

Group-Evolving Agents: Open-Ended Self-Improvement via Experience Sharing

Zhaotian Weng¹ Antonis Antoniadis¹ Deepak Nathani¹ Zhen Zhang¹
Xiao Pu¹ Xin Eric Wang¹

¹University of California, Santa Barbara

Abstract

Open-ended self-improving agents can autonomously modify their own structural designs to advance their capabilities and overcome the limits of pre-defined architectures, thus reducing reliance on human intervention. We introduce **Group-Evolving Agents (GEA)**, a new paradigm for open-ended self-improvements, which treats a group of agents as the fundamental evolutionary unit, enabling explicit experience sharing and reuse within the group throughout evolution. Unlike existing open-ended self-evolving paradigms that adopt tree-structured evolution, GEA overcomes the limitation of inefficient utilization of exploratory diversity caused by isolated evolutionary branches. We evaluate GEA on challenging coding benchmarks, where it significantly outperforms state-of-the-art self-evolving methods (71.0% vs. 56.7% on SWE-bench Verified, 88.3% vs. 68.3% on Polyglot) and matches or exceeds top human-designed agent frameworks (71.8% and 52.0% on two benchmarks, respectively). Analysis reveals that GEA more effectively converts early-stage exploratory diversity into sustained, long-term progress, achieving stronger performance under the same number of evolved agents. Furthermore, GEA exhibits consistent transferability across different coding models and greater robustness, fixing framework-level bugs in 1.4 iterations on average, versus 5 for self-evolving methods.

Code:  [UCSB-AI/GEA](https://github.com/UCSB-AI/GEA)




Figure 1: **Overview of Group-Evolving Agents (GEA) vs. tree-structured self-evolution for open-endedness.** GEA treats a *group of agents*, rather than an individual agent, as the fundamental unit of evolution. At each iteration, a parent group jointly gives rise to an offspring group through explicit intra-group experience sharing and reuse.

1 Introduction

Open-endedness and cumulative progress are key characteristics of scientific breakthroughs (Stanley et al., 2017; Zhang et al., 2025a; Stanley & Lehman, 2015). However, most existing AI systems rely on pre-defined model architectures designed by humans. Although such systems can accumulate experience through training, they often struggle to

Correspondence: zhaotian@ucsb.edu, ericxwang@ucsb.edu

transcend the capability boundaries imposed by their initial designs, as they lack the ability to modify their own structural configurations (Yin et al., 2025). Thus, progress remains heavily dependent on continuous human intervention.

Existing open-ended self-improving systems are always designed around individual-centric evolutionary processes (Zhang et al., 2025a; Yin et al., 2025; Schmidhuber, 2003; Robeyns et al., 2025; Novikov et al., 2025), many of which are largely inspired by biological evolution. At each iteration, a single agent is selected as the parent and refined to produce one or more offspring (Figure 1a). The overall structure follows chain- or tree-structured evolution, where different branches remain strictly isolated (e.g., agents in the green subtree never acquire experience from those in the purple subtree in Figure 1a). Although such systems often exhibit substantial exploratory diversity, this diversity rarely serves as effective stepping stones (Mouret & Clune, 2015; Pugh et al., 2016). Instead, many agents provide only temporary diversity, producing short-lived variants that fail to contribute to long-term cumulative progress.

It is time to rethink agent evolution. *AI agents are not biological individuals; why should their evolution remain constrained by biological paradigms or proceed at the individual level?* In fact, AI agents can directly share trajectories, tools, and learned artifacts, and they can aggregate complementary skills without the constraints of reproduction or lineage.

Therefore, we introduce **Group-Evolving Agents (GEA)**, a new paradigm for open-ended self-improvement that treats a *group of agents*, rather than an individual agent, as the fundamental unit of evolution (Figure 1b). This shift enables explicit experience sharing and reuse across agents within a group, naturally allowing exploratory discoveries from different agents to be consolidated and accumulated into long-term progress rather than remaining as short-lived variants. At each iteration, GEA first selects a parent group of agents using a Performance-Novelty criterion that balances immediate performance gains with evolutionary diversity. The parent agents then jointly produce a child group through a shared pool of aggregated experience from all members.

We evaluate GEA on challenging coding benchmarks, achieving success rates of 71.0% on SWE-bench Verified and 88.3% on Polyglot, **significantly outperforming state-of-the-art open-ended self-evolving methods** (56.7% and 68.3%, respectively). Analysis reveals that GEA more effectively consolidates the diversity generated during open-ended exploration, yielding sustained progress and stronger performance given the same number of evolved agents. By leveraging experience from better-performing agents, GEA also exhibits stronger robustness to framework-level perturbations. Furthermore, its improvements stem from workflow and tool enhancements rather than model-specific optimizations, thus transferring consistently across GPT- and Claude-series models.

Additionally, by leveraging meta-learning for self-improvement in open-ended exploration, without any human intervention, GEA achieves performance comparable to or even surpassing human-designed state-of-the-art frameworks on both benchmarks (71.0% vs. 71.8% on SWE-bench Verified, 88.3% vs. 52.0% on Polyglot).

In summary, we propose **Group-Evolving Agents**, a new paradigm for open-ended self-improvement that:

1. Overcomes the limitation of inefficient utilization of exploratory diversity caused by branch isolation in existing tree-structured evolution, by enabling explicit experience sharing and reuse within the group during evolution.
2. More effectively consolidates and reuses experience and evolutionary diversity from other agents, achieving significant performance gains and stronger robustness over state-of-the-art open-ended self-evolving methods, with improvements that transfer consistently across different coding models.
3. Matches or surpasses human-designed state-of-the-art frameworks through meta-learning-based self-improvement without human intervention.

2 Related Work

Recent years have witnessed growing interest in how AI systems can continuously improve themselves without human intervention (Fang et al., 2025; Gao et al., 2025; Wang et al., 2025b). Most existing self-improving approaches mainly focus on continuous, iterative refinement of the given agent system (Shinn et al., 2023; Song et al.; Gao et al., 2025; Zelikman et al., 2024; Singh et al.; He et al., 2025), typically evolving toward a specific optimization objective and following a linear, chain-based evolutionary structure (Zhang et al., 2025b; Wang et al., 2025b; Robeyns et al., 2025). Such systems achieve self-improvement through mechanisms such as self-play against historical versions or self-generated verification (Setlur et al., 2024; Wang et al., 2024; Huang et al., 2025; Wang et al., 2025a; Wei et al., 2025), supervised fine-tuning (Zelikman et al., 2022; Gou et al., 2023; Ni et al., 2024) or reinforcement learning on selectively filtered feedback (Xin et al., 2024; Zhao et al., 2025; Yuan et al., 2024), and reflection-based methods (Wang et al., 2025c; Yin et al., 2025; Shinn et al., 2023; Rosser & Foerster, 2026; Lin et al., 2026; Liu et al., 2026) or in-context learning (Hua et al., 2024; Sun et al., 2023). While this goal-oriented, chain-based evolutionary paradigm enables autonomous improvement along a particular direction, it inherently limits the ability of self-evolving systems to explore diverse evolutionary directions in open-ended solution spaces.

A line of work has pointed out that one of the key challenges in enabling unbounded improvement and innovation lies in developing open-ended AI systems that can continuously produce both novel and learnable artifacts (Stanley et al., 2017; Zhang et al., 2025a; Stanley & Lehman, 2015; Shapere, 1964). Building on this insight, open-endedness has been characterized as the capability of systems to continuously generate artifacts that are novel, interesting, and learnable from a human perspective (Zhang et al., 2025a; Clune, 2019; Hughes et al., 2024; Jiang et al., 2023; Zhang et al.; Faldor et al.).

Motivated by the potential of enabling unbounded evolution through open-ended exploration in self-evolving agents, more recent studies adopt lineage-based, tree-structured evolutionary strategies (Zhang et al., 2025a; Novikov et al., 2025; Fernando et al., 2024; Huang et al., 2023; Antoniadis et al.). In these frameworks, individual parent agents are selected at each iteration to independently produce offspring, enabling various branching exploration across multiple evolutionary directions and helping avoid local optima. However, the strict isolation between evolutionary branches prevents effective information and experience sharing and reuse across lineages. As a result, many promising directions discovered early in evolution persist only as temporary diversity and fail to contribute to long-term cumulative progress. To overcome this limitation, we introduce a group-centric evolutionary paradigm, **Group-Evolving Agents (GEA)**, which explicitly enables intra-group experience sharing and reuse throughout the evolutionary process. By consolidating complementary discoveries across agents, GEA more effectively leverages the diversity generated by open-ended exploration to support sustained cumulative progress.

3 Method

We propose **Group-Evolving Agents (GEA)**, a framework that treats a *group of agents* as the fundamental unit of evolution. GEA maintains an archive that stores all discovered agents throughout the evolutionary process. Each agent consists of three modules: a **reflection module** that derives improvement objectives from shared group-level experience, an **evolving module** that updates the agent by generating model patches and applying them to the agent itself, and an **action module** that evaluates the updated agent on downstream tasks (Appendix A.4). At each iteration, GEA proceeds in two core stages (Figure 1):

(1) Parent Group Selection (§3.1): GEA first selects K parent agents from the archive using a Performance–Novelty selection strategy (Mouret & Clune, 2015; Pugh et al., 2016; Chatzilygeroudis et al., 2021) that balances immediate task-solving competence with long-term evolutionary diversity and potential.

Algorithm 1 Parent Group Selection with KNN Novelty

```

1: Input: Archive of agents  $\mathcal{A}$ ; agent-representation vectors  $\{z_i \in \{0,1\}^D\}_{i \in \mathcal{A}}$ ; performance scores  $\{\alpha_i\}_{i \in \mathcal{A}}$ ; parent group size  $K$ ; KNN size  $M$ .
2: Output: Parent agent group  $\mathcal{G}$  with  $|\mathcal{G}| = K$ .
3: /* Compute novelty for each agent */
4: for  $i \in \mathcal{A}$  do
5:   Initialize empty list  $\mathcal{D}_i$ 
6:   for  $j \in \mathcal{A}, j \neq i$  do
7:      $d_{ij} \leftarrow 1 - \frac{z_i^\top z_j}{\|z_i\|_2 \|z_j\|_2 + \varepsilon}$ 
8:     Append  $d_{ij}$  to  $\mathcal{D}_i$ 
9:   end for
10:  Let  $\mathcal{N}_M(i)$  be the indices of the  $M$  smallest values in  $\mathcal{D}_i$ 
11:   $\text{nov}(i) \leftarrow \frac{1}{M} \sum_{j \in \mathcal{N}_M(i)} d_{ij}$ 
12: end for
13: /* Rank agents by Performance–Novelty score */
14: for  $i \in \mathcal{A}$  do
15:    $\text{score}(i) \leftarrow \alpha_i \cdot \sqrt{\text{nov}(i)}$ 
16: end for
17:  $\mathcal{G} \leftarrow$  the top- $K$  agents in  $\mathcal{A}$  ranked by  $\text{score}(\cdot)$ 
18: return  $\mathcal{G}$ 

```




Figure 2: **Detailed illustration of group-level evolution in GEA.** Aggregated evolutionary traces from the parent group are shared across all agents to generate evolution directives and framework-level patches.

(2) Open-ended Group Evolution (§3.2): The selected agents form a parent group that jointly produces an offspring group of the same size through explicit experience sharing and reuse across parent agents.

3.1 Parent Group Selection

Inspired by Mouret & Clune (2015); Pugh et al. (2016); Chatzilygeroudis et al. (2021), parent group selection in GEA balances two key principles: *performance* and *novelty*. We prioritize agents with strong task performance, as performance reflects an agent’s immediate competence and its likelihood of producing effective offspring, since evolution in GEA proceeds through iterative modifications of the agent’s implementation, which itself constitutes a form of solving coding problems. At the same time, we also encourage exploration beyond currently well-optimized regions of the search space, as agents that exhibit novel evolutionary directions may contribute to long-term cumulative progress even when their current performance is not optimal.

We represent each agent i using a task-success vector $z_i \in \{0,1\}^D$, where each dimension indicates whether the agent successfully solves a corresponding probe task. Similar binary task–response representations of this form have been widely used to characterize an agent’s coding capabilities and to better understand how these capabilities are distributed across various tasks Polo et al. (2024); Rodriguez et al. (2021). Using this representation, we measure the dissimilarity between two agents via cosine distance:

$$d(i, j) = 1 - \frac{z_i^\top z_j}{\|z_i\|_2 \|z_j\|_2 + \varepsilon}. \quad (1)$$

We define the novelty of agent i as the average cosine distance to its M most similar neighbors:

$$\text{nov}(i) = \frac{1}{M} \sum_{j \in \mathcal{N}_M(i)} d(i, j), \quad (2)$$

where $\mathcal{N}_M(i)$ denotes the set of M agents with the smallest cosine distance to agent i .

To construct the parent group, we rank agents according to a combined score

$$\text{score}(i) = \alpha_i \cdot \sqrt{\text{nov}(i)}, \quad (3)$$

where α_i denotes the performance of agent i on downstream coding tasks, and $\sqrt{\text{nov}(i)}$ moderates the influence of novelty. Finally, we select the top- K agents according to this score to form the parent group. Performance serves as the primary selection criterion, while novelty is incorporated as a mild bias without dominating performance, enabling a balanced trade-off between exploitation and exploration. The full procedure is summarized in Algorithm 1.

3.2 Open-Ended Group Evolution

Unlike conventional approaches where parent agents evolve independently without information and experience exchange, GEA explicitly enables group-level experience sharing and reuse during evolution, allowing agents to integrate complementary evolutionary directions explored by different agents while maintaining open-ended exploration. As a result, exploratory diversity is transformed from transient variations into long-term useful experience, supporting sustained evolutionary progress.

Given a selected parent group $\mathcal{G} = \{a_1, a_2, \dots, a_K\}$, GEA generates a new group $\mathcal{G}'$ of size K , where each agent evolves by leveraging both its own evolutionary history and experience aggregated from other members of the parent group (Figure 2). For each agent $a_i \in \mathcal{G}$, we collect a set of evolutionary traces consisting of: a) **the code modification patches** applied to the agent’s framework; b) **a predicted task patch** generated by a_i for a randomly sampled unsolved task during evaluation; c) the corresponding **task execution logs**, including the tool invocation history and execution workflow; d) **the evaluation outcome** of the same task, exposing failure modes and potential directions for framework-level improvement.

The aggregated traces are shared with all agents. For each agent a_i , this shared group-level experience is fed into and analyzed by its **reflection module** to generate evolution directives over agent’s workflow, tool usage, and prompting strategies. These directives are then passed to the **evolving module** to generate framework-level patches. Finally, the updated agent is evaluated on downstream tasks via the **acting module** (more implementation details in Appendix A.4). Agents that compile successfully and exhibit basic coding functionality are retained in the archive for future evolution (Figure 2). Each agent adapts its own codebase based on shared experience, enabling diverse yet coordinated evolution across the group. Applying this process to each agent in the parent group yields an off-spring group, and the evolution iterates in an open-ended manner. We summarize this process in Algorithm 2, with additional implementation details including agent initialization (Appendix A.2) and the prompts used in group evolution provided in the Appendix A.5.

4 Experiments

4.1 Benchmarks

Following the evaluation protocol established by the state-of-the-art open-ended self-evolving system, Darwin Gödel Machine (DGM) (Zhang et al., 2025a), we evaluate GEA on two structurally distinct benchmarks to assess its coding capabilities in both repository-level software engineering and multi-language code synthesis settings. To mitigate the substantial cost of evaluating every evolved agent on the full benchmarks, we adopt a staged evaluation strategy where agents must pass smaller subsets before advancing.

SWE-bench. We evaluate on **SWE-bench Verified** (OpenAI, 2024), a curated, human-validated subset of SWE-bench in which every task is confirmed to be solvable. We employ a three-stage evaluation process. First, agents undergo a sanity check on a small set of 10 tasks to discard those with framework-level failures (i.e., cannot compile or solve 0 out of 10 tasks). Agents that demonstrate basic coding functionality are then evaluated on the 50-task *Verified-Mini* set (Hobbhahn, 2025), which is designed to preserve a similar difficulty and pass-rate distribution as the full benchmark with fewer samples. Finally, the top-2 performing agents from this stage are evaluated on the full SWE-bench Verified test suite.




Figure 3: **Performance comparison between GEA and DGM (self-evolving baseline) on two coding benchmarks.** Under the same number of evolved agents, GEA exhibits substantially larger performance gains than DGM on both SWE-bench and Polyglot, demonstrating the improved efficiency of group-level evolution.

Polyglot. Polyglot (Gauthier, 2024a;b) assesses algorithmic code generation across diverse languages (C++, Rust, Java, etc.) and serves as an out-of-domain generalization test, since it is rarely used for training or fine-tuning coding models (Zhang et al., 2025a). We report pass@1 performance. Unlike SWE-bench Verified (OpenAI, 2024), group-level evolution is conducted exclusively on a 10-task small set. Agents achieving a success rate above 40% are subsequently evaluated on a separate, unseen 50-task medium set, ensuring that the larger set remains a strict out-of-distribution test for generalization.

4.2 Experimental Settings

For both SWE-bench Verified and Polyglot, we set the group size to $K = 2$ and use $M = 4$ nearest neighbors in the KNN-based parent group selection. At each iteration, a parent group of size 2 produces an offspring group of the same size.

For SWE-bench Verified, we run group-evolution for 30 iterations. Due to computational budget constraints, during the first 20 iterations, the evolving and acting modules are powered by Claude Haiku 4.5 (Anthropic, 2025a), while the final 10 iterations use Claude Sonnet 4.5 (Anthropic, 2025b). The reflection module is consistently powered by GPT-o1 across all iterations (Jaech et al., 2024). For Polyglot, we run group-evolution for 20 iterations. The first 10 iterations use Claude Haiku 4.5 (Anthropic, 2025a) for evolving and acting module, and the remaining 10 iterations use Claude Sonnet 4.5 (Anthropic, 2025b). The reflection module is again consistently powered by GPT-o1 (Jaech et al., 2024).

4.3 Baselines

We use state-of-the-art open-ended self-evolving agents as our primary baseline for comprehensive comparison, to systematically examine how experience sharing affects diversity utilization and performance improvement during open-ended evolution, as well as robustness to framework-level perturbations. Additionally, we compare the final performance of GEA, which uses a meta-learning approach without any human intervention, against state-of-the-art human-designed coding agents.

Open-Ended Self-Evolving Baseline. We compare against the current *state-of-the-art open-ended self-evolving framework*, implemented following DGM (Zhang et al., 2025a). Unlike GEA, this baseline enforces a strict tree-structured and individual-centric evolution, where a single agent a_i is selected as the parent and independently evolves into one child agent at each iteration. Evolutionary experience is not shared across different branches.

For SWE-bench Verified (OpenAI, 2024), we run this baseline for 60 iterations in total: the evolution and coding modules are powered by Claude Haiku 4.5 (Anthropic, 2025a) for the first 40 iterations and Claude Sonnet 4.5 (Anthropic, 2025b) for the final 20 iterations. For Polyglot, we run the baseline for 40 iterations: Claude Haiku 4.5 (Anthropic, 2025a) for the first 20 iterations and Claude Sonnet 4.5 (Anthropic, 2025b) for the remaining 20. In all baseline experiments, the reflection module is consistently powered by GPT-o1 (Jaech et al., 2024). To ensure fair comparison, we intentionally run the baseline for twice as many iterations as GEA so that the total number of evolved agents is comparable across methods, ensuring all comparisons are conducted under matched model schedules.

Human-Designed Frameworks. We additionally compare against *state-of-the-art human-designed frameworks* on both benchmarks. The top-performing, open-scaffold, checked entry on SWE-bench Verified is OpenHands + GPT-5 (Wang et al.; Singh et al., 2025), achieving 71.8%, where “checked” indicates that the SWE-bench team successfully reproduced the reported patch generations (Jimenez et al., 2024; Zhang et al., 2025a). For Polyglot, which was originally used to evaluate Aider (Gauthier, 2024b;a) by its developers, we compare against Aider, a widely adopted coding agent under continuous development and testing by human developers. The state-of-the-art performance is Aider + GPT-5 (high) (Gauthier, 2024b; Singh et al., 2025), achieving a 52.0% pass@1 success rate.

5 Results and Analysis

5.1 Main Results

GEA vs. State-of-the-Art Open-Ended Self-Evolving Systems. As shown in Figure 3, GEA demonstrates substantial performance improvements over the DGM (self-evolving baseline) on both SWE-bench Verified and Polyglot. On SWE-bench Verified, GEA improves performance from 20.0% to 71.0%, while under the same number of evolved agents, the DGM baseline achieves only 56.7%. On Polyglot, GEA boosts performance from 38.2% to 88.3%, significantly outperforming the DGM baseline (68.3%). Notably, GEA exhibits faster and more pronounced improvement in the mid-to-late stages of evolution compared to DGM (Zhang et al., 2025a), one potential reason is that the archive has accumulated sufficient diverse evolutionary directions by this point, that can be progressively consolidated and reused, leading to more rapid and pronounced performance gain, aligning with Figure 4. We discuss this phenomenon in detail in Section 5.2.

GEA vs. State-of-the-Art Human-Designed Agents. GEA achieves performance comparable to or exceeding state-of-the-art human-designed agents on both benchmarks: 71.0% vs. 71.8% on SWE-bench Verified, and 88.3% vs. 52.0% on Polyglot. Using meta-learning without any human intervention, GEA automatically evolves agent frameworks that match or surpass carefully engineered human designs (Gauthier, 2024a; Singh et al., 2025; Wang et al.), demonstrating the potential and viability of fully autonomous agent improvement.

Analysis of evolutionary patterns on two benchmarks. GEA achieves state-of-the-art open-ended self-improvement on both benchmarks, with particularly strong performance on Polyglot (Gauthier, 2024a;b). We further analyze and observe that the two benchmarks differ in task complexity: Polyglot (Gauthier, 2024a;b) primarily requires agents to modify a single file from scratch to resolve tasks, involving lower editing complexity without multi-file coordination. In contrast, SWE-bench Verified (OpenAI, 2024; Jimenez et al., 2024) typically requires coordinated modifications across multiple files, demanding that agents understand inter-file dependencies and locate relevant files for coordinated edits.

This difference in complexity leads to distinct evolutionary patterns. On Polyglot (Gauthier, 2024a;b), meta-learning produces larger, more concentrated patches: each iteration yields substantial performance gains, reaching 88.3% in just 4 iterations with a total of 8,677 lines of code added. On SWE-bench Verified (OpenAI, 2024; Jimenez et al., 2024), the evolved patches are smaller and more distributed, requiring 8 iterations to reach 71.0% with 9,663 lines of code added. These observations suggest that *GEA adapts its evolutionary behavior*




Figure 4: **Evolution analysis of tool discovery and integration over iterations.** Each row (T1–T9) corresponds to a key tool-level functionality. Blue markers indicate tools that have been discovered but not yet integrated into the current best agent, while red markers indicate tools integrated into the best-performing agent.




Figure 5: **Model transfer results on both benchmarks.** Across all coding models, the GEA best agent consistently outperforms the corresponding initial (iteration-0) agent, demonstrating that the improvements induced by group-level evolution generalize across different underlying model backbones.

to varying task complexity, demonstrating the flexibility and generality of group-level meta-learning across different problem settings.

5.2 Evolution Analysis

Overall, our analysis shows that GEA can efficiently consolidate tool-level innovations discovered across the agents, rather than letting them remain isolated in separate evolutionary branches. Figure 4 summarizes nine key tool-level modifications on agents’ framework that drove improvements. GEA integrated eight of these functionalities into its best agent, whereas the best DGM agent integrated only five. Crucially, the four tools missing from the DGM agent *were* explored in isolated branches (e.g., T4 at iteration 9) but failed to propagate due to lineage isolation. In contrast, GEA systematically consolidated these dispersed capabilities; five of its integrated tools originated from different parent agents, confirming that explicit experience sharing prevents beneficial innovations from dying out.

To quantify this consolidation, we track the number of unique *ancestor agents* contributing to the final solutions (Table 1). The best GEA agent integrates experiences from 17 unique ancestors (28.3% of the population)—nearly double that of the best DGM agent (9 ancestors).

This broader integration correlates directly with population-wide quality. As shown in Table 1, we report the *worst-case* performance among the top- k agents. Notably, the **worst** of GEA’s top-5 agents achieves 58.3%, which strictly outperforms the **single best** agent produced by DGM (56.7%). This confirms that GEA does not merely produce outliers, but systematically elevates the entire population by effectively consolidating complementary and diverse evolutionary paths. More broadly, this efficient experience consolidation suggests that GEA may exhibit stronger evolutionary capabilities in environments with greater diversity. Since open-ended evolution naturally leads to a monotonically growing diversity in the archive, GEA not only achieves stronger performance as evolution progresses but also gains enhanced self-improvement capabilities, consistent with the results in Figure 3.

5.3 Transferability

We evaluate the generalization and model transferability of the best-discovered agent produced by GEA. Specifically, we replace the coding model used in the acting module with different GPT-series and Claude-series models during benchmark evaluation. We

Rank	Worst-Case Success $\uparrow$		Ancestor Count $\uparrow$	
	DGM	GEA	DGM	GEA
Top-1	56.7%	71.0%	9 (15.0%)	17 (28.3%)
Top-3	48.3%	63.3%	13 (21.7%)	17 (28.3%)
Top-5	45.0%	58.3%	16 (26.7%)	18 (30.0%)

Table 1: **Comparison of performance (Success Rate) and ancestor integration across the Top-k agents on SWE-bench Verified.** Performance is reported as the *worst-case (minimum) success rate* among the top-k agents. *Ancestor Count* denotes the count of unique historical agents integrated into the solution. Notably, the worst-case performance of GEA’s top-5 agents (58.3%) exceeds the single best agent from DGM (56.7%).

Method	E1	E2	E3	E4	E5	Avg. $\downarrow$
DGM	5	4	5	6	5	5.0
GEA (Ours)	1	1	2	1	2	1.4

Table 2: **Robustness to framework-level bugs.** We report the number of evolution iterations required to repair injected bugs across five independent trials (E1–E5). Each entry denotes the number of iterations needed for the agent to successfully repair the injected framework-level bug in a given trial. **Lower means** that GEA repairs bugs significantly **faster** than DGM.

then compare the performance of the initial(iteration=0) agent and the GEA best-discovered agent under each coding model.

As illustrated in Figures 5, GEA’s best-discovered agent driven by GPT-series and Claude-series models consistently achieve higher performance than their corresponding initial agents on both SWE-bench and Polyglot benchmarks. This indicates that the improvements induced by group-evolving persist across different backbone models.

Further analysis reveals that all performance-improving patches discovered during GEA evolution, including those from the best agent and the top-3 performing agents, primarily target the agent’s workflow and tool usage rather than model-specific prompting, details can be found in Table 3 in Appendix A.6.

These findings together with Figures 5 demonstrate that although GEA leverages a specific backbone model to drive evolution, it discovers agent-level improvements that are largely model-agnostic and the evolved agents could generalize across different coding models.

5.4 Robustness

To evaluate the robustness of GEA, we introduce framework-level bugs by manually injecting errors into agent implementations. Specifically, we randomly select an agent from the GEA archive and manually inject framework-level bugs into its implementation. We then form a group consisting of this faulty agent and another bug-free agent from its original parent group, and perform group evolution to assess whether GEA can leverage experience from the better-performing agent (i.e., the one without framework-level bugs) to repair the faulty one. For comparison, under the self-evolution setting, the bugged agent evolves independently without access to external experience sharing. In both settings, we measure the number of iterations required to successfully repair the bug.

As shown in Table 2, across five independent trials, GEA requires only 1.4 iterations on average to repair injected bugs, whereas the self-evolving baseline (DGM) requires 5. This gap demonstrates that group-evolving agents benefit from intra-group experience sharing, enabling successful framework-level experiences from better-performing agents to guide the repair of faulty ones, confirming the robustness of the group-evolving paradigm.

6 Conclusion

We introduce **Group-Evolving Agents (GEA)**, a new paradigm for open-ended self-improvement that treats a *group of agents*, rather than an individual agent, as the fundamental unit of evolution. By enabling explicit group-level experience sharing and reuse, agents can learn from each other’s evolutionary experiences and adaptively integrate complementary improvements throughout evolution, which efficiently transforms transient diversity

into long-term useful experience. Therefore, group-level evolution achieves substantially stronger performance given the same number of evolved agents than self-evolving methods.

Further analysis shows that GEA’s improvements stem from enhancements to agent workflows and tool usage, rather than model-specific overfitting, enabling consistent transfer across different models, including both GPT-series and Claude-series. In addition, group-level experience reuse allows better-performing agents to guide the repair of faulty ones, enabling GEA to recover from framework-level bugs with fewer evolution iterations and exhibits stronger robustness than individual-centric self-evolving approaches.

Ethics statement

GEA demonstrates the potential and viability of group-evolving open-ended systems to autonomously modify their own implementation for continuous improvement. While this potential aligns with the goal of building AI that benefits humanity, open-ended exploration also carries inherent considerations worth noting. For instance, the evolutionary process may inadvertently introduce directions misaligned with human intent while consuming substantial computational resources, or produce patches that lack structural clarity, leading to increasingly complex systems that are difficult to fully understand. Therefore, it is essential to establish appropriate boundaries and guide the system to preserve exploratory diversity while ensuring alignment with human intent. Following [Zhang et al. \(2025a\)](#), all experiments in this work are conducted in isolated sandbox environments, thereby limiting potential impacts on host systems.

On the other hand, although we focus on evolving agents’ coding capabilities in this work, this paradigm has broader potential applications, for example, enabling systems to mitigate biases through self-improvement, thereby becoming more trustworthy and beneficial for social good.

References

- Anthropic. Claude haiku 4.5. <https://www.anthropic.com/news/claude-haiku-4-5>, 2025a. Accessed: 2025-11-10.
- Anthropic. Claude sonnet 4.5. <https://www.anthropic.com/news/claude-sonnet-4-5>, 2025b. Accessed: 2025-11-10.
- Antonis Antoniadis, Albert Örwall, Kexun Zhang, Yuxi Xie, Anirudh Goyal, and William Yang Wang. Swe-search: Enhancing software agents with monte carlo tree search and iterative refinement. In *The Thirteenth International Conference on Learning Representations*.
- Konstantinos Chatzilygeroudis, Antoine Cully, Vassilis Vassiliades, and Jean-Baptiste Mouret. Quality-diversity optimization: a novel branch of stochastic optimization. In *Black box optimization, machine learning, and no-free lunch theorems*, pp. 109–135. Springer, 2021.
- Jeff Clune. Ai-gas: Ai-generating algorithms, an alternate paradigm for producing general artificial intelligence. *arXiv preprint arXiv:1905.10985*, 2019.
- Maxence Faldor, Jenny Zhang, Antoine Cully, and Jeff Clune. Omni-epic: Open-endedness via models of human notions of interestingness with environments programmed in code. In *The Thirteenth International Conference on Learning Representations*.
- Jinyuan Fang, Yanwen Peng, Xi Zhang, Yingxu Wang, Xinhao Yi, Guibin Zhang, Yi Xu, Bin Wu, Siwei Liu, Zihao Li, et al. A comprehensive survey of self-evolving ai agents: A new paradigm bridging foundation models and lifelong agentic systems. *arXiv preprint arXiv:2508.07407*, 2025.
- Chrisantha Fernando, Dylan Banarse, Henryk Michalewski, Simon Osindero, and Tim Rocktäschel. Promptbreeder: self-referential self-improvement via prompt evolution. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 13481–13544, 2024.

- Huan-ang Gao, Jiayi Geng, Wenyue Hua, Mengkang Hu, Xinzhe Juan, Hongzhang Liu, Shilong Liu, Jiahao Qiu, Xuan Qi, Yiran Wu, et al. A survey of self-evolving agents: On path to artificial super intelligence. *arXiv preprint arXiv:2507.21046*, 2025.
- Paul Gauthier. Aider: Ai pair programming in your terminal. <https://github.com/Aider-AI/aider>, 2024a. Accessed: 2025-11-15.
- Paul Gauthier. o1 tops aider’s new polyglot leaderboard. <https://aider.chat/2024/12/21/polyglot.html>, December 2024b. Accessed: 2025-11-10.
- Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Minlie Huang, Nan Duan, and Weizhu Chen. Tora: A tool-integrated reasoning agent for mathematical problem solving. *arXiv preprint arXiv:2309.17452*, 2023.
- Yufei He, Juncheng Liu, Yue Liu, Yibo Li, Tri Cao, Zhiyuan Hu, Xinxing Xu, and Bryan Hooi. Evotest: Evolutionary test-time learning for self-improving agentic systems. *arXiv preprint arXiv:2510.13220*, 2025.
- Marius Hobbhahn. Swe-bench verified mini. <https://github.com/mariushobbhahn/SWEBench-verified-mini>, April 2025. Accessed: 2025-11-16.
- Wenyue Hua, Xianjun Yang, Mingyu Jin, Zelong Li, Wei Cheng, Ruixiang Tang, and Yongfeng Zhang. Trustagent: Towards safe and trustworthy llm-based agents through agent constitution. In *Trustworthy Multi-modal Foundation Models and AI Agents (TiFA)*, 2024.
- Chengsong Huang, Wenhao Yu, Xiaoyang Wang, Hongming Zhang, Zongxia Li, Ruosen Li, Jiaxin Huang, Haitao Mi, and Dong Yu. R-zero: Self-evolving reasoning llm from zero data. *arXiv preprint arXiv:2508.05004*, 2025.
- Jiaxin Huang, Shixiang Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve. In *Proceedings of the 2023 conference on empirical methods in natural language processing*, pp. 1051–1068, 2023.
- Edward Hughes, Michael Dennis, Jack Parker-Holder, Feryal Behbahani, Aditi Mavalankar, Yuge Shi, Tom Schaul, and Tim Rocktaschel. Open-endedness is essential for artificial superhuman intelligence. *arXiv preprint arXiv:2406.04268*, 2024.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.
- Minqi Jiang, Tim Rocktäschel, and Edward Grefenstette. General intelligence requires rethinking exploration. *Royal Society Open Science*, 10(6):230539, 2023.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *12th International Conference on Learning Representations, ICLR 2024*, 2024.
- Minhua Lin, Juncheng Wu, Zijun Wang, Zhan Shi, Yisi Sang, Bing He, Zewen Liu, Tianxin Wei, Zongyu Wu, Zhiwei Zhang, Dakuo Wang, Xiang Zhang, Benoit Dumoulin, Cihang Xie, Yuyin Zhou, Suhang Wang, and Hanqing Lu. Harness updating is not harness benefit: Disentangling evolution capabilities in self-evolving llm agents, 2026. URL <https://arxiv.org/abs/2605.30621>.
- Zewen Liu, Zhan Shi, Yisi Sang, Bing He, Minhua Lin, Tianxin Wei, Dakuo Wang, Benoit Dumoulin, Wei Jin, and Hanqing Lu. Adaptive auto-harness: Sustained self-improvement for agentic system deployment on open-ended task streams. *arXiv preprint arXiv:2606.01770*, 2026.
- Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.

- Ansong Ni, Miltiadis Allamanis, Arman Cohan, Yinlin Deng, Kensen Shi, Charles Sutton, and Pengcheng Yin. Next: Teaching large language models to reason about code execution. *arXiv preprint arXiv:2404.14662*, 2024.
- Alexander Novikov, Ngân Vũ, Marvin Eisenberger, Emilien Dupont, Po-Sen Huang, Adam Zsolt Wagner, Sergey Shirobokov, Borislav Kozlovskii, Francisco JR Ruiz, Abbas Mehrabian, et al. Alphaevolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131*, 2025.
- OpenAI. Introducing swe-bench verified. <https://openai.com/index/introducing-swe-bench-verified/>, August 2024. Accessed: 2025-11-10.
- Felipe Maia Polo, Lucas Weber, Leshem Choshen, Yuekai Sun, Gongjun Xu, and Mikhail Yurochkin. tinybenchmarks: evaluating llms with fewer examples. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 34303–34326, 2024.
- Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.
- Maxime Robeyns, Martin Szummer, and Laurence Aitchison. A self-improving coding agent. *arXiv preprint arXiv:2504.15228*, 2025.
- Pedro Rodriguez, Joe Barrow, Alexander Miserlis Hoyle, John P Lalor, Robin Jia, and Jordan Boyd-Graber. Evaluation examples are not equally informative: How should that change nlp leaderboards? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 4486–4503, 2021.
- J Rosser and Jakob Foerster. Agentbreeder: Mitigating the ai safety risks of multi-agent scaffolds via self-improvement. *Advances in Neural Information Processing Systems*, 38: 149564–149594, 2026.
- Jürgen Schmidhuber. Gödel machines: self-referential universal problem solvers making provably optimal self-improvements. *arXiv preprint cs/0309048*, 2003.
- Amrith Setlur, Chirag Nagpal, Adam Fisch, Xinyang Geng, Jacob Eisenstein, Rishabh Agarwal, Alekh Agarwal, Jonathan Berant, and Aviral Kumar. Rewarding progress: Scaling automated process verifiers for llm reasoning. *arXiv preprint arXiv:2410.08146*, 2024.
- Dudley Shapere. The structure of scientific revolutions. *The Philosophical Review*, 73(3): 383–394, 1964.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652, 2023.
- Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, et al. Openai gpt-5 system card. *arXiv preprint arXiv:2601.03267*, 2025.
- Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, et al. Beyond human data: Scaling self-training for problem-solving with language models. *Transactions on Machine Learning Research*.
- Yuda Song, Hanlin Zhang, Carson Eisenach, Sham M Kakade, Dean Foster, and Udaya Ghai. Mind the gap: Examining the self-improvement capabilities of large language models. In *The Thirteenth International Conference on Learning Representations*.
- Kenneth O Stanley and Joel Lehman. Why greatness cannot be planned: The myth of the objective. (*No Title*), 2015.

- Kenneth O Stanley, Joel Lehman, and Lisa Soros. Open-endedness: The last grand challenge you’ve never heard of. *While open-endedness could be a force for discovering intelligence, it could also be a component of AI itself*, 2017.
- Haotian Sun, Yuchen Zhuang, Ling kai Kong, Bo Dai, and Chao Zhang. Adaplaner: Adaptive planning from feedback with language models. *Advances in neural information processing systems*, 36:58202–58245, 2023.
- Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9426–9439, 2024.
- Shaobo Wang, Zhengbo Jiao, Zifan Zhang, Yilang Peng, Xu Ze, Boyu Yang, Wei Wang, Hu Wei, and Linfeng Zhang. Socratic-zero: Bootstrapping reasoning via data-free agent co-evolution. *arXiv preprint arXiv:2509.24726*, 2025a.
- Xingyao Wang, Boxuan Li, Yufan Song, Frank F Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, et al. Openhands: An open platform for ai software developers as generalist agents. In *The Thirteenth International Conference on Learning Representations*.
- Yingxu Wang, Siwei Liu, Jinyuan Fang, and Zaiqiao Meng. Evoagentx: An automated framework for evolving agentic workflows. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 643–655, 2025b.
- Zhenhailong Wang, Haiyang Xu, Junyang Wang, Xi Zhang, Ming Yan, Ji Zhang, Fei Huang, and Heng Ji. Mobile-agent-e: Self-evolving mobile assistant for complex tasks. *arXiv preprint arXiv:2501.11733*, 2025c.
- Yuxiang Wei, Zhiqing Sun, Emily McMilin, Jonas Gehring, David Zhang, Gabriel Synnaeve, Daniel Fried, Lingming Zhang, and Sida Wang. Toward training superintelligent software agents through self-play swe-rl. *arXiv preprint arXiv:2512.18552*, 2025.
- Huajian Xin, ZZ Ren, Junxiao Song, Zhihong Shao, Wan jia Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiushi Du, et al. Deepseek-prover-v1. 5: Harnessing proof assistant feedback for reinforcement learning and monte-carlo tree search. *arXiv preprint arXiv:2408.08152*, 2024.
- Xunjian Yin, Xinyi Wang, Liangming Pan, Li Lin, Xiaojun Wan, and William Yang Wang. Gödel agent: A self-referential agent framework for recursively self-improvement. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 27890–27913, 2025.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason E Weston. Self-rewarding language models. In *Forty-first International Conference on Machine Learning*, 2024.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488, 2022.
- Eric Zelikman, Eliana Lorch, Lester Mackey, and Adam Tauman Kalai. Self-taught optimizer (stop): Recursively self-improving code generation. In *First Conference on Language Modeling*, 2024.
- Jenny Zhang, Joel Lehman, Kenneth Stanley, and Jeff Clune. Omni: Open-endedness via models of human notions of interestingness. In *The Twelfth International Conference on Learning Representations*.
- Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. Darwin godel machine: Open-ended evolution of self-improving agents. *arXiv preprint arXiv:2505.22954*, 2025a.

Yanzhi Zhang, Yitong Duan, Zhaoxi Zhang, Jiyan He, and Shuxin Zheng. Population-evolve: a parallel sampling and evolutionary method for llm math reasoning. *arXiv preprint arXiv:2512.19081*, 2025b.

Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. Absolute zero: Reinforced self-play reasoning with zero data. *arXiv preprint arXiv:2505.03335*, 2025.

A Appendix

Appendix Contents

- [Cost Estimate](#)
- [Initial Agent](#)
- [Algorithm Details](#)
- [Details about the Agent](#)
- [Prompt Usage for Open-Ended Group-Evolution and Self-Evolution](#)
- [Case Study](#)

A.1 Cost Estimate

The primary cost for both GEA and DGM arises from benchmark evaluation. Since we generate the same number of agents for both methods, their overall costs are very similar. Following the settings described in Sections 4.2 and 4.3, the estimated cost of completing a full run is approximately USD 13,000 per method on SWE-bench and USD 1,500 on Polyglot. A more detailed estimated cost breakdown is provided below:

Coding Model	Benchmark	Number of Tasks	Cost Estimate (USD)
Claude Sonnet 4.5	SWE-bench	60	\$370
Claude Sonnet 4.5	Polyglot	60	\$60
Claude Haiku 4.5	SWE-bench	60	\$120
Claude Haiku 4.5	Polyglot	60	\$20

A.2 Initial Agent

We follow the setup of (Zhang et al., 2025a). The initial agent is equipped with two basic tools: a Bash tool that executes commands in a shell and returns the results, and an Edit tool for viewing, creating, and modifying files. We detail these tools below:

```

1 def tool_info():
2     return {
3         "name": "bash",
4         "description": """Run commands in a bash shell\n
5 * When invoking this tool, the contents of the "command" parameter does
6   NOT need to be XML-escaped.\n
7 * You don't have access to the internet via this tool.\n
8 * You have access to a mirror of common linux and python packages via
9   apt and pip.\n
10 * State is persistent across command calls and discussions with the
11   user.\n
12 * To inspect a particular line range of a file, e.g. lines 10-25, try
13   'sed -n 10,25p /path/to/the/file'.\n
14 * Please avoid commands that may produce a very large amount of output.\n
15 * Please run long lived commands in the background, e.g. 'sleep 10 &' or
16   start a server in the background."""
17     },
18     "input_schema": {

```

```

13         "type": "object",
14         "properties": {
15             "command": {
16                 "type": "string",
17                 "description": "The bash command to run."
18             }
19         },
20         "required": ["command"]
21     }
22 }

```

The information of the given Bash tool

```

1 async def tool_function_call(command):
2     """Execute a command in the bash shell."""
3     try:
4         bash_session = BashSession()
5
6         if not bash_session._started:
7             await bash_session.start()
8
9         output, error = await bash_session.run(command)
10        error = filter_error(error)
11        result = ""
12        if output:
13            result += output
14        if error:
15            result += "\nError:\n" + error
16        return result.strip()
17    except Exception as e:
18        return f"Error: {str(e)}"
19
20 def tool_function(command):
21     return asyncio.run(tool_function_call(command))

```

Details about the given Bash tool

```

1 def tool_info():
2     return {
3         "name": "editor",
4         "description": """Custom editing tool for viewing, creating, and
5 editing files\n
6 * State is persistent across command calls and discussions with the
7 user.\n
8 * If `path` is a file, `view` displays the entire file with line
9 numbers. If `path` is a directory, `view` lists non-hidden files and
10 directories up to 2 levels deep.\n
11 * The `create` command cannot be used if the specified `path` already
12 exists as a file.\n
13 * If a `command` generates a long output, it will be truncated and
14 marked with `<response clipped>`.\n
15 * The `edit` command overwrites the entire file with the provided
16 `file_text`.\n
17 * No partial/line-range edits or partial viewing are supported.""",
18         "input_schema": {
19             "type": "object",
20             "properties": {
21                 "command": {
22                     "type": "string",
23                     "enum": ["view", "create", "edit"],
24                     "description": "The command to run: `view`,
25 `create`, or `edit`."
26                 },
27                 "path": {

```

```

20         "description": "Absolute path to file or directory,
21             e.g. `/repo/file.py` or `/repo`.",
22         "type": "string"
23     },
24     "file_text": {
25         "description": "Required parameter of `create` or
26             `edit` command, containing the content for the
27             entire file.",
28         "type": "string"
29     }
30 }

```

Information about the given Edit tool.

```

1 def tool_function(command: str, path: str, file_text: str = None) -> str:
2     """
3     Main tool function that handles:
4     - 'view' : View the entire file or directory listing
5     - 'create': Create a new file with the given file_text
6     - 'edit' : Overwrite an existing file with file_text
7     """
8     try:
9         path_obj = validate_path(path, command)
10
11         if command == "view":
12             return view_path(path_obj)
13
14         elif command == "create":
15             if file_text is None:
16                 raise ValueError("Missing required `file_text` for
17                     'create' command.")
18             write_file(path_obj, file_text)
19             return f"File created successfully at: {path}"
20
21         elif command == "edit":
22             if file_text is None:
23                 raise ValueError("Missing required `file_text` for
24                     'edit' command.")
25             write_file(path_obj, file_text)
26             return f"File at {path} has been overwritten with new
27                 content."
28
29         else:
30             raise ValueError(f"Unknown command: {command}")
31
32     except Exception as e:
33         return f"Error: {str(e)}"

```

Details about the given Edit tool.

A.3 Algorithm Details

Here is the pseudocode for Open-Ended Group-Evolving(Algorithm 2).

Algorithm 2 Open-Ended Group-Evolving

```

1: Input: Parent group  $\mathcal{G} = \{a_1, \dots, a_K\}$ ; archive  $\mathcal{A}$ ; coding benchmark  $\mathcal{T}$ 
2: Output: Offspring group  $\mathcal{G}'$  with  $|\mathcal{G}'| = K$ ; updated archive  $\mathcal{A}$ 
3: Initialize offspring group  $\mathcal{G}' \leftarrow \emptyset$ 
4: for  $a_i \in \mathcal{G}$  do
5:   /* Collect evolutionary traces */
6:    $\mathcal{P}_i^{\text{Applied}} \leftarrow \text{GETAPPLIEDPATCHES}(a_i)$ 
7:    $t_i^{\text{fail}} \leftarrow \text{SAMPLEUNSOLVEDTASK}(a_i)$ 
8:    $\mathcal{P}_i^{\text{pred}} \leftarrow \text{GETPREDICTEDTASKPATCH}(a_i, t_i^{\text{fail}})$ 
9:    $\mathcal{L}_i \leftarrow \text{GETEXECUTIONLOGS}(a_i, t_i^{\text{fail}})$ 
10:   $\mathcal{O}_i \leftarrow \text{GETOUTCOMELOG}(a_i, t_i^{\text{fail}})$ 
11:   $T_i \leftarrow \{\mathcal{P}_i^{\text{Applied}}, \mathcal{P}_i^{\text{pred}}, \mathcal{L}_i, \mathcal{O}_i\}$ 
12: end for
13: /* Aggregate and share group-level experience */
14:  $S \leftarrow \bigcup_{a_j \in \mathcal{G}} T_j$ 
15: for  $a_i \in \mathcal{G}$  do
16:   /* Reflection: analyze shared experience */
17:    $\Delta_i \leftarrow \text{REFLECT}(a_i; S)$  // evolution directives
18:   /* Evolution: generate framework-level patches */
19:    $\pi'_i \leftarrow \text{EVOLVE}(a_i; \Delta_i)$ 
20:   /* Acting: evaluate updated agent */
21:    $a'_i \leftarrow \text{APPLYPATCH}(a_i, \pi'_i)$ 
22:    $\text{ACTANDEVALUATE}(a'_i; \mathcal{T})$ 
23:    $\mathcal{G}' \leftarrow \mathcal{G}' \cup \{a'_i\}$ 
24:   /* Archive update */
25:   if  $\text{COMPILES}(a'_i)$  and  $\text{BASICCODINGFUNC}(a'_i)$  then
26:      $\mathcal{A} \leftarrow \mathcal{A} \cup \{a'_i\}$ 
27:   end if
28: end for
29: return  $\mathcal{G}', \mathcal{A}$ 

```

A.4 Details about the Agent

Here we provide details of the **Reflection**, **Evolving**, and **Acting** modules in GEA. Each module is powered by a dedicated foundation model, as described in Section 4.2.

The **Reflection module** takes as input the aggregated experience from the parent agent group (Section 3.2) and outputs a framework-level evolution directives. The prompt used is provided in Appendix A.5, and an example output is shown in Example A.4.

The evolution directives is then passed to the **Evolving module**, which generates a model patch to update the agent’s codebase. The prompt used for the Evolving module is provided in Appendix A.4.

The updated agent is subsequently evaluated on downstream tasks via the **Acting module**. For each task, the Acting module follows the prompt in Appendix A.4 to generate a predicted patch, which is then submitted for evaluation.

Example: Framework-Level Evolution Directives Proposed by the Reflection Module

To Implement

Within `coding_agent.py`, introduce a new function that analyzes diffs before finalizing them. The function should verify if any relevant non-test file has been altered. If no such changes are detected, the agent should be prompted to generate a revised patch, ensuring meaningful modifications are proposed.

We need a way to detect empty or trivial patches before concluding a patch cycle. Specifically, it should reject changes that do not modify the core source files and prompt the agent to create a meaningful patch. The coding agent currently has no mechanism to detect if the proposed patch is empty, modifies only test files, or is irrelevant. We should incorporate a validation layer (in `coding_agent.py`) that compares changed files against test file patterns. If the changes do not touch core source code, the agent is instructed to attempt another patch.

Acting Module Prompt

```

1 I have uploaded a Python code repository in the directory /testbed/.
   Help solve the following problem.
2 <problem_description>
3 {self.problem_statement}
4 </problem_description>
5 <test_description>
6 {self.test_description}
7 </test_description>
8 Your task is to make changes to the files in the /testbed/ directory
   to address the <problem_description>. I have already taken care
   of the required dependencies.
```

Evolving Module Prompt

```

1 I have uploaded a Python code repository in the directory
   {self.git_tempdir}. Help solve the following problem.
2
3 <problem_description>
4 {self.problem_statement}
5 </problem_description>
6
7 <test_description>
8 {self.test_description}
9 </test_description>
10
11 Your task is to make changes to the files in the {self.git_tempdir}
   directory to address the <problem_description>. I have already
   taken care of the required dependencies.
12
13 # Coding Agent Summary
14
15 - **Main File**: `coding_agent.py`
16 - Primary Class: `AgenticSystem`
17 - The `forward()` function is the central entry point.
18 - Prompts are located either within the `forward()` function or in
   the `prompts/` directory.
19 - **Tools**: `tools/`
20 - The `tools/` directory contains various tools that LLMs can use
   to perform specific tasks.
21 - Each tool must have a `tool_info()` function that returns a JSON
   object containing 'name', 'description', and 'input_schema'.
```

```

    The 'input_schema' should be a JSON object containing 'type',
    'properties', and 'required'.
22 - Each tool must have a `tool_function()` function that takes the
    arguments defined in input_schema, performs the tool's task,
    and returns a string.
23 - See other tools for reference.
24 - **Utilities**: `utils/`
25 - The `utils/` directory contains utility functions used across the
    codebase.
26
27 - **Additional Details**:
28 - The agent is very good at automatically utilizing the right
    available tools at the right time. So do not have an agentic
    flow that explicitly forces a tool's usage.
29 - Common tools, such as file editing and bash commands, are easy
    for the agent to recognize and use appropriately. However, more
    complex and niche tools may require explicit instructions in
    the prompt.
30 - Tools should be designed to be as general as possible, ensuring
    they work across any GitHub repository. Avoid hardcoding
    repository-specific details or behaviors (e.g., paths).
31 - Do not use 'while True' loops in the agent's code. This can cause
    the agent to get stuck and not respond.
32 - Verify the implementation details of helper functions prior to
    usage to ensure proper integration and expected behavior.
33 - Do not install additional packages or dependencies directly.
    Update `requirements.txt` if new dependencies are required and
    install them using `pip install -r requirements.txt`.
34
35 # To Implement\n\n
36 {implementation_suggestion}\n\n
37 {problem_description}

```

A.5 Prompt usage for open-ended group-evolution and self-evolution

Here is an example of the prompt used to guide the Reflection Module in generating improvement requirements for the agent framework including workflow, tool usage, and prompt strategy during the open-ended group-evolving stage:

Group-Evolution Prompt

```

1 Coding Agent Summary
2
3 - **Main File**: `coding_agent.py`
4 - Primary Class: `AgenticSystem`
5 - The `forward()` function is the central entry point.
6 - Prompts are located either within the `forward()` function or in
    the `prompts/` directory.
7 - **Tools**: `tools/`
8 - The `tools/` directory contains various tools that LLMs can use
    to perform specific tasks.
9 - Each tool must have a `tool_info()` function that returns a JSON
    object containing 'name', 'description', and 'input_schema'.
    The 'input_schema' should be a JSON object containing 'type',
    'properties', and 'required'.
10 - Each tool must have a `tool_function()` function that takes the
    arguments defined in input_schema, performs the tool's task,
    and returns a string.
11 - See other tools for reference.
12 - **Utilities**: `utils/`

```

```

13 - The `utils/` directory contains utility functions used across the
    codebase.
14 - **Additional Details**:
15 - The agent is very good at automatically utilizing the right
    available tools at the right time. So do not have an agentic
    flow that explicitly forces a tool's usage.
16 - Common tools, such as file editing and bash commands, are easy
    for the agent to recognize and use appropriately. However, more
    complex and niche tools may require explicit instructions in
    the prompt.
17 - Tools should be designed to be as general as possible, ensuring
    they work across any GitHub repository. Avoid hardcoding
    repository-specific details or behaviors (e.g., paths).
18 - Do not use 'while True' loops in the agent's code. This can cause
    the agent to get stuck and not respond.
19 - Verify the implementation details of helper functions prior to
    usage to ensure proper integration and expected behavior.
20 - Do not install additional packages or dependencies directly.
    Update `requirements.txt` if new dependencies are required and
    install them using `pip install -r requirements.txt`.
21
22 Here is the implementation of the coding agent.
23
24 # Coding Agent Implementation
25 ----- Coding Agent Implementation Start -----
26 {code}
27 ----- Coding Agent Implementation End -----
28 Your task is to identify ONE detailed plan that would improve the
    agent's coding ability. The improvement should not be specific to
    any particular GitHub issue or repository.
29 Here is the log for the coding agent trying to solve the GitHub
    issues but failed.
30 # Agent Running Log
31 ----- Agent Running Log Start -----
32 {md_log}
33 ----- Agent Running Log End -----
34 # GitHub Issue
35 The GitHub issue that the agent is trying to solve.
36 ----- GitHub Issue Start -----
37 {github_issue}
38 ----- GitHub Issue End -----
39 # Predicted Patch
40 The agent's predicted patch to solve the issue.
41 ----- Predicted Patch Start -----
42 {predicted_patch}
43 ----- Predicted Patch End -----
44 # Private Test Patch
45 SWE-bench's official private tests to detect whether the issue is
    solved. This is not available to the agent during evaluation. The
    agent should try to implement its own tests.
46 ----- Private Test Patch Start -----
47 {test_patch}
48 ----- Private Test Patch End -----
49 # Issue Test Results
50 The test results from SWE-bench using the above official private
    tests.
51 ----- Issue Test Results Start -----
52 {eval_log}
53 ----- Issue Test Results End -----
54
55 and Here is the log of another coding agent --- not the current one
    --- that attempted to solve the GitHub issues but failed. You can
    use it as a reference to gain complementary insights.

```

```

56
57 # Agent Running Log
58 ----- Agent Running Log Start -----
59 {md_log1}
60 ----- Agent Running Log End -----
61
62 # GitHub Issue
63 The GitHub issue that the agent is trying to solve.
64 ----- GitHub Issue Start -----
65 {github_issue1}
66 ----- GitHub Issue End -----
67
68 # Predicted Patch
69 The agent's predicted patch to solve the issue.
70 ----- Predicted Patch Start -----
71 {predicted_patch1}
72 ----- Predicted Patch End -----
73
74 # Private Test Patch
75 SWE-bench's official private tests to detect whether the issue is
    solved. This is not available to the agent during evaluation. The
    agent should try to implement its own tests.
76 ----- Private Test Patch Start -----
77 {test_patch1}
78 ----- Private Test Patch End -----
79
80 # Issue Test Results
81 The test results from SWE-bench using the above official private
    tests.
82 ----- Issue Test Results Start -----
83 {eval_log1}
84 ----- Issue Test Results End -----
85
86 Respond precisely in the following format including the JSON start
    and end markers:
87
88 ```json
89 <JSON>
90 ```
91
92 In <JSON>, provide a JSON response with the following fields:
93 - "log_summarization": Analyze the above logs and summarize how the
    agent tried to solve the GitHub issue. Note which tools and how
    they are used, the agent's problem-solving approach, and any
    issues encountered.
94 - "potential_improvements": Identify potential improvements to the
    coding agent that could enhance its coding capabilities. Focus on
    the agent's general coding abilities (e.g., better or new tools
    usable across any repository) rather than issue-specific fixes
    (e.g., tools only usable in one framework). All necessary
    dependencies and environment setup have already been handled, so
    do not focus on these aspects.
95 - "improvement_proposal": Choose ONE high-impact improvement from the
    identified potential improvements and describe it in detail. This
    should be a focused and comprehensive plan to enhance the agent's
    overall coding ability.
96 - "implementation_suggestion": Referring to the coding agent's
    summary and implementation, think critically about what feature
    or tool could be added or improved to best implement the proposed
    improvement. If the proposed feature can be implemented by
    modifying the existing tools, describe the modifications needed,
    instead of suggesting a new tool.

```

```

97 - "problem_description": Phrase the improvement proposal and
    implementation suggestion as a GitHub issue description. It
    should clearly describe the feature so that a software engineer
    viewing the issue and the repository can implement it.
98
99 Your response will be automatically parsed, so ensure that the string
    response is precisely in the correct format. Do NOT include the
    `<JSON>` tag in your output.

```

Here is an example of the prompt used to guide the Reflection Module in generating improvement requirements for the open-ended self-evolving method:

Self-Evolution Prompt

```

1 Coding Agent Summary
2
3 - **Main File**: `coding_agent.py`
4   - Primary Class: `AgenticSystem`
5   - The `forward()` function is the central entry point.
6   - Prompts are located either within the `forward()` function or in
    the `prompts/` directory.
7 - **Tools**: `tools/`
8   - The `tools/` directory contains various tools that LLMs can use
    to perform specific tasks.
9   - Each tool must have a `tool_info()` function that returns a JSON
    object containing 'name', 'description', and 'input_schema'.
    The 'input_schema' should be a JSON object containing 'type',
    'properties', and 'required'.
10  - Each tool must have a `tool_function()` function that takes the
    arguments defined in input_schema, performs the tool's task,
    and returns a string.
11  - See other tools for reference.
12 - **Utilities**: `utils/`
13   - The `utils/` directory contains utility functions used across the
    codebase.
14 - **Additional Details**:
15   - The agent is very good at automatically utilizing the right
    available tools at the right time. So do not have an agentic
    flow that explicitly forces a tool's usage.
16   - Common tools, such as file editing and bash commands, are easy
    for the agent to recognize and use appropriately. However, more
    complex and niche tools may require explicit instructions in
    the prompt.
17   - Tools should be designed to be as general as possible, ensuring
    they work across any GitHub repository. Avoid hardcoding
    repository-specific details or behaviors (e.g., paths).
18   - Do not use 'while True' loops in the agent's code. This can cause
    the agent to get stuck and not respond.
19   - Verify the implementation details of helper functions prior to
    usage to ensure proper integration and expected behavior.
20   - Do not install additional packages or dependencies directly.
    Update `requirements.txt` if new dependencies are required and
    install them using `pip install -r requirements.txt`.
21
22 Here is the implementation of the coding agent.
23
24 # Coding Agent Implementation
25 ----- Coding Agent Implementation Start -----
26 {code}
27 ----- Coding Agent Implementation End -----

```

```

28 Your task is to identify ONE detailed plan that would improve the
    agent's coding ability. The improvement should not be specific to
    any particular GitHub issue or repository.
29 Here is the log for the coding agent trying to solve the GitHub
    issues but failed.
30 # Agent Running Log
31 ----- Agent Running Log Start -----
32 {md_log}
33 ----- Agent Running Log End -----
34 # GitHub Issue
35 The GitHub issue that the agent is trying to solve.
36 ----- GitHub Issue Start -----
37 {github_issue}
38 ----- GitHub Issue End -----
39 # Predicted Patch
40 The agent's predicted patch to solve the issue.
41 ----- Predicted Patch Start -----
42 {predicted_patch}
43 ----- Predicted Patch End -----
44 # Private Test Patch
45 SWE-bench's official private tests to detect whether the issue is
    solved. This is not available to the agent during evaluation. The
    agent should try to implement its own tests.
46 ----- Private Test Patch Start -----
47 {test_patch}
48 ----- Private Test Patch End -----
49 # Issue Test Results
50 The test results from SWE-bench using the above official private
    tests.
51 ----- Issue Test Results Start -----
52 {eval_log}
53 ----- Issue Test Results End -----
54
55 Respond precisely in the following format including the JSON start
    and end markers:
56
57 ```json
58 <JSON>
59 ```
60
61 In <JSON>, provide a JSON response with the following fields:
62 - "log_summarization": Analyze the above logs and summarize how the
    agent tried to solve the GitHub issue. Note which tools and how
    they are used, the agent's problem-solving approach, and any
    issues encountered.
63 - "potential_improvements": Identify potential improvements to the
    coding agent that could enhance its coding capabilities. Focus on
    the agent's general coding abilities (e.g., better or new tools
    usable across any repository) rather than issue-specific fixes
    (e.g., tools only usable in one framework). All necessary
    dependencies and environment setup have already been handled, so
    do not focus on these aspects.
64 - "improvement_proposal": Choose ONE high-impact improvement from the
    identified potential improvements and describe it in detail. This
    should be a focused and comprehensive plan to enhance the agent's
    overall coding ability.
65 - "implementation_suggestion": Referring to the coding agent's
    summary and implementation, think critically about what feature
    or tool could be added or improved to best implement the proposed
    improvement. If the proposed feature can be implemented by
    modifying the existing tools, describe the modifications needed,
    instead of suggesting a new tool.

```

```

66 - "problem_description": Phrase the improvement proposal and
    implementation suggestion as a GitHub issue description. It
    should clearly describe the feature so that a software engineer
    viewing the issue and the repository can implement it.
67
68 Your response will be automatically parsed, so ensure that the string
    response is precisely in the correct format. Do NOT include the
    `<JSON>` tag in your output.

```

A.6 Case Study

Agent	Patch Description	Δ Score
Top-1	docutils_debug	+0.10
	Modified core agent logic	+0.10
	Updated tests logic	+0.22
	Add snippet_extract tool	+0.10
	Updating edit tools (v1)	+0.01
	Updating edit tools (v2)	+0.05
	Updating edit tools (v3)	+0.02
Top-2	docutils_debug	+0.10
	Modified core agent logic	+0.10
	Updated tests logic	+0.22
	Add snippet_extract tool	+0.10
	Add bash tools	+0.01
	Updating edit tools	+0.07
Top-3	docutils_debug	+0.10
	Modified core agent logic	+0.10
	Updated tests logic	+0.22
	Add multi-file tools (v1)	+0.01
	Add multi-file tools (v2)	+0.07

Table 3: **Evolutionary trajectories of the top-3 performing agents discovered by GEA on SWE-bench Verified.** Across all three agents, the performance-improving patches primarily focus on enhancing the agents’ workflows and tool usage, rather than relying on model-specific prompting strategies.