rather than random task separation. Under matched external conditions, automatic methods differ only in how they encode and reuse training experience; the benchmark’s own Ability-grounded skills serve as a diagnostic reference with curator-side routing.

Using EvoAgentBench, we expose a substantial gap between high-quality procedural content and current automatic self-evolution methods. The Ability-grounded reference condition improves held-out performance in each of the four domain-specific evaluations, but automatic methods remain brittle, with gains varying across settings. We also identify instances of negative transfer, where injected artifacts hurt task performance. Finally, interaction cost is an unreliable proxy for accuracy, motivating joint reporting of performance and overhead.

In summary, this paper makes three contributions:

- **Evaluation framework.** We formulate agent self-evolution as a controlled experience-to-transfer evaluation problem and provide a matched evaluation substrate so that methods differ only in how they encode and reuse training experience.
- **Benchmark.** We introduce EvoAgentBench, a multi-domain benchmark whose trace-grounded Ability Graphs enable separate procedure-level transfer measurement within each of four agentic domains.
- **Empirical findings.** Evaluation across two

scaffolds and multiple backbones shows that reusable procedural content transfers when correctly delivered, but current automatic methods remain brittle as the evaluated domain, scaffold, and cost regime change.

2 Related Work

Agent self-evolution via reusable artifacts.

Methods for non-parametric self-evolution can be organized by artifact type. Reflexion (Shinn et al., 2023) introduced verbal reinforcement via transient reflections; ExpeL (Zhao et al., 2024) generalized this across tasks into natural-language insights; Voyager (Wang et al., 2024) built executable skill libraries. AWM (Wang et al., 2025) targets workflows, while prompt-evolution systems such as Promptbreeder and GEPA (Fernando et al., 2024; Agrawal et al., 2026) target evolved prompts. Continuously updated or structured memories include Dynamic Cheatsheet (Suzgun et al., 2026), streaming reasoning pools (Ouyang et al., 2026b), Memento (Zhou et al., 2025), EverMemOS (Hu et al., 2026a), and HyperMem (Yue et al., 2026). A recent wave converges on curated skill packages that distill trajectories into structured, persistent files (Zhou et al., 2026; Ouyang et al., 2026a; Mi et al., 2026; Yang et al., 2026). Yet no method explicitly models which reusable abilities are present in training and whether they transfer. The gap is not artifact diversity but transfer awareness: without modeling which abilities are present at training time, a method’s improvement could stem from its extracted procedures, retrieval, or incidental task-level overlap. EvoAgentBench grounds transfer evaluation in trace-derived Ability units and evaluates automatic self-evolution methods under matched task, tool, and scoring conditions.

Benchmarks for self-evolving agents.

Existing benchmarks are largely streaming: LifelongAgentBench (Zheng et al., 2025), SEA-Eval (Jiang et al., 2026), StuLife (Cai et al., 2025), and EvoMemory (Wei et al., 2025) study how memory or skill libraries evolve as tasks arrive sequentially, conflating artifact quality with scheduling and retrieval dynamics. AMA-Bench (Zhao et al., 2026) and EverMemBench (Hu et al., 2026b) target in-episode agentic or multi-party conversational memory rather than cross-task procedural transfer. SkillsBench (Li et al., 2026), SkillFlow (Zhang et al., 2026), and SkillLearnBench (Zhong et al., 2026) probe skill-artifact quality: SkillsBench re-

Benchmark	Eval Target	Cross-Task	Abil.-Aware Split	#D	#Tasks
LongMemEval	Info. retention	✗	✗	1	500
AMA-Bench	Traj. memory	✗	✗	6	3,696
StreamBench	Online adapt.	✓	✗	5	9,702
LifelongAgentBench	Online adapt.	✓	✗	3	1,396
SEA-Eval	Online adapt.	✓	✗	2	90
SkillsBench	Skill quality	✗	✗	11	84
SkillLearnBench	Skill learning	✗	✗	6	100
SkillFlow	Skill evolution	✓	✗	5	166
EvoAgentBench	Proc. transfer	✓	✓	4	528 / 267*

Table 1: Comparison with related benchmarks. #D: number of domains. * Train / test split; other benchmarks report undivided evaluation pools or sequential streams.

ports that curated skills substantially outperform self-generated ones, and even SkillLearnBench’s three-level evaluation operates within individual tasks. Yet when a method fails, none can pinpoint whether the bottleneck lies in missing training-side ability, flawed extraction, or retrieval breakdown, a distinction essential to method development.

EvoAgentBench closes this gap (Table 1) with an Ability Graph that guarantees training-side Ability support for every test task, enabling diagnosable transfer evaluation.

3 EvoAgentBench

EvoAgentBench is built around an *Ability Graph*. Nodes are tasks, and edges indicate shared reusable Abilities extracted from agent executions. An Ability is not a dataset label, topic, or repository name; it is a reusable operation such as a search strategy, debugging procedure, or validation workflow. Task relatedness, data splitting, and diagnostic construction all derive from these operations rather than surface metadata. Figure 1 overviews the three-stage construction pipeline.

3.1 Benchmark Formulation

Let D_{train} and D_{test} denote the training and test splits. A self-evolution method m receives training-side evidence only—task prompts, verifier outcomes, associated metadata, and training trajectories, either the benchmark’s construction traces or rollouts the method itself performs on D_{train} (the protocol in our experiments; Section 4)—and produces an evolution state z_m encoding this experience into a reusable form such as a skill library, case bank, workflow set, or evolved prompt. The benchmark is not tied to a particular representation; future methods may instantiate z_m as lightweight

updates, provided the update derives only from training-side evidence.

During evaluation, the agent is run on D_{test} with z_m available through the method’s own interface. Let $\mathcal{D}$ denote the evaluated domains and D_{test}^d the test tasks in domain d . For a fixed scaffold–backbone configuration s , let $r_{0,s}(x)$ be the no-evolution baseline score on test task x , and $r_{m,s}(x)$ the score after applying method m . The primary observable is the domain-balanced average transfer gain:

$$\delta_{m,s}^d = \frac{1}{|D_{\text{test}}^d|} \sum_{x \in D_{\text{test}}^d} (r_{m,s}(x) - r_{0,s}(x)),$$

$$\Delta_{m,s} = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \delta_{m,s}^d.$$

Thus, taskwise gains are averaged within each domain before the domain means receive equal weight. When reporting an aggregate across a set $\mathcal{S}$ of evaluation configurations, we additionally give each configuration equal weight: $\Delta_m^{\text{overall}} = |\mathcal{S}|^{-1} \sum_{s \in \mathcal{S}} \Delta_{m,s}$. In our experiments, $|\mathcal{D}| = 4$ and $|\mathcal{S}| = 6$. We additionally report cost changes when available.

Within each agent–backbone setting, we fix the task statement, tool set, scoring contract, timeout, and base agent configuration. Methods differ only in what evolution state they construct and how that state is applied at test time, and EvoAgentBench measures whether each chosen form of self-evolution improves test-time behavior on tasks with verified train-side Ability support.

3.2 Ability-Guided Data Construction

EvoAgentBench covers four long-horizon agentic domains: web research, algorithmic reasoning, software engineering, and knowledge




Figure 1: EvoAgentBench construction pipeline. **Stage 1:** no-skill executions from multiple backbones across four domains yield a trace pool with reusable spans. **Stage 2:** raw Ability cards are merged into canonical families (Method, Guard, or Workflow) via embedding blocking, LLM adjudication, and expert review. **Stage 3:** the task–Ability graph is partitioned into communities, and constrained splitting guarantees train-side Ability support for every test task.

work, instantiated with BrowseComp-Plus, LiveCodeBench, SWE-Bench Verified, and GDPVal respectively (Chen et al., 2026; Jain et al., 2025; Jimenez et al., 2024; Patwardhan et al., 2026). Each domain favors distinct reusable procedures: search and verification strategies for web research, algorithm design and debugging patterns for algorithmic reasoning, repository-level repair procedures for software engineering, and artifact construction or validation workflows for knowledge work. Rather than treating source benchmark labels or domain metadata as transfer units, EvoAgentBench builds transfer structure from **trace-grounded Abilities**.

3.2.1 Multi-Backbone Trace Collection

Ability construction is grounded in agent behavior rather than task metadata alone. For each task, we collect executions from a set of construction backbones $\mathcal{B}_{\text{trace}}$. In our implementation, $\mathcal{B}_{\text{trace}}$ contains three construction backbones: Kimi-K2.5, GLM-5.1, and DeepSeek-V3.2 (Team et al., 2026; GLM-5-Team et al., 2026; DeepSeek-AI et al., 2025).

For each task x and backbone $b \in \mathcal{B}_{\text{trace}}$, we collect executions across construction scaffolds (Appendix A.2); $\tau_{x,b}$ denotes the resulting trace set, with $o_{x,b}$ and $\nu_{x,b}$ the corresponding outputs and verifier scores (binary, scalar, or structured depend-

ing on the source benchmark). The task-level multi-backbone trace pool is $\mathcal{T}(x) = \{(\tau_{x,b}, o_{x,b}, \nu_{x,b}) : b \in \mathcal{B}_{\text{trace}}\}$. The extractor compares successful and failed attempts across backbones, identifying recurring procedures, failure modes, and discriminative verification steps at the task level. This cross-backbone aggregation reduces sensitivity to any single model’s shortcuts, tool-use habits, or idiosyncratic failures.

3.2.2 Trace-Grounded Ability Extraction

EvoAgentBench defines transfer over *Abilities*: trace-grounded abstractions of reusable operations with explicit applicability conditions, specifying when an agent should apply a procedure, correction, or validation step.

For each task x , the extractor observes the task query q_x , the ground-truth target y_x^* , and the multi-backbone trace pool $\mathcal{T}(x)$, and produces task-local raw Ability cards $\mathcal{A}_{\text{raw}}(x) = E_\phi(q_x, y_x^*, \mathcal{T}(x))$.

Each raw Ability card is a tuple $a = (\gamma_a, \pi_a, \mathcal{E}_a, \partial_a, \rho_a)$, where γ_a is a trigger condition, π_a is the reusable procedure, and $\mathcal{E}_a$ is supporting evidence linking the card to trajectory spans, outputs, verifier results, and the ground-truth target. The boundary ∂_a restricts when the procedure should transfer, preventing broad instructions from forming spurious Ability links, and ρ_a assigns one of three roles: **Method** for primary task-solving

procedures, **Guard** for corrections against recurring invalid behavior, and **Workflow** for execution-control procedures such as verification. These roles support analysis and auditing but do not define separate evaluation metrics.

The extractor operates jointly over multi-backbone evidence for each task rather than processing individual runs independently, letting E_ϕ compare successful and failed executions, identify operations that distinguish reliable completions from brittle attempts, and abstract recurring failure patterns into corrective procedures.

3.2.3 Ability Canonicalization

The raw Ability cards produced by E_ϕ are task-local: different tasks may describe the same operation with different wording, while superficially similar cards may correspond to different procedures. EvoAgentBench therefore canonicalizes raw cards through conservative adjudication targeting operational equivalence rather than semantic clustering.

For each domain d , we use embedding similarity as a recall-oriented blocking step over the domain’s raw cards $\mathcal{A}_{\text{raw}}^d = \bigcup_{x \in D_d} \mathcal{A}_{\text{raw}}(x)$. We embed each card using its trigger, procedure, boundary, and role fields, and take as candidates all pairs whose embedding cosine similarity meets a domain-specific threshold θ_d . The threshold serves only to select pairs for adjudication; embedding similarity is never a merge criterion.

Each candidate pair is judged under an operational-equivalence rubric. A merge requires the same role type, compatible triggers, equivalent reusable procedures, the same success mechanism or correction target, and compatible applicability boundaries. Shared topic, lexical overlap, or generic verbs (“search”, “debug”, “validate”) are insufficient.

To reduce single-judge variance, each pair is evaluated by three independent LLM adjudicators; unanimously approved pairs are accepted automatically, and all others are reviewed by domain experts under the same rubric, yielding pair-level merge and cannot-link decisions.

Merge decisions are not transitively closed: $a_i \sim a_j$ and $a_j \sim a_k$ does not imply that a_i, a_j, a_k form one Ability. Accepted links define a compatibility graph, but a connected component becomes a canonical Ability unit only after a group-level consistency check confirms that every internal pair is merge-compatible and none carries a cannot-link decision. Domain experts split components that vi-

olate this condition, assigning broad bridge cards to their most specific compatible unit or downgrading them to annotation-only status.

Each accepted unit $u = (\mathcal{R}_u, \Gamma_u, \Pi_u, \mathcal{E}_u, \partial_u, \rho_u, \mathcal{X}_u)$ aggregates the member raw cards $\mathcal{R}_u$ into canonical fields (Γ_u trigger, Π_u procedure, $\mathcal{E}_u$ aggregated evidence, ∂_u boundary, ρ_u role), with supporting task set $\mathcal{X}_u = \{x : \mathcal{R}_u \cap \mathcal{A}_{\text{raw}}(x) \neq \emptyset\}$. Units whose procedure becomes too generic after merging are retained for annotation rather than used as transfer-defining Abilities. The resulting canonical units form the shared Ability vocabulary for constructing the Ability Graph.

3.2.4 Ability Graph Construction

We map canonical units back to tasks through their supporting raw cards. Let $\mathcal{U}_d$ denote the canonical Ability units in domain d ; the Abilities assigned to task x are $A(x) = \{u \in \mathcal{U}_d : x \in \mathcal{X}_u\}$.

Not every canonical unit defines a transfer edge. A unit is *edge-eligible* if at least two distinct tasks support it and its procedure remains operationally specific after merging; singleton, overly broad, and annotation-only units are retained for audit but cannot create graph edges. Let $A^+(x) \subseteq A(x)$ denote the edge-eligible Abilities of task x .

For each domain d , we retain tasks with at least one edge-eligible Ability and construct an undirected Ability Graph $G_d = (V_d, E_d)$, where $V_d = \{x \in D_d : A^+(x) \neq \emptyset\}$. Across the four domains, this filtering retains 1,128 task nodes. An edge $(x_i, x_j) \in E_d$ is added exactly when two tasks share at least one edge-eligible Ability, i.e., $A^+(x_i) \cap A^+(x_j) \neq \emptyset$. Each edge stores the shared Ability identifiers and their roles, enabling later analysis of whether a relation is induced by Method, Guard, or Workflow Abilities. Tasks without an edge-eligible Ability are excluded before graph construction.

We then partition each domain graph into Ability communities via Louvain community detection (Blondel et al., 2008): local regions of procedural overlap, not external labels, used only to localize train-side support during split construction.

3.2.5 Ability-Aware Transfer Split

We sample the supported evaluation subset from the retained graph nodes V_d within each domain by community, so that train-side support comes from the same local Ability neighborhood: every test task must share at least one edge-eligible Abil-

Domain	Eval. Tasks		Full Ability Graph		
	Train	Test	Abl.	Comm.	Abl./T
BrowseComp-Plus	154	65	41	14	2.34
SWE-Bench Verified	87	56	37	15	2.16
LiveCodeBench	182	86	61	23	2.19
GDPVal	105	60	34	7	2.57
Total	528	267	173	59	2.29

Table 2: Dataset statistics. Train/Test report the supported evaluation split used in experiments. The Ability Graph columns summarize the full retained graph before evaluation-subset sampling. Abl.: edge-eligible canonical Ability units; Comm.: Ability communities; Abl./T: average edge-eligible Ability units per retained graph task. Zero test tasks are unsupported.

ity with a training task in its own community. We perform constrained sampling under this invariant: moving a task into the test split is allowed only if it and all previously selected test tasks retain train-side support. Among feasible assignments, no-evolution verifier scores are used only as a secondary stratification signal to balance baseline difficulty, preventing the test split from being dominated by either ceiling or floor cases.

The resulting evaluation split contains 528 training tasks and 267 test tasks, with zero unsupported test tasks. The splits are instance-disjoint: $D_{\text{train}} \cap D_{\text{test}} = \emptyset$; training and test tasks may share canonical Abilities but never the same source-task instance.

3.3 Dataset Statistics and Construction Checks

Table 2 separates graph-level task counts from the supported evaluation subset used in the main experiments. The source pool contains 2,605 tasks across four domains, and trace-grounded extraction produces 7,326 raw Ability cards from 2,516 tasks. Of the adjudicated canonical Ability units, 173 satisfy the edge-eligibility criterion and induce graph edges. Retaining tasks with at least one such unit yields a full Ability Graph of 1,128 transferable task nodes. The main experiments use a supported subset of 795 graph nodes, containing 528 training tasks and 267 test tasks with zero unsupported test tasks; the remaining 333 graph nodes are not part of the evaluation split.

4 Experiments

4.1 Experimental Setup

Evaluation settings. We evaluate two agent scaffolds, **OpenClaw** (OpenClaw Contributors, 2026) and **Nanobot** (Ren and Nanobot Contributors, 2026), with three backbones: **Qwen3.5-27B**, **Qwen3.5-397B**, and **Gemma-4-31B** (Qwen Team, 2026; Gemma Team et al., 2026). The resulting six combinations of scaffold and backbone span two scaffold architectures and two model families, including two scales of Qwen3.5. Within each setting, tasks, MCP tools, timeouts, scoring contracts, and base agent configuration are fixed across methods; results are averaged over three independent runs per instance, and we report the standard error across tasks.

Evolution conditions. We compare **Vanilla**, which has no evolution state, with three automatic self-evolution methods and one diagnostic reference. **Memento** retrieves the closest retained training case (Zhou et al., 2025); **Reasoning-Bank** retrieves from a distilled reasoning-memory pool (Ouyang et al., 2026b); and **GEPA** evolves one prompt on D_{train} and broadcasts it to all test tasks, making it the only no-retrieval evolution condition (Agrawal et al., 2026). **Anchor Skill** loads Ability-grounded skills produced by our extraction pipeline and retrieved by the test task’s curator-side Ability cluster. Anchor Skill is not a deployable method because it uses Ability labels provided by curators, which are unavailable to automatic methods. Its skill content is constructed exclusively from evidence in the training split. Canonical family descriptions and skill text use only raw cards extracted from training tasks, while raw cards from test tasks contribute to neither. Test information therefore enters Anchor only through the routing label. Its construction backbones (Kimi-K2.5, GLM-5.1, DeepSeek-V3.2) are disjoint from the evaluation backbones, allowing Anchor to test whether the resulting Ability content remains useful across model families.

Protocol and metrics. All methods build their evolution state from D_{train} only, before any test evaluation begins. Each automatic method runs its own ingestion on D_{train} with the same combination of scaffold and backbone used at evaluation, with task prompts and verifier outcomes available; test answers, successful test trajectories, and test-specific shortcuts are excluded. Do-

Method	OpenClaw					Nanobot				
	Web	Algo	SWE	KW	Avg.	Web	Algo	SWE	KW	Avg.
Qwen3.5-27B										
Vanilla	10.8±3.4	51.9±5.2	42.3±5.3	43.6±4.6	37.2	8.7±2.6	52.7±5.0	45.8±5.5	43.9±4.4	37.8
+ Memento	11.3±3.0 +0.5	53.9±5.0 +2.0	45.2±5.4 +2.9	44.3±4.3 +0.7	38.7 +1.5 ±10.5	12.3±3.1 +3.6	57.0±5.2 +4.3	9.5±3.4 -36.3	47.6±4.8 +3.7	31.6 -6.2 ±8.5
+ RB	12.8±2.8 +2.0	57.4±5.1 +5.5	36.9±5.1 -5.4	48.2±4.5 +4.6	38.8 +1.7 ±6.4	9.2±3.2 +0.5	62.4±5.1 +9.7	53.0±5.6 +7.2	49.2±4.7 +5.3	43.5 +5.7 ±3.6
+ GEPA	14.9±3.6 +4.1	53.9±5.0 +2.0	41.7±5.8 -0.6	31.3±4.0 -12.3	35.5 -1.7 ±13.8	11.8±3.4 +3.1	60.5±5.0 +7.8	48.8±5.7 +3.0	46.6±4.6 +2.7	41.9 +4.2 ±13.7
+ Anchor [†]	14.9±2.9 +4.1	60.5±4.8 +8.6	52.4±5.2 +10.1	48.6±5.9 +5.0	44.1 +7.0 ±4.3	13.3±2.9 +4.6	63.2±5.1 +10.5	49.4±5.4 +3.6	57.7±5.2 +13.8	45.9 +8.1 ±7.7
Qwen3.5-397B										
Vanilla	31.8±4.6	52.7±5.2	66.7±5.7	55.8±4.9	51.8	12.8±2.9	57.8±5.1	62.5±5.8	53.4±4.4	46.6
+ Memento	33.8±4.8 +2.0	60.9±5.1 +8.2	63.7±5.3 -3.0	50.0±4.6 -5.8	52.1 +0.3 ±6.7	16.9±3.4 +4.1	63.5±5.1 +5.7	64.9±5.7 +2.4	52.2±4.9 -1.2	49.4 +2.8 ±11.5
+ RB	39.0±5.0 +7.2	59.3±5.2 +6.6	61.3±5.8 -5.4	57.3±4.5 +1.5	54.2 +2.5 ±4.5	15.9±3.4 +3.1	64.3±5.0 +6.5	63.1±5.5 +0.6	52.5±4.9 -0.9	49.0 +2.3 ±4.6
+ GEPA	32.3±4.4 +0.5	61.6±5.1 +8.9	67.9±5.6 +1.2	56.5±4.7 +0.7	54.6 +2.8 ±9.5	14.9±3.3 +2.1	65.1±4.8 +7.3	71.4±5.3 +8.9	52.2±4.5 -1.2	50.9 +4.3 ±10.0
+ Anchor [†]	45.3±3.7 +13.5	64.0±4.9 +11.3	76.8±4.9 +10.1	59.6±4.4 +3.8	61.4 +9.7 ±3.1	20.0±4.1 +7.2	69.8±4.6 +12.0	79.8±4.7 +17.3	62.0±4.4 +8.6	57.9 +11.3 ±2.9
Gemma-4-31B										
Vanilla	13.3±3.3	72.1±4.6	25.0±4.9	44.1±4.8	38.6	9.2±2.7	69.0±4.8	17.3±4.4	43.2±4.5	34.7
+ Memento	16.4±4.0 +3.1	67.8±4.8 -4.3	20.2±3.9 -4.8	33.2±4.6 -10.9	34.4 -4.2 ±9.6	10.3±3.4 +1.1	72.4±4.6 +3.4	21.4±4.6 +4.1	46.2±5.0 +3.0	37.6 +2.9 ±4.9
+ RB	13.3±3.3 +0.0	61.2±5.3 -10.9	31.5±5.1 +6.5	42.7±4.9 -1.4	37.2 -1.4 ±7.4	4.6±2.2 -4.6	72.5±4.7 +3.5	24.4±4.6 +7.1	46.5±4.7 +3.3	37.0 +2.3 ±1.2
+ GEPA	17.4±4.1 +4.1	69.8±4.6 -2.3	42.3±5.7 +17.3	44.8±4.6 +0.7	43.6 +5.0 ±3.1	9.7±2.7 +0.5	70.2±4.7 +1.2	37.5±5.5 +20.2	47.9±4.7 +4.7	41.3 +6.7 ±6.3
+ Anchor [†]	17.9±3.1 +4.6	72.5±4.6 +0.4	32.7±4.8 +7.7	54.8±4.3 +10.7	44.5 +5.8 ±3.8	12.3±3.3 +3.1	73.3±4.5 +4.3	26.8±4.8 +9.5	49.3±4.6 +6.1	40.4 +5.8 ±3.7

Table 3: **Main results on EvoAgentBench.** Per-domain accuracy (%), averaged over three runs per task; $\pm$ denotes the standard error across tasks. Colored values show Δ vs. Vanilla (positive, negative). Avg. is the equal-weighted four-domain mean; blue $\pm$ values are 95% paired-bootstrap CI half-widths for Δ . Web = web research; Algo = algorithmic reasoning; SWE = software engineering; KW = knowledge work. [†]Curator-side diagnostic reference with deterministic Ability-cluster routing.

mains use their native metrics and evaluation prompts: BrowseComp-Plus uses LLM-as-judge with the official evaluation prompt (Chen et al., 2026), SWE-Bench Verified uses hidden repository test suites and resolve rate (Jimenez et al., 2024), LiveCodeBench uses hidden test cases and pass@1 (Jain et al., 2025), and GDPVal uses reference-based LLM judging (Patwardhan et al., 2026). Within each combination of scaffold and backbone, Avg. is the equally weighted mean of the four domain scores. For uncertainty in Avg. Δ , we average the three trials per task, compute paired differences relative to Vanilla, bootstrap tasks within each domain ($B = 10,000$), and combine the four domain effects with equal weight; aggregate effects additionally average the six configurations with equal weight. Intervals for individual configurations are descriptive and unadjusted for multiple comparisons. Cost is reported as the percentage change in agent turns relative to Vanilla.

4.2 Main Results

Table 3 reports accuracy by domain and average transfer gain (Δ) across six combinations of scaffold and backbone. We focus on the consistency of transfer across evaluation configurations rather than ranking methods by a single average.

Curated Ability content transfers consistently across model families. Anchor Skill yields pos-

itive average Δ in all six configurations, with every paired confidence interval excluding zero. Its average gains by backbone are +7.5, +10.5, and +5.8 points, and all 24 domain comparisons have positive point estimates. Averaging the six configurations gives an aggregate gain of +7.9 points with a 95% confidence interval of [+5.6, +10.3]. The disjoint construction and evaluation backbones show that curated Ability content remains useful across model families.

Automatic gains vary across configurations.

Memento, GEPA, and ReasoningBank have negative point estimates in 7, 4, and 6 of the 24 domain comparisons, respectively. The paired intervals reinforce this heterogeneity, with positive evidence concentrated in selected configurations. No automatic method maintains positive transfer across all evaluation conditions, which motivates evaluation across multiple configurations rather than reliance on a single aggregate score.

The observed gap motivates a controlled routing analysis.

Anchor combines curated content construction with curator-label routing, whereas automatic methods construct and retrieve reusable artifacts from training trajectories. We next isolate test-time routing by holding the curated skill library and one-skill injection procedure fixed.

Routing	Algo ($n = 86$)		SWE ($n = 56$)	
	Hit@1	Δ	Hit@1	Δ
Curator label	—	+11.3	—	+10.1
Reranker top-1	91.9%	+7.8	89.3%	+7.3
Embedding top-1	68.6%	+5.6	62.5%	+7.2
Random top-1	0.0%	+0.0	1.8%	+3.7

Table 4: **Routing ablation on OpenClaw $\times$ Qwen3.5-397B.** The curated skill library and one-skill injection procedure are fixed across conditions. Hit@1 is agreement with the skill selected by Anchor’s curator-label routing; it is not applicable to the curator-label reference itself. Δ is the absolute performance change in percentage points relative to Vanilla, averaged over three independent trials per task.

4.3 Routing Ablation

We compare curator-label routing with embedding top-1 retrieval, a cascade of embedding retrieval and reranking, and random top-1 selection. We run the ablation on OpenClaw $\times$ Qwen3.5-397B in algorithmic reasoning (LiveCodeBench; $n = 86$) and software engineering (SWE-Bench Verified; $n = 56$), the two domains with deterministic verifiers. Automatic routing uses the test task statement as the query and skill triggers as retrieval documents; the reranking condition retrieves five candidates with Qwen3-Embedding-8B and selects one with Qwen3-Reranker-8B (Zhang et al., 2025). We count a top-1 hit when the selected skill matches the skill chosen by Anchor’s curator-label routing.

Table 4 shows an ordered pattern in the Algo point estimates: curator-label routing obtains +11.3 points, followed by reranking (+7.8), embedding top-1 (+5.6), and random selection (+0.0). Reranking also increases Hit@1 over embedding top-1 from 68.6% to 91.9%. On SWE, the corresponding point estimates are +10.1, +7.3, +7.2, and +3.7 points; reranking increases Hit@1 from 62.5% to 89.3%, while its downstream gain differs from embedding top-1 by only 0.1 points. Thus, agreement with curator-label routing has a clearer relationship with downstream performance on Algo than on SWE in this setting.

4.4 Error Analysis

Self-evolution depends on four observable stages: training experience must be *encoded* into a reusable artifact, *routed* to a relevant task, *exposed* in the agent’s working context, and *taken up* in agent behavior. We inspect artifact structure, a fixed-library routing ablation, and session-level traces to

identify failure signatures at each stage.

Encoding. Automatic methods suggest distinct granularity trade-offs. GEPA compresses the training experience into one prompt and broadcasts it to every test task, removing task-level selectivity. ReasoningBank instead distills trajectories into abstract strategies, which may omit executable details. Their contrasting behavior is consistent with a trade-off between broad coverage and task-specific procedural content.

Routing. The fixed-library ablation in Table 4 shows that embedding retrieval often disagrees with curator-label routing. Qualitative inspection of embedding-retrieval misses suggests that surface similarity can obscure procedural mismatch: the same few software-engineering skills recur across multiple retrieval misses. This pattern is also consistent with Memento’s vulnerability to lexically similar but operationally mismatched cases. Offline reranking substantially improves Hit@1, while its downstream effect remains domain-dependent.

Exposure. Because the evaluated methods expose their artifacts directly in the agent context, we additionally inspect Trace2Skill (Ni et al., 2026) as a skill-based method whose executable procedures are stored in auxiliary reference files rather than injected inline. Across 426 sessions on the same 142 Algo and SWE tasks, with three trials per task, agents did not access any reference file. Consequently, the file-resident procedures never entered the agents’ working context, revealing a gap between stored skill content and the content actually exposed during execution.

Uptake. The effect of relevant inline content can be observed in agent behavior. On *Tile Distance 2* (AtCoder ABC359 C, abc359_c) from LiveCodeBench (Jain et al., 2025), which requires mapping grid coordinates to 2×1 tiles, the curator-routed skill provides the key mapping $L = S_x - ((S_x + S_y) \% 2)$, which appears verbatim in all three passing programs. Under an unrelated retrieved skill, the generated programs contain no trace of this procedure. This case illustrates how relevant inline procedural content can be taken up when correctly routed.

4.5 Cost Analysis

Evolution does not uniformly reduce cost. Table 5 reports turn cost change relative to Vanilla on all tasks, on tasks solved after evolution, and on

Method	OpenClaw			Nanobot		
	All	Solv.	Uns.	All	Solv.	Uns.
Qwen3.5-27B						
+ Memento	+39.9	+38.6	+24.7	-2.3	+5.7	-2.4
+ RB	-5.5	-4.2	-14.7	+4.2	+4.4	+6.4
+ GEPA	-4.5	+4.4	-17.0	-7.8	+8.6	-13.4
+ Anchor [†]	-1.4	-12.4	-3.9	-9.1	-9.0	-9.4
Qwen3.5-397B						
+ Memento	-5.2	+3.1	-20.2	-8.9	-9.9	-22.7
+ RB	-11.1	-0.4	-26.0	-1.7	+1.9	+5.2
+ GEPA	+24.0	+27.5	+13.5	+15.1	+24.5	-7.4
+ Anchor [†]	-6.6	-6.1	-10.3	-3.3	+1.4	+8.3
Gemma-4-31B						
+ Memento	+13.1	+40.7	+12.2	+44.7	+12.5	-17.5
+ RB	+33.6	+31.3	+43.0	+37.6	+28.9	+40.5
+ GEPA	+44.7	+40.8	+53.9	+40.4	+10.5	-7.1
+ Anchor [†]	+6.3	+15.9	+26.0	+33.5	+13.0	-8.2

Table 5: **Turn cost change** (ΔT %) relative to Vanilla, averaged across four domains per scaffold. **Negative** (fewer turns) is desirable; **positive** (more turns) indicates overhead. Solv.=tasks solved after evolution; Uns.=tasks unsolved. All, Solv., and Uns. are per-domain means averaged with equal domain weights, so All need not lie between Solv. and Uns. [†]Diagnostic reference.

tasks remaining unsolved. Across the six combinations of scaffold and backbone, Anchor reduces overall cost in four (savings up to -9.1%); ReasoningBank reduces overall cost in three; GEPA increases cost in four of six settings, while Memento is evenly split between increases and reductions. Gemma-4-31B is the most costly backbone: Memento, GEPA, and ReasoningBank all add double-digit overhead on Gemma settings, and even Anchor incurs $+6.3\%$ to $+33.5\%$ on this backbone.

The relation between cost and accuracy depends on the method. When evolution provides a procedure well-matched to the test task, the agent can bypass early-stage exploration and follow a directed strategy, reducing turn count. This pattern is clearest for Anchor on Qwen backbones, which saves turns on solved tasks in three of the four Qwen settings (ReasoningBank in two of four). In contrast, Memento and GEPA frequently increase turns even on tasks the agent eventually solves (e.g., Memento $+38.6\%$ on solved / OpenClaw / Qwen3.5-27B), suggesting that generic or poorly

matched artifacts add overhead without shortcutting exploration. Cost overhead is not a stable function of method type or accuracy direction; reporting accuracy alone hides substantial heterogeneity in how methods reshape the agent’s behavioral distribution.

5 Conclusion

We introduced EvoAgentBench, a benchmark for agent self-evolution with guaranteed training-side Ability support within each of four domain-specific graphs. Evaluation across two scaffolds and three backbones shows a positive aggregate effect for curated Ability content, while gains from automatic methods vary across configurations and remain uncertain in aggregate. Transfer gain depends jointly on the evaluated domain, scaffold, and backbone, which makes evaluation across multiple configurations essential for reliable method comparison. The critical question is not *whether* an agent improves with experience, but *where* improvement breaks down: in experience encoding, routing, exposure, or uptake.

Limitations

We evaluate within-domain, cross-task transfer across four long-horizon text-based agentic domains; benchmark construction uses three open-source backbones. The evaluation does not cover multimodal or embodied agents, or larger or proprietary model families. Ability extraction uses a single LLM, so the resulting vocabulary may depend on the extractor; we do not evaluate alternative extractors. Each evaluation cell aggregates three independent runs over per-domain test sets of 56 to 86 tasks, limiting finer-grained per-domain analysis.

Ethical Considerations

EvoAgentBench is intended for research evaluation of text-based agents. It draws on public benchmark releases, and reuse of source-task content remains subject to the licenses and terms of the corresponding source benchmarks. The construction pipeline combines LLM-based extraction and adjudication with domain-expert review; consequently, the resulting Ability vocabulary may inherit biases or omissions from both the source benchmarks and the construction models. Benchmark performance should not be interpreted as evidence that an agent is safe for deployment in unconstrained settings.

Acknowledgments

We thank the domain experts who contributed to Ability canonicalization and group-consistency review, and the anonymous reviewers for their constructive feedback.

References

- Lakshya A. Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziemis, 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. 2026. **GEPA: Reflective prompt evolution can outperform reinforcement learning**. In *International Conference on Learning Representations*.
- Anthropic. 2026. Claude Sonnet 4.6 system card. <https://www-cdn.anthropic.com/bbd8ef16d70b7a1665f14f306ee88b53f686aa75.pdf>.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, et al. 2024. LongBench: A bilingual, multitask benchmark for long context understanding. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 3119–3137.
- Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. **Fast unfolding of communities in large networks**. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008.
- Yuxuan Cai, Yipeng Hao, Jie Zhou, Hang Yan, Zhikai Lei, Rui Zhen, Zhenhua Han, Yutao Yang, Junsong Li, Qianjun Pan, Tianyu Huai, Qin Chen, Xin Li, Kai Chen, Bo Zhang, Xipeng Qiu, and Liang He. 2025. **Building self-evolving agents via experience-driven lifelong learning: A framework and benchmark**. *arXiv preprint arXiv:2508.19005*.
- Zijian Chen, Xueguang Ma, Shengyao Zhuang, Ping Nie, Kai Zou, Sahel Sharifmoghammad, Andrew Liu, Joshua Green, Kshama Patel, Ruoxi Meng, Mingyi Su, Yanxi Li, Haoran Hong, Xinyu Shi, Xuye Liu, Hosna Oyarhoseini, Nandan Thakur, Crystina Zhang, Luyu Gao, and 2 others. 2026. **BrowseCompPlus: A fair and disentangled evaluation benchmark for deep search agents**. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 22349–22370.
- DeepSeek-AI, Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, et al. 2025. **Deepseek-v3.2: Pushing the frontier of open large language models**. *Preprint*, arXiv:2512.02556.
- Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. **The faiss library**. *Preprint*, arXiv:2401.08281.
- Chrisantha Fernando, Dylan Sunil Banarse, Henryk Michalewski, Simon Osindero, and Tim Rocktäschel. 2024. Promptbreeder: Self-referential self-improvement via prompt evolution. In *International Conference on Machine Learning*, pages 13481–13544. PMLR.
- Gemma Team et al. 2026. **Gemma 4 technical report**. *Preprint*, arXiv:2607.02770.
- GLM-5-Team, Aohan Zeng, Xin Lv, Zhenyu Hou, Zhengxiao Du, Qinkai Zheng, Bin Chen, et al. 2026. **Glm-5: from vibe coding to agentic engineering**. *Preprint*, arXiv:2602.15763.
- Chuanrui Hu, Xingze Gao, Zuyi Zhou, Dannong Xu, Yi Bai, Xintong Li, Hui Zhang, Tong Li, Chong Zhang, Lidong Bing, et al. 2026a. Evermemos: A self-organizing memory operating system for structured long-horizon reasoning. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 45836–45853.
- Chuanrui Hu, Tong Li, Xingze Gao, Hongda Chen, Yi Bai, Dannong Xu, Tianwei Lin, Xiaohong Li, Yunyun Han, Jian Pei, et al. 2026b. Evaluating long-horizon memory for multi-party collaborative dialogues. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 9082–9090.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2025. **Live-CodeBench: Holistic and contamination free evaluation of large language models for code**. In *International Conference on Learning Representations*, pages 58791–58831.
- Sihang Jiang, Lipeng Ma, Zhonghua Hong, Keyi Wang, Zhiyu Lu, Shisong Chen, Jinghao Zhang, Tianjun Pan, Weijia Zhou, Jiaqing Liang, and Yanghua Xiao. 2026. **SEA-Eval: A benchmark for evaluating self-evolving agents beyond episodic assessment**. *arXiv preprint arXiv:2604.08988*.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. 2024. SWE-bench: Can language models resolve real-world GitHub issues? In *International Conference on Learning Representations (ICLR)*.
- Jinhyuk Lee, Feiyang Chen, Sahil Dua, Daniel Cer, Madhuri Shanbhogue, Iftekhar Naim, Gustavo Hernández Ábrego, Zhe Li, Kaifeng Chen, Henrique Schechter Vera, Xiaoqi Ren, Shanfeng Zhang, Daniel Salz, Michael Boratko, Jay Han, Blair Chen, Shuo Huang, Vikram Rao, Paul Suganthan, and 28

- others. 2025. [Gemini embedding: Generalizable embeddings from gemini](#). *Preprint*, arXiv:2503.07891.
- Xiangyi Li, Yimin Liu, Wenbo Chen, Bingran You, Zonglin Di, Yifeng He, Shenghan Zheng, Kyoung Whan Choe, Jiankai Sun, Shuyi Wang, Chujun Tao, Binxu Li, Xuandong Zhao, Hejia Geng, Xiaojun Wu, Junwei Zhou, Xiaokun Chen, Hanwen Xing, Yubo Li, and 59 others. 2026. [SkillsBench: Benchmarking how well agent skills work across diverse tasks](#). *Preprint*, arXiv:2602.12670.
- Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. Evaluating very long-term conversational memory of LLM agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13851–13870.
- Qirui Mi, Zhijian Ma, Mengyue Yang, Haoxuan Li, Yisen Wang, Haifeng Zhang, and Jun Wang. 2026. [Skill-Pro: Learning reusable skills from experience via non-parametric PPO for LLM agents](#). *arXiv preprint arXiv:2602.01869*.
- Grégoire Mialon, Clémentine Fourrier, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2024. GAIA: A benchmark for general AI assistants. In *International Conference on Learning Representations*.
- Jingwei Ni, Yihao Liu, Xinpeng Liu, Yutao Sun, Mengyu Zhou, Pengyu Cheng, Dexin Wang, Erchao Zhao, Xiaoxi Jiang, and Guanjuan Jiang. 2026. [Trace2Skill: Distill trajectory-local lessons into transferable agent skills](#). *arXiv preprint arXiv:2603.25158*.
- OpenAI. 2024. GPT-4o system card. <https://cdn.openai.com/gpt-4o-system-card.pdf>.
- OpenClaw Contributors. 2026. OpenClaw: Personal AI assistant. <https://github.com/openclaw/openclaw>. GitHub repository.
- Siru Ouyang, Jun Yan, Yanfei Chen, Rujun Han, Zifeng Wang, Bhavana Dalvi Mishra, Rui Meng, Chun-Liang Li, Yizhu Jiao, Kaiwen Zha, Maohao Shen, Vishy Tirumalashetty, George Lee, Jiawei Han, Tomas Pfister, and Chen-Yu Lee. 2026a. [Skill-IO: Learning skill curation for self-evolving agents](#). *arXiv preprint arXiv:2605.06614*.
- Siru Ouyang, Jun Yan, I-Hung Hsu, Yanfei Chen, Ke Jiang, Zifeng Wang, Rujun Han, Long T. Le, Samira Daruki, Xiangru Tang, Vishy Tirumalashetty, George Lee, Mahsan Rofouei, Hangfei Lin, Jiawei Han, Chen-Yu Lee, and Tomas Pfister. 2026b. [ReasoningBank: Scaling agent self-evolving with reasoning memory](#). In *International Conference on Learning Representations*.
- Tejal Patwardhan, Rachel Dias, Elizabeth Proehl, Grace Kim, Michele Wang, Olivia Watkins, Simon Fishman, Marwan Aljubeh, Phoebe Thacker, Laurance Fauconnet, Natalie Kim, Samuel Miserendino, Gildas Chabot, David Li, Patrick Chao, Michael Sharman, Alexandra Barr, Amelia Glaese, and Jerry Tworek. 2026. [GDPval: Evaluating AI model performance on real-world economically valuable tasks](#). In *International Conference on Learning Representations*, pages 24005–24040.
- Qwen Team. 2026. [Qwen3.5: Towards native multi-modal agents](#).
- Xubin Ren and Nanobot Contributors. 2026. Nanobot: Lightweight, open-source AI agent for tools, chats, and workflows. <https://github.com/HKUDS/nanobot>. GitHub repository.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflexion: Language agents with verbal reinforcement learning](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 8634–8652.
- Mirac Suzgun, Mert Yuksekgonul, Federico Bianchi, Dan Jurafsky, and James Zou. 2026. [Dynamic cheat-sheet: 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)*, pages 7080–7106.
- Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, S. H. Cai, Yuan Cao, Y. Charles, H. S. Che, Cheng Chen, Guanduo Chen, Huarong Chen, et al. 2026. [Kimi k2.5: Visual agentic intelligence](#). *Preprint*, arXiv:2602.02276.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2024. [Voyager: An open-ended embodied agent with large language models](#). *Transactions on Machine Learning Research*.
- Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. 2025. Agent workflow memory. In *International Conference on Machine Learning*, pages 63897–63911. PMLR.
- Tianxin Wei, Noveen Sachdeva, Benjamin Coleman, Zhankui He, Yuanchen Bei, Xuying Ning, Mengting Ai, Yunzhe Li, Jingrui He, Ed H. Chi, Chi Wang, Shuo Chen, Fernando Pereira, Wang-Cheng Kang, and Derek Zhiyuan Cheng. 2025. [Evo-Memory: Benchmarking LLM agent test-time learning with self-evolving memory](#). *arXiv preprint arXiv:2511.20857*.
- Cheng-Kuang Wu, Zhi R Tam, Chieh-Yen Lin, Yun-Nung Chen, and Hung-yi Lee. 2024. StreamBench: Towards benchmarking continuous improvement of language agents. *Advances in Neural Information Processing Systems*, 37:107039–107063.
- Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, and Dong Yu. 2025. Longmemeval: Benchmarking chat assistants on long-term interactive memory. In *The Thirteenth International Conference on Learning Representations*.

- Yutao Yang, Junsong Li, Qianjun Pan, Bihao Zhan, Yuxuan Cai, Lin Du, Jie Zhou, Kai Chen, Qin Chen, Xin Li, Bo Zhang, and Liang He. 2026. [AutoSkill: Experience-driven lifelong learning via skill self-evolution](#). *Preprint*, arXiv:2603.01145.
- Juwei Yue, Chuanrui Hu, Jiawei Sheng, Zuyi Zhou, Wenyan Zhang, Tingwen Liu, Li Guo, and Yafeng Deng. 2026. Hypermem: Hypergraph memory for long-term conversations. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 35237–35254.
- Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. 2025. [Qwen3 embedding: Advancing text embedding and reranking through foundation models](#). *Preprint*, arXiv:2506.05176.
- Ziao Zhang, Kou Shi, Shiting Huang, Avery Nie, Yu Zeng, Yiming Zhao, Zhen Fang, Qishen Su, Haibo Qiu, Wei Yang, Qingnan Ren, Shun Zou, Wenxuan Huang, Lin Chen, Zehui Chen, and Feng Zhao. 2026. [SkillFlow: Benchmarking lifelong skill discovery and evolution for autonomous agents](#). *arXiv preprint arXiv:2604.17308*.
- Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. 2024. ExpeL: LLM agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19632–19642.
- Yujie Zhao, Boqin Yuan, Junbo Huang, Haocheng Yuan, Zhongming Yu, Haozhou Xu, Lanxiang Hu, Abhilash Shankarampeta, Zimeng Huang, Wentao Ni, Yuan-dong Tian, and Jishen Zhao. 2026. [AMA-Bench: Evaluating long-horizon memory for agentic applications](#). In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*.
- Junhao Zheng, Xidi Cai, Qiuqi Li, Duzhen Zhang, Zhongzhi Li, Yingying Zhang, Le Song, and Qianli Ma. 2025. [LifelongAgentBench: Evaluating LLM agents as lifelong learners](#). *arXiv preprint arXiv:2505.11942*.
- Shanshan Zhong, Yi Lu, Jingjie Ning, Yibing Wan, Lihan Feng, Yuyi Ao, Leonardo F. R. Ribeiro, Markus Dreyer, Sean Ammirati, and Chenyan Xiong. 2026. [SkillLearnBench: Benchmarking continual learning methods for agent skill generation on real-world tasks](#). *arXiv preprint arXiv:2604.20087*.
- Huichi Zhou, Yihang Chen, Siyuan Guo, Xue Yan, Kin Hei Lee, Zihan Wang, Ka Yiu Lee, Guchun Zhang, Kun Shao, Linyi Yang, and Jun Wang. 2025. [Memento: Fine-tuning LLM agents without fine-tuning LLMs](#). *arXiv preprint arXiv:2508.16153*.
- Huichi Zhou, Siyuan Guo, Anjie Liu, Zhongwei Yu, Ziqin Gong, Bowen Zhao, Zhixun Chen, Menglong Zhang, Yihang Chen, Jinsong Li, Runyu Yang, Qiangbin Liu, Xinlei Yu, Jianmin Zhou, Na Wang, Chunyang Sun, and Jun Wang. 2026. [Memento-Skills: Let agents design agents](#). *arXiv preprint arXiv:2603.18743*.

A Construction Pipeline Details

A.1 Source Task Selection

EvoAgentBench draws from the full public releases of four source benchmarks: BrowseComp-Plus (830 tasks), SWE-Bench Verified (500 instances), LiveCodeBench V6 (1,055 problems), and GDP-Val gold subset (220 tasks), yielding a source pool of 2,605 tasks. No task-level filtering is applied at this stage; all publicly available instances enter the trace collection phase.

A.2 Trace Collection

For each task, we collect no-skill executions using two agent scaffolds, OpenClaw (OpenClaw Contributors, 2026) and Nanobot (Ren and Nanobot Contributors, 2026), each run twice per construction backbone. The three construction backbones are Kimi-K2.5, GLM-5.1, and DeepSeek-V3.2 (Team et al., 2026; GLM-5-Team et al., 2026; DeepSeek-AI et al., 2025), yielding four trials per backbone per task. Each execution uses only the scaffold’s default tool set for the target domain, with a uniform timeout of 1,800 seconds. No evolution artifacts are injected during trace collection; these are strictly no-skill baselines.

A.3 Ability Extraction

The extractor E_ϕ is instantiated with Claude Sonnet 4.6 (Anthropic, 2026). For each task, E_ϕ jointly reads the task query, reference target, and all construction-backbone trajectories to produce task-local raw Ability cards. A raw card is retained only when it satisfies three conditions: (1) it is grounded in at least one trajectory span, (2) it specifies an actionable operation rather than a topic label or generic advice, and (3) its applicability boundary is narrow enough to support meaningful transfer.

Table 6 reports the raw card statistics. Extraction produces 7,326 raw Ability cards from 2,516 tasks (of the 2,605 source tasks, 89 produced no extractable cards). The per-task Ability count ranges from 1 to 8, with a median of 3.

A.4 Canonicalization

Embedding blocking. We embed each raw card using Gemini Embedding 001 (Lee et al., 2025) and construct candidate pairs within each domain using cosine similarity thresholds: $\theta = 0.85$ for BrowseComp-Plus and GDPVal, and $\theta = 0.82$ for SWE-Bench Verified and LiveCodeBench. These

Domain	Tasks	Raw	Avg.	1	2	3	4+
Web Research	814	2,486	3.05	69	221	260	264
Algo. Reasoning	993	2,554	2.57	120	430	268	175
Software Eng.	490	1,402	2.86	45	185	130	130
Knowledge Work	219	884	4.04	5	25	45	144
Total	2,516	7,326	2.91	239	861	703	713

Table 6: Raw Ability card statistics over source tasks that produce at least one raw card. Raw: total cards; Avg.: cards per task; 1–4+: tasks with exactly 1, 2, 3, or ≥ 4 cards.

thresholds serve only as a recall-oriented prefilter; embedding similarity is never a merge criterion.

LLM adjudication. Each candidate pair is independently judged by the three construction backbones (Kimi-K2.5, GLM-5.1, DeepSeek-V3.2), providing model-diverse adjudication. Each adjudicator assigns one of the following labels:

- **Accept** (merge): Same_Tactic, Same_Strategy, or Same_Diagnostic, i.e., the two cards describe operationally equivalent procedures.
- **Reject**: Related_Only (shared topic but different operation), Different (unrelated), Conflict (contradictory procedures), or Invalid (malformed card).

Pairs with unanimous positive agreement ($v = 3$) are accepted automatically. All non-unanimous pairs are reviewed by a domain expert under the same rubric (see Appendix C).

Group consistency and canonical families. Accepted merge links define a compatibility graph over raw cards. Connected components are converted into canonical Ability families only after a group-level consistency check: all internal pairs must be merge-compatible, and no internal pair may carry a cannot-link decision. Components that violate this condition are split by domain experts. Broad bridge cards are assigned to their most specific compatible family or downgraded to annotation-only status.

A.5 Graph Construction

Table 7 reports the full retained Ability Graph statistics per domain. A canonical Ability unit is edge-eligible if it is supported by at least two distinct tasks and its procedure remains operationally specific after merging. Retaining tasks with at least

Domain	Nodes	Edges	Dens.	Deg.
Web Research	337	12,318	.218	73.1
Algo. Reasoning	402	7,593	.094	37.8
Software Eng.	213	2,864	.127	26.9
Knowledge Work	176	2,217	.144	25.2
Total	1,128	24,992	.039	44.3

Table 7: Ability Graph statistics. Dens.: density per domain subgraph; Deg.: average degree.

one edge-eligible unit yields 1,128 graph nodes, and graph edges are induced by 173 such units. Of these nodes, 795 form the supported evaluation subset reported in Table 2 (528 train and 267 test); the remaining 333 are not part of the evaluation split.

BrowseComp-Plus exhibits the highest graph density (0.218), reflecting extensive procedural overlap among information-seeking tasks (e.g., shared search and verification strategies).

A.6 Split Algorithm

The evaluation split is constructed independently per domain after Ability extraction and canonicalization. The inputs are the retained graph nodes V_d , their community assignments $c(t)$, the unit assignments $F(t)$ per task, and the no-evolution baseline reward for each task. The split is a supported evaluation subset of the retained graph rather than a partition of every graph node.

Procedure.

1. **Initialize the candidate pool.** Use the retained graph nodes V_d , each of which has at least one edge-eligible canonical Ability unit.
2. **Balance baseline difficulty.** Stratify feasible candidates by no-evolution verifier score and sample across baseline-performance ranges. Secondary objectives include preserving Ability-unit coverage and avoiding test-only units.
3. **Select test tasks under support constraints.** A task t enters the test split only if at least one unit in $F(t)$ retains train-side support within the same community: $\forall t \in D_{\text{test}}, \exists f \in F(t)$ such that $|\{t' \in D_{\text{train}} : c(t') = c(t) \wedge f \in F(t')\}| > 0$.
4. **Validate unit-level support.** Every Ability unit that appears in test must also appear in train.

Domain	Test	Supp@10	Zero
Web Research	65	9.72	0
Algo. Reasoning	86	7.17	0
Software Eng.	56	7.11	0
Knowledge Work	60	8.83	0
Total	267	8.15	0

Table 8: Train-side support per test task. Supp@10: average supporting training tasks (capped at 10). Zero: number of unsupported test tasks.

5. **Select skill-train evidence.** For each test-relevant unit, select train-side evidence tasks for skill construction, prioritizing failure evidence, success contrast, informative traces, and task diversity within the unit.

Guarantees. The resulting split satisfies three invariants: (1) every test task shares at least one edge-eligible Ability unit with a training task in the same community, (2) every test-side Ability unit has train-side task support, and (3) every test-side Ability unit has selected train evidence for skill construction.

Support distribution. Table 8 reports the train-side support available to each test task. Supp@10 denotes the effective support count capped at 10, matching the retrieval budget used by skill- and case-based evolution methods.

B Information Access and Leakage Control

Table 9 specifies the information available to each party in the EvoAgentBench protocol. The curator uses all task data (including test answers, test traces, and Ability labels) during benchmark construction; this information is never exposed to self-evolution methods or the test-time agent.

Automatic self-evolution methods receive training prompts and training verifier outcomes, and generate their own training trajectories by rolling out on D_{train} with the same scaffold-backbone pair used at evaluation. They do not receive any test-side information during evolution-state construction. At evaluation time, methods receive only the test prompt and apply their evolution state through their own routing mechanism.

The Anchor Skill diagnostic reference uses curator-side test Ability labels to retrieve skills from the training skill library. It does not access test answers or test traces, and skill content is con-

	Curator	Auto.	Anchor [†]	Agent
Train prompts	✓	✓	✓	—
Train traces	✓	✓	✓	via z_m
Train verifier	✓	✓	✓	—
Train Ability labels	✓	—	✓	—
Test prompts	✓	at eval	at eval	at eval
Test answers	✓	—	—	—
Test traces	✓	—	—	—
Test raw cards	✓	—	—	—
Test Ability labels	✓	—	✓	—

Table 9: Information access by party. ✓: full access; —: no access; at eval: available only during test execution; via z_m : through the method’s evolution state. Automatic methods’ train traces are their own rollouts on D_{train} with the evaluation scaffold–backbone; Anchor skill content derives from train-side raw cards only. [†]Diagnostic reference.

structed exclusively from train-side evidence: raw cards extracted from test tasks contribute neither to canonical family text nor to skill content. Because it uses curator-side labels for retrieval, it is a diagnostic reference rather than a deployable method.

C Human Review Protocol

Canonicalization quality depends on the expert review of non-unanimous adjudication pairs. We employ four domain experts, one per domain, each with research or professional experience in their respective area: web research for BrowseComp-Plus, software engineering for SWE-Bench Verified, algorithmic reasoning for LiveCodeBench, and knowledge work for GDPVal.

Workflow. For each non-unanimous pair ($v \neq 3$), the assigned domain expert reviews the two raw Ability cards alongside their trajectory evidence, following the same operational-equivalence rubric used by the LLM adjudicators. The expert assigns a final merge or reject decision. Additionally, experts review all group-level consistency violations: when transitive closure produces components with incompatible internal pairs, the expert splits the component and reassigns broad bridge cards.

Rubric. A merge requires:

1. Same role type (Method, Guard, or Workflow).
2. Compatible trigger conditions.
3. Equivalent reusable procedures (same mechanism, not merely same topic).

4. Same success mechanism or correction target.

5. Compatible applicability boundaries.

Shared topic, lexical overlap, or generic verbs (“search”, “debug”, “validate”) are explicitly insufficient for a merge.

D Ability Examples

We illustrate the Ability concept with a concrete example from the algorithmic reasoning domain. Figure 2 shows the full extraction flow: from the source task and baseline failure to the extracted raw Ability card.

E Evaluation and Method Details

E.1 Domain-Specific Evaluation

Each domain uses its native evaluation protocol. We adopt the official evaluation prompts and harnesses without modification. For the two judge-based benchmarks, BrowseComp-Plus and GDPVal, we use GPT-4o (OpenAI, 2024) as the fixed judge across all scaffold–backbone settings.

Web Research (BrowseComp-Plus). The agent’s answer is compared against the ground-truth reference using an LLM-as-judge with the fixed evaluation prompt provided by the benchmark (Chen et al., 2026). The judge outputs a binary correct/incorrect decision. We use the same prompt template with GPT-4o.

Algorithmic Reasoning (LiveCodeBench). Each generated program is executed against the benchmark’s hidden test cases (averaging ~ 17 per problem) in an isolated sandbox. The metric is pass@1: a problem is solved only if all test cases pass. No LLM judge is involved; evaluation is purely programmatic (Jain et al., 2025).

Software Engineering (SWE-Bench Verified). The candidate patch is applied to the target repository and evaluated by running the instance’s designated FAIL_TO_PASS and PASS_TO_PASS test suites. An instance is resolved only when all tests pass. No LLM judge is involved; evaluation is purely programmatic (Jimenez et al., 2024).

Knowledge Work (GDPVal). The model’s deliverable is compared against a human expert’s reference using LLM-as-judge with task-specific evaluation prompts from the benchmark. The judge assigns a score based on pairwise comparison (model

Domain: Algorithmic Reasoning		Task: String Transformation (Cyclic Shift Counting)	
Task description. Count how many ways string s can be transformed into string t after exactly k cyclic suffix-prepend operations. Input length up to 5×10^5 .			
Baseline failure. The agent correctly identifies that the operation is a cyclic shift, but enumerates shifts by slicing the doubled string at every position: <pre>for p in range(n): if (s+s)[p:p+n] == t: ...</pre> This creates an $O(n)$ substring per shift, so the shift-counting phase is $O(n^2)$. Public examples pass; a large hidden test times out.			
Extracted Raw Ability Card			
<i>Name</i>	Linear-Time Substring Matching for Cyclic Shift Counting		
<i>Trigger</i>	Use when a task requires finding all cyclic rotations or shift positions of a long string, especially under large-length constraints.		
<i>Procedure</i>	(1) Construct $S = s + s$. (2) Treat t as the pattern. (3) Run a linear-time matching algorithm (KMP, Z-algorithm, or rolling hash) over S . (4) Keep only positions $p < n$. (5) Split counts by $p = 0$ vs. $p \neq 0$ if needed. (6) Sanity-test on large n .		
<i>Boundary</i>	Must not contain an $O(n)$ -length substring extraction inside a loop over all shifts.		
<i>Role</i>	Guard (corrects a recurring $O(n^2)$ anti-pattern).		
Canonical Ability Family: <i>String Pattern Matching for Large-Scale Cyclic Transformations</i>			
The reusable intervention is not a memorized answer but a procedural correction: when cyclic string matching is large-scale, replace repeated slicing with a linear pattern-matching algorithm. This family links multiple tasks that share the same anti-pattern and the same fix.			

Figure 2: Ability extraction example from the algorithmic reasoning domain. Top: source task and baseline failure. Middle: extracted raw Ability card with trigger, procedure, boundary, and role. Bottom: the canonical Ability family this card belongs to.

vs. expert reference). We use the official evaluation prompts with GPT-4o (Patwardhan et al., 2026).

E.2 Agent Scaffold Configuration

We evaluate on two scaffolds of different scale and design philosophy:

- **OpenClaw** (OpenClaw Contributors, 2026): a full-featured, modular agent framework with 26 built-in tools spanning filesystem operations, shell execution, web search and fetching, browser automation, memory management, and session control.
- **Nanobot** (Ren and Nanobot Contributors, 2026): an ultra-lightweight agent framework ($\sim 4K$ lines of Python) with 7 core tools covering filesystem, shell execution, web search, and communication.

Both scaffolds are MCP-native and support extensibility through external tool servers. In all experiments, each scaffold’s full default tool set is enabled across all four domains. The one domain-specific addition is BrowseComp-Plus, where we attach an MCP search server that exposes a FAISS (Douze et al., 2024) index over the official BrowseComp-Plus corpus (embedded with

Qwen3-Embedding-8B (Zhang et al., 2025)). The agent queries this server via a search tool that returns top- k snippets by embedding similarity. This tool replaces open-web access: `web_search` and `web_fetch` are disabled for this domain so that all evidence comes from the controlled corpus. Other domains use only the scaffold’s built-in tools with no additional servers.